

The Dark Side of the Language: Pre-trained Transformers in the DarkNet

Leonardo Ranaldi^(*), Aria Nourbakhsh^(*), Arianna Patrizi^(*), Elena Sofia Ruzzetti^(*),
Dario Onorati^(•), Michele Mastromattei^(*), Francesca Fallucchi^(*), Fabio Massimo Zanzotto^(*)

^(*) Human-Centric ART Group, Department of Enterprise Engineering, University of Rome Tor Vergata.

^(•) Department of Computer, Control and Management Engineering, University of Rome "Sapienza".

[first name].[last name]@uniroma2.it

Abstract

Pre-trained Transformers are challenging human performances in many NLP tasks. The massive datasets used for pre-training seem to be the key to their success on existing tasks. In this paper, we explore how a range of pre-trained Natural Language Understanding models perform on definitely unseen sentences provided by classification tasks over a DarkNet corpus. Surprisingly, results show that syntactic and lexical neural networks perform on par with pre-trained Transformers even after fine-tuning. Only after what we call extreme domain adaptation, that is, retraining with the masked language model task on all the novel corpus, pre-trained Transformers reach their standard high results. This suggests that huge pre-training corpora may give Transformers unexpected help since they are exposed to many of the possible sentences.

1 Introduction

Transformers (Zhang et al., 2019; Radford and Narasimhan, 2018) have been rocking the field of NLP. These Transformers are outperforming all previous methods and, sometimes, even humans in many NLP tasks (Wang et al., 2018, 2020; Kalyan et al., 2021; Guo et al., 2022).

Pre-training on large corpora seems to be the key that boosts performances, as these Transformers may induce clear models of target languages. Indeed, BERT is pre-trained on an English corpus of 3,300M words consisting of books (Zhu et al., 2015a) and Wikipedia. The English version of the last ERNIE (Sun et al., 2021a) is trained on an even more extensive corpus. MEGATRON-LM (Shoeybi et al., 2019) utilizes an incredible corpus of 174 GB, and the Chinese version of ERNIE breaks the records by exploiting a 4TB corpus (Sun et al., 2021b). Therefore, the challenge is training over always more massive corpora.

Huge pre-training corpora may give unexpected help to Transformers: their successes in downstream tasks can be because Transformers have seen large parts of possible sentences. This could be a sort of overfitting. This possible shortcoming is sometimes considered when novel Transformers are introduced (Radford et al., 2019; Shoeybi et al., 2019). For this reason, Radford et al. (2019) have excluded Wikipedia pages for pre-training as it is a common data source for other downstream datasets. Yet, when using off-the-shelf pre-trained models, this caution is generally disregarded. For example, the discovering ongoing conversation (DOC) task, introduced by Zanzotto and Ferrone (2017) was found challenging for humans, but the BERT baseline model achieved the astonishing 88.4 F1 score (Wang et al., 2020). DOC consists of determining if two utterances are contiguous in classical theatrical plays. These plays may be included in the book dataset (Zhu et al., 2015a) used for pre-training BERT. This may explain the superhuman performance in such a challenging task.

Corpora and related tasks derived from the DeepWeb and DarkWeb (Ranaldi et al., 2022b,a; Avarikioti et al., 2018; Choshen et al., 2019) offer a tremendous opportunity to study Transformers and other natural language models on definitely unseen sentences.

Performances on these tasks cannot depend on overfitting over-seen sentences.

Indeed, it is extremely unlikely that texts extracted from these sources are included in pre-training corpora. Moreover, language on the DarkNet may have very different characteristics with respect to the one accessible from the surface web (Choshen et al., 2019).

In this paper, we aim to explore how pre-trained Natural Language Understanding models behave on definitely unseen sentences. These definitely unseen sentences are provided by the DarkNet corpus

	Onion Drugs		Onion Forums		Surface Web	
	Legal	Illegal	Legal	Illegal	eBay drugs	DBpedia sample
#tokens	41,683	67,506	41,683	43,654	114,817	46,792
#types	8,576	12,334	8,576	9,411	18,405	7,941
types/token ratio	4,86	5,47	4,86	4,36	6,23	5,88
BERT’s types pieces/OOVs ratio	2,94	3,96	3,67	3,15	5,03	2,40

Table 1: Lexical description of the corpus subsets and their lexical coverage with respect to the vocabulary of BERT compared with a sample of the DbPedia dataset (Zhang et al., 2015).

along with a classification task. We experimented with Stylistic Classifiers based on the bleaching text model (van der Goot et al., 2018), with Lexical Neural Networks based on GloVe (Pennington et al., 2014) and word2vec (Mikolov et al., 2013), with Syntactic-based neural networks based on KERMIT (Zanzotto et al., 2020), and with holistic Transformers such as BERT (Devlin et al., 2019), XLNet (Yang et al., 2019), ERNIE (Zhang et al., 2019) and Electra (Clark et al., 2020). Results show that syntactic and lexical neural networks surprisingly outperform pre-trained Transformers even after fine-tuning (Wei et al., 2021). Only when pre-training is extended to the novel corpus and Transformers see these definitely unseen sentences do their performances increase to the expected level. This seems to suggest that huge pre-training corpora may give Transformers the unexpected help of showing them many possible sentences.

