

# DaLA: Danish Linguistic Acceptability Evaluation Guided by Real World Errors

Gianluca Barmina<sup>1</sup>, Nathalie Carmen Hau Norman<sup>2</sup>,  
Peter Schneider-Kamp<sup>1</sup>, Lukas Galke Poech<sup>1</sup>

<sup>1</sup>Department of Mathematics and Computer Science, University of Southern Denmark

<sup>2</sup>Department of Nordic Studies and Linguistics, University of Copenhagen

{gbarmina, petersk, galke}@imada.sdu.dk  
aha@hum.ku.dk

## Abstract

We present an enhanced benchmark for evaluating linguistic acceptability in Danish. We first analyze the most common errors found in written Danish. Based on this analysis, we introduce a set of fourteen corruption functions that generate incorrect sentences by systematically introducing errors into existing correct Danish sentences. To ensure the accuracy of these corruptions, we assess their validity using both manual and automatic methods. The results are then used as a benchmark for evaluating Large Language Models on a linguistic acceptability judgement task. Our findings demonstrate that this extension is both broader and more comprehensive than the current state of the art. By incorporating a greater variety of corruption types, our benchmark provides a more rigorous assessment of linguistic acceptability, increasing task difficulty, as evidenced by the lower performance of LLMs on our benchmark compared to existing ones. Our results also suggest that our benchmark has a higher discriminatory power which allows to better distinguish well-performing models from low-performing ones.

**Keywords:** Linguistic Acceptability, Evaluation, Low Resource Languages, LLM, LLMs Evaluation

**Datasets:** [default](#)<sup>1</sup>, [medium](#)<sup>2</sup>, [large](#)<sup>3</sup>

**Code:** [DaLA GitHub](#)<sup>4</sup>

## 1. Introduction

Judging linguistic acceptability involves evaluating the well-formedness, nativeness, or naturalness of linguistic forms, such as sentences, typically regarded as a binary classification task. Linguistic acceptability has had several applications in the past (Myers, 2017), including: assessment of the learning of second-language learners (Sorace, 1996), justification of grammatical analyses (Sprouse, 2018), study of language development in cognitive processes (McKercher and Jaswal, 2011). Moreover, in recent years, several datasets and benchmarks for linguistic acceptability have been developed to assess the natural language understanding capabilities of large language models (LLMs), some of which we briefly review later.

A model's success in this task suggests that it is able to capture the underlying structural rules of language, indicating deeper linguistic capabilities. However, existing linguistic acceptability resources have substantially been developed for a

limited number of widely spoken languages. There is less research and fewer datasets available for low-resource languages, such as Danish.

This scarcity hinders the training and development of language models, particularly multilingual ones. Furthermore, the responses of multilingual models can be influenced by multiple languages, particularly when the languages are closely related (e.g., Danish and Norwegian), which may lead to errors that, while sometimes subtle, remain salient and relevant for human evaluation.

The most relevant prior work on Danish is ScaLA (Nielsen, 2023). This dataset begins with existing Danish sentences, introduces errors, and creates pairs of sentences, one correct and one incorrect. However, this approach has some limitations: it includes only two corruptions (word removal and word swap); the introduced corruptions are not representative of errors typically made in Danish and are not based on language rules; despite strategies to avoid creating a corrupted sentence that is still grammatically correct, there is no data regarding the accuracy of the corruptions applied; and, finally, the method is restricted to the Universal Dependencies dataset (Buch-Kromann et al., 2003; Johannsen et al., 2015; Keson, 1998) (or datasets with the same format), as it relies on POS tags annotations from that dataset.

In this work, introducing a new linguistic acceptability dataset for Danish: DaLA – addressing each

<sup>1</sup>[huggingface.co/datasets/giannor/dala](https://huggingface.co/datasets/giannor/dala)

<sup>2</sup>[huggingface.co/datasets/giannor/dala\\_medium](https://huggingface.co/datasets/giannor/dala_medium)

<sup>3</sup>[huggingface.co/datasets/giannor/dala\\_large](https://huggingface.co/datasets/giannor/dala_large)

<sup>4</sup>[github.com/N-essuno/DaLA](https://github.com/N-essuno/DaLA)

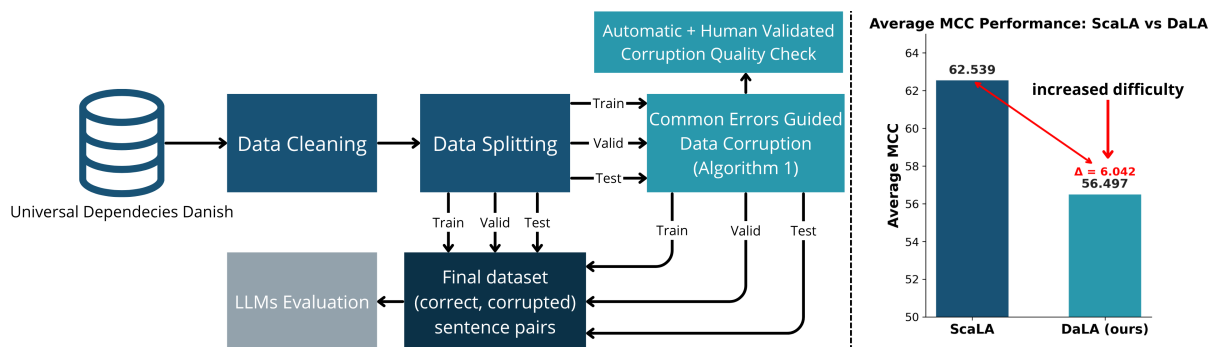


Figure 1: Left: Overview of DaLA creation method, including automatic and human corruption-quality validation, as well as LLM evaluation. Right: Comparison of LLM performance on ScaLA and DaLA, measured with the Matthews Correlation Coefficient (higher is better). LLMs perform worse on DaLA, indicating increased task difficulty.

of these limitations. Specifically, we introduce 14 new types of corruptions. These corruption types are designed to reflect common frequent grammatical errors made by Danish speakers. We focus on grammar because it has been shown to exhibit a strong relationship with linguistic acceptability judgements (Christensen and Nyvad, 2024). We evaluate the accuracy of the corruptions, i.e., whether their result is indeed incorrect, using both automatic and manual methods. Our approach does not require any annotations and does not make any assumptions about the initial data format, other than the sentence being in Danish and acceptable. An overview of this approach is shown in Figure 1.

In order to compare both methods during the model evaluation phase, we follow the same train, validation, and test proportions of previous work by Nielsen (2023). Hence, the DaLA dataset consists of 1,664 sentence pairs, with 512 for training, 1,024 for validation, and 128 for testing, for a total of 3,328 samples. Each pair includes a grammatically correct sentence and its systematically corrupted counterpart. Bigger versions of this dataset are also made available, up to a total of 7,656 samples.

To facilitate DaLA, we introduce multiple types of corruption, providing both a binary acceptability label and a classification of the specific error type. Each type of corruption is generated using predefined functions that automatically modify the original sentences. This same method can be applied to new sentences enabling a straight-forward expansion of the dataset. Moreover, it can be easily extended by adding new corruption types. To ensure the corruption actually produces a grammatically incorrect sentence, we implemented measures to minimize the cases where a sentence remains unintentionally uncorrupted. Additionally, we individually evaluate each function, that would potentially lead to unintentionally uncorrupted sentences, to

assess its effectiveness in introducing errors and maintaining data integrity. For this, as detailed in Section 4, we employ both automatic and human evaluation. Finally, we evaluate several of the best performing open-source and open-weights large language models on the DaLA corpus and compare their performance on DaLA with the currently most widely used Danish linguistic acceptability corpus, ScaLA (Nielsen, 2023).

