

An investigation of structures responsible for gender bias in BERT and DistilBERT

T. Leteno¹, A. Gourru¹, C. Laclau², and C. Gravier¹

¹ Université Jean Monnet Saint-Étienne, CNRS, Institut d'Optique Graduate School, Laboratoire Hubert Curien UMR 5516, F-42023, SAINT-ÉTIENNE, FRANCE

`firstname.lastname@univ-st-etienne.fr`

² Télécom Paris, Institut Polytechnique de Paris

`charlotte.laclau@telecom-paris.fr`

Abstract. In recent years, large Transformer-based Pre-trained Language Models (PLM) have changed the Natural Language Processing (NLP) landscape, by pushing the performance boundaries of the state-of-the-art on a wide variety of tasks. However, this performance gain goes along with an increase in complexity, and as a result, the size of such models (up to billions of parameters) represents a constraint for their deployment on embedded devices or short-inference time tasks. To cope with this situation, compressed models emerged (e.g. DistilBERT), democratizing their usage in a growing number of applications that impact our daily lives. A crucial issue is the fairness of the predictions made by both PLMs and their distilled counterparts. In this paper, we propose an empirical exploration of this problem by formalizing two questions: (1) Can we identify the neural mechanism(s) responsible for gender bias in BERT (and by extension DistilBERT)? (2) Does distillation tend to accentuate or mitigate gender bias (e.g. is DistilBERT more prone to gender bias than its *uncompressed version*, BERT)? Our findings are the following: (I) one cannot identify a specific layer that produces bias; (II) every attention head uniformly encodes bias; except in the context of underrepresented classes with a high imbalance of the sensitive attribute; (III) this subset of heads is different as we re-fine tune the network; (IV) bias is more homogeneously produced by the heads in the distilled model.

Keywords: Language Models · Fairness · Imbalance · Compression.

1 Introduction

The introduction of large Pre-trained Language Models (PLM) has marked an important paradigm shift in Natural Language Processing (NLP). It leads to unprecedented progress in tasks such as machine translation, document classification [10], and multitasks text generation [29]. The strength of these approaches lies in their ability to produce contextual representations. They have been initially based on Recurrent Neural Networks (RNN) [5] and they have gradually integrated the Transformers model [34] as is the case for GPT3 [29] or BERT [10], for example. Compared to RNNs, Transformers can be parallelized, which

opens the way, on one hand, to the use of ever-increasing training corpus (for example, GPT3 is trained on 45TB of data - almost the entire public web), and on the other hand, to the design of increasingly complex architectures (e.g., BERT large comprises 345 million parameters, BERT base 110 million). In a nutshell, Transformers [34] are founded on three key innovations: positional encoding, scaled dot product attention, and multi-head attention (we will come back to these elements in more detail in Section 2). As a result of a combination of all these elements, Transformers can learn an internal understanding of language automatically from the data. Despite their good performance on many different tasks, the use of these models in so-called sensitive applications or areas raised concerns over the past couple of years. Indeed, when decisions have an impact on individuals, for example in the medical and legal domains [9] or human resources [20], it becomes crucial to study the fairness of these models.

The core definition of fairness is still a hotly debated topic in the scientific community. In our work, we adopt the following commonly accepted definition [26]: *fairness refers to the absence of any prejudice or favoritism towards an individual or a group based on their intrinsic or acquired traits*. In machine learning, we assume that *unfairness* is the result of biased predictions (prejudice or favoritism), which are defined as elements that conduct a model to treat groups of individuals conditionally on some particular *protected* attributes, such as gender, race, or sexual orientation.

As an example, in human resources, the NLP-based recruitment task consists in analyzing and then selecting the relevant candidates. A lack of diversity inherent to the data, for instance, a corpus containing a large majority of male profiles (i.e. sample bias), will cause the model to maintain and accentuate a gender bias [33]. When handling simple linear models trained on reasonable size corpora, creating safeguards to avoid this type of bias is conceivable. With PLM, the characteristics that allow them to perform so well are numerous: the size of their training corpus, the number of parameters, and their ability to infer a fine-grained semantic from the data. However, they are also what make it difficult to prevent them from encoding societal biases [2].