The rest of the paper is organized as follows: Material and Methods, Results and Discussion, Conclusions, and Limitations. The code and data are publicly available at: https://github.com/ART-Group-it/Transformers_DarkWeb.

2 Material and Methods

This section describes the corpus used to challenge Transformers that is a source of definitely unseen sentences (Section 2.1) and the investigated transformers along with other neural network methods (Section 2.2).

2.1 Material: A Dark Web Dataset

Corpora scraped from DarkWeb to fight illegal actions offer a tremendous opportunity for studying how large pre-trained models behave on definitely unseen sentences. Indeed, as discussed in section 2.2, pre-trained Transformers as well as “pre-trained” syntactic neural networks have not used any of these corpora for training. Hence, corpora collected for totally different reasons can help to

shed light on an important research question: is pre-training on large datasets a sort of overfitting on many of the possible sentences?

Classifying legal and illegal actions is a key task in the DarkWeb (also called Onion Web). Hence, it is also crucial to try to understand whether or not Transformers or other neural network-based models add value to models for this challenging task.

2.1.1 Using an Existing Corpus: DUTA-10k

For our experiments, we used the “Darknet Usage Text Addresses” (DUTA-10k) (Nabki et al., 2019) as exploited in Choshen et al. (2019). The corpus DUTA-10k (Nabki et al., 2019) contains onion web data manually tagged in legal and illegal samples. Choshen et al. (2019) selected only the drug subdomain and used four different subsets of 571 samples each: (1) legal onion drugs, (2) illegal onion drugs, (3) onion forums discussing legal activities, and (4) onion forums discussing illegal topics. Additionally, to compare with the data from surface web, Choshen et al. (2019) have extracted item descriptions from eBay as well. These descriptions were selected by searching the keywords, which are, ‘marijuana’, ‘weed’, ‘grass’, and ‘drug’. The resulting DUTA-10K used in Choshen et al. (2019) and in our experiments contains 5 subsets (see Table 5 for some examples): the four manually annotated onion web sets and the eBay drugs set.

Choshen et al. (2019) propose to use data from DUTA-10k for the task of classifying legal and illegal activities. Hence, the five subsets are used to produce four different classification datasets: (1) eBay vs. legal drugs; (2) legal vs. illegal drugs; (3) legal vs. illegal forums; and, finally, (4) legal and illegal drugs as training data vs. legal and illegal forums as testing dataset. The last task is the most important and complex as it tests knowledge transferability as training is on the drug dataset and

the testing is on a totally different domain.

2.1.2 Retrieving DUTA-10K and Final Corpus

Using the corpus proposed in Choshen et al. (2019) is not straightforward as DUTA-10k has to be re-constructed by scraping the DarkWeb again. Indeed, for legal reasons, the dataset contains only manually tagged links to the DarkWeb¹.

In our experiments, we then extracted the corpus and prepared the dataset by using links and the tools provided by Choshen et al. (2019). The preprocessing consists of removing: HTML tags, non-linguistic content such as buttons, encryption keys, metadata, and common words such as “Show more results”. All the 571 samples for all the subsets still exist in the DarkWeb and, thus, have been used to create our dataset. The resulting subsets (see Tab. 1) are not extremely large in terms of tokens.

Since there is not a standard split in the provided datasets, we prepared 5 different 70%-30% splits in training and testing of each subtask. Results presented in Choshen et al. (2019) are obtained on a single 70%-30% split. However, in our experiments, we opted for these 5 different 70%-30% splits to evaluate the stability of experimental results. Results may vary from split to split. Indeed, our results differ from the experiments conducted in Choshen et al. (2019) when we tried to replicate their models.

In addition to the provided datasets, we also used a random subset of the DBpedia dataset (Zhang et al., 2015) for comparing the language of the datasets with a more general language.

2.2 Methods: Classification Models

Our goal is to investigate how Transformers and other pre-trained models perform on definitely unseen sentences. This section introduces the models used in this study: lexical-based neural networks (Sec. 2.2.2), syntax-based neural networks (Section 2.2.3), and Transformers (Section 2.2.4). The description of the pre-trained models discusses the size of the corpora used to train each model as well.