In sum, the key contributions of this work are as follows:

- We study common errors in Danish and derive 14 corruption functions, applicable to sentences from arbitrary sources datasets.
- We create DaLA, a dataset for Danish linguistic acceptability representative of real-world common errors, consisting of 3,328 sentences (extended versions are also available), through applying these corruption functions to Universal Dependencies Data (Buch-Kromann et al., 2003; Johannsen et al., 2015; Keson, 1998).
- We evaluate the quality of the corruption functions with automatic and manual methods.
- We supply baseline results for some of the best models on previous linguistic acceptability datasets in Danish, and show that DaLA is on average 6.04% harder, going up to 14.55% for some models. The results suggest also higher discriminatory power between under-performing and high-performing models.

## 2. Related Work

### 2.1. Linguistic Acceptability

Early discussions on Linguistic Acceptability primarily focused on grammatical aspects, emphasizing the significance of such judgments from the outset (Chomsky, 1957; Chomsky, 1965). Over time, the need for more innovative methods to assess lin-

guistic acceptability became evident, expanding the range of available approaches (Bard and Robertson, 1996). The introduction of these new methods led to efforts to systematically review and evaluate their effectiveness (Sprouse et al., 2013). Subsequent research delved deeper into these judgments, examining various methodologies and reinforcing their importance despite the inherent complexity of acceptability assessments (Schütze, 2016). With the advent of big data and recent technological advancements, machine learning techniques and neural networks have increasingly been explored for predicting linguistic acceptability judgments (Wagner et al., 2013; Lawrence et al., 2000).

More recent discussions have shifted the focus from purely grammatical and syntactic factors to the role of cognitive processes, perception, and an individual's familiarity with a given sentence (Myers, 2017). Recent studies have also investigated the influence of structural complexity and frequency on linguistic acceptability (Christensen and Nyvad, 2024), particularly in Danish, which is of primary interest in this context. Findings from these studies suggest that linguistic acceptability can be shaped by the frequency and everyday use of words and phrases. However, this influence is not absolute. That is, while usage patterns may affect judgments, there is not always a direct correlation between them. Instead, a stronger relationship exists between acceptability judgments and grammatical rules, which appear to have a more consistent impact. In other words, well-established grammatical and structural rules remain among the most reliable predictors of linguistic acceptability judgments. Further evidence for this comes from research on language acquisition, which demonstrates that language learners tend to focus more on sentence structure than on less formal factors such as word frequency or usage patterns when acquiring a language (Culbertson and Adger, 2014).

## 2.2. Existing Resources

Regarding linguistic acceptability, two main types of resources can be considered: those aligned with classical Linguistic Acceptability, which consist of sentence sets annotated with binary judgments (e.g. CoLA (Warstadt et al., 2019)), and those following a syntactic minimal pair approach, where sentence pairs are provided one acceptable and the other differing by a single error that renders it unacceptable (e.g. BLiMP (Warstadt et al., 2020)).

**CoLA** (Warstadt et al., 2019; Warstadt et al. (2018)) consists of a collection of English linguistic acceptability examples, each annotated as either acceptable or unacceptable. These examples are primarily sourced from scientific linguistic resources, such as books and journal articles. The dataset is divided into two subsets: "in-domain" and

"out-of-domain". The in-domain subset, after appropriate splits, is used for both training and testing, whereas the out-of-domain subset (disjoint from the in-domain set) is used exclusively for testing.

**NoCoLA** (Jentoft and Samuel, 2023; Jentoft and Samuel (2023)) consists of two subsets: one following a binary classification approach for linguistic acceptability, and another following a pairwise approach, where each pair contains one acceptable and one unacceptable sentence. All sentences in NoCoLA are sourced from the ASK Corpus (Tenfjord et al., 2006; Tenfjord et al. (2006)), which includes both correct and incorrect sentences.

**ScaLA** (Nielsen, 2023) provides datasets for multiple languages, including Danish. The original sentences are derived from the respective Universal Dependencies Dataset and are assumed to be linguistically acceptable. Each sentence undergoes one of two types of automatic corruption to generate an unacceptable counterpart.

**JCoLA** (Someya et al., 2024; Someya et al. (2024)) and **ItaCoLA** (Trotta et al., 2021; Trotta et al. (2021)) follow a structure closely resembling CoLA, maintaining similar sources and pre-existing acceptability judgments. A notable difference in JCoLA is that more complex examples are included exclusively in the out-of-domain test set, enabling an assessment of a model's generalization ability after training on simpler or more generic cases.

**CoLAC** (Hu et al., 2023; Hu et al. (2023)) similarly follows CoLA's structure but does not differentiate between in-domain and out-of-domain subsets. Instead of relying on acceptability judgments from linguists and academic sources, CoLAC employs judgments from native speakers, arguing that this approach provides a more reliable and usage-based representation of linguistic acceptability.

**MELA** (Zhang et al., 2024; Zhang et al. (2024)) covers multiple languages. It incorporates existing linguistic acceptability resources while creating new datasets for some languages that previously lacked them, namely German, French, Spanish, Arabic, Icelandic, and Japanese. The methodology closely follows that of CoLA, sourcing sentences primarily from books in the fields of linguistics and syntax.

**BLiMP** (Warstadt et al., 2020; Warstadt et al. (2020)) adopts a different approach, generating sentences automatically using grammar templates designed by linguists to model common syntactic structures in English. Given templates for acceptable sentences, the approach also generates corresponding unacceptable sentences, enabling controlled dataset construction.

**JBLiMP** (Someya and Oseki, 2023; Someya and Oseki (2023)), for Japanese, and **BLiMP-NL** (Suijkerbuijk et al., 2025; Suijkerbuijk et al. (2025)), for Dutch, follow a structure similar to BLiMP. However, unlike BLiMP, the former sources sentences

from books and scientific articles on syntax and linguistics, incorporating pre-existing acceptability judgments, whereas the latter employs a hybrid approach, manually creating 10 minimal pairs per grammatical paradigm and automatically generating the remaining 90 pairs by prompting a large language model.

**MultiBLiMP** (Jumelet et al., 2025; Jumelet et al. (2025)), is a multilingual resource covering 101 languages, including Danish. It introduces two types of subject-verb agreement errors for creating linguistic minimal pairs through an automated pipeline.

**DaLAJ** (Volodina et al., 2021; Volodina et al., 2023; Text (2021); Volodina and Mohammed (2025)) for Swedish is based on the processing of existing datasets of sentences produced by second language learners, annotated with errors, and correct sentences extracted from Swedish language teaching books. The resulting dataset comprises an almost equal number of correct and incorrect sentences, with each incorrect sentence containing a single error.

### 2.3. Comparison with Related Resources

**DaLA** occupies an intermediate position between CoLA-like and BLiMP-like resources. Unlike CoLA and similarly to BLiMP, it introduces automatic corruption, allowing it to be applied to a broader set of sentences even in the absence of explicit acceptability judgments. This enables the creation of a balanced dataset, provided that the sentences are sourced from reliable resources and can be assumed to be acceptable. Similarly to BLiMP, DaLA allows the generation of minimal pair sets by applying a single corruption function to each sentence. However, it has been prepared to be easily extendable and permit multiple corruptions to be applied to the same sentence, further increasing the variability of the dataset. In contrast to BLiMP, and similarly to CoLA, DaLA does not generate sentences automatically from pre-defined patterns but operates on pre-existing sentences. Unlike DaLAJ (Volodina et al., 2021), DaLA includes an automatic and expandable sentence corruption process based on an analysis of the most common errors and therefore does not require a corpus of annotated sentences. Unlike MultiBLiMP (Jumelet et al., 2025), DaLA introduces 12 additional types of corruption covering a range of grammatical phenomena, compared to the single one addressed in MultiBLiMP (subject-verb agreement). Moreover, the results of LLM evaluations reported on the Danish part of MultiBLiMP show 100% accuracy for 9 out of 11 models, suggesting that the task is not particularly challenging. Finally, while MultiBLiMP relies on annotations from Universal Dependencies (Nivre et al., 2016, 2020; De Marneffe et al., 2021) and UniMorph (Batsuren et al., 2022), DaLA does not

require such annotations to be provided, as they are instead derived automatically. In contrast to all previous approaches DaLA relies on functions derived from real-world errors applicable to arbitrary Danish sentences, making it more representative of real-world scenarios than other resources relying on acceptability judgements derived from scientific sources.

