

I Wish I Would Have Loved This One, But I Didn't – A Multilingual Dataset for Counterfactual Detection in Product Reviews

James O'Neill^{†,*}
James.O-Neill@liverpool.ac.uk

Polina Rozenshtein^{†,*}
prrozens@amazon.co.jp

Ryuichi Kiryo[†]
kiryor@amazon.co.jp

Motoko Kubota[†]
kubmotok@amazon.co.jp

Danushka Bollegala^{†,‡}
danubol@amazon.com

Amazon[†], University of Liverpool[‡]

Abstract

Counterfactual statements describe events that did not or cannot take place. We consider the problem of counterfactual detection (CFD) in product reviews. For this purpose, we annotate a multilingual CFD dataset from Amazon product reviews covering counterfactual statements written in English, German, and Japanese languages. The dataset is unique as it contains counterfactuals in multiple languages, covers a new application area of e-commerce reviews, and provides high quality professional annotations. We train CFD models using different text representation methods and classifiers. We find that these models are robust against the selectional biases introduced due to cue phrase-based sentence selection. Moreover, our CFD dataset is compatible with prior datasets and can be merged to learn accurate CFD models. Applying machine translation on English counterfactual examples to create multilingual data performs poorly, demonstrating the language-specificity of this problem, which has been ignored so far.

1 Introduction

Counterfactual statements are an essential tool of human thinking and are often found in natural languages. Counterfactual statements may be identified as statements of the form – *If p was true, then q would be true* (i.e. assertions whose antecedent (p) and consequent (q) are known or assumed to be false) (Milmed, 1957). In other words, a counterfactual statement describes an event that may not, did not, or cannot take place, and the subsequent consequence(s) or alternative(s) did not take place. For example, consider the counterfactual statement – *I would have been content with purchasing this iPhone, if it came with a warranty!*. Counterfactual statements can be broken into two parts: a statement about the event (*if it came with a warranty*), also referred to as the **antecedent**, and

the consequence of the event (*I would have been content with purchasing this iPhone*), referred to as the **consequent**. Counterfactual statements are ubiquitous in natural language and have been well-studied in fields such as philosophy (Lewis, 2013), psychology (Markman et al., 2007; Roese, 1997), linguistics (Ippolito, 2013), logic (Milmed, 1957; Quine, 1982), and causal inference (Höfler, 2005).

Accurate detection of counterfactual statements is beneficial to numerous applications in natural language processing (NLP) such as in medicine (e.g., clinical letters), law (e.g., court proceedings), sentiment analysis, and information retrieval. For example, in information retrieval, counterfactual detection (CFD) can potentially help to remove irrelevant results to a given query. Revisiting our previous example, we should not return the iPhone in question for a user who is searching for *iPhone with warranty* because that iPhone does not come with a warranty. A simple bag-of-words retrieval model that does not detect counterfactuals would return the iPhone in question because all the tokens in the query (i.e. *iPhone, with, warranty*) occur in the review sentence. Detecting counterfactuals can also be a precursor to capturing causal inferences (Wood-Doughty et al., 2018) and interactions, which have shown to be effective in fields such as health sciences (Höfler, 2005). Janocko et al. (2016) and Son et al. (2017) studied CFD in social media for automatic psychological assessment of large populations.

CFD is often modelled as a binary classification task (Son et al., 2017; Yang et al., 2020a). A manually annotated sentence-level counterfactual dataset was introduced in SemEval-2020 (Yang et al., 2020a) to facilitate further research into this important problem. However, successful developments of classification methods require extensive high quality labelled datasets. To the best of our knowledge, currently there are only two labelled datasets for counterfactuals: (a) the pio-

[†]The two first authors contributed equally

neering small dataset of tweets (Son et al., 2017) and (b) a recent larger corpus covering the area of the finance, politics, and healthcare domains (Yang et al., 2020a). However, these datasets are limited to the English language.

In this paper, we contribute to this emerging line of work by annotating a novel CFD dataset for a new domain (i.e. product reviews), covering languages in addition to English, such as Japanese and German, ensuring a balanced representation of counterfactuals and the high quality of the labelling. Following prior work, we model counterfactual statement detection as a binary classification problem, where given a sentence extracted from a product review, we predict whether it expresses a counterfactual or a non-counterfactual statement. Specifically, we annotate sentences selected from Amazon product reviews, where the annotators provided sentence-level annotations as to whether a sentence is counterfactual with respect to the product being discussed. We then represent sentences using different encoders and train CFD models using different classification algorithms.

The percentage of sentences that contain a counterfactual statement in a random sample of sentences has been reported to be low as 1-2% (Son et al., 2017). Therefore, all prior works annotating CFD datasets have used clue phrases such as *I wished* to select candidate sentences that are likely to be true counterfactuals, which are then subsequently annotated by human annotators (Yang et al., 2020a). However, this selection process can potentially introduce a selection bias towards the clue phrases used.

To the best of our knowledge, while the data selection bias is a recognised problem in other NLP tasks (e.g., Larson et al. (2020)), this selection bias on CFD classifiers has not been studied previously. Therefore, we train counterfactual classifiers with and without masking the clue phrases used for candidate sentence selection. Furthermore, we experiment with enriching the dataset with sentences that do not contain clue phrases but are semantically similar to the ones that contain clue phrases. Interestingly, our experimental results reveal that compared to the lexicalised CFD such as bag-of-words representations, CFD models trained using contextualised masked language models such as BERT are robust against the selection bias (Devlin et al., 2019). Our contributions in this paper are as follows:

First-ever Multilingual Counterfactual Dataset:

We introduce the first-ever multilingual CFD dataset containing manually labelled product review sentences covering English, German, and Japanese languages.¹ As already mentioned above, counterfactual statements are naturally infrequent. We ensure that the positive (i.e. counterfactual) class is represented by at least 10% of samples for each language. Distinguishing between a counterfactual and non-counterfactual statements is a fairly complex task even for humans. Unlike previous works, which relied on crowdsourcing, we employ professional linguists to produce a high quality annotation. We follow the definition of counterfactuals used by Yang et al. (2020a) to ensure that our dataset is compatible with the SemEval-2020 CFD dataset (SemEval). We experimentally verify that by merging our dataset with the SemEval CFD dataset, we can further improve the accuracies of counterfactual classifiers. Moreover, applying machine translation on the English CFD dataset to produce multilingual CFD datasets results in poor CFD models, indicating the language-specificity of the problem that require careful manual annotations.

Accurate CFD Models: Using the annotated dataset we train multiple classifiers using (a) lexicalised word-order insensitive bag-of-words representations as well as (b) contextualised sentence embeddings. We find that there is a clear advantage to using contextualised embeddings over non-contextualized embeddings, indicating that counterfactuals are indeed context-sensitive.

2 Related Work

Counterfactuals have been studied in various contexts such as for problem solving (Markman et al., 2007), explainable machine learning (Byrne, 2019), advertisement placement (Joachims and Swaminathan, 2016) and algorithmic fairness (Kusner et al., 2017). Kaushik et al. (2020) proposed an annotation scheme whereby the original data is augmented in a counterfactual manner to overcome spurious associations that a classifier heavily relies upon, thus failing to perform well on test data distributions that are not identical. Unlike Kaushik et al. (2020) and closely related work by Gardner et al. (2020), we are interested in identifying exist-

¹<https://github.com/amazon-research/amazon-multilingual-counterfactual-dataset>

ing counterfactuals and filtering these statements to improve search performance.

A CFD task was presented in SemEval-2020 Challenge (Yang et al., 2020b). The provided dataset contains counterfactual statements from news articles. However, the dataset does not cover counterfactuals in e-commerce product reviews, which is our focus in this paper. One of the earliest CFD datasets was annotated by Son et al. (2017) and covers counterfactual statements extracted from social media. Both datasets are labelled for binary classification by crowdsourcing and contain only sentences in English. We will compare our dataset to these previous works in § 3.4. To summarise, our dataset is unique as it contains counterfactuals in multiple languages, covers a new application area of e-commerce reviews, and provides high quality annotations.