To discard the idea that the chosen task has strong stylistic signals where Transformers perform poorly (see results for the corpus linguistic acceptability task in Warstadt et al. (2019)), this section utilizes two style-oriented classifiers (Section 2.2.1)

¹Data and code are available in Choshen et al. (2019)’s GitHub repository <https://github.com/huji-nlp/cyber>

to investigate whether determining legal and illegal activities is indeed only a stylistic task. The code can be found in the file *Code The Dark Side of the Language.zip* and will be made publicly available.

2.2.1 Style-oriented Classifiers

Legal and illegal activities may be described with different styles of language: a formal vs. a more informal style of writing. For this reason, we use two style-oriented classifiers as detectors to understand if this task is merely stylistic.

The first is a family of classifiers that exploits *part-of-speech (POS) tags*, treating them as bag-of-POSs. These are very simple models which mainly capture the distribution of POSs in target texts. In line with Choshen et al. (2019), we tested two non-neural models, namely Naive Bayes NB(POS) and Support Vectors Machines. In addition, we tested BoPOS, a simple feed-forward neural network using bag-of-POS as input representation.

Bleaching text (van der Goot et al., 2018) is a model proposed to capture the style of writing. Originally, it has been applied to cross-lingual authors’ gender prediction. This model converts sequences of tokens, e.g., ‘1x Pcs Mobile Case!? US\$65’, into abstract sequences according to these rules presented with the effect on the example:

- each token is replaced by its length (effect: ‘02 03 06 06 05’)
- alphanumeric characters are merged into one single letter and other characters are kept (effect: ‘W W W W!? W\$W’)
- punctuation marks are transformed into a unified character (effect: ‘W W W WPP W’)
- upper case letters are replaced with ‘u’, lower case letters with ‘l’, digits with ‘d’, and the rest to ‘x’ (effect: ‘dl ull ull ullxx uuxdd’)
- consonants are replaced with ‘c’, vowels to ‘v’ and the rest to ‘o’ (effect: ‘oc ccc cvcvcv cvcvoo vcooo’)

Finally, a sample is represented by the concatenation of all the above transformations. For classification, we use a linear SVM classifier with a binary bag of word representation.

2.2.2 Lexical-based Neural Networks

To investigate the role of pre-trained word embeddings, we used a classifier based on vanilla feed-forward neural networks (FFNN) over bag-of-word-embedding (BoE) representations. In BoE(GloVe), sentence representations are computed as the sum of word embeddings representing

their words. BoE(GloVe) uses GloVe word embeddings (Pennington et al., 2014) trained on 2014 Wikipedia dumps and Giga5 (see Table 6). The supporting FFNN of BoE(GloVe) consists of an input layer of dimension 300 and 2 hidden layers of 150 and 50 dimensions with the *ReLU* activation function.

2.2.3 Syntactic-based Neural Networks

To evaluate the role of “pre-trained” universal syntactic models, we used the Kernel-inspired Encoder with Recursive Mechanism for Interpretable Trees (KERMIT) (Zanzotto et al., 2020). This model positively exploits parse trees in neural networks as it increases the performances of pre-trained Transformers when it is used in combined models.

The version used in the experiments encodes parse trees in vectors of 4,000 dimensions. The rest of the feed-forward network is composed of 2 hidden layers of dimension 4,000 and 2,000 respectively, and finally the output layer of dimension 2. Between each layer, the *ReLU* activation function and a dropout of 0.1 is used to avoid overfitting on the train data.

Even in this case, the model is somehow ‘pre-trained’. In fact, KERMIT exploits parse trees produced by a traditional parser. In our experiments, we used the English constituency-based parser in CoreNLP (Zhu et al., 2013). The parser is trained on the standard WSJ Penn Treebank (Marcus et al., 1993), which contains only around 1M words.

2.2.4 Holistic Transformers

We tested the following Transformers to cover the majority of cases of pre-training size (see Table 6) and models:

BERT (Devlin et al., 2019), the architecture Bidirectional Encoder Representations from Transformers, trained on the BooksCorpus (Zhu et al., 2015b) and English Wikipedia ($BERT_{base}$).

XLNet (Yang et al., 2019), a generalized autoregressive pre-training technique that allows the learning of bidirectional contexts by maximizing the expected likelihood over all permutations of the factorization order and to its autoregressive formulation. XLNet is trained on 32.89 billion tokens, taken from datasets gathered from the surface web or publicly available datasets, such as Wikipedia, Bookcorpus, Giga5, Clueweb, and Common Crawl.

ERNIE (Sun et al., 2021a), an improved language model that addresses the inadequacy of

BERT and utilizes an external knowledge graph for named entities. ERNIE is pre-trained on Wikipedia corpus and Wikidata knowledge base.