## 3. DaLA: Danish Corpus of Linguistic Acceptability

The total number of corrupted sentences is 3,828. The final dataset therefore consists of up to 7,656 sentences, as each corrupted sentence is paired with its corresponding acceptable version. However, in order to compare DaLA with ScaLA during model evaluation (Section 5), we keep the same proportions used in ScaLA, reducing the number of samples for each split. We provide also the extended versions using all the 7,656 sentences.

### 3.1. Starting Dataset

The dataset we propose was constructed from sentences in the Universal Dependencies Danish (UDD) corpus (Buch-Kromann et al., 2003; Johansson et al., 2015; Keson, 1998; Johansen et al. (2015)). The sentences in this resource are extracted from original Danish texts; therefore, we assume they are correct and linguistically acceptable, consistent with the approach of Nielsen (2023). Following Nielsen (2023), the dataset was cleaned by removing duplicates, excluding sentences that are too short (<5 tokens), removing overly simple sentences (<5 POS tags), fixing lower and upper bounds on the number of characters per sentence (between 2 and 5,000) and filtering out sentences with ambiguous punctuation.

### 3.2. Corruption Type Selection

Previous studies have aimed to create linguistic acceptability datasets based on examples provided by linguistics resources, existing annotated datasets of correct and incorrect sentences or general linguistic patterns. In contrast, DaLA adopts an approach driven by automatic corruptions of existing sentences and mistakes commonly made by humans. This method aims to generate acceptability samples that better reflect real-world scenarios, thereby providing a language model evaluation that closely aligns with practical use cases. Our selection is based on very common and most noticeable errors made by native Danish speakers: the error selection was informed by linguistic expert judgements and established resources. In particular, we

used the list *Typiske Problemer*<sup>5</sup> (common problems) which has the purpose of addressing the most common needs for advice on the Danish language. The list is jointly written by experts from The Danish Language Council and the Society of Danish Language and Literature based on the official spelling rules and other works on the Danish language and grammar. We manually examined each common error and gave them a score from 0 to 5 based on implementation complexity (how difficult is it to automatically corrupt), error type (grammatical, semantic, writing mechanic), and how likely a corruption would produce an incorrect sentence (or how often a corruption is conflated with an acceptable form). After implementing the highest ranking errors, we used an additional criteria of the frequency of applying the corruption to our starting dataset. For instance, the corruption of switching the adverbs *af* (to, along) and *ad* (of, with) would only yield 18 instances. The lowest-ranking problems were almost exclusively informational guides on basic grammar terminology such as what constitutes a word class, sentence, or subordinate clause. We also saw a group of errors related to writing mechanics, e.g. when to use lower or upper case, how to split words at line breaks, and how to use punctuation. Since we focus primarily on grammatical errors, as they are among the most frequent and accurate predictors of linguistic acceptability judgments (Christensen and Nyvad, 2024), we gave this group a lower ranking. In total, 14 corruption types (and corresponding errors) were considered most relevant for implementation alongside with a basic generic corruption derived from Nielsen (2023), resulting in 15 implemented corruptions overall, 14 of which have been introduced by us. For more details on each specific corruption, we refer to Appendix A

### 3.3. Introduced Corruptions and Implementation

For each selected error, a dedicated function was developed to inject the error into an input sentence<sup>6</sup>. The main focus during implementation was to minimize the generation of corrupted sentences that remain grammatically correct, thereby ensuring the highest possible quality and precision. To achieve this, we adopted a conservative approach: for each corruption, an expert analyzed in which cases its application would most likely result in sentences that were still grammatically correct. In such cases, we preferred applying stricter filtering to the pool of corruptible sentences rather than risk introducing a large number of still-grammatical sentences. Through this conservative strategy, we achieved a

<sup>5</sup>[www.sproget.dk/typiske-problemer/](http://www.sproget.dk/typiske-problemer/)

<sup>6</sup>The main library used for this purpose was SpaCy (Honnibal et al., 2020)

balanced representation of each error, as further detailed in Section 3.6. The types of errors introduced are:

- Corruption of indefinite determiners.
- Corruption of determiners that can be appended to the end of a noun as a suffix.
- Confusion between the pronouns *nogle* and *nogen*.
- Confusion between the suffixes *-ende* and *-ene*.
- Corruption of subject and object pronouns.
- Corruption of the pronouns *som* and *der*.
- Corruption of personal and impersonal pronouns (*han/hun/det*).
- Commonly observed spelling errors.
- Confusion between the verbs *ligge* and *lægge*.
- Introduction of the “R problem” for verbs, nouns, and adjectives, related to confusion over certain types of words containing suffixes with the letter “r”.
- Corruption of different types of genitives.
- Corruption of the verb *får* and the conjunction *for*.
- The two basic corruptions implemented in Nielsen (2023).

### 3.4. Iterative Corruption Algorithm

The individual corruption functions were combined to produce an algorithm applied to existing Danish sentences for generating the final dataset. Following the method introduced by BLIMP (Warstadt et al., 2020), a single error is introduced per sentence. The errors considered may vary in representation depending on the distribution of the starting sentence dataset. For example, if the starting dataset contains few sentences with pronouns, errors applicable to pronouns may be less represented. To mitigate this issue, our proposed algorithm (Algorithm 1) applies corruptions in ascending order according to their frequency of applicability in the starting dataset. This means that the corruption applicable to the fewest sentences is applied first, corrupting all possible sentences, followed by subsequent corruptions in order.

Any change in the dataset or in the set of applied corruptions requires recalculating the applicability frequency to account for a potentially different error distribution or data imbalance. However, this process is quick and inexpensive.

### 3.5. Output Dataset

The final corpus is divided into training, validation, and test splits, and is ready to be used as a fine-tuning dataset for LLMs in subsequent evaluations. Similar to Warstadt et al. (2020), each split consists of minimal pairs of a correct sentence and a corrupted version of it with a single error. For each

### Algorithm 1 DaLA Iterative Corruption

```
1: Input: List of corruption functions cor sorted by applicability frequency on starting dataset; Starting sentence dataset sents
2: Output: List of tuples (corrupted sentence, corruption type) res
3: for s in sents do
4:   if s is not corrupted then
5:     for f in cor do
6:       candidate = f(sent)
7:       if candidate is corrupted then
8:         Mark s as corrupted
9:         Append tuple (candidate, error type from s) to res
10:      end if
11:    end for
12:  end if
13: end for
14: return res
```

corrupted sentence the label and corruption type are also indicated. The final dataset sizes are:

- Training: 512 pairs = 1024 sentences
- Validation: 128 pairs = 256 sentences
- Test: 1024 pairs = 2048 sentences

These dimensions align with those used in the previous ScaLA work (Nielsen, 2023) and have been maintained for the sake of comparability. However, these data represent only a subset of the original potential dataset. Therefore, we present two additional versions of DaLA that utilize all available examples. Below, we report the proportions and the number of examples for the training, validation and test splits, respectively:

- DaLA Medium
  - Proportions: 0.6 / 0.05 / 0.35
  - Samples: 4592 / 386 / 2678
- DaLA Large
  - Proportions: 0.8 / 0.05 / 0.15
  - Samples: 6124 / 384 / 1148

### 3.6. Corruption Diversity

From Figure 2, it can be observed that the “Basic corruptions”, derived from ScaLA (Nielsen, 2023), have the potential to cover the entire starting dataset, as they are highly generic. However, they are less representative of common real-world errors. Furthermore, we can see that Algorithm 1 has successfully limited the overrepresentation of specific corruptions, enabling a more balanced distribution of all error types. This approach favors less represented errors while constraining the proportion of “Basic corruptions”, which, although less significant, remain useful as “rollback corruptions” and for improving generalization. Although we prioritized underrepresented errors, it is not possible to avoid the distribution of the original dataset to influence the final proportions. However, the initial

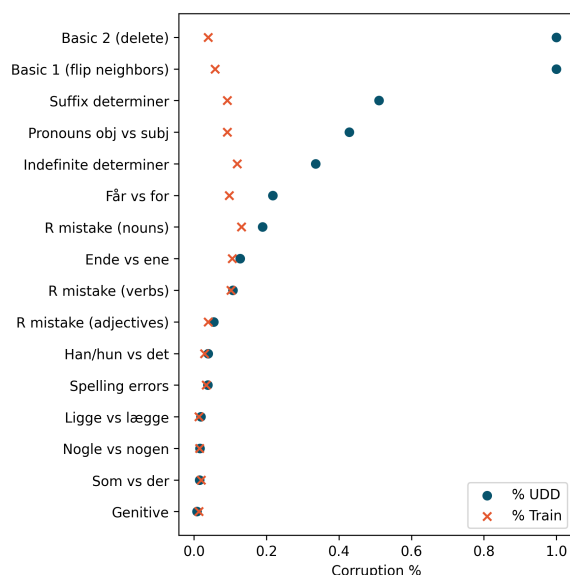


Figure 2: Proportion of corruptible examples among all Universal Dependencies samples (blue circles) vs proportion of actually corrupted examples in training set (red crosses).

dataset was collected with a focus on text diversity, and it may as well be caused by the natural distribution of the various grammatical phenomena. Nonetheless, to prevent excessive underrepresentation, we ensured that less frequent errors are included in the dataset as much as possible. Finally, we verified the distance between the training and test datasets, ensuring a similar representation of the different errors. We observed a Jensen-Shannon (JS) divergence of less than 0.01, full results can be seen in Appendix B. Therefore, the training and test distributions are sufficiently similar and appropriate for use as such.

### 4. Validating the Corruption Method

For all automatic corruptions that could potentially result in a still-acceptable sentence, we conducted a hybrid validation comprising automatic classification and manual validation to ensure high quality of the final dataset. The only corruption excluded from the validation was *spelling errors*, as these were derived from a list of spelling errors identified during a work on finding word candidates for the Danish Dictionary (Sørensen et al., 2023; Appel et al., 2024). Importantly, each spelling error occurs in corpus data and is manually categorized by lexicographers to follow common spelling error patterns for Danish and, therefore, does not represent a random typing mistake. As the primary aim of the method is to generate unacceptable sentences, we use precision as the evaluation metric, defined

as the proportion of truly unacceptable sentences (post-corruption) among all corrupted sentences. Thus, a key concern is that the corrupted sentences are in fact unacceptable rather than modified but still acceptable. The validation procedure consisted of two phases: an initial automatic validation phase (see Section 4.1), followed by a manual validation phase by an expert (see Section 4.2).

#### 4.1. Automatic Validation

The automatic system employed for validating corruptions is called *writeassistant*<sup>7</sup>. This is a free platform provided by a Danish language technology company. The system accepts Danish text as input and produces grammaticality judgments, identifying which words or parts of the sentence contain errors, the types of errors, explanations, and automatic corrections. According to internal evaluations conducted by the company and made available to us, the system is among the leading grammatical error correction tools for the Danish language. We also performed a further evaluation of the system’s precision in identifying truly corrupted sentences. We randomly extracted a sample of sentences (all, if the sentences were few, 50 otherwise) detected as incorrect for each error type and manually verified the correctness of the system’s classification. The average precision (proportion of correctly classified incorrect sentences among the total number of corrupted sentences considered across all errors) is 99.7%. For each corruption type, we analyzed which internal error categories the system detected and subsequently automated this process across all corrupted sentences. This allowed us to obtain a precision score for each corruption. Despite the generally high performance of the system, for certain error types, it showed low recall, meaning that some corrupted and actual incorrect sentences were still classified as correct or were assigned wrong error categories. Subsequent linguistic analysis revealed that many such sentences were indeed erroneous but went undetected by the system, prompting a second round of manual evaluation for these cases.

#### 4.2. Manual Validation

Following the automatic evaluation of each error type, a second manual evaluation was conducted on the corrupted sentences that were either not classified as erroneous or were assigned a different error category by the automatic system. The manual evaluation was performed by two native Danish speakers with experience in the fields of linguistics and lexicography. The agreement between the two raters was 100%. This was insofar expected as

the considered error types are highly frequent and verified as errors against the official orthographical rules in Danish (*Dansk Sprognævn*, 2024) and descriptions of Danish grammar (*Hansen and Heltoft*, 2011). Thus, the annotation can be considered an application of well-established linguistic rules for Danish, which leaves little room for annotator subjectivity. The annotators’ task was to check whether the corruption did indeed contain the intended grammatical error following the aforementioned linguistic literature. In cases of unsuccessful corruption, a comment was made on why the corruption failed. Some of the comments covered examples like error in the original sentence, the corruption conflated with another viable form, and a corruption on a named entity. For all corruptions except two, the manual evaluation covered all false positives, resulting in an accurate update of the precision. The other two corruptions, namely *pronoun corruption* and *suffix determiner corruption*, are relatively common and thus generated a higher number of examples for evaluation compared to the others. Consequently, we annotated a subset of 304 examples: 140 out of 877 examples for *pronoun corruption* and 164 out of 1199 examples for *suffix determiner corruption*. Keeping in mind that a positive is a sentence which after corruption is grammatically wrong, we then estimate the new precision based on this subset through the following equation:

$$prec_{new} = \frac{tp_{new}}{\#corruptions} \quad (1)$$

$$tp_{new} = \frac{tp_{man}}{tp_{man} + fp_{man}} fp_{auto} + tp_{auto}$$

where  $tp_{new}$  represents the newly estimated total true positives.  $fp_{auto}$  and  $tp_{auto}$  refer to the totals of false positives (classified as still correct after corruption) and true positives (classified as grammatically corrupted), respectively, detected by the automatic evaluation.  $tp_{man}$  and  $fp_{man}$  correspond to the total number of samples extracted from  $fp_{auto}$ , manually annotated as true positives or false positives, respectively. Since during the annotation we observed recurrent linguistic error patterns, we argue that our estimate is reliable.

#### 4.3. Corruption Quality Results

The final results of the corruption quality validation are reported in Table 1. For each error we report the precision (proportion of post-corruption unacceptable sentences among all corrupted sentences for that error) calculated using the automatic validation method and then adjusted with manual validation according to Equation 1. The average precision of the corruptions is approximately 0.957, which

<sup>7</sup>[www.writeassistant.com](http://www.writeassistant.com)

successfully validates our method for introducing errors into initially correct sentences.