A range of CFD methods was recently proposed in response to the SemEval-2020 challenge (Yang et al., 2020b). Most of the high performing methods (Ding et al., 2020; Fajcik et al., 2020; Lu et al., 2020; Ojha et al., 2020; Yabloko, 2020) use state-of-the-art pretrained language models (Devlin et al., 2019; Liu et al., 2019; Lan et al., 2020; Radford et al., 2019; Yang et al., 2019). Traditional ML methods, such as SVM and random forests were also used but with less success (Ojha et al., 2020).

To achieve the best prediction quality, ensemble strategies are employed. The top performing systems use an ensemble of transformers (Ding et al., 2020; Fajcik et al., 2020; Lu et al., 2020), while others include Convolutional Neural Networks (CNNs) with Global Vectors (GloVe; Pennington et al., 2014) embeddings (Ojha et al., 2020). Various structures are used on top of transformers. For example, Lu et al. (2020); Ojha et al. (2020) use a CNN as the top layer, while Bai and Zhou (2020) use a Bi-GRUs and Bi-LSTMs. Some other proposed methods use additional modules, such as constituency and dependency parsers, in the lower layers of the architecture (Yabloko, 2020).

CFD datasets tend to be highly imbalanced because counterfactual statements are less frequent in natural language texts. Prior work has used techniques such as pseudo-labelling (Ding et al., 2020) and multi sample dropout (Chen et al., 2020) to address the data imbalance and overfitting problems.

3 Dataset Curation

We adopt the definition of a counterfactual statement proposed by Janocko et al. (2016) where they define it as *a statement which looks at how a hypothetical change in past experience could have affected the outcome of that experience*. Their definition is based on linguistic structures of 6 types of counterfactuals as following.

Conjunctive Normal: The antecedent is followed by the consequent. The antecedent consists of a conditional conjunction followed by a past tense subjunctive verb or past modal verb. The consequent contains a past or present tense modal verb. (Example: *If everyone **got** along, it would be more enjoyable.*)

Conjunctive Converse: The consequent is followed by the antecedent. The consequent consists of a modal verb and past or present tense verb. The antecedent consists of a conditional conjunction followed by a past tense subjunctive verb or past tense modal. (Example: *I would be stronger, if I **had lifted** weights.*)

Modal Normal: The antecedent is followed by the consequent. The antecedent consists of a modal verb and past participle verb. The consequent consists of a past/present tense modal verb. (Example: *We **should have gone** bowling, that would have been better.*)

Wish/Should Implied: The antecedent is present, the consequent is implied. The antecedent is the independent clause following ‘wish’ or ‘should’. The consequent is implied and can be paraphrased as “would be better off”. (Examples: *I **wish I had been** richer. I **should have** revised my rehearsal lines.*)

Verb Inversion: No specific order of the antecedent and consequent. The antecedent uses the subjunctive mood by inverting the verbs ‘had’ and ‘were’ to create a hypothetical conditional statement along with a past tense verb. The consequent consists of a modal verb and past or present tense verb. (Example: ***Had I listened** to your advice, I **may have got** the job.*)

Modal Propositional, Would/Could Have: The consequent is followed by the antecedent. The antecedent consists of a past/present modal verb. The consequent consists of a prepositional phrase (only certain types). (Examples: *I **would have been better off** not reading this. I **would have been happier** without John.*)

Note that, while Yang et al. (2020a) explicitly

mention only 5 types of counterfactual and [Son et al. \(2017\)](#) work with 7 types, their definitions and clue words used for data collection effectively cover the same 6 types defined by [Janocko et al. \(2016\)](#). We worked with professional linguists to extend these counterfactual definitions for the German and Japanese languages. While the extension of the definition from English to German is relatively straightforward, the extension to syntactically and orthographically different structure of Japanese sentences was challenging ([Jacobsen, 2011](#)) and required re-writing the annotation guidelines including additional examples. The annotation guidelines are included in the dataset release.

3.1 Data Collection

The main step of data collection in the previous works ([Son et al., 2017](#); [Yang et al., 2020a](#)) is filtering of the data using a pre-compiled list of clue words/phrases. Because the exact list of clue phrases used by [Janocko et al. \(2016\)](#) was not publicly available, we created a new list of clue phrases following the definitions of counterfactual types. In addition, we compiled similar clue phrase lists for German and Japanese languages. [Yang et al. \(2020a\)](#) applied a more complex procedure, where they match Part of Speech (PoS)-tagged sentences against lexico-syntactic patterns. In our work, we do not consider PoS-based patterns, which are difficult to generalise across languages.

We use the Amazon Customer Reviews Dataset,² which contains over 130 million customer reviews collected and released by Amazon to the research community. To create an annotated dataset, we select reviews in different categories as detailed in the Supplementary. Next, we sample candidate sentences for annotation in two iterations.

In the first iteration, we consider reviews written by customers with a verified purchase (i.e., the customer has bought the product about which he or she is writing the review). Given that counterfactual statements are infrequent, all prior works ([Son et al., 2017](#); [Yang et al., 2020a](#)) have used clue phrase lists for selecting data for human annotation. Following this practice, we select sentences that contain exactly one clue phrase from our pre-compiled clue phrase lists for each language. We remove sentences that are exceedingly long (more than 512 tokens) or short (less than 10 tokens).

Shorter sentences might not contain sufficient information for a human annotator to decide whether it is a counterfactual statement, whereas longer sentences are likely to contain various other information besides counterfactuals.

The above-mentioned first iteration might produce a biased dataset in the sense that all sentences contain counterfactual clues from the predefined lists. There are two possible drawbacks in this selection method. First, the manually compiled clue phrase lists might not cover all the different ways in which we can express a counterfactual in a particular language. Therefore, the sentences selected using the clue phrase lists might have coverage issues. Second, a counterfactual classification model might assign high confidence scores for some high precision clue phrases (e.g., “wish” for English). Such a classifier is likely to perform poorly on test data that do not use clue phrases for expressing counterfactuality. On the contrary, adding sentences with no clue words to the dataset might result in a greater bias: those additional sentences are likely to be negative examples, and thus discriminatory power of the clue phrases can get amplified. Later in our experiments, we empirically evaluate the effect of selection bias due to the reliance on clue phrases.

To address the selection bias, in addition to the sentences selected in the first iteration, we conduct a second iteration where we select sentences that *do not* contain counterfactual clues from our lists. For this purpose, we create sentence embeddings for each sentence selected in the first iteration. We use a pretrained multilingual BERT model³. We then use *k*-means clustering to cluster these sentences into $k = 100$ clusters. We assume each cluster represents some aspect of a product, and represented by its centroid. Next, we pick sentences that do not contain the clue phrases, compute their sentence embeddings, and measure the similarity to each of the centroids. For each centroid we select the top n most similar sentences for manual annotation. We set n such that we obtain an approximately equal number of sentences to the number of sentences that contain clue phrases selected in the first iteration. All selected sentences are manually annotated for counterfactuality as described in § 3.2.

²<https://s3.amazonaws.com/amazon-reviews-pds/readme.html>

³<https://huggingface.co/bert-base-multilingual-uncased>

3.2 Annotation

The annotators were provided guidelines with definitions, extensive examples and counterexamples. Briefly, counterfactual statements were identified if they belong to any of the counterfactual types described in § 3. If any part of a sentence contains a counterfactual, then we consider the entire sentence to be a counterfactual. This annotation process increases the number of counterfactual examples and the coverage across the counterfactual types in the dataset, thereby improving the class imbalance. We require that at least 90% of the sentences have agreement of 2 professional linguists (2 out of 2 agreement), the rest at most 10% cases had a third linguist to resolve the disagreement (2 out of 3 agreement).

3.3 Dataset Statistics

The basic dataset statistics can be found in Table 1. We present two versions of the English dataset: *EN* contains only sentences filtered by the clue words, *EN-ext* is a superset of *EN* enriched by sentences with no clue words as described above. The clue-based dataset *EN* contains about 1/5-th of positive examples, while its extended version contains 1/10-th of counterfactuals. Only 76 out of 4977 added sentences were labelled positively. *DE* dataset contains 69.1% and *JP* contains 9.5% of counterfactuals.