ELECTRA (Clark et al., 2020), an improved BERT where, instead of masking input tokens, these are “corrupted” with replacement tokens that potentially fit the place. The training procedure is a classification of each token on if it is a corrupted input or not. To make its performance comparable to BERT, they have trained the model on the same dataset that BERT was trained on.

The models for all the proposed Transformers were implemented using the Transformers library from Huggingface and the pre-trained version of AutoModelForSequenceClassification (Wolf et al., 2020). For each model, we chose the Adam optimizer (Kingma and Ba, 2015) with a batch size of 32 and fine-tuned the model for 5 epochs, following the original paper (Devlin et al., 2019). For hyperparameter tuning, the best learning rate is different for each task, and all original authors choose one between 1×10^{-5} and 5×10^{-5} . All the other settings are the same as those used in the original papers.

3 Results and Discussion

This section investigates how pre-trained Transformers and other pre-trained language models behave on definitely unseen sentences (Sec. 3.3). Yet, to exclude that the nature of the task biases our study, we have performed additional analyses: 1) to understand whether the onion language is really different from the surface web language (Sec. 3.1), and 2) to determine the nature of the proposed classification task (Sec. 3.2).

3.1 Onion Language and Surface Web Language

One important concern is whether Onion texts have some specific features that make it hard to analyze with *surface-web-oriented* language analyzers. For this purpose, we compared the onion subsets with the surface web subset, that is, the *eBay drugs* set (see Table 1 and Figure 1).

Surface web and onion web texts seem to have a language with similar basic features. For example, type/token ratio of the different onion subsets is similar to the one of *eBay drugs* (Table 1). Indeed, all these datasets contain many unique tokens as different drugs have different names (see Table 5). Moreover, the frequency of tokens vs. their

	eBay/Legal Drugs	Drugs	Forums	Drugs/Forums
<i>NB (POS)</i> (Choshen et al., 2019)	91.4	77.6	74.1	78.4
<i>SVM (POS)</i> (Choshen et al., 2019)	63.8	63.8	85.3	62.1
(our) <i>NB (POS)</i>	90.4(± 1.3)	83.7(± 0.7)	64.3(± 2.2)	48.2(± 1.8)
(our) <i>SVM (POS)</i>	86.6(± 0.6)	83.6(± 1.2)	61.6(± 1.5)	43.2(± 2.3)
<i>Bleaching text</i>	84.73(± 0.8)	81.3(± 0.6)	56.65(± 1.7)	55.68(± 1.3)

Table 2: Accuracies of style-oriented classifiers over 5 different splits on the Legal vs. Illegal Classification Task and accuracies of NB (POS) and SVM (POS) as reported in (Choshen et al., 2019)

		eBay/Legal Drugs	Drugs	Forums	Drugs→Forums
<i>Freeze</i>	<i>BERT</i>	66.62(± 3.1)	64.35(± 2.7)	53.12(± 1.2)	49.2(± 1.8)
	<i>Electra</i>	71.36(± 2.9)	61.56(± 3.1)	54.22(± 1.9)	50.33(± 2.2)
	<i>XLNet</i>	59.92(± 3.4)	56.77(± 2.7)	50.32(± 2.6)	51.86(± 1.8)
	<i>Ernie</i>	74.56(± 2.4)	60.33(± 1.9)	60.41(± 3.2)	50.33(± 2.3)
	<i>BoE(GloVe)</i>	91.50(± 0.5)	81.60(± 1.4)	54.60(± 1.4)	53.50(± 1.5)
	<i>KERMIT</i>	90.50(± 1.0)	79.00(± 1.0)	66.60(± 1.4)	58.37 (± 1.26)
	<i>BoE(GloVe) + KERMIT</i>	93.54(± 1.46)	83.10(± 1.4)	66.20(± 1.4)	54.30(± 2.34)
<i>Fine-Tuning</i>	<i>Electra</i>	93.47(± 2.5)	85.15(± 1.33)	64.22(± 1.38)	48.96(± 2.22)
	<i>XLNet</i>	92.58(± 1.99)	84.95(± 2.13)	65.32(± 1.69)	48.22(± 2.46)
	<i>Ernie</i>	94.97(± 2.11)	85.19(± 2.32)	66.43(± 1.79)	50.12(± 2.23)
	<i>BERT</i>	94.36(± 3.20)	84.35(± 3.16)	65.18(± 2.28)	50.68(± 2.49)
	<i>BERT</i> _{with DomA}	95.43(± 2.17)	83.76(± 1.70)	70.95(± 2.56)	51.7(± 2.23)
	<i>BERT</i> _{with ExtremeDomA}	97.4 (± 2.30)	89.7 (± 3.10)	72.4 (± 3.30)	55.6(± 2.90)