Corruption Type	Precision
Som vs der	0.983
Ligge vs lægge	1.0
Nogle vs nogen	0.968
Genitive corruption	0.8
Ende vs ene	0.998
R problem (verbs)	0.995
Får vs for	1.0
R problem (nouns)	0.983
R problem (adjectives)	1.0
Han/Hun vs det	0.824
Indefinite determiner	0.992
Pronouns obj vs subj	0.935
Suffix determiner	0.962

Table 1: Corruption precision for each error, i.e., the fraction of sentences being corrupted in a way that the output is grammatically incorrect. Evaluation is done semi-automatically with iterative human checks on a subset and subsequent adjustments.

## 5. Large Language Models Evaluation

Using the EuroEval evaluation framework (Nielsen, 2023), we evaluated some of the best-performing large language models on the current Danish linguistic acceptability judgment task, ScaLA. Specifically, we selected two of the highest-performing models for each typology tracked by the benchmark: encoder-only, base generative, instruction-tuned, and reasoning. We focused exclusively on open-source or open-weight models to ensure reproducibility. Given that the best-performing models are generally large and relatively expensive to run, we included two smaller models to provide a broader overview of model performances on DaLA. These models still demonstrate high performance on ScaLA, and in some cases, their performance is very close to that of the bigger models.

### 5.1. Evaluation Setting

According to the EuroEval framework’s methodology (Nielsen, 2023), encoder-only models are fine-tuned on the training set and subsequently evaluated on the test set, whereas for encoder-decoder and decoder-only model types, a few-shot in-context learning approach is applied. After initial experiments showed that the default evaluation setting was not allowing the LLMs to properly adapt to new data, ending in similar performances despite the higher quality data, we adjusted some default parameters of the framework, resulting in a longer but more accurate and reliable evaluation process

for both datasets. In the fine-tuning setting, we increased the patience parameter from 2 to 20, while in the few-shot evaluation setting, we increased the number of evaluation runs from 10 to 50. The primary evaluation metric is the Matthews Correlation Coefficient (MCC)<sup>8</sup>, which is particularly well-suited for imbalanced classes (Chicco and Jurman, 2020). It ranges from -100% to +100%, with 0% indicating a random guess. Our secondary metric is the unweighted macro-F1 score, calculated as the average of the F1-scores for each class across all evaluation runs. For more details on the default settings, refer to Appendix D and Nielsen (2023).

### 5.2. Results

The results are presented in terms of MCC performance, full results are available in Table 2. Nine out of the ten models considered exhibit a performance decrease when evaluated with our DaLA dataset compared to the most widely used current benchmark (ScaLA). The average performance decrease across these models is 6.04%, ranging from 1.71% to 14.85%. These results indicate that our new dataset is generally more challenging for some of the best-performing models for the linguistic acceptability task according to the previous evaluation method (ScaLA). The only model showing a performance increase is the `DFM Sentence Encoder Large`. This may be attributed to the fact that it is the biggest Danish-only model among those considered, providing it with both enough generalization power and more specialized knowledge of the Danish language compared to models trained on multilingual data or models trained only on Danish data but too small to have enough Danish language knowledge (e.g. `Electra Small Discriminator`). This supports the conclusion that our dataset is more representative of typical Danish errors and favors models with deeper knowledge of the language. The worst performance on DaLA was observed in reasoning models, which showed an average decline of 14.32%. When evaluated with ScaLA, these models performed similarly to other models assessed using the few-shot method (i.e., instruction-finetuned and base generative models). However, under DaLA evaluation, while the other few-shot models showed a comparable decline in performance, the reasoning models experienced a substantially sharper decrease. This phenomenon suggests that our dataset, being more representative and diverse than previous benchmarks, is more discriminative regarding model performance and may better capture differences in linguistic acceptability judgments of language models.

<sup>8</sup>[wikipedia.org/wiki/Phi\\_coefficient](https://wikipedia.org/wiki/Phi_coefficient)

Model Type	Model (Hugging Face id)	Danish part of ScaLA		DaLA (ours)	
		MCC	F1	MCC	F1
Small	sarnikowski/electra-small-discriminator-da-256-cased	70.37 ± 2.73	84.82 ± 1.44	<b>65.36 ± 2.09</b>	<b>81.97 ± 1.20</b>
	jonfd/electra-small-nordic	66.25 ± 3.95	82.24 ± 2.22	<b>64.54 ± 4.99</b>	<b>80.78 ± 2.78</b>
Encoder	ai-sweden-models/roberta-large-1160k	75.04 ± 2.84	87.19 ± 1.58	<b>71.50 ± 2.09</b>	<b>85.13 ± 1.07</b>
	kennethenevoldsen/dfm-sentence-encoder-large	<b>70.29 ± 4.63</b>	<b>84.60 ± 2.52</b>	74.19 ± 7.17	86.18 ± 3.78
Instruction-Finetuned	google/gemma-3-27b-it	60.10 ± 1.54	79.67 ± 0.84	<b>56.77 ± 1.65</b>	<b>78.14 ± 0.83</b>
	google/gemma-3-12b-it	56.74 ± 1.63	77.59 ± 0.86	<b>50.66 ± 1.75</b>	<b>75.02 ± 0.89</b>
Base Generative	google/gemma-3-27b-pt	62.46 ± 1.60	80.97 ± 0.80	<b>59.78 ± 2.03</b>	<b>79.56 ± 1.04</b>
	google/gemma-3-12b-pt	57.32 ± 2.05	77.28 ± 1.38	<b>53.93 ± 1.90</b>	<b>75.80 ± 1.18</b>
Reasoning	qwen/qwen3-32b	60.07 ± 1.74	79.10 ± 0.89	<b>45.22 ± 1.72</b>	<b>72.43 ± 0.86</b>
	qwen/qwen3-14b	54.50 ± 1.77	74.71 ± 0.99	<b>40.71 ± 1.79</b>	<b>69.26 ± 0.97</b>

Table 2: ScaLA and DaLA benchmark results (Matthews Correlation Coefficient and F1 with confidence intervals). **Scores are highlighted in bold, where DaLA is below ScaLA**, indicating that DaLA is harder.

## 6. Conclusion

In this work, we introduced DaLA, a real-world-data driven method for evaluating linguistic acceptability judgments in the Danish language. Our proposed approach, based on observation of real errors and leveraging existing sentences, enables the creation of a dataset that closely reflects real-world scenarios, thereby providing a more reliable, precise and authentic evaluation of large language models. This was accomplished through a combination of linguistic analysis of the most common errors made by Danish speakers and rigorous automatic and manual evaluation of the quality of the introduced corruptions. The results demonstrate that DaLA is overall more challenging for some of the best-performing open-source and open-weight large language models according to the previous state-of-the-art evaluation for Danish. Furthermore, the findings suggest that DaLA possesses greater discriminatory power, allowing for a clearer distinction between truly high-performing and underperforming models during the evaluation process.

## 7. Limitations

Although the work presented here focuses on Danish, future work could transfer the proposed methodology to other languages by developing corruption functions based on common errors in each language. Another extension would be to broaden the types of errors, including not only grammatical errors but also semantic, pragmatic and aesthetics ones, thus providing broader coverage of linguistic acceptability. We observed that some sentences in the employed starting (UDD) dataset used may be considered outdated or sound odd, however they were a negligible amount with respect to the dataset size and were filtered out during data cleaning. We have chosen UDD to be able to compare our results with previous work. However, our methodology is applicable to sentences from arbitrary source datasets and does not rely on any annotations.

## 8. Ethical Considerations

We do not see any potential harm that could arise from our work focusing on evaluating grammatical competences of language models in Danish.

## 9. Acknowledgments

This work was partially supported by the Danish Foundation Models<sup>9</sup> project.