Related Works Several recent studies highlight fairness issues raised by models based on the Transformer architecture. These issues are observed in different levels of the NLP pipeline: text encoding [1, 23], during the fine-tuning process [8], or simply as the potential harm caused on downstream tasks [23], with dedicated studies on language generation [32], document classification [3], toxicity detection, and sentiment analysis [19]. Besides measuring the fairness issue, *locating* the neural mechanism responsible for these issues is largely understudied and unsolved – locating such mechanisms would unlock the possibility for countermeasures in neural architectures. At the same time, a segment of research focused on compressing these large pre-trained models to attain similar performances with fewer parameters, so that running these models is more sustainable and more cost-effective. Several model compression techniques have been proposed, as discussed in [12] and namely the following compression families : pruning, quantization and distillation. The primer [24] increases the speed and general-

ization capacities by removing the less important model’s weights with regard to the task, while quantization approximates the model’s weights to reduce its complexity (e.g. reducing the numerical precision of the weights [22]). Finally, distillation [16] consists in training a smaller model (called *student model*) to mimic the predictions of the large PLM to distill (*teacher model*). In the present work, we focus on this latter, approach. One of the earliest model, DistilBERT [31] is able to reduce the number of parameters of BERT by 40% while maintaining 96% of accuracy in document classification. Looking at the impact of model distillation through fairness lenses has started to be investigated, mainly in the context of computer vision [18, 17, 25]. To summarize, their findings: i) compressed models impact underrepresented visual features directly related to bias, and ii) distilled models tend to accentuate discrimination already made by the teacher model. In NLP, fewer works have been conducted, and the conclusions are sometimes contradictory. While some works have shown that distilled versions of PLMs can exacerbate bias [28, 7], other articles seem to reach an opposite conclusion [35]; in this latter, authors state that model distillation acts as a regularization technique allowing bias reduction.

Contribution Based on existing results, we start from the postulate that PLMs, and more specifically BERT, encode undesirable bias. With a focus on the task of document classification on the Bias in Bio dataset, our objective is to identify the inner structure of the neural network architecture that produce bias, both for BERT and its distilled version DistilBERT. To this end, we design and conduct a series of experiments to verify the relation between models’ fairness and their intermediate representation or the attention they carry to the embedding in different data balance setting.

Organisation Section 2 provides background knowledge about BERT and DistilBERT. Section 3 presents the empirical protocols that we design. Section 4 details the technical setting and shows the obtained results of our experiments. Finally, we conclude in Section 5 and provide several perspectives unlocked by our experiments.

2 Preliminaries and Background

We study two PLM, BERT [10] and its distilled version, DistilBERT [31]. BERT is a general-purpose language model trained on masked language modeling task³. A small fraction of the words of each training document is masked, and the model is trained to reconstruct those masked words on a large amount of textual data. More precisely, the encoder part of a Transformer architecture takes a sentence, or short document, as input, and maps each token (word or subwords) to an initial representation space in \mathbb{R}^{768} .

³ The model is also trained on a next sentence prediction task, but that is irrelevant in our work and therefore not presented here.