The summary of clue phrase distributions in positive and negative classes is shown in Table 2. Interestingly, English and German lists have approximately the same number of clues, but the precision for German clues is much higher, resulting in more counterfactual statements being extracted using those clue phrases. On the contrary, the Japanese list has the largest number of clues, yet results in the lowest precision. The specification of counterfactual clue phrases for Japanese is a linguistically hard problem because the meaning of the clues is highly context dependent. The large number of Japanese clue phrases is due to the orthographic variations present in Japanese where the same phrase can be written using kanji, hiragana, katakana characters or a mixture of them. Because we were able to select a sufficiently large datasets for German and Japanese using the clue phrases, we did not consider the second iteration step described in § 3.1 for those languages.

Dataset	Positive	Negative	Total	CF %
<i>EN</i>	954	4069	5023	18.9
<i>EN-ext</i>	1030	8970	10000	10.0
<i>DE</i>	4840	2160	7000	69.1
<i>JP</i>	667	6333	7000	9.5

Table 1: Dataset statistics: the number of positive (counterfactual) and negative (non-counterfactual) examples, total sizes of the datasets, percentage of counterfactual (CF) examples.

Dataset	N	f_P	f_N	f_{data}
<i>EN</i>	29	100.	100.	100.
<i>EN-ext</i>	29	92.6	45.3	50.2
<i>DE</i>	27	100.	100.	100.
<i>JP</i>	70	100.	100.	100.

Table 2: Clue phrases summary for the datasets: N is the total number of clue phrases in each clue phrase list. f_P and f_N are the percentages of examples containing clue phrases respectively in counterfactual and non-counterfactual classes. f_{data} is the percentage of sentences containing a clue phrase in a dataset.

3.4 Comparison with Existing Datasets

We compare the multilingual counterfactual dataset we create against existing datasets in Table 3. Our dataset is well-aligned with the two other existing datasets in the sense that we use the same definition of a counterfactual, keep a similar percentage of positive examples, and use similar keywords for dataset construction. These properties ensure that our dataset of product reviews can be used on its own, as well as organically combined with the existing datasets from other domains. A distinctive feature of our dataset is its coverage of a novel domain, e-commerce reviews, which is not covered by any of the existing counterfactual datasets. Furthermore, our dataset is available for three languages: English, German, and Japanese. This is the first counterfactual dataset not limited to English language. Unlike previous works, which relied on crowdsourcing, we employ professional linguists to produce the lists of clue words and supervise the annotation. This ensures the high quality of the labelling.

4 Evaluations

We conduct a series of experiments to systematically evaluate several important factors related to counterfactuality such as (a) selection bias due to clue phrases (§ 4.1), (b) effect of merging multiple counterfactual datasets (§ 4.2), (c) use of machine

Dataset	Language	Size	CF %
Son et al. (2017)	English	1637 (2137)	10.1 (31.2)
Yang et al. (2020a)	English	20000	11.0
This work	English / German / Japanese	10000 (5023) / 7000 / 7000	10.0 (18.9) / 69.1 / 9.5

Dataset	CF definition	Domain	Construction	Annotation
Son et al. (2017)	Janocko et al. (2016)	Twitter	keywords filtering	mixed: manual (unknown), automatic pattern matching
Yang et al. (2020a)	Janocko et al. (2016)	News: finance, politics, healthcare	keywords filtering, pattern matching	manual (crowdsourcing, strong agreement)
This work	Janocko et al. (2016)	Amazon Reviews	keywords filtering	manual (curated by linguists)

Table 3: Dataset comparisons. The numbers in parenthesis for Son et al. (2017) correspond to the union of manually and automatically labelled datasets. The numbers in parenthesis for this work correspond to clue-based English dataset *EN*.

translation (MT) to translate counterfactual statements (§ 4.3), and (d) effect of different sentence encoders and classifiers for training CFD models (§ 4.4).

For evaluations in (a), (b), and (c), we fine-tune a widely used multilingual transformer model BERT (mBERT) (Devlin et al., 2019) to train a CFD model. The model is pretrained for the tasks of masked language modelling and next sentence prediction for 104 languages⁴ and is used with the default parameter settings. The model is implemented using the Transformer⁵ library. We fine-tune a linear layer on top of these pretrained language models for the CFD task using the training process as described next.⁶

We use an 80%-20% train-test data split and tune hyperparameters via 5-fold cross-validation. Hyperparameters in the already pretrained transformer models are kept fixed. F1, Matthew’s Correlation Coefficient (MCC; Boughorbel et al., 2017), and accuracy are used as evaluation metrics. MCC ($\in [-1, 1]$) accounts for class imbalance and incorporates all correlations within the confusion matrix (Chicco and Jurman, 2020). Accuracy may be misleading in highly imbalanced datasets because a simple classification of all instances to the majority class has a high accuracy. However, for consistency with prior work, we report all three evaluation metrics in this paper. All the reported results are averaged over at least 3 independently trained models initialised with the same hyperparameter values. For tokenisation, unless the tokeniser is pre-specified for the model, we use `word.tokenize`

from `nltk.tokenize.punkt`⁷ for English and German languages; and MeCab⁸ as the morphological analyser for Japanese.

4.1 Selection Bias due to Clue Phrases

To evaluate the effectiveness of clue phrases for selecting sentences for human annotation and any selection bias due to this process, we fine-tune mBERT with and without masking the clue phrases. Classification performance values are shown in Table 4. Overall, we see that **no mask** (training without masking) returns slightly better performance than **mask** (training with masking), however the differences are not statistically significant. This is reassuring because it shows that the sentence embeddings produced by mBERT generalise well beyond the clue phrases used to select sentences for manual annotation. On the other hand, if a CFD model had simply *memorised* the clue phrases and was classifying based on the occurrences of the clue phrases in a sentence, we would expect a drop in classification performance in **no mask** setting due to overfitting to the clue phrases that are not observed in the test data. Indeed for *EN* where all sentences contain clue phrases, we see a slight drop in all evaluation measure for **no mask** relative to **mask**, which we believe is due to this overfitting effect. The performance on *JP* is the lowest among all languages compared. This could be attributed to the tokenisation issues and lack of Japanese coverage in mBERT. Many counterfactual clues in Japanese are parts of verb/adjective inflections, which can get split/removed during the tokenisation.

Table 5 shows recall (R) and precision (P) on

⁴<https://huggingface.co/bert-base-multilingual-uncased>

⁵<https://github.com/huggingface/transformers>

⁶See Supplementary for the details on fine-tuning.

⁷<https://www.nltk.org/api/nltk.tokenize.html>

⁸<https://pypi.org/project/mecab-python3/>

Dataset	Mask	mBERT		
		F1	MCC	Acc
<i>EN</i>	mask	0.92	0.76	0.92
	no mask	0.89	0.73	0.89
<i>EN-ext</i>	mask	0.93	0.69	0.93
	no mask	0.94	0.74	0.94
<i>DE</i>	mask	0.86	0.68	0.86
	no mask	0.90	0.79	0.90
<i>JP</i>	mask	0.86	0.48	0.84
	no mask	0.85	0.49	0.82

Table 4: F1, MCC and Accuracy (Acc) for CFD models trained with and without masking the clue phrases.

Metric	<i>EN</i>	<i>EN-ext</i>	<i>DE</i>	<i>JP</i>
R_{nm}	0.93	0.94	0.92	0.85
P_{nm}	0.71	0.59	0.94	0.30
R_m	0.87	0.79	0.86	0.88
P_m	0.68	0.66	0.93	0.37

Table 5: Precision and Recall for mBERT trained with (m) and without (nm) masking the clue phrases.

masked (subscript m) and non-masked (subscript nm) settings. In all datasets the recall is higher than precision for both masked and non-masked versions due to dataset imbalance with an underrepresented positive class. The number of positive examples misclassified under masked and non-masked settings are typically very small. We see that the CFD model trained on *EN-ext* has a higher recall, but lower precision than the one on *EN*. Most of the added examples in *EN-ext* are negatives, which makes it hard to maintain a high precision.

4.2 Cross-Dataset Adaptation

To study the compatibility of our dataset with existing datasets, we train a CFD model on one dataset and test the trained model on a different dataset. Prior work on domain adaptation (Ben-David et al., 2009) has shown that the classification accuracy of such a cross-domain classifier is upper-bounded by the similarity between the train and test datasets. Further, we merge our *EN-ext* dataset with the *SemEval* dataset (Yang et al., 2020a) to create a dataset denoted by *Comb*. Specifically, we separately pool the the counterfactual and noncounterfactual instances in each dataset to create *Comb*.