## 10. Bibliographical References

- Kirsten Appel, Nathalie Hau Sørensen, and Jonas Jensen. 2024. Jagten på hverdagssproget—brugen af tekster fra internetfora i arbejdet med den danske ordbog. *I: LexicoNordica*, 31.
- Ellen Bard and Dan Robertson. 1996. [Magnitude estimation of linguistic acceptability](#). *Language*, 72.
- Khuyagbaatar Batsuren, Omer Goldman, Salam Khalifa, Nizar Habash, Witold Kieraś, Gábor Bella, Brian Leonard, Garrett Nicolai, Kyle Gorman, Yustinus Ghanggo Ate, et al. 2022. Uni-morph 4.0: universal morphology. *arXiv preprint arXiv:2205.03608*.
- Matthias T. Buch-Kromann, Line Mikkelsen, and Stine Kern Lyng. 2003. [Danish dependency treebank](#). In *Proceedings of the 2nd International Workshop on Treebanks and Linguistic Theories (TLT)*, Växjö, Sweden. Växjö University Press.
- Davide Chicco and Giuseppe Jurman. 2020. [The advantages of the matthews correlation coefficient \(mcc\) over f1 score and accuracy in binary classification evaluation](#). *BMC Genomics*, 21(1):6.

<sup>9</sup>[www.foundationmodels.dk](http://www.foundationmodels.dk)

- Noam Chomsky. 1957. *Syntactic Structures*. Mouton and Co., The Hague.
- Noam Chomsky. 1965. *Aspects of the Theory of Syntax*. MIT Press, Cambridge, MA, USA.
- Ken Ramshøj Christensen and Anne Mette Nyvad. 2024. [Complexity, frequency, and acceptability](#). 9(1).
- Jennifer Culbertson and David Adger. 2014. [Language learners privilege structured meaning over surface frequency](#). *Proceedings of the National Academy of Sciences*, 111(16):5842–5847.
- Dansk Sprognævn. 2024. *Retskrivningsordbogen*, 5. udgave. edition. Alinea.
- Marie-Catherine De Marneffe, Christopher D Manning, Joakim Nivre, and Daniel Zeman. 2021. Universal dependencies. *Computational linguistics*, 47(2):255–308.
- Erik Hansen and Lars Heltoft. 2011. *Grammatik over det Danske Sprog*, 1 edition, volume 1-3. Syddansk Universitetsforlag. Bd.1: Indledning og oversigt bd.2: Syntaktiske og semantiske helheder bd.3: Sætningen og dens konstruktion.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spacy: Industrial-strength natural language processing in python](#).
- Hai Hu, Ziyin Zhang, Weifang Huang, Jackie Yan-Ki Lai, Aini Li, Yina Patterson, Jiahui Huang, Peng Zhang, Chien-Jer Charles Lin, and Rui Wang. 2023. [Revisiting acceptability judgements](#).
- Matias Jentoft and David Samuel. 2023. [NoCoLA: The norwegian corpus of linguistic acceptability](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 610–617. University of Tartu Library.
- Anders Johannsen, Héctor Martínez Alonso, and Barbara Plank. 2015. [Universal dependencies for danish](#). In *Proceedings of the 14th International Workshop on Treebanks and Linguistic Theories (TLT14)*, pages 157–167, Warsaw, Poland. Association for Computational Linguistics.
- Jaap Jumelet, Leonie Weissweiler, Joakim Nivre, and Arianna Bisazza. 2025. Multiblmp 1.0: A massively multilingual benchmark of linguistic minimal pairs. *arXiv preprint arXiv:2504.02768*.
- Britt Keson. 1998. [Documentation of the danish morpho-syntactically tagged parole corpus](#). Technical report, Society for Danish Language and Literature (DSL).
- S. Lawrence, C.L. Giles, and S. Fong. 2000. [Natural language grammatical inference with recurrent neural networks](#). *IEEE Transactions on Knowledge and Data Engineering*, 12(1):126–140.
- David A. McKercher and Vikram K. Jaswal. 2011. [Using Judgment Tasks to Study Language Knowledge](#), chapter 10. John Wiley & Sons, Ltd.
- Vladislav Mikhailov, Tatiana Shamardina, Max Ryabinin, Alena Pestova, Ivan Smurov, and Ekaterina Artemova. 2022. [RuCoLA: Russian corpus of linguistic acceptability](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5207–5227. Association for Computational Linguistics.
- James Myers. 2017. [Acceptability Judgments](#). Oxford University Press.
- Dan Nielsen. 2023. [ScandEval: A benchmark for scandinavian natural language processing](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 185–201. University of Tartu Library.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajič, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. Universal dependencies v1: A multilingual treebank collection. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 1659–1666.
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Jan Hajič, Christopher D Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. Universal dependencies v2: An evergrowing multilingual treebank collection. *arXiv preprint arXiv:2004.10643*.
- Carson T. Schütze. 2016. [The empirical base of linguistics](#). Number 2 in Classics in Linguistics. Language Science Press, Berlin.
- Taiga Someya and Yohei Oseki. 2023. [JBLiMP: Japanese benchmark of linguistic minimal pairs](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1581–1594. Association for Computational Linguistics.
- Taiga Someya, Yushi Sugimoto, and Yohei Oseki. 2024. [JCoLA: Japanese corpus of linguistic acceptability](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 9477–9488. ELRA and ICCL.

- Antonella Sorace. 1996. *The use of acceptability judgements in second language acquisition research*, 1 edition, volume 1, pages 375–409. Academic Press, United States.
- Nathalie Hau Sørensen, Nicolai Hartvig Sørensen, Kirsten Lundholm Appel, and Sanni Nimb. 2023. Trawling the corpus for the overlooked lemmas. In *Electronic lexicography in the 21st century (eLex 2023): Invisible Lexicography. Proceedings of the eLex 2023 conference. Brno, 27–29 June 2023*.
- Jon Sprouse. 2018. Acceptability judgments and grammaticality, prospects and challenges. *Syntactic structures after*, 60:195–224.
- Jon Sprouse, Carson T. Schütze, and Diogo Almeida. 2013. [A comparison of informal and formal acceptability judgments using a random sample from linguistic inquiry 2001-2010](#). *Lingua*, 134:219–248.
- Michelle Suijkerbuijk, Zoë Prins, Marianne de Heer Kloots, Willem Zuidema, and Stefan L Frank. 2025. Blimp-nl: A corpus of dutch minimal pairs and acceptability judgments for language model evaluation. *Computational Linguistics*, pages 1–35.
- Kari Tenfjord, Paul Meurer, and Knut Hofland. 2006. [The ASK corpus - a language learner corpus of Norwegian as a second language](#). In *Proceedings of the Fifth International Conference on Language Resources and Evaluation (LREC'06)*, Genoa, Italy. European Language Resources Association (ELRA).
- Daniela Trotta, Raffaele Guarasci, Elisa Leonardelli, and Sara Tonelli. 2021. [Monolingual and cross-lingual acceptability judgments with the italian CoLA corpus](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2929–2940. Association for Computational Linguistics.
- Elena Volodina, Yousuf Ali Mohammed, Aleksandrs Berdicevskis, Gerlof Bouma, and Joey Öhman. 2023. [DaLAJ-GED - a dataset for grammatical error detection tasks on Swedish](#). In *Proceedings of the 12th Workshop on NLP for Computer Assisted Language Learning*, pages 94–101, Tórshavn, Faroe Islands. LiU Electronic Press.
- Elena Volodina, Yousuf Ali Mohammed, and Julia Klezl. 2021. [Dalaj - a dataset for linguistic acceptability judgments for swedish: Format, baseline, sharing](#).
- Joachim Wagner, Jennifer Foster, and Josef Genabith. 2013. [Judging grammaticality: Experiments in sentence classification](#). *Wagner, Joachim and Foster, Jennifer and van Genabith, Josef (2009) Judging grammaticality: experiments in sentence classification. CALICO Journal, 26 (3). pp. 474-490. ISSN 0742-7778, 26*.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for english](#). 8:377–392.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Ziyin Zhang, Yikang Liu, Weifang Huang, Junyu Mao, Rui Wang, and Hai Hu. 2024. [MELA: Multilingual evaluation of linguistic acceptability](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2658–2674. Association for Computational Linguistics.