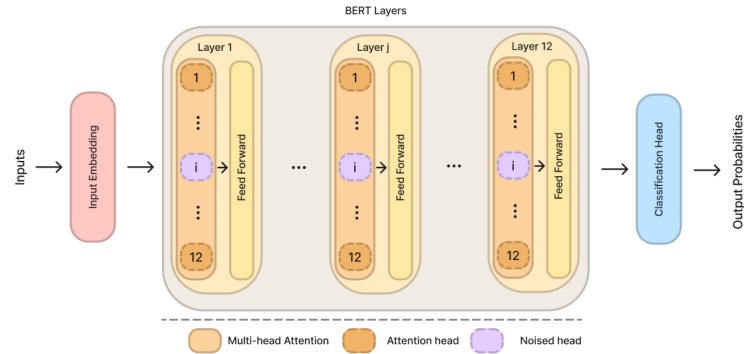


Fig. 1: Scheme of head's noising experiments

The encoder then *contextualizes* these representations using a multi-head attention mechanism. The attention mechanism builds attention weights between each pair of words, based on their similarity in a latent space. These are then used to build new word representations by a simple weighted average operation. More precisely, this attention operation is computed for slices of dimensions in parallel before concatenation, this is why it is called multi-head. This is followed by a pass through a feed-forward neural network with residual connections. See Figure 1 for an illustration of this operation. The output of the Transformer encoder is a matrix of size $L \times 768$, where L is the maximum document length of the model (512 for BERT) and 768 the number of dimensions. After pre-training on this task, the model can be *fine-tuned* on a downstream task, such as classification, by adding a linear layer on top of this model, that either inputs the final hidden state's "CLS" token (a special token corresponding to a representation of the sequence) or by pooling the representations of the words.

BERT, in its base version, has 110 M of parameters, 12 layers, and 12 attention heads. DistilBERT is a shallow version of BERT, trained with half the number of layers using distillation [16]. The principle of this compression method is to train a student model to replicate the behavior of the teacher. To do so, one feeds a dataset to the teacher to retrieve its predictions for each sample (the outputs of the teacher are soft targets, i.e. the probabilities over each class instead of the predicted label). On the other hand, the student receives the same input as training data and the predicted soft targets as training labels. The objective of the student is then to match the (soft) predictions of the teacher. The data used to train the student can either be unlabeled or labeled; in the second case, the true labels can also be fed to the student and a regularization term is added to the objective function to improve the student's performance. Using this approach, DistilBERT obtains up to 95% of the performance of BERT.

In our experiments, we use the pre-trained models from Hugging Face⁴, more specifically, the Transformer models for sequence classification with a linear layer as classification head on top of the pooled output.

3 A gender-bias neural exploration protocol for language models

3.1 Fine-tuning Scenarios for Fairness Evaluation

In many real-life datasets, we observe gender ratio imbalance, making ML models outcomes prone to unfairness. For example, in the Bias in Bios dataset [6], more than 90% of nurses are women, while they are less than 2% to be a surgeon. In [11], authors show that training with imbalanced data allows the model to learn a correlation between the target label and sensitive attributes, therefore inducing bias. These observations are unsurprising: this kind of bias is considered to be extrinsic, and is caused by the data used during the fine-tuning. However, with pre-trained models, another type of bias emerges: intrinsic bias. They are encoded during the pre-training and are out of the control of the practitioner. The understanding of the neural behavior that leads to them remains unclear. In our work, we are primarily interested in understanding and exploring the inner operations of the Transformer architecture that are at stake in these findings.

To study in detail the effect of those biases, we fine-tune both BERT and DistilBERT on two sub-samples of the initial dataset: a balanced and an imbalanced one with regard to the sensitive groups (class imbalance remains identical for both datasets). These models will be referred to as **Mi** and **Mb** when fine-tuned with the imbalanced or the balanced dataset respectively. We believe that starting from the same models, but with different fine-tuning strategies will make it possible to make comparisons of the fairness of these models.

3.2 Attention and Hidden States comparison (E1)

We first investigate the inner differences, induced by the fine-tuning process, between **Mi** and **Mb** models. (cf. section 3.1). More precisely, we focus on the attention weights and the hidden states (the intermediate representations) between both models for similar input. For this first set of experiments, we propose to investigate fairness through the lens of the learning dynamic of the PLMs. Recent works [30] show that the first layers of deep architectures capture low-level information about the input data, and that the learned representations tend to become more abstract and finer when moving through the body of the network towards its heads. In [13], the authors specifically studied the dynamic of BERT fine-tuning and conclude that mainly the last layers are significantly changing, both their attention mode and the hidden representation that they produce. Our objective is then to verify the two following hypotheses.