As can be seen from Table 6, the models trained on *EN* and *EN-ext* perform poorly on *SemEval*, while the model trained on *SemEval* has relatively high values of F1, MCC, and Accuracy on *EN* and *EN-ext*. This implies that the product reviews we

Train	Test	mBERT		
		F1	MCC	Acc
<i>EN</i>	<i>EN</i>	0.89	0.73	0.89
	<i>EN-ext</i>	0.96	0.85	0.96
	<i>SemEval</i>	0.65	0.28	0.59
	<i>Comb</i>	0.68	0.31	0.62
<i>EN-ext</i>	<i>EN</i>	0.92	0.80	0.92
	<i>EN-ext</i>	0.94	0.74	0.94
	<i>SemEval</i>	0.50	0.19	0.42
	<i>Comb</i>	0.49	0.19	0.42
<i>SemEval</i>	<i>EN</i>	0.82	0.56	0.80
	<i>EN-ext</i>	0.86	0.48	0.83
	<i>SemEval</i>	0.93	0.71	0.92
	<i>Comb</i>	0.96	0.84	0.96
<i>Comb</i>	<i>EN</i>	0.95	0.86	0.95
	<i>EN-ext</i>	0.94	0.72	0.94
	<i>SemEval</i>	0.93	0.70	0.92
	<i>Comb</i>	0.96	0.84	0.96

Table 6: Classification quality, combining datasets for training and evaluation.

use cover a narrow subdomain compared to the domains in *SemEval*. Interestingly, the CFD model trained on *Comb* reports the best performance across all measures, indicating that our dataset is compatible with *SemEval* and can be used in conjunction with existing datasets to train better CFD models.

4.3 Cross-Lingual Transfer via MT

Considering the costs involved in manually annotating counterfactual statements for each language, a frugal alternative would be to train a model for English and then apply it on test sentences in a target language of interest, which are translated into English using a machine translation (MT) system. To evaluate this possibility, we first translate the German and Japanese CFD datasets into English (denoted respectively by *DE-EN* and *JP-EN*) using Amazon MT.⁹ Next, we train separate English CFD models using *EN*, *EN-ext* and *SemEval* datasets, and apply those models on *DE-EN* and *JP-EN*.

As shown in Table 7, the MCC values for the MT-based CFD model are significantly lower than that for the corresponding in-language baseline, which is trained using the target language data. Therefore, simply applying MT on test data is *not* an alternative to annotating counterfactual datasets from scratch for a novel target language. This result shows the importance of developing counterfactual datasets for languages other than English, which has not been done prior to this work. Moreover,

⁹<https://aws.amazon.com/translate/>

Train	Test	mBERT		
		F1	MCC	Acc
<i>EN</i>	<i>DE-EN</i>	0.65	0.41	0.64
<i>EN-ext</i>	<i>DE-EN</i>	0.73	0.49	0.72
<i>SemEval</i>	<i>DE-EN</i>	0.58	0.35	0.58
<i>DE</i>	<i>DE</i>	0.90	0.79	0.90
<i>EN</i>	<i>JP-EN</i>	0.80	0.26	0.78
<i>EN-ext</i>	<i>JP-EN</i>	0.80	0.28	0.76
<i>SemEval</i>	<i>JP-EN</i>	0.86	0.22	0.86
<i>JP</i>	<i>JP</i>	0.85	0.49	0.82

Table 7: Classification quality of English translations.

the performance for German, which belongs to the same Germanic language group as English, is better than for Japanese. The model trained on *SemEval* performs the worst on *DE-EN* dataset, and has the lowest MCC on *JP-EN*. This experimental result indicates the importance of introducing new languages to the counterfactual dataset family.

4.4 Sentence Encoders and Classifiers

We evaluate the effect of the sentence encoding and binary classification methods on the performance of CFD using multiple settings.

Bag-of-N-grams (BoN): We represent a sentence using tf-idf weighted unigrams and bi-grams and ignore n -grams with a frequency less than 2 or more than 95% of the frequency distribution. Next, Principal Component Analysis (PCA; [Wold et al., 1987](#)) is used to create 600-dimensional sentence embeddings.

Word Embeddings (WE): We average the 300-dimensional fastText embeddings trained on Common Crawl and Wikipedia¹⁰ for the words in a sentence to create its sentence embedding. We note that there have been meta-embedding methods ([Bollegala and Bao, 2018](#); [Bollegala et al., 2018](#)) proposed to combine multiple word embeddings to further improve their accuracy. However, their consideration for CFD is beyond the scope of current work.

BoN and WE representations are used to train binary CFD models using different classification methods such as a Support Vector Machine (SVM; [Cortes and Vapnik, 1995](#)) with a Radial Basis function, an ID3 Decision Tree (DT; [Breiman et al., 1984](#)), a Random Forest (RF; [Breiman, 2001](#)) with 20 trees.

¹⁰<https://fasttext.cc/docs/en/crawl-vectors.html>

Method	Mask	Dataset			
		<i>EN</i>	<i>EN-ext</i>	<i>DE</i>	<i>JP</i>
mBERT	mask	0.76	0.69	0.68	0.48
	no mask	0.73	0.74	0.79	0.49
XLM-RoBERTa	mask	0.75	0.68	0.59	0.42
	no mask	0.79	0.76	0.80	0.38
XLM-w/o-Emb	mask	0.71	0.64	0.67	0.47
	no mask	0.76	0.70	0.79	0.47
SVM (BoN)	mask	0.50	0.44	0.47	0.58
	no mask	0.74	0.70	0.76	0.58
DT (BoN)	mask	0.36	0.28	0.37	0.43
	no mask	0.64	0.58	0.70	0.48
RF (BoN)	mask	0.16	0.11	0.20	0.14
	no mask	0.40	0.34	0.60	0.11
SVM (WE)	mask	0.42	0.32	0.40	0.49
	no mask	0.56	0.49	0.67	0.49
DT (WE)	mask	0.23	0.25	0.28	0.42
	no mask	0.37	0.37	0.56	0.40
RF (WE)	mask	0.20	0.08	0.17	0.16
	no mask	0.26	0.14	0.39	0.14

Table 8: MCC for the different CFD Models.

Pretrained Language Models Along with mBERT, we fine-tune a linear layer for CFD task on top of two following pretrained transformer models: XLM model ([Conneau and Lample, 2019](#))¹¹ and base XLM-RoBERTa model ([Conneau et al., 2020](#)).¹² Both models were trained for the task of masked language modelling for 100 languages.

Results Here we extend our experiment with clue word masking. For the transformer-based models we mask the clue words similar to mBERT. For the traditional ML methods we remove the clue words from the sentences before tokenization.

The results with and without masking are reported in [Table 8](#) (F1 and Accuracy are reported in the Supplementary). First, we note that masking decreases the performance of all classifiers on all datasets. Transformer-based classifiers are the least affected by masking: they are able to learn semantic dependencies from the remaining text. We could also say that transformers are the least affected by the data-selection bias as they do not rely on the clue words. Traditional ML methods with BoN features are affected by masking the most: they seem to use clue words for discrimination. Interestingly, for these methods the performance drops equally for clue-based *EN* and enriched *EN-ext* datasets. This could indicate that in both cases the classifier relies on the clue words.

Overall transformer-based models (especially

¹¹<https://huggingface.co/xlm-mlm-100-1280>

¹²<https://huggingface.co/xlm-roberta-base>

XLM-RoBERTa) perform the best across all datasets except for *JP*. For *JP* the best performance is obtained by an SVM model with BoN features. This could indicate that for Japanese, a language-specific tokenisation works for the lexicalised (BoN) models better than the language-independent subtokenisation methods such as Byte Pair Encoding (BPE; Sennrich et al., 2016) that are used when training contextualised transformer-based sentence encoders. The former preserves more information than the latter at the expense of a sparser and larger feature space (Bollegala et al., 2020). Transformer-based masked language models on the other hand require subtokenisation as they must use a smaller vocabulary to make the token prediction task efficient (Yang et al., 2018; Li et al., 2019).