Table 3: Accuracies of pre-trained models on the Legal vs. Illegal Classification Task on the DarkWeb Corpus (Choshen et al., 2019). BERT, XLNet, Ernie, and Electra are those specific for SequenceClassification (Wolf et al., 2019). Where applicable, models are: (1) without or with fine-tuning; (2) with domain-adaptation using only the training sets (*DomA*); (3) and, with extreme domain-adaptation using training and testing sets (*ExtremeDomA*). Experiments are obtained over 5 runs over the 5 different splits with 5 different seeds for initializing weights of neural networks.

length is quite similar in the surface web and the onion web (see Figure 1a). There are no strange peaks or outliers. Finally, when analyzed with a symbolic parser (Zhu et al., 2013), these subsets have the same syntactic characteristics. In fact, frequencies of POS tags and constituent types are similar across dataset subsets (see Figure 1b). Indicators of not analyzed sentences such as FRAG have a similar distribution. Also, the indicators of unknown words, which are NNP and NNPS, are similar across all the subsets. The only differences are tags for pronouns (PRP) and cardinal numbers (CD). It seems that pronouns are more frequent for legal activities on the onion web. Overall, onion and surface web data are mostly parsed in a similar manner.

In conclusion, the difference in the performance of the different models on the different datasets is not due to an inherent difference in the languages of the onion and surface web domains. Moreover,

syntactic-based neural networks are not taking advantage of some hidden syntactic bias.

3.2 Illegal vs. Legal is only a Stylistic Task?

Determining whether or not a piece of text is illegal seems to be a stylistic task. The hypothesis is that illegal texts are written in a way that is *stylistically* different. Indeed, experiments of Choshen et al. (2019) seem to suggest that this is the case. Simple models based on POS tags using SVM or Naive Bayes have very high performances (lines 1 and 2 in Table 2).

According to our experiments, the task is not only stylistic. We repeated the measures on the 5 splits we proposed (see Section 2.1.2) and results are quite different with respect to those of Choshen et al. (2019) on a single split (lines 3 and 4 vs. lines 1 and 2 in Table 2). Two datasets of eBay/Legal Drugs and Drugs seem to be strongly correlated with style and simple features. Yet, no strong POS

tag features emerged from the Naive Bayes models and, thus, there are no apparent artifacts in these datasets. Instead, Forum and Drugs/Forums datasets are less correlated, if not unrelated.

Moreover, our tests with *Bleaching text* model confirm that the proposed task is not solely a stylistic task. This model has been designed to capture only stylistic features (van der Goot et al., 2018) (see Section 2.2.1). Its results are quite high for eBay/Legal Drugs and Drugs and relatively low for Forum and Drugs/Forums. This is in line with the shown findings on POS-tag-based models.

3.3 Investigating pre-trained Transformers, Lexical, and Syntactic Models

We can now focus on the performance of the different pre-trained models on the novel, unexplored task – classifying legal and illegal texts in the onion web – taking into account that there is not a real difference between the language of onion and the one of the surface web, but it is very unlikely that these definitely unseen sentences of the onion corpus, or very similar sentences, have been used for pre-training models. Lexical-based and syntactic-based neural networks outperform Holistic transformers when all these models are considered universal linguistic knowledge embedders and general parameters are not fine-tuned (*Freeze* in Tab. 3). Indeed, the accuracy of BoE(Glove), KERMIT, and their combination are well above the results of Holistic Transformers. It is common knowledge that Transformers need fine-tuning. However, this is the fairest comparison with other models, which cannot benefit from fine-tuning. Syntactic parsers aim to capture the general structure of language and cannot and should not be adapted to a particular task. In fact, language and knowledge about language are general, hence it is not clear why this general knowledge should be adapted to tasks with fine-tuning.

Fine-tuning boosts the performance of holistic transformers except for the task *Drugs*→*Forums* (see Tab. 3). In fact, performance for the other three tasks *eBay/Legal Drugs*, *Drugs*, and *Forums* have a dramatic increase in accuracy. To obtain these results, all layers of transformers should be fine-tuned (see Tab. 8). Apparently, there is not a predominant set of layers that help performance to have a big increase in accuracy suggesting that some kind of linguistic knowledge is more important than another. The absence of an increase in

performance for the task *Drugs*→*Forums* is extremely interesting. Indeed, this task asks to learn a legal/illegal classifier in an environment and apply it in another environment. Fine-tuning is definitely not helping for this out-of-domain task.

Fine-tuned holistic transformers do not have an important increase in performance with respect to lexical-based and syntactic-based neural networks on these datasets with definitely unseen sentences. BoE(GLove)+KERMIT, which cannot be fine-tuned, are basically on par with fine-tuned transformers for the three tasks. This suggests that fine-tuning is not helping transformers to grab additional knowledge on definitely unseen sentences, but it seems to adapt weights to solve final tasks better. Moreover, BoE(GLove)+KERMIT and KERMIT alone still outperform transformers on the out-of-domain task *Drugs*→*Forums*.