## 11. Language Resource References

- Hu, Hai and Zhang, Ziyin and Huang, Weifang and Lai, Jackie Yan-ki and Li, Aini and Patterson, Yina and Huang, Jiahui and Zhang, Peng and Lin, Chien-Jer Charles and Wang, Rui. 2023. [CoLAC: Corpus of Linguistic Acceptability in Chinese](#). Shanghai Jiao Tong University, Computational Linguistics Group. PID <https://github.com/huhailinguist/CoLAC>.
- Jentoft, Matias and Samuel, David. 2023. [NoCoLA: The Norwegian Corpus of Linguistic Acceptability](#). Language Technology Group, University of Oslo. PID <https://github.com/lgtoslo/nocola>.
- Johannsen, Anders and Martínez Alonso, Héctor and Plank, Barbara. 2015. [UD Danish DDT: Universal Dependencies for Danish](#). Universal Dependencies. PID [https://github.com/UniversalDependencies/UD\\_Danish-DDT](https://github.com/UniversalDependencies/UD_Danish-DDT).
- Jumelet, Jaap and Weissweiler, Leonie and Nivre, Joakim and Bisazza, Arianna. 2025. [MultiBLiMP 1.0: A Massively Multilingual Benchmark of Linguistic Minimal Pairs](#). University of Groningen. PID <https://github.com/jumelet/multibliimp>.
- Someya, Taiga and Oseki, Yohei. 2023. [JBLiMP: Japanese Benchmark of Linguistic Minimal Pairs](#). Oseki Laboratory, The University of Tokyo. PID <https://github.com/osekilab/JBLiMP>.

- Someya, Taiga and Sugimoto, Yushi and Oseki, Yohei. 2024. *JCoLA: Japanese Corpus of Linguistic Acceptability*. Oseki Laboratory, The University of Tokyo. PID <https://github.com/osekilab/JCoLA>.
- Suijkerbuijk, Michelle and Prins, Zoë and Kloots, Marianne de Heer and Zuidema, Willem and Frank, Stefan L. 2025. *BLiMP-NL: BLiMP-NL: The Benchmark of Linguistic Minimal Pairs for Dutch*. Radboud University. PID <http://hdl.handle.net/2066/322008>.
- Tenfjord, Kari and Meurer, Paul and Hofland, Knut. 2006. *ASK: The Norwegian Second Language Corpus*. CLARINO Bergen Centre, University of Bergen. PID <http://hdl.handle.net/11495/DA23-DEB6-9EE5-2>.
- Språkbanken Text. 2021. *DaLAJ v.1.0*. Språkbanken Text. PID <https://spraakbanken.gu.se/resurser/dalaj-1-0>.
- Trotta, Daniela and Guarasci, Raffaele and Leonardelli, Elisa and Tonelli, Sara. 2021. *ItaCoLA: The Italian Corpus of Linguistic Acceptability*. Fondazione Bruno Kessler / FBK Digital Humanities Group. PID <https://github.com/dhfbk/ItaCoLA-dataset>.
- Volodina, Elena and Mohammed, Yousuf Ali. 2025. *DaLAJ-GED-Superlim 2.0*. Språkbanken Text. PID <https://spraakbanken.gu.se/resurser/dalaj-ged-superlim>.
- Warstadt, Alex and Singh, Amanpreet and Bowman, Samuel R. 2018. *CoLA: The Corpus of Linguistic Acceptability*. New York University. PID <http://hdl.handle.net/2451/60441>.
- Warstadt, Samuel R. and Parrish, Alicia and Liu, Haokun and Mohananey, Anhad and Peng, Wei and Wang, Sheng-Fu and Bowman, Samuel R. 2020. *BLiMP: The Benchmark of Linguistic Minimal Pairs for English*. New York University. PID <http://hdl.handle.net/2451/61422>.
- Zhang, Ziyin and Liu, Yikang and Huang, Weifang and Mao, Junyu and Wang, Rui and Hu, Hai. 2024. *MELA: Multilingual Evaluation of Linguistic Acceptability*. Shanghai Jiao Tong University, Computational Linguistics Group. PID <https://github.com/sjtu-compling/MELA>.

## A. Details on Implemented Corruptions

In this section we explain what each corruption does and show one example pair for each of them. We report first the translated sentence from Danish to English, then the original Danish sentence and finally the corrupted one. The latter is marked with †. We highlight the differences between the correct and incorrect sentence by marking the affected part in blue for the original sentence and in red for the corrupted one.

**Indefinite determiner:** This error involves swapping the articles *en* and *et*. The implementation detects nouns in the sentence, identifies their associated article, if present, via dependency labels, and swaps it (from *en* to *et* or vice versa).

It is best done in a comprehensive school with differentiated teaching.

Det gøres bedst i *en* enhedsskole med differentieret undervisning.

Det gøres bedst i *et* enhedsskole med differentieret undervisning.†

**Determiner as suffix:** When a noun is singular and definite in Danish, the articles *en* and *et* may appear as suffixes attached to the noun, which can also be erroneously swapped. This case is handled by identifying specific singular definite nouns using morphological labels.

Finally, the inspection ship Fylla is also for sale.

Endelig er også inspektionsskib*et* Fylla til salg.

Endelig er også inspektionsskib*en* Fylla til salg.†

**Nogle vs Nogen:** These pronouns are pronounced similarly but are used in different contexts and frequently confused in written Danish too. *Nogle* is used with plural nouns, while *nogen* is used with singular nouns. The implementation locates the noun immediately following the pronoun and, based on pronoun-noun agreement, performs the swap. Since inverting *nogle* to *nogen* can sometimes yield grammatically acceptable sentences, particularly in interrogative or negative constructions, we implemented checks to identify and exclude such cases.

Should anyone think that it is because the Queen is overweight, I can dismiss this.

Skulle *nogen* tro, at det er fordi Dronningen er overvægtig, kan jeg afvise dette.

Skulle *nogle* tro, at det er fordi Dronningen er overvægtig, kan jeg afvise dette.†

**-ende vs -ene suffixes:** These suffixes are commonly confused due to "silent d mistakes" in Danish. The suffix *-ene* typically appears in nouns, whereas *-ende* occurs in adjectives and verbs. The implementation identifies and swaps these suffixes accordingly.

His satisfaction recently manifested itself in his offering the former welfare recipient permanent employment.

Hans tilfredshed udmøntede sig for nylig i, at han tilbød den forhenværende bistandsklient fast ansættelse.

Hans tilfredshed udmøntede sig for nylig i, at han tilbød den forhenværene bistandsklient fast ansættelse.†

**Subject vs Object pronouns:** Another common issue, especially among second language learners (L2), is confusion between subject and object pronouns. The implementation maps and swaps subject pronouns accordingly. In addition to pronoun identification, dependency relations are used to determine the syntactic position: subject pronouns in subject or conjunct positions and object pronouns in object or oblique nominal positions are swapped appropriately.

However, I am fond of this principle.

*Jeg* holder imidlertid af dette princip.