In general, unlike the simpler word embedding and bag of words approaches, large pretrained contextualized embeddings maintain high test performance according to the reported evaluation metrics. We note that these also converged after a few epochs using a relatively small number of labelled instances, based on the model with the best 5-fold validation accuracy. Hence, contextualized embeddings can identify various context-dependent counterfactuals from a diverse range of reviews using a small number of mini-batch gradient updates of a single linear layer. Among the different sentence embedding methods compared, the best performance is reported by XLM-RoBERTa.

Between the two baselines, we see that using word embeddings to represent the sentences does not offer clear benefits for traditional ML methods and BoN features are sufficient. However, embedding based methods suffer generally a smaller performance drop when clues are masked. This suggests that embeddings provide a more general and robust representation of counterfactuals in the semantic space than BoN features.

5 Conclusion

We annotated a multilingual counterfactual dataset using Amazon product reviews for English, German and Japanese languages. Experimental results show that our English dataset is compatible with the previously proposed SemEval-2020 Task 5 dataset. Moreover, the CFD models trained using our dataset are relatively robust against selection bias due to clue phrases. Simply applying MT on test data results in poor cross-lingual classifica-

tion performance, indicating the need for language-specific CFD datasets.

6 Ethical Considerations

In this work, we annotated a multilingual dataset covering counterfactual statements. Moreover, we train CFD models using different sentence representation methods and binary classification algorithms. In this section, we discuss the ethical considerations related to these contributions.

With regard to the dataset being released, all sentences that are included in the dataset were selected from a publicly available Amazon product review dataset. In particular, we do *not* collect or release any additional product reviews as part of this paper. Moreover, we have manually verified that the sentences in our dataset do not contain any customer sensitive information. However, product reviews do often contain subjective opinions, which can sometimes be socially biased. We do not filter out any such biases.

We use two pretrained sentence encoders, mBERT and XLM-RoBERTa, when training the CFD models. It has been reported that pretrained masked language model encode unfair social biases such as gender, racial and religious biases (Bommasani et al., 2020). Although we have evaluated ourselves the mBERT and XLM-RoBERTa based CFD models that we use in our experiments, we suspect any social biases encoded in these pretrained masked language models could propagate into the CFD models that we train. In particular, these social biases could be further amplified during the CFD model training process, if the counterfactual statements in the training data also contain such biases. Debiasing masked language models is an active research field (Kaneko and Bollegala, 2021) and we plan to evaluate the social biases in CFD models in our future work.

References

- Yang Bai and Xiaobing Zhou. 2020. Byteam at semeval-2020 task 5: Detecting counterfactual statements with bert and ensembles. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 640–644.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. 2009. A theory of learning from different domains. *Machine Learning*, 79:151–175.

- Danushka Bollegala and Cong Bao. 2018. Learning word meta-embeddings by autoencoding. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1650–1661, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Danushka Bollegala, Kohei Hayashi, and Ken-ichi Kawarabayashi. 2018. Think globally, embed locally — locally linear meta-embedding of words. In *Proc. of IJCAI-EACL*, pages 3970–3976.
- Danushka Bollegala, Ryuichi Kiryo, Kosuke Tsujino, and Haruki Yukawa. 2020. Language-independent tokenisation rivals language-specific tokenisation for word similarity prediction. In *Proc. of LREC*.
- Rishi Bommasani, Kelly Davis, and Claire Cardie. 2020. Interpreting Pretrained Contextualized Representations via Reductions to Static Embeddings. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4758–4781, Online. Association for Computational Linguistics.
- Sabri Boughorbel, Fethi Jarray, and Mohammed El-Anbari. 2017. Optimal classifier for imbalanced data using matthews correlation coefficient metric. *PLoS one*, 12(6):e0177678.
- Leo Breiman. 2001. Random forests. *Machine learning*, 45(1):5–32.
- Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. 1984. *Classification and regression trees*. CRC press.
- Ruth MJ Byrne. 2019. Counterfactuals in explainable artificial intelligence (xai): evidence from human reasoning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 6276–6282.
- Weilong Chen, Yan Zhuang, Peng Wang, Feng Hong, Yan Wang, and Yanru Zhang. 2020. Ferryman at semeval-2020 task 5: Optimized bert for detecting counterfactuals. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 653–657.
- 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.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. Unsupervised cross-lingual representation learning at scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*.
- Alexis Conneau and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*.
- Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning*, 20(3):273–297.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- Xiao Ding, Dinghui Hao, Yuewei Zhang, Kuo Liao, Zhongyang Li, Bing Qin, and Ting Liu. 2020. Hitscrl at semeval-2020 task 5: Training pre-trained language model with pseudo-labeling data for counterfactuals detection. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 354–360.
- Martin Fajcik, Josef Jon, Martin Docekal, and Pavel Smrz. 2020. BUT-FIT at SemEval-2020 task 5: Automatic detection of counterfactual statements with deep pre-trained language representation models. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 437–444, Barcelona (online). International Committee for Computational Linguistics.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khachabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. Evaluating models’ local decision boundaries via contrast sets. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Corrado Gini. 1912. Variabilità e mutabilità (variability and mutability). *Tipografia di Paolo Cuppini, Bologna, Italy*, page 156.
- M Höfler. 2005. Causal inference based on counterfactuals. *BMC medical research methodology*, 5(1):28.
- Michela Ippolito. 2013. Counterfactuals and conditional questions under discussion. In *Semantics and Linguistic Theory*, volume 23, pages 194–211.
- Wesley M. Jacobsen. 2011. The interrelationship of time and realis in japanese – in search of the semantic roots of hypothetical meaning. *NINJAL project review*, 1(5).
- Anthony Janocko, Allegra Larche, Joseph Raso, and Kevin Zembroski. 2016. Counterfactuals in the language of social media: A natural language processing project in conjunction with the world well being

- project. Technical report, University of Pennsylvania.
- Thorsten Joachims and Adith Swaminathan. 2016. Counterfactual evaluation and learning for search, recommendation and ad placement. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 1199–1201.
- Masahiro Kaneko and Danushka Bollegala. 2021. Debiasing pre-trained contextualised embeddings. In *Proc. of the 16th European Chapter of the Association for Computational Linguistics (EACL)*.
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2020. Learning the difference that makes a difference with counterfactually-augmented data. In *International Conference on Learning Representations*.
- Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. In *Advances in Neural Information Processing Systems*, pages 4066–4076.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. ALBERT: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- Stefan Larson, Anthony Zheng, Anish Mahendran, Rishi Tekriwal, Adrian Cheung, Eric Guldan, Kevin Leach, and Jonathan K Kummerfeld. 2020. Iterative feature mining for constraint-based data collection to increase data diversity and model robustness. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8097–8106.
- David Lewis. 2013. *Counterfactuals*. John Wiley & Sons.
- Liunian Harold Li, Patrick H. Chen, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Efficient contextual representation learning with continuous outputs. *Transactions of the Association for Computational Linguistics*, 7:611–624.
- 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. *arXiv preprint arXiv:1907.11692*.
- Yaojie Lu, Annan Li, Hongyu Lin, Xianpei Han, and Le Sun. 2020. IsCAS at semeval-2020 task 5: Pre-trained transformers for counterfactual statement modeling. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 658–663.
- Keith D Markman, Matthew J Lindberg, Laura J Kray, and Adam D Galinsky. 2007. Implications of counterfactual structure for creative generation and analytical problem solving. *Personality and Social Psychology Bulletin*, 33(3):312–324.
- Bella K Milmed. 1957. Counterfactual statements and logical modality. *Mind*, 66(264):453–470.
- Anirudh Anil Ojha, Rohin Garg, Shashank Gupta, and Ashutosh Modi. 2020. Iitk-rsa at semeval-2020 task 5: Detecting counterfactuals. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 458–467.
- Jeffery Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: global vectors for word representation. In *Proc. of EMNLP*, pages 1532–1543.
- John Platt et al. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74.
- Willard Van Orman Quine. 1982. *Methods of logic*. Harvard University Press.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Neal J Roese. 1997. Counterfactual thinking. *Psychological bulletin*, 121(1):133.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Youngseo Son, Anneke Buffone, Joe Raso, Allegra Larche, Anthony Janocko, Kevin Zembroski, H Andrew Schwartz, and Lyle Ungar. 2017. Recognizing counterfactual thinking in social media texts. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 654–658.
- Svante Wold, Kim Esbensen, and Paul Geladi. 1987. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52.
- Zach Wood-Doughty, Ilya Shpitser, and Mark Dredze. 2018. Challenges of using text classifiers for causal inference. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing*, volume 2018, page 4586. NIH Public Access.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.