Extreme domain adaptation produces a real change in the performance of transformers with respect to lexical-based and syntactic-based neural networks (see last line of Tab. 3). Classical domain adaptation is not improving for *Drugs* and it is improving only a little for *eBay/Legal Drugs* and *Drugs*→*Forums*. Hence, when transformers see definitely unseen sentences with MLM, they seem to incorporate the knowledge needed to treat sentences in final tasks better. This may suggest that, in other classical tasks, pre-training plays a crucial role as sentences may at least have been partially seen during pre-training.

Despite all the domain adaptation and fine-tuning, holistic transformers are not gaining real clues on the difference between legal and illegal language. The best accuracy in the out-of-domain task *Drugs*→*Forums* remains that of KERMIT, the syntax-based neural network. Hence, the compelling question is: what are these models really learning?

3.4 Qualitative analysis

Transformers confirm to have astonishing results if considered in a task within a single dataset and if they have partially seen sentences, that is, *BERT* with *ExtremeDomA*. Then, the compelling question of what they are learning should at least be addressed. For this reason, we performed a qualitative analysis of a very small part of the dataset. Focusing on examples with the most frequent range of lengths (see Fig. 1a), we selected 12 examples to better analyze the results (see Tab. 7).

Acknowledgements

This paper has been supported by “Knowledge At the Tip of Your fingers: Clinical Knowledge for Humanity” (KATY) project funded by the European Union’s Horizon 2020 research and innovation program under grant agreement No. 101017453.

We would also like to thank all the anonymous reviewers that helped to strengthen this paper.

Limitations

We believe that the main results obtained in this paper are convincing: Transformers behave better if they see in advance, at least, part of the corpus. Yet, our paper leaves some open avenues to explore besides the two future research lines described in the conclusions.

One limitation is due to the fact that we explored only one possible corpus with definitely unseen sentences. To assess these results better, additional corpora should be taken into consideration. For researchers outside big companies, retrieving such corpora is extremely difficult. As a possible solution, these corpora should be retrieved where they naturally and publicly occur.

Ethics Statement

Navigating the Dark Web may be extremely dangerous. Indeed, it may contain offensive content or illegal content, and tasks themselves could potentially have harmful content. We really need to describe how this is navigated and, more importantly, what are the real pros of using such resources.

We propose to navigate the Dark Web in a textual manner so that there is no exposure and no need to download sensitive visual material that may result in a crime. In this textual version of the Dark Web, systems may be exposed to offensive or illegal content, but offensive content is everywhere, and then it is not a real problem.

However, using Dark Web material is a way to access really unseen text, which has never been used in Pre-trained Transformers. Unlike what we believe can happen in large companies, public researchers hardly have the possibility to access user-produced data that are not seen. This is a democratic way to obtain such truly unseen data.

References

- Georgia Avarikioti, Roman Brunner, Aggelos Kiayias, Roger Wattenhofer, and Dionysis Zindros. 2018. [Structure and content of the visible darknet](#).
- Jamie Callan, Mark Hoy, Changkuk Yoo, and Le Zhao. 2009. Clueweb09 data set.
- Leshem Choshen, Dan Eldad, Daniel Hershcovich, Elior Sulem, and Omri Abend. 2019. [The language of legal and illegal activity on the Darknet](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4271–4279, Florence, Italy. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training text encoders as discriminators rather than generators](#). In *ICLR*.
- Common Crawl. 2019. Common crawl. URL: <http://commoncrawl.org>.
- 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)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Rob van der Goot, Nikola Ljubešić, Ian Matroos, Malvina Nissim, and Barbara Plank. 2018. [Bleaching text: Abstract features for cross-lingual gender prediction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 383–389, Melbourne, Australia. Association for Computational Linguistics.
- Shangwei Guo, Chunlong Xie, Jiwei Li, Lingjuan Lyu, and Tianwei Zhang. 2022. [Threats to pre-trained language models: Survey and taxonomy](#).
- Katikapalli Subramanyam Kalyan, Ajit Rajasekharan, and Sivanesan Sangeetha. 2021. [Ammus : A survey of transformer-based pretrained models in natural language processing](#).
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- Mitchell P. Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. [Building a large annotated corpus of English: The Penn Treebank](#). *Computational Linguistics*, 19(2):313–330.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.