⁴ BERT: https://huggingface.co/docs/Transformers/model_doc/bert,
DistilBERT: https://huggingface.co/docs/Transformers/model_doc/distilbert

Hypothesis 1 *As layers specialize on the different granularity of the textual content, from grammatical to semantic aspects, we assume that monitoring the attention scores and hidden states of the successive layers of the models allows determining which one(s) is encoding bias.*

Hypothesis 2 *The distillation process implies that the student model will reproduce the behavior of the teacher, including biases in predictions. We assume that by reducing the depth, and hence the expressive power of the model, compression encourages amplified bias in the hidden representations.*

To proceed, we adapt the protocol of [13] to verify both Hypothesis 1 and 2. We first take a look at the modification of the similarity between tokens, where the attention is computed as a function of the similarity. In a second step, we look at how much hidden word representations are impacted by the model independently of the pairwise similarities.

Attention values comparison. The Jensen-Shannon divergence is a symmetrized version of the Kullback-Leibler divergence. It allows comparing two probability distributions P and Q . We propose to compare the attention of the two models layer-wise. Formally, we evaluate the divergence for each sample and for each head between the layers of two models (e.g., \mathbf{Mi} and \mathbf{Mb}). Let N be the number of examples in the evaluation set, H the number of attentions head, and W the number of tokens in a sequence. $A_i^h(token_t)$ and $A_b^h(token_t)$ are the attention scores for $token_t$ on head h respectively for models \mathbf{Mi} and \mathbf{Mb} . In this context, the JS divergence is defined as follows:

$$D_{JS}(\mathbf{Mi}||\mathbf{Mb}) = \frac{1}{N} \frac{1}{H} \sum_{n=1}^N \sum_{h=1}^H \frac{1}{W} \sum_{t=1}^W D_{JS}(A_i^h(token_t)||A_b^h(token_t)) \quad (1)$$

$D_{JS}(\cdot||\cdot) \in [0, 1]$, where 0 indicates that the distributions are identical.

Hidden states comparison. We compute the Singular Vector Canonical Correlation Analysis distance (SVCCA) [30] to observe the evolution of hidden states. SVCCA allows analyzing and comparing representations in deep learning models; in our case, the hidden representations produced by each Transformer layer. When computing SVCCA, we first perform a Singular Value Decomposition (SVD) of the representations produced by the two models for each input observation. Then, we compute the Canonical Correlation Analysis (CCA) [14] between the two subspaces created by the SVD to evaluate the correlation between the two representations and finally, condense the correlations obtained for each dimension into a distance.

Let c be the hidden size of the model and $\rho \in [0, 1]$ the CCA. The SVCCA distance is defined as follows:

$$D_{SVCCA}(\mathbf{Mi}||\mathbf{Mb}) = 1 - \frac{1}{c} \sum_{j=1}^c \rho^{(j)} \quad (2)$$

$D_{SVCCA}(\cdot||\cdot) \in [0, 1]$ with 0 meaning identical representations.

3.3 Head’s ablation (E2)

With multi-head attention, Transformers build for each head a different representation of the input embedding. We make the hypothesis that some representations might induce more biases than others. Thus, complementary to the previous experiments and in continuity with the goal of finding where are biases encoded in PLM, we successively ablate heads of the model to infer if some of them are responsible for biases in the model. The ablation is done by setting all its attention weights to 0 through all the layers, as shown in Figure 1.

Hypothesis 3 *By ablating attention heads, we aim at removing the bias due to a given head and identify the ones contributing to unfairness. In other words, we expect that when ablating a head responsible for bias, the model will obtain a better fairness score and reciprocally.*

In practice, we first fine-tune the model so that it learns the weights as in real-world applications. Then, we successively ablate heads and evaluate the performance and fairness of the model on new data to evaluate the bias encoded by the aforementioned heads. We are interested in the results of our models following two criteria: their predicting performance and their fairness.

Performance. The model performance is evaluated using a weighted version of the F-Score (since our target variable is multivalued).

More precisely, we compute the F-Score for each class, then compute the average weighted by the number of samples per class.