Len Yabloko. 2020. Ethan at semeval-2020 task 5: Modelling causal reasoning in language using neuro-symbolic cloud computing. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 645–652.

Xiaoyu Yang, Stephen Obadinma, Huasha Zhao, Qiong Zhang, Stan Matwin, and Xiaodan Zhu. 2020a. SemEval-2020 task 5: Counterfactual recognition. In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 322–335, Barcelona (online). International Committee for Computational Linguistics.

Xiaoyu Yang, Stephen Obadinma, Huasha Zhao, Qiong Zhang, Stan Matwin, and Xiaodan Zhu. 2020b. [SemEval-2020 Task 5: Counterfactual Recognition](#).

Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W. Cohen. 2018. [Breaking the softmax bottleneck: A high-rank RNN language model](#). In *International Conference on Learning Representations*.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *Advances in neural information processing systems*, pages 5754–5764.

Supplementary Materials

A Fine-tuned multilingual BERT for counterfactual classification

Given that we select mBERT (Devlin et al., 2019) as the main classification method in the paper, we describe how the original BERT architecture is adapted for fine-tuned for CF classification.

Consider a dataset $D = \{(X_i, y_i)\}_{i=1}^m$ for $D \in \mathcal{D}$ and a sample $s := (X, y)$ where the sentence $X := (x_1, \dots, x_n)$ with n being the number of words $x \in X$. We can represent a word as an input embedding $\mathbf{x}_w \in \mathbb{R}^d$, which has a corresponding target vector \mathbf{y} . In the pre-trained transformer models we use, X_i is represented by 3 types of embeddings; word embeddings ($\mathbf{X}_w \in \mathbb{R}^{n \times d}$), segment embeddings ($\mathbf{X}_s \in \mathbb{R}^{n \times d}$) and position embeddings ($\mathbf{X}_p \in \mathbb{R}^{n \times d}$), where d is the dimensionality of each embedding matrix. The self-attention block in a transformer mainly consists of three sets of parameters: the query parameters $\mathbf{Q} \in \mathbb{R}^{d \times l}$, the key parameters $\mathbf{K} \in \mathbb{R}^{d \times l}$ and the value parameters $\mathbf{V} \in \mathbb{R}^{d \times o}$. For 12 attention heads (as in BERT-base), we express the forward pass as follows:

$$\vec{\mathbf{X}} = \mathbf{X}_w + \mathbf{X}_s + \mathbf{X}_p \quad (1)$$

$$\vec{\mathbf{Z}} := \bigoplus_{i=1}^{12} \text{softmax}(\vec{\mathbf{X}} \mathbf{Q}_{(i)} \mathbf{K}_{(i)}^T \vec{\mathbf{X}}^T) \vec{\mathbf{X}} \mathbf{V}_{(i)} \quad (2)$$

$$\vec{\mathbf{Z}} = \text{Feedforward}(\text{LayerNorm}(\vec{\mathbf{Z}} + \vec{\mathbf{X}})) \quad (3)$$

$$\overleftarrow{\mathbf{Z}} = \text{Feedforward}(\text{LayerNorm}(\overleftarrow{\mathbf{Z}} + \overleftarrow{\mathbf{X}})) \quad (4)$$

The last hidden representations of both directions are then concatenated $\mathbf{Z}' := \overleftarrow{\mathbf{Z}} \oplus \vec{\mathbf{Z}}$ and projected using a final linear layer $\mathbf{W} \in \mathbb{R}^d$ followed by a sigmoid function $\sigma(\cdot)$ to produce a probability estimate \hat{y} , as shown in (5). As in the original BERT paper, WordPiece embeddings (Wu et al., 2016) are used with a vocabulary size of 30,000. Words from (step-3) that are used for filtering the sentences are masked using a [PAD] token to ensure the model does not simply learn to correctly classify some samples based on the association of these tokens with counterfactuals. A linear layer is then fine-tuned on top of the hidden state, $\mathbf{h}_{X, [\text{CLS}]}$ emitted corresponding to the [CLS] token. This fine-tunable linear layer is then used to predict whether the sentence is counterfactual or not, as shown in Equation 5, where $\mathcal{B} \subset \mathcal{D}$ is a mini-batch and \mathcal{L}_{ce} is the cross-entropy loss.

$$\mathcal{L}_{ce} := \frac{1}{|\mathcal{B}|} \sum_{(X,y) \in \mathcal{B}} \mathbf{y} \log(\sigma(\mathbf{h}_{X, [\text{CLS}]} \cdot \mathbf{W})) \quad (5)$$

Configurations For the mBERT counterfactual model we use BERT-base, which uses 12 Transformer blocks, 12 self-attention heads with a hidden size of 768. The default size of 512 is used for the sentence length and the sentence representation is taken as the final hidden state of the first [CLS] token. This model is already pre-trained and we fine-tune a linear layer \mathbf{W} on top of BERT, which is fed to through a sigmoid function σ as $p(c|h) = \sigma(\mathbf{W}h)$ where c is the binary class label and we maximize the log-probability of correctly predicting the ground truth label.

B Matthews Correlation Coefficient

Unlike metrics such as F1, MCC accounts for class imbalance and incorporates all correlations within the confusion matrix (Chicco and Jurman, 2020). For MCC, the range is [-1, 1] where 1 represents a

perfect prediction, 0 an average random prediction and -1 an inverse prediction.

$$\text{MCC} = \frac{\text{tp} \times \text{tn} - \text{fp} \times \text{fn}}{\sqrt{(\text{tp} + \text{fp})(\text{tp} + \text{fn})(\text{tn} + \text{fp})(\text{tn} + \text{fn})}} \quad (6)$$

C Extended version of Table 8

We report F1, MCC, and accuracy in Table 9.

D Examples of Incorrect Predictions

Table 10 shows examples of misclassifications given by transformer models. The second column indicates which of the remaining transformer models misclassified each review where B=mBERT, XR=XLM-RoBERTa, X=XLM without embedding.

E Hardware Used

All transformer, RNN and CNN models were trained using a GeForce NVIDIA GTX 1070 GPU which has 8GB GDDR5 Memory.

F Model Configuration and Hyperparameter Settings

BERT-base uses 12 Transformer blocks, 12 self-attention heads with a hidden size of 768. The default size of 512 is used for the sentence length and the sentence representation is taken as the final hidden state of the first [CLS] token. A fine-tuned linear layer \mathbf{W} is used on top of BERT-base, which is fed to through a sigmoid function σ as $p(c|h) = \sigma(\mathbf{W}h)$ where c is used to calibrate the class probability estimate and we maximize the log-probability of correctly predicting the ground truth label.

Table 11 shows the pretrained model configurations that were already predefined before our experiments. The number of (Num.) hidden groups here are the number of groups for the hidden layers where parameters in the same group are shared. The intermediate size is the dimensionality of the feed-forward layers of the the Transformer encoder. The ‘Max Position Embeddings’ is the maximum sequence length that the model can deal with.

We now detail the hyperparameter settings for transformer models and the baselines. We note that all hyperparameter settings were performed using a manual search over development data.

F.1 Transformer Model Hyperparameters

We did not change the original hyperparameter settings that were used for the original pre-training of each transformer model. The hyperparameter settings for these pretrained models can be found in the class arguments python documentation in each configuration python file in the https://github.com/huggingface/transformers/blob/master/src/transformers/e.g_configuration_py and are also summarized in Table 11.

For fine-tuning transformer models, we manually tested different combinations of a subset of hyperparameters including the learning rates $\{50^{-4}, 10^{-5}, 50^{-5}\}$, batch sizes $\{16, 32, 128\}$, warmup proportion $\{0, 0.1\}$ and ϵ which is a hyperparameter in the adaptive momentum (adam) optimizer. Please refer to the huggingface documentation at <https://github.com/huggingface/transformers> for further details on each specific model e.g at https://github.com/huggingface/transformers/blob/master/src/transformers/modeling_bert.py, and also for the details of the architecture for BertForSequenceClassification pytorch class that is used for our sentence classification and likewise for the remaining models.