- Mhd Wesam Al Nabki, Eduardo FIDALGO, Enrique Alegre, and Laura Fernández-Robles. 2019. Torank: Identifying the most influential suspicious domains in the tor network. *Expert Syst. Appl.*, 123:212–226.
- R Parker, D Graff, J Kong, K Chen, and K Maeda. 2011. English gigaword fifth edition ldc2011t07 (tech. rep.). Technical report, Technical Report. Linguistic Data Consortium, Philadelphia.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. **GloVe: Global vectors for word representation**. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Alec Radford and Karthik Narasimhan. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Leonardo Ranaldi, Aria Nourbakhsh, Francesca Fallucchi, and Fabio Massimo Zanzotto. 2022a. **C-OSINT: COVID-19 open source artificial intelligence framework**. In *Proceedings of the Italian Conference on Cybersecurity (ITASEC 2022), Rome, Italy, June 20-23, 2022*, volume 3260 of *CEUR Workshop Proceedings*, pages 219–235. CEUR-WS.org.
- Leonardo Ranaldi, Federico Ranaldi, Francesca Fallucchi, and Fabio Massimo Zanzotto. 2022b. **Shedding light on the dark web: Authorship attribution in radical forums**. *Information*, 13(9).
- Mohammad Shoeybi, Mostofa Ali Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-LM: Training multi-billion parameter language models using model parallelism. *ArXiv*, abs/1909.08053.
- Yu Sun, Shuohuan Wang, Shikun Feng, Siyu Ding, Chao Pang, Junyuan Shang, Jiayang Liu, Xuyi Chen, Yanbin Zhao, Yuxiang Lu, Weixin Liu, Zhihua Wu, Weibao Gong, Jianzhong Liang, Zhizhou Shang, Peng Sun, Wei Liu, Xuan Ouyang, Dianhai Yu, Hao Tian, Hua Wu, and Haifeng Wang. 2021a. **Ernie 3.0: Large-scale knowledge enhanced pre-training for language understanding and generation**. *ArXiv*, abs/2107.02137.
- Yu Sun, Shuohuan Wang, Shikun Feng, Siyu Ding, Chao Pang, Junyuan Shang, Jiayang Liu, Xuyi Chen, Yanbin Zhao, Yuxiang Lu, Weixin Liu, Zhihua Wu, Weibao Gong, Jianzhong Liang, Zhizhou Shang, Peng Sun, Wei Liu, Xuan Ouyang, Dianhai Yu, Hao Tian, Hua Wu, and Haifeng Wang. 2021b. **Ernie 3.0: Large-scale knowledge enhanced pre-training for language understanding and generation**.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2020. **Superglue: A stickier benchmark for general-purpose language understanding systems**.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. **GLUE: A multi-task benchmark and analysis platform for natural language understanding**. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. **Neural network acceptability judgments**. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Colin Wei, Sang Michael Xie, and Tengyu Ma. 2021. Why do pretrained language models help in downstream tasks? an analysis of head and prompt tuning. In *NeurIPS*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. HuggingFace’s Transformers: State-of-the-art Natural Language Processing. *ArXiv*, abs/1910.0.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. **Transformers: State-of-the-art natural language processing**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. **XLNet: Generalized Autoregressive Pretraining for Language Understanding**. Curran Associates Inc., Red Hook, NY, USA.
- Fabio Massimo Zanzotto and Lorenzo Ferrone. 2017. **Have you lost the thread? discovering ongoing conversations in scattered dialog blocks**. *ACM Trans. Interact. Intell. Syst.*, 7(2).
- Fabio Massimo Zanzotto, Andrea Santilli, Leonardo Ranaldi, Dario Onorati, Pierfrancesco Tommasino, and Francesca Fallucchi. 2020. **KERMIT: Complementing transformer architectures with encoders of explicit syntactic interpretations**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 256–267, Online. Association for Computational Linguistics.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. **Character-level convolutional networks for text classification**. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [Ernie: Enhanced language representation with informative entities](#).

Muhua Zhu, Yue Zhang, Wenliang Chen, Min Zhang, and Jingbo Zhu. 2013. [Fast and accurate shift-reduce constituent parsing](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 434–443, Sofia, Bulgaria. Association for Computational Linguistics.

Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015a. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#). In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 19–27.

Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015b. [Aligning books and movies: Towards story-like visual explanations by watching movies and reading books](#).