Gender fairness. We are interested in group fairness, and several metrics have been proposed in the literature [4]. We choose the commonly used Equalized Odds (EO) [15], defined as follows $\mathbb{P}(\bar{y} = 1|y, S = 0) = \mathbb{P}(\bar{y} = 1|y, S = 1)$. To ease the interpretation, we compute the difference version of EO given by

$$EO = |\mathbb{P}(\bar{y} = 1|y, S = 0) - \mathbb{P}(\bar{y} = 1|y, S = 1)|, \quad (3)$$

where \bar{y} are the predictions, $y \in 0, 1$ are the true labels and S corresponds to the sensitive attributes (0 and 1 indicating the belonging to a sensitive group). $EO \in [0, 1]$ where a score closer to 0 indicates fairer predictions.

4 Experiments

4.1 Task and Dataset

In our experiments, we focus on a classification task and use a subset of the Bias in Bios dataset [6] called the Curriculum Vitae dataset⁵. It contains a set of short biographies associated with an occupation and a gender. The dataset is composed of 217,197 entries, and 28 professional occupations. The distribution

⁵ Dataset: <https://www.kaggle.com/competitions/defi-ia-insa-toulouse/data>

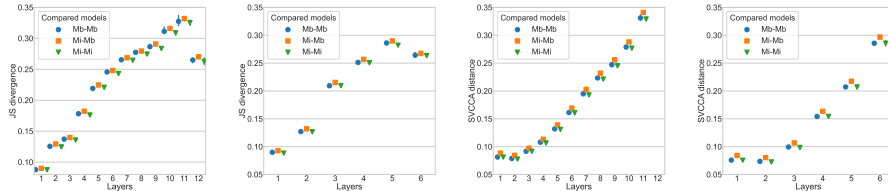


Fig. 2: JS divergence comparison for BERT (left) and DistilBERT (center-left) and SVCCA distance for BERT (center-right) and DistilBERT (right).

of classes (occupations) and groups (genders) within each occupation is highly imbalanced. For example, class 19 corresponding to ‘professor’ represents 32.23% of the dataset and within the class, women represent 44.88% of the entries; class 23 corresponding to ‘paralegal’ represents 0.44% of the dataset, and women 84.17% of the class entries.

We create two versions of this dataset, a balanced version, and an imbalanced version, depending on the gender attribute. The former one is a subset, where for each class the largest sensitive group is truncated to equalize the proportion of individuals of each gender. The latter is a subset, where we reproduce the imbalance between gender observed in the initial dataset, but both groups are truncated to ensure that the number of samples in both subsets are equal.

Based on these two versions, we exploit the relationship between fairness and gender imbalance (cf. section 3.1) to build two models **Mi** and **Mb** to further explore the mechanisms of bias. We evaluate the EO of BERT and DistilBERT fine-tuned on 70% of the samples, for both versions of the original dataset. In the sequel, we refer to these models as **BERT_{Mb}** and **DistilBERT_{Mb}** for the balanced versions, and **BERT_{Mi}** and **DistilBERT_{Mi}** for the imbalanced ones. To confirm our premise, we perform classification and observe an average EO over all classes three times higher for the imbalanced versions (0.13 vs. 0.42). Following this first experience, one might think that balancing the fine-tuning data is a sufficient and satisfactory solution to ensure fairness. However, before proceeding further, two remarks are in order. First, balancing the data is a first step in reducing bias, but it does not guarantee a fair model (EO above 0). Second, in many real-world scenarios, where multiple protected attributes can be observed simultaneously (e.g. women of color), this solution appears to be shortsighted as one cannot slice the data into more sub-population infinitely to rebalance classes.

4.2 Attention and Hidden States comparison (E1)

We present the results of this experiment averaged over five random seeds in Figure 2.

If we observe higher divergence between models **Mi** and **Mb** than between **Mi** and **Mi** or **Mb** and **Mb** on a given layer, we can assume this layer to be responsible for encoding bias. For the JS divergence, first we can note a similar pattern for