Fine-tuning all language models with a sentence classifier took less than two and half hours for all models. For example, for the largest transformer model we used, BERT, the estimated average runtime for a full epoch with batch size 16 (of 2, 682 training samples) is 184.13 seconds. In the worst case, if the model does not already converge early and all 50 training epochs are carried out, training lasts for 2 hour and 30 minutes.

F.2 Baseline Hyperparameters

SVM Classifier: A radial basis function was used as the nonlinear kernel, tested with an ℓ_2 regularization term settings of $C = \{0.01, 0.1, 1\}$, while the kernel coefficient γ is autotuned by the scikit-learn python package and class weights are used inversely proportional to the number of samples in each class. To calibrate probability estimates for AUC scores, we use Platt’s scaling (Platt et al., 1999).

Decision Tree and Random Forest Classifiers: We use 20 decision tree classifiers with no restriction on tree depth and the minimum number of samples required to split an internal node is set to 2. The criterion for splitting nodes is the Gini importance (Gini, 1912).

Dataset	Mask	F1	MCC	Acc	F1	MCC	Acc	F1	MCC	Acc
		mBERT			XLM-RoBERTa			XLM-w/o-Emb		
EN	mask	0.92	0.76	0.92	0.91	0.75	0.90	0.89	0.71	0.89
	no mask	0.89	0.73	0.89	0.92	0.79	0.92	0.91	0.76	0.90
EN-ext	mask	0.93	0.69	0.93	0.93	0.68	0.92	0.91	0.64	0.91
	no mask	0.94	0.74	0.94	0.95	0.76	0.95	0.93	0.70	0.92
DE	mask	0.86	0.68	0.86	0.82	0.59	0.82	0.85	0.67	0.85
	no mask	0.90	0.79	0.90	0.91	0.80	0.91	0.91	0.79	0.90
JP	mask	0.86	0.48	0.84	0.81	0.42	0.78	0.84	0.47	0.81
	no mask	0.85	0.49	0.82	0.86	0.38	0.85	0.83	0.47	0.79
EN	mask	SVM (BoN)			DT (BoN)			RF (BoN)		
	no mask	0.83	0.50	0.83	0.80	0.36	0.82	0.73	0.16	0.81
EN-ext	mask	0.89	0.44	0.88	0.87	0.28	0.89	0.85	0.11	0.89
	no mask	0.94	0.70	0.94	0.92	0.58	0.93	0.87	0.34	0.90
DE	mask	0.76	0.47	0.75	0.73	0.37	0.75	0.66	0.20	0.71
	no mask	0.89	0.76	0.89	0.87	0.70	0.87	0.82	0.60	0.83
JP	mask	0.91	0.58	0.91	0.90	0.43	0.91	0.85	0.14	0.89
	no mask	0.91	0.58	0.91	0.90	0.48	0.92	0.85	0.11	0.89
EN	mask	SVM (WE)			DT (WE)			RF (WE)		
	no mask	0.78	0.42	0.77	0.77	0.23	0.79	0.74	0.20	0.81
EN-ext	mask	0.80	0.32	0.76	0.87	0.25	0.90	0.84	0.08	0.89
	no mask	0.86	0.49	0.84	0.89	0.37	0.90	0.85	0.14	0.89
DE	mask	0.71	0.40	0.70	0.70	0.28	0.72	0.65	0.17	0.70
	no mask	0.84	0.67	0.84	0.81	0.56	0.82	0.73	0.39	0.76
JP	mask	0.87	0.49	0.84	0.90	0.42	0.91	0.85	0.16	0.90
	no mask	0.86	0.49	0.84	0.89	0.40	0.91	0.85	0.14	0.89

Table 9: F1, Matthew’s Correlation Coefficient & Accuracy for the different CFD Models.

G Further Details on the Datasets

Review categories represented in the datasets and clue words breakdown:

Below we list the breakdown of product categories for each dataset in the format “category (total number of review sentences from the category / number of counterfactual examples/ number of non-counterfactual examples)”.

Dataset *EN-ext* contains review sentences from 4 product categories: Apparel (2500 / 297 / 2203), Digital_Ebook_Purchase (2500 / 287 / 2213), Electronics (2500 / 213 / 2287), Home (2500 / 233 / 2267).

Dataset *DE* contains review sentences from 20 categories: Automotive (47 / 31 / 16), Baby (99 / 80 / 19), Camera (816 / 597 / 219), Digital_Ebook_Purchase (426 / 259 / 167), Digital_Video_Download (1297 / 961 / 336), Electronics (7 / 5 / 2), Home Entertainment (94 / 62 / 32), Home Improvement (87 / 54 / 33), Kitchen (20 / 10 / 10), Lawn and Garden (47 / 34 / 13), Luggage (21 / 9 / 12), Music (1297 / 909 / 388), Musical Instruments (162 / 113 / 49), Office Products (40 / 25 / 15),

PC (1297 / 873 / 424), Personal_Care_Appliances (56 / 36 / 20), Sports (5 / 3 / 2), Toys (378 / 216 / 162), Watches (186 / 126 / 60), Wireless (618 / 437 / 181).

Dataset *JP* contains review sentences from 18 categories: Automotive (191 / 19 / 172), Baby (182 / 6 / 176), Camera (490 / 67 / 423), Digital_Ebook_Purchase (490 / 22 / 468), Digital_Video_Download (490 / 49 / 441), Electronics (490 / 43 / 447), Home (102 / 16 / 86), Home Entertainment (227 / 34 / 193), Home Improvement (221 / 29 / 192), Kitchen (221 / 23 / 198), Music (490 / 21 / 469), Musical Instruments (490 / 42 / 448), PC (490 / 61 / 429), Shoes (490 / 52 / 438), Sports (466 / 39 / 427), Toys (490 / 53 / 437), Watches (490 / 37 / 453), Wireless (490 / 54 / 436).

The clue phrases for English, German and Japanese are shown respectively in [Table 13](#), [Table 14](#) and [Table 15](#).

Misclassifications of Reviews Containing No Counterfactuals		Models
If you workout regularly, an extra set of 'expendable' earbuds like these is a must-have.		B XR X
I put over 500 songs on it the first day and still have around 17 GB left, probably could have done with a much smaller one.		B XR X
If you have a similar build compared to mine, buy this shirt without hesitation.		B XR X
If they ever need replacing I would definitely buy these again.		B XR X
If this device for whatever reason fails within a year or two, I think I would look to buy the same machine again.		B XR X
I was hoping she would be able to grow in it but it fits her now with no room to grow.		B XR X
I must have read reviews on about 20 different models.		B XR X
Because it is fleece, if you are in the US, I would suggest a second cool water rinse with a touch of fabric softener.		B XR X
There are ways to get it like you want it but its not as easy as it could have been.		B XR X
Could be improved with a size adjustment and chin strap.		B XR X
If you need more desk space and have a location where you can use a wall mount for your monitors, this thing is the way to go.		B XR X
It should be about \$20 cheaper to make it worth while.		B XR X

Misclassifications of Counterfactual Reviews		Models
At the end of a series like The Wheel of Time, it might be appropriate to lament the loss of familiar characters.		B XR X
You would have to be 5' 10 and super thin to fit into these.		B XR X
From the picture the dress looks like it should be long enough for someone at least 5' 6.		B XR X
To say "the usual awesome Stephen King novel" would be an understatement.		B XR X
I don't like to go into the plot a lot unless the blurb doesn't represent the book fairly.		B XR X
I've thought about it, and I guess that's because what happened to the characters in Missing are stuff that I could imagine happening to me as well.		B XR
For the price that this particular seller charged for this T-shirt, the material SHOULD be HEAVY-DUTY.		B X
If one can put aside their religious beliefs about heaven and hell I think they will find this to be something they've always known deep inside about the afterlife.		XR X
If you think leakage is a problem it really isn't they are as bad as a pair of ear-buds.		XR X

Table 10: Qualitative Examples of Incorrect Predictions from Fine-tuned BERT

Hyperparameters	mBERT	XLm-RoBERTa	XLm-w/o-Emb
Vocab Size	119547	250002	119547
Max Pos. Embeddings	512	514	514
Hidden Size	3072	3072	3072
Encoder Size	768	768	768
Num. Hidden Layers	12	12	12
Num. Hidden Groups	1	1	1
Num. Attention Heads	12	12	12
Hidden Activations	GeLU	GeLU	GeLU
Layer Norm. Epsilon	10^{-12}	10^{-12}	10^{-12}
Fully-Connected Dropout Prob.	0.1	0.1	0.1
Attention Dropout Prob.	0	0	0