A Appendix

eBay drugs	Legal Onion	Illegal Onion
Asclepias (butterfly weed) seeds. This E-Z grower is a butterfly magnet & host plant for the Queen & Monarch butterfly. Loves full sun, is drought tolerant and prefers sandy, dry soil or gravel soil. It is a favorite of hummingbirds. 20 seeds per package.	Generic Synthroid is used for treating low thyroid activity and treating or suppressing different types of goiters. It is also used with surgery and other medicines for managing certain types of thyroid cancer.	Known as: Clomid / Clofi / Fertomid / Milophene / Wellfert Generic Clomid is used for treating female infertility. Select dosage 100mg Package Price Per pill Savings Order 25mg x 30 pills Price: \$29.95 Per pill: \$1.00
Big book-opening drugs bag 2 parts with adjustable elastics to feature 60 different sized ampoules Central padding with transparent pocket for ampoules list and expiry date. Made in red, tear-resistant, water-resistant.	Propecia is used for treating certain types of male pattern hair loss (androgenic alopecia) in men. It is also used to treat symptoms of benign prostatic hyperplasia (BPH) in men with an enlarged prostate.	TAMOXIFEN blocks the effects of estrogen. It is commonly used to treat breast cancer. It is also used to decrease the chance of breast cancer coming back in women who have received treatment for the disease.

Table 5: Example paragraphs (data instances) taken from the Legal Onion and Illegal Onion subsets training sets of the drug-related corpus. Each paragraph is reduced to the first 50 characters for space reasons.

B Appendix

<i>Corpus</i>	<i>Size</i>	BERT		<i>Electra</i>	<i>XLNet</i>	<i>Ernie</i>	KERMIT (Parser)
		<i>base</i>	<i>multi</i>				
BooksCorpus (Zhu et al., 2015b)	800M words	✓		✓	✓		
2010-and-2014-English Wikipedia dump	2,500M words	✓	✓	✓	✓	✓	
Giga5 (Parker et al., 2011)	16GB	✓		✓	✓		
Common Crawl (Crawl, 2019)	110GB				✓		
ClueWeb (Callan et al., 2009)	19GB				✓		
Penn Treebank (Marcus et al., 1993)	1M words						✓

Table 6: Pre-training corpora with their size. All corpora are derived from the surface web.

C Appendix

Text	Oracle	$BERT_{\text{with ExtremeDomA}}$	
		Runs 1 and 3	Run 2
You should never take more than one dose more than once a day.	illegal	illegal	illegal
4. Fill in the order information required	illegal	illegal	illegal
All Items Ship Via First Class Air Mail - registered or unregistered, add an extra \$25 and we ship express mail EMS	legal	legal	legal
All Major Credit, Debit, Gift, and Prepaid Cards Accepted	legal	legal	legal
All prices are in Australian dollars. (AUD) Weight: 3.5g 20 Clear	illegal	illegal	illegal
All times are UTC	illegal	illegal	illegal
auince September 15, 2016 Super fast shipping. Great product as always.	illegal	illegal	illegal
Balkan Pharmaceuticals Ltd. (Moldova) Turinabol	legal	legal	legal
Do you have a coupon code?	legal	legal	legal
free shipping On all orders \$50.00 or more	legal	legal	legal
Generic Clomid is used for treating female infertility. More Info	legal	legal	legal
Known as: Clomid / Clofi / Fertomid / Milophene / Ovamid / Serophene / Wellfert Generic Clomid is used for treating female infertility. Select dosage 100mg Package Price Per pill Savings Order 25mg x 30 pills Price:\$ 29.95 Per pill:\$ 1.00 Order:	illegal	legal	illegal

Table 7: Dataset Drugs: Oracle classifications along with classifications of three runs of $BERT_{\text{with ExtremeDomA}}$ with three different seeds.

D Appendix

Model	eBay/Legal Drugs	Drugs	Forums	Drugs/Forums
$BERT$	94.36(± 3.20)	84.35(± 3.16)	65.18(± 2.28)	50.68(± 2.49)
$BERT_{\text{last } 2}$	73.25(± 2.6)	68.26(± 3.4)	59.94(± 2.8)	49.93(± 3.6)
$BERT_{\text{last } 4}$	80.02(± 2.2)	70.62(± 2.8)	59.11(± 2.9)	50.75(± 1.9)
$BERT_{\text{last } 6}$	86.89(± 2.9)	71.47(± 2.6)	60.56(± 1.9)	52.17(± 2.9)
$BERT_{\text{last } 8}$	77.62(± 2.6)	75.22(± 3.2)	65.08(± 2.7)	50.81(± 2.4)
$BERT_{\text{last } 10}$	89.95(± 2.4)	79.37(± 2.7)	62.96(± 4.2)	50.11(± 2.7)
$BERT_{\text{with DomA}}$	95.43(± 2.17)	83.76(± 1.70)	70.95(± 2.56)	51.7(± 2.23)
$BERT_{\text{with ExtremeDomA}}$	97.4 (± 2.30)	89.7 (± 3.10)	72.4 (± 3.30)	55.6 (± 2.90)

Table 8: Accuracies for $BERT_{\text{base}}$: (1) fine-tuned on the *last n* layers; (2) domain-adapted (DomA) without and with fine-tuning; (3) sentence-adapted (SenA) without and with fine-tuning. Experiments are obtained over 5 runs with different seeds.