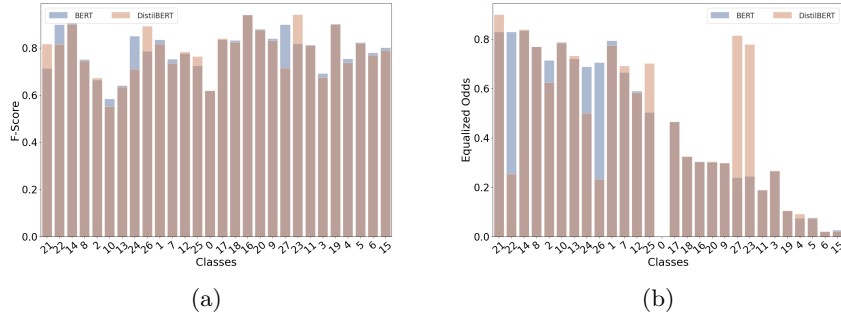


Fig. 3: (a) F-Score and (b) Equalized Odds for model M_i per class, without ablation and ordered by ascending ratio W/M .

both BERT and DistilBERT: the divergence increases as we move forward in the architecture, with a peak on the penultimate layer. For the SVCCA distance, the trend is similar, and we reach the highest value on the last layers. These findings are perfectly in line with the results of [13] and [27] claiming that fine-tuning mainly affects top layers. These first observations allow us to state that DistilBERT follows the same learning dynamic than BERT during fine-tuning. Now, comparing $[M_i-M_i, M_b-M_b]$ vs. M_i-M_b , we observe that the values for both metrics are slightly above on M_i-M_b , but not significantly. In addition, this difference is consistent over all layers. Finally, we observe the exact same behavior for DistilBERT. These particular results are counterintuitive with our Hypotheses 1 and 2, and we cannot conclude that extrinsic bias makes some layers different with regard to internal representations and attention scores for both architectures.

4.3 Head’s ablation (E2)

Let us now look at the relationship between the heads and bias. Once the model is fine-tuned, we neutralize the heads in turn at inference time to estimate the bias they encode. Since the distribution of the classes is highly imbalanced, we evaluate the fairness of the model for each class individually.

Figure 3 shows the EO and F-Score for each class on the original model (M_i) (classes are sorted by ascending ratio women/men). Three observations can be drawn: firstly, for the highly imbalanced class (left-learning) EO is significantly higher than average; secondly, comparing BERT and DistilBERT, we see that for the majority of classes, we obtain an equivalent level of fairness, except for a few classes, where either one or the other is more biased. However, for the most imbalanced classes, DistilBERT is reaching a better level of fairness in comparison with BERT (lower or equal EO scores); thirdly, both architectures obtain comparable F-Score.

Now, we reproduce this evaluation twelve times (one time for each ablation), and observe different levels of variations for the EO depending on both the class

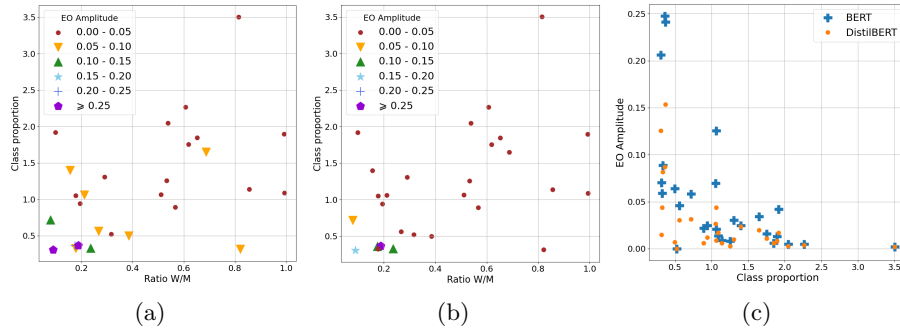


Fig. 4: EO amplitude for BERT’s Mi (a), DistilBERT’s Mi (b), BERT’s and DistilBERT’s Mb (c).

and the head that is ablated. This implies that the representations produced by the attention head are different enough to impact the fairness of the models as assumed in Hypothesis 3. However, when the network is re-fine tuned we do not observe the same variations for a given head. For each class, we compute the difference of EO obtained when neutralizing a head gives the fairest score vs. the most unfair score and call it the amplitude. Let $EO_{class} = \{EO_{class}^1, \dots, EO_{class}^{12}\}$ with EO_{class}^{head} the EO computed for a given class after noising a head. We defined the amplitude for a class as :

$$\text{amplitude}_{class} = \max_{head}(EO_{class}) - \min_{head}(EO_{class}) \quad (4)$$