Table 11: Final Transformer Model Hyperparameter Settings

Hyperparameters	mBERT	XLm-RoBERTa	XLm-w/o-Emb
Seed	1234	1234	1234
Learning rates	10^{-5}	10^{-5}	50^{-5}
Max Seq. Length	256	256	256
Max Train Epochs	50	50	50
Warmup Proportion	0.1	0.1	0.1
Classifier Dropout Prob.	0.2	0.2	0.2
Adam eps	10^{-8}	10^{-8}	10^{-8}

Table 12: Final Transformer Model Hyperparameter Settings

Clue phrase	N_P	N_N	f_P	f_N	f_{data}
without	25	571	2.42	6.36	5.96
doesn't	11	569	1.06	6.34	5.8
wanted	18	512	1.74	5.70	5.3
would be	100	389	9.70	4.33	4.89
would have	281	114	27.2	1.27	3.95
wish	305	81	29.6	0.90	3.86
couldn't	11	289	1.06	3.22	3.0
won't	9	274	0.87	3.05	2.83
must	13	258	1.26	2.87	2.71
haven't	5	208	0.48	2.31	2.13
instead of	18	132	1.74	1.47	1.5
should be	20	126	1.94	1.40	1.46
came with	9	128	0.87	1.42	1.37
could have	100	36	9.70	0.40	1.36
should have	106	19	10.2	0.21	1.25
miss	6	116	0.58	1.29	1.22
could be	19	103	1.84	1.14	1.22
except	7	115	0.67	1.28	1.22
comes with	1	80	0.09	0.89	0.81
none	2	68	0.19	0.75	0.7
missing	3	56	0.29	0.62	0.59
if it was	21	30	2.03	0.33	0.51
might have	10	20	0.97	0.22	0.3
wished	18	4	1.74	0.04	0.22
had not	3	13	0.29	0.14	0.16
if it were	13	3	1.26	0.03	0.16
if it had	10	3	0.97	0.03	0.13
wishing	5	4	0.48	0.04	0.09
had thought	3	4	0.29	0.04	0.07
Total	954	4069	92.6	45.3	50.2

Table 13: English clue words (statistics for *EN-ext* dataset). N_P (and N_N) is the number of positive (and negative) examples with the clue word. f_P (and f_N) is the percent of positive (negative) examples with the clue word. f_{data} is the frequency of the clue word in the dataset.

Clue phrase	N_P	N_N	f_P	f_N	f_{data}
hätte	1804	11	37.2	0.50	25.9
wäre	1397	22	28.8	1.01	20.2
könnte	1143	28	23.6	1.29	16.7
müssen	122	479	2.52	22.1	8.58
fehlt	111	429	2.29	19.8	7.71
wenn es	296	227	6.11	10.5	7.47
statt	107	200	2.21	9.25	4.38
außer	52	184	1.07	8.51	3.37
wünschen	115	80	2.37	3.70	2.78
müsste	174	13	3.59	0.60	2.67
wird nicht	15	167	0.30	7.73	2.6
eigentlich nicht	51	119	1.05	5.50	2.42
dürfen	34	63	0.70	2.91	1.38
vermisse	10	55	0.20	2.54	0.92
gewollt	4	33	0.08	1.52	0.52
wünschte	25	11	0.51	0.50	0.51
verpassen	13	22	0.26	1.01	0.5
konnte nicht	4	27	0.08	1.25	0.44
hatte nicht	6	11	0.12	0.50	0.24
haben nicht	1	14	0.02	0.64	0.21
könnte sein	6	0	0.12	0.0	0.08
hatte gedacht	2	4	0.04	0.18	0.08
nicht hätte	5	0	0.10	0.0	0.07
anstelle von	1	2	0.02	0.09	0.04
hätte haben können	0	0	0.0	0.0	0.0
sollte haben	0	0	0.0	0.0	0.0
wenn es hatte	0	0	0.0	0.0	0.0
Total	4840	2160	100.	100.	100.

Table 14: German clue words. N_P (and N_N) is the number of positive (and negative) examples with the clue word. f_P (and f_N) is the percent of the clue word in the dataset.

Clue phrase	N_P	N_N	f_P	f_N	f_{data}
思います	84	1556	12.5	24.5	23.4
れば	258	805	38.6	12.7	15.1
なら	84	841	12.5	13.2	13.2
でしょう	28	633	4.19	9.99	9.44
良かった	63	324	9.44	5.11	5.52
思う	32	344	4.79	5.43	5.37
いいです	17	282	2.54	4.45	4.27
よかった	53	239	7.94	3.77	4.17
思った	12	251	1.79	3.96	3.75
良いです	9	233	1.34	3.67	3.45
思いました	22	204	3.29	3.22	3.22
だろう	12	211	1.79	3.33	3.18
もっと	94	116	14.0	1.83	3.0
もう少し	131	64	19.6	1.01	2.78
良いと	21	154	3.14	2.43	2.5
ほうが	18	155	2.69	2.44	2.47
いいと	14	143	2.09	2.25	2.24
助か	7	119	1.04	1.87	1.8
べき	13	102	1.94	1.61	1.64
さらに	6	108	0.89	1.70	1.62
欲しかった	42	44	6.29	0.69	1.22
としても	2	67	0.29	1.05	0.98
いいかも	6	58	0.89	0.91	0.91
ならば	8	52	1.19	0.82	0.85
更に	7	51	1.04	0.80	0.82
たかった	10	46	1.49	0.72	0.8
思っていました	5	48	0.74	0.75	0.75
できれば	24	28	3.59	0.44	0.74
良いのですが	6	35	0.89	0.55	0.58
だったら	15	26	2.24	0.41	0.58
いいかも	2	39	0.29	0.61	0.58
おもいます	3	35	0.44	0.55	0.54
ほしかった	25	11	3.74	0.17	0.51
よいと	5	29	0.74	0.45	0.48
いいのですが	6	26	0.89	0.41	0.45
いいかな	7	22	1.04	0.34	0.41
思ったた	1	25	0.14	0.39	0.37
よいです	0	25	0.0	0.39	0.35
いいな	7	18	1.04	0.28	0.35
たらな	3	20	0.44	0.31	0.32
いいのに	8	8	1.19	0.12	0.22

Continued on next column

Continued from previous page					
Clue phrase	N_P	N_N	f_P	f_N	f_{data}
良いのでは	3	13	0.44	0.20	0.22
良いかな	2	10	0.29	0.15	0.17
いいのでは	1	10	0.14	0.15	0.15
したかった	0	10	0.0	0.15	0.14
おもう	1	7	0.14	0.11	0.11
いいんですか	0	7	0.0	0.11	0.1
良いのだが	4	3	0.59	0.04	0.1
だったのに	1	6	0.14	0.09	0.1
おもった	0	6	0.0	0.09	0.08
おもいました	1	4	0.14	0.06	0.07
良いな	2	3	0.29	0.04	0.07
良かったのに	3	2	0.44	0.03	0.07
よいかな	2	2	0.29	0.03	0.05
よいのでは	0	4	0.0	0.06	0.05
よいのですか	2	2	0.29	0.03	0.05
たすかり	1	2	0.14	0.03	0.04
よいかも	0	3	0.0	0.04	0.04
よかったのに	2	0	0.29	0.0	0.02
よいのに	2	0	0.29	0.0	0.02
ところだった	0	1	0.0	0.01	0.01
よいな	1	0	0.14	0.0	0.01
良いのに	0	1	0.0	0.01	0.01
おもっていました	0	0	0.0	0.0	0.0
おもってた	0	0	0.0	0.0	0.0
たすかった	0	0	0.0	0.0	0.0
たすかる	0	0	0.0	0.0	0.0
よいんですか	0	0	0.0	0.0	0.0
あれば	0	0	0.0	0.0	0.0
Total	667	6333	100.	100.	100.

Table 15: Japanese clue words. N_P (and N_N) is the number of positive (and negative) examples with the clue word. f_P (and f_N) is the fraction of positive (negative) examples with the clue word. f_{data} is the frequency of the clue word in the dataset.