*Mig* holder imidlertid af dette princip.†

**Som vs Der:** Both are relative pronouns but used in different syntactic positions, fulfilling different roles. The most frequent error is confusing *der*, which is used in subject position, with *som*, used in object position, partly because *der* is considerably more frequent. Dependency labels are leveraged to detect relevant instances and perform the swap between *som* and *der*.

A wish that the industry apparently shares with several students.

Et ønske *som* industrien åbenbart deler med flere studerende.

Et ønske *der* industrien åbenbart deler med flere studerende.†

**Personal pronouns:** This error consists of swapping the personal pronouns *han/hun* with the impersonal pronoun *det*. The most common and clearly erroneous case is substituting *det* for *han/hun*. When a sentence contains a pronoun without an explicit antecedent, such as sentences extracted from texts, evaluating acceptability can be challenging. Thus, the implementation includes a check for the presence of a potential referent noun.

Bjørn did as she said, and then they stood facing each other.

Bjørn gjorde, som *hun* sagde, og så stod de over for hinanden.

Bjørn gjorde, som *det* sagde, og så stod de over for hinanden.<sup>†</sup>

**Spelling errors:** Frequently observed spelling mistakes in a corpus from the Danish Dictionary are addressed by replacing words with their commonly occurring incorrect variants.

He collected a year's junk mail on a spike in the basement, just as people used to keep track of receipts in the old days.

Han samlede et års junkmail på et spyd i kælderens, *ligesom* man i gamle dage holdt styr på kvitteringer.

Han samlede et års junkmail på et spyd i kælderens, *ligsom* man i gamle dage holdt styr på kvitteringer.<sup>†</sup>

**Ligge vs Lægge:** These verbs are often confused due to their similar meanings: *Ligge* is used when there is no object, whereas *Lægge* is used when an object is present. The implementation maps and swaps correct usages of these verbs accordingly through the use of dependency labels.

If you have developed sores from sitting, you must lie down continuously until it is gone.

Har du fået sår af at sidde, må du *ligge* uafbrudt, til det er væk.

Har du fået sår af at sidde, må du *lægge* uafbrudt, til det er væk.<sup>†</sup>

**R problem:** Another common error arises from confusion involving the letter *r* in certain contexts. This occurs because the pronunciation and written forms of the alternatives can be very similar.

**R problem (verbs):** For verbs, the issue involves confusing the present tense suffix *-rer* with *-re*, and conversely, confusing infinitive verbs ending in *-re* with those ending in *-rer*. These distinctions are implemented using morphological features and POS tags.

At the same time as we learn about the world, we learn our first language.

Samtidig med at vi *lærer* verden at kende, lærer vi vores første sprog.

Samtidig med at vi *lære* verden at kende, lærer vi vores første sprog.<sup>†</sup>

**R problem (nouns):** The cases considered include plural definite nouns with the suffix *-erne* confused with *-ene*; plural indefinite nouns with the

suffix *-ere* confused with *-er*; and singular definite nouns with the suffix *-eren* confused with *-erne*. These cases are handled via morphological features and POS tags.

But despite the payment, the creditors have still suffered losses of millions in the case, haven't they?

Men trods betalingen har *kreditorerne* jo stadig haft milliontab på sagen?

Men trods betalingen har *kreditorene* jo stadig haft milliontab på sagen?<sup>†</sup>

**R problem (adjectives):** This error consists of confusing adjectives ending with the suffix *-ere* with those ending in *-er*, and vice versa.

I must have hit harder than I thought and been unlucky.

Jeg må have slået *hårdere*, end jeg troede, og været uheldig.

Jeg må have slået *hårder*, end jeg troede, og været uheldig.<sup>†</sup>

**Genitive:** In Danish, several genitive forms exist depending on specific cases ('s, s, '), which can be confused or misused. The corruption implementation detects and swaps these different genitive types. Detection relies on morphological features or POS tags, supplemented by suffix checks when morphological features are unavailable.

The appeal appeared in a telegram from the council to the UN's Secretary-General, Boutros Boutros-Ghali.

Opfordringen fremgik af et telegram fra rådet til FN's generalsekretær, Boutros Boutros-Ghali.

Opfordringen fremgik af et telegram fra rådet til FN's generalsekretær, Boutros Boutros-Ghali.<sup>†</sup>

**Får vs For:** This error involves confusing the present tense of the verb *få* with the conjunction *for*. Although this is perhaps the least significant error among those considered, it remains sufficiently relevant for inclusion.

The motivation for signing was not hard to find.

Motivationen *for* at skrive under var ikke svær at finde.

Motivationen *får* at skrive under var ikke svær at finde.<sup>†</sup>

**Basic corruptions:** Following (Nielsen, 2023), two simple and broadly applicable generic corruptions have been retained as fallback mechanisms

to enable the corruption of virtually any sentence. These corruptions involve deleting a token from the sentence or swapping two tokens, with safeguards to ensure the resulting sentences are ungrammatical (e.g., corruptions are not applied to adjectives).

Two questions arise in this connection.

To spørgsmål melder *sig* i denne forbindelse.

To spørgsmål melder *sig* i denne forbindelse.<sup>†</sup>

Lotte wanted to know how much older Søren was than Peter.

Lotte *ville vide* hvor meget ældre Søren var end Peter.

Lotte *vide ville* hvor meget ældre Søren var end Peter.<sup>†</sup>

## B. Train-Test Distribution Distance

Metric	Train - Test distance
KL Divergence (Tr    Te)	0.018222
KL Divergence (Te    Tr)	0.018126
JS Divergence	0.004530
Total Variation Distance	0.082031
Wasserstein Distance	0.003418
Hellinger Distance	0.067354

Table 3: Distribution distance metrics between train and test splits

## C. Corruption Proportions

Table 4 presents the maximum proportion of corruptible sentences for each corruption type among all sentences in the starting dataset after the cleaning process (3,828). Notably, the “Basic corruptions”, derived from ScaLA (Nielsen, 2023), have the potential to cover the entire dataset, as they are highly generic, but also less representative of real-world common errors. Table 5 shows the proportions of each error in the total of corrupted sentences in the training set, which consists of 512 out of 1,024 sentences, as half of the dataset contains correct sentences and the other half contains incorrect sentences.

<sup>†</sup>Corrupted sentence

Corruption Type	% UDD
Basic corruptions	1.000000
Suffix determiner	0.510449
Pronouns obj vs subj	0.428422
Indefinite determiner	0.335946
Får vs for	0.217085
R problem (nouns)	0.189107
Ende vs ene	0.127249
R problem (verbs)	0.107612
R problem (adjectives)	0.054118
Han/hun vs det	0.038674
Spelling errors	0.037909
Ligge vs lægge	0.019073
Nogle vs nogen	0.016469
Som vs der	0.015147
Genitive	0.007839

Table 4: Frequency of each single corruption among the entire starting dataset of correct sentences

Corruption Type	Proportion
R problem (nouns)	0.130859
Indefinite determiner	0.119141
Ende vs ene	0.105469
R problem (verbs)	0.101562
Får vs for	0.097656
Pronouns obj vs subj	0.091797
Suffix determiner	0.091797
Basic 1 (flip neighbors)	0.058594
R problem (adjectives)	0.039062
Basic 2 (delete)	0.039062
Spelling errors	0.033203
Han/hun vs det	0.029297
Som vs der	0.019531
Nogle vs nogen	0.015625
Genitive	0.013672
Ligge vs lægge	0.013672

Table 5: Corruption distribution on training set

## D. EuroEval Default Evaluation Setting

In the case of few-shot evaluation, the number of few-shot examples was set to 12. The label assigned by the model was determined based on the token (containing one of the possible labels) with the highest logit value. For more details on the default settings, we refer the reader to the original evaluation framework paper Nielsen (2023).