Figure 4 shows the amplitude depending on the class imbalance and the ratio of women/men within the classes. The minority group is not the same for every class, thus, we compute the ratio as follows $\text{ratio W/M} = \min(\frac{\%women}{\%men}, \frac{\%men}{\%women}) \in [0, 1]$. For better readability, we rescale the class proportion and amplitude vectors by taking $\log(\text{vector} + 1_n)$, n being the dimension of the vectors. We note, in Figures 4a and 4b that the more a class suffers from double imbalance (under-represented class and highly sensitive group imbalance) the more the heads will produce different representations, some being more biased than others. Also, when comparing the values of EO amplitude for BERT and DistilBERT, we observe that BERT is more sensitive to those scenarios than DistilBERT. According to Figure 4c, where sensitive groups are balanced within the classes, the less a class is represented the higher the amplitude is, meaning that **BERT is generally more sensitive to class imbalance than DistilBERT with regard to the homogeneity of head representations**. On the other hand, we evaluate the correlation between F-Score and EO, when ablating each head, and have not been able to establish a relation caused by the process.

5 Conclusion

This paper investigated the implication of inner elements of BERT-based models’ architecture in bias encoding through empirical experiments on the Transformers’ layers and attention heads. We also studied the attention carried by the ”CLS” token to the words of the sequence, specifically the pronouns ’he’ and ’she’, but also to the ones receiving the most attention from the aforesaid token, in an attempt to understand what the model was focusing on; this study did not lead us to any convincing results. Similarly, investigating the JS divergence and SVCCA distance between different layers (e.g. the 1 and 2) was not conclusive, we suspect it might since layers specialize on different aspects of the input text [21]. To summarize, we show that gender bias is not encoded in a specific layer or head. We also demonstrate that the distilled version of BERT, DistilBERT, is more robust to double imbalance of classes and sensitive groups than the original model. Even more specifically, we observe that the representations generated by the attention heads in such a context are more homogeneous for DistilBERT than for BERT in which some attention heads will be fair while others are very unfair. Thus, we advise giving special care to such patterns in the data but do not recommend ablating the heads producing more unfair representations since it could seriously harm the performance of the model. Finally, we recommend DistilBERT to the practitioner using datasets containing underrepresented classes with a high imbalance between sensitive groups, while cautiously evaluating class independently, using the protocol that we propose in this paper.

Acknowledgment

This work is funded by the french National Agency for Research (ANR) in the context of the Diké project (ANR-21-CE23-0026).

References

1. Basta, C., Costa-jussà, M.R., Casas, N.: Evaluating the underlying gender bias in contextualized word embeddings. In: the First Workshop on Gender Bias in NLP. pp. 33–39. ACL (2019)
2. Bender, E.M., Gebru, T., McMillan-Major, A., Shmitchell, S.: On the dangers of stochastic parrots: Can language models be too big? In: FAccT. pp. 610–623 (2021)
3. Bhardwaj, R., Majumder, N., Poria, S.: Investigating gender bias in bert. Cognitive Computation (2020)
4. Caton, S., Haas, C.: Fairness in machine learning: A survey. arXiv preprint arXiv:2010.04053 (2020)
5. Dai, A.M., Le, Q.V.: Semi-supervised sequence learning. In: NeurIPS. p. 3079–3087 (2015)
6. De-Arteaga, M., Romanov, A., Wallach, H., Chayes, J., Borgs, C., Chouldechova, A., Geyik, S., Kenthapadi, K., Kalai, A.T.: Bias in bios: A case study of semantic representation bias in a high-stakes setting. In: FaccT. pp. 120–128 (2019)

7. Delobelle, P., Berendt, B.: Fairdistillation: Mitigating stereotyping in language models. *ECML-PKDD* (2022)
8. Delobelle, P., Tokpo, E., Calders, T., Berendt, B.: Measuring fairness with biased rulers: A comparative study on bias metrics for pre-trained language models. In: *Proceedings of NAACL*. pp. 1693–1706 (2022)
9. Demner-Fushman, D., Chapman, W.W., McDonald, C.J.: What can natural language processing do for clinical decision support? *Journal of biomedical informatics* **42**(5), 760–772 (2009)
10. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of NAACL-HLT*. pp. 4171–4186 (2019)
11. Dixon, L., Li, J., Sorensen, J., Thain, N., Vasserman, L.: Measuring and mitigating unintended bias in text classification. In: *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. pp. 67–73 (2018)
12. Gupta, M., Agrawal, P.: Compression of deep learning models for text: A survey. *ACM Trans. Knowl. Discov. Data* **16**(4) (2022)
13. Hao, Y., Dong, L., Wei, F., Xu, K.: Investigating learning dynamics of BERT fine-tuning. In: *Proceedings of the Conference of the Asia-Pacific Chapter of the ACL and the IJCNLP*. Suzhou, China (2020)
14. Hardoon, D.R., Szedmak, S., Shawe-Taylor, J.: Canonical correlation analysis: An overview with application to learning methods. *Neural computation* **16**(12), 2639–2664 (2004)
15. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. *NeurIPS* **29** (2016)
16. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. *arxiv pre-print* (2015)
17. Hooker, S., Courville, A., Clark, G., Dauphin, Y., Frome, A.: What do compressed deep neural networks forget? *arxiv pre-print* (2021)
18. Hooker, S., Moorosi, N., Clark, G., Bengio, S., Denton, E.: Characterising bias in compressed models. *arXiv preprint arXiv:2010.03058* (2020)
19. Hutchinson, B., Prabhakaran, V., Denton, E., Webster, K., Zhong, Y., Denuyl, S.: Social biases in NLP models as barriers for persons with disabilities. In: *ACL*. pp. 5491–5501 (2020)
20. Jatobá, M., Santos, J., Gutierriz, I., Moscon, D., Fernandes, P.O., Teixeira, J.P.: Evolution of artificial intelligence research in human resources. *Procedia Computer Science* **164**, 137–142 (2019)
21. Jawahar, G., Sagot, B., Seddah, D.: What does bert learn about the structure of language? In: *ACL* (2019)
22. Kim, S., Gholami, A., Yao, Z., Mahoney, M.W., Keutzer, K.: I-bert: Integer-only bert quantization (2021)
23. Kurita, K., Vyas, N., Pareek, A., Black, A.W., Tsvetkov, Y.: Measuring bias in contextualized word representations. In: *the First Workshop on Gender Bias in Natural Language Processing*. pp. 166–172 (2019)
24. LeCun, Y., Denker, J., Solla, S.: Optimal brain damage. *Advances in neural information processing systems* **2** (1989)
25. Lukasik, M., Bhojanapalli, S., Menon, A.K., Kumar, S.: Teacher’s pet: understanding and mitigating biases in distillation. *arXiv preprint arXiv:2106.10494* (2021)
26. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. *ACM Computing Surveys* **54**(6), 1–35 (2021)

27. Merchant, A., Rahimtoroghi, E., Pavlick, E., Tenney, I.: What happens to BERT embeddings during fine-tuning? In: Proceedings of the Third BlackboxNLP Workshop. pp. 33–44. ACL (2020)
28. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training. OpenAI blog (2018)
29. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. OpenAI blog **1**(8), 9 (2019)
30. Raghu, M., Gilmer, J., Yosinski, J., Sohl-Dickstein, J.: Svcca: Singular vector canonical correlation analysis for deep learning dynamics and interpretability. NeurIPS **30** (2017)
31. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arxiv pre-print (2020)
32. Sheng, E., Chang, K., Natarajan, P., Peng, N.: The woman worked as a babysitter: On biases in language generation. In: EMNLP/IJCNLP. pp. 3405–3410 (2019)
33. Swinger, N., De-Arteaga, M., Heffernan IV, N.T., Leiserson, M.D., Kalai, A.T.: What are the biases in my word embedding? In: Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society. pp. 305–311 (2019)
34. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. NeurIPS **30** (2017)
35. Xu, G., Hu, Q.: Can model compression improve nlp fairness. arXiv preprint arXiv:2201.08542 (2022)