

MARKUS BAYER, MARC-ANDRÉ KAUFHOLD, and CHRISTIAN REUTER,

PEASEC*, Technical University of Darmstadt

Data augmentation, the artificial creation of training data for machine learning by transformations, is a widely studied research field across machine learning disciplines. While it is useful for increasing a model's generalization capabilities, it can also address many other challenges and problems, from overcoming a limited amount of training data to regularizing the objective, to limiting the amount of data used to protect privacy. Based on a precise description of the goals and applications of data augmentation and a taxonomy for existing works, this survey is concerned with data augmentation methods for textual classification and aims at providing a concise and comprehensive overview for researchers and practitioners. Derived from the taxonomy, we divide more than 100 methods into 12 different groupings and give state-of-the-art references expounding which methods are highly promising by relating them to each other. Finally, research perspectives that may constitute a building block for future work are provided.

CCS Concepts: • Computing methodologies \rightarrow Supervised learning by classification; Natural language processing; *Neural networks*; Adversarial learning; Regularization;

Additional Key Words and Phrases: Data augmentation, low data regimes, small data analytics

ACM Reference format:

Markus Bayer, Marc-André Kaufhold, and Christian Reuter. 2022. A Survey on Data Augmentation for Text Classification. *ACM Comput. Surv.* 55, 7, Article 146 (December 2022), 39 pages. https://doi.org/10.1145/3544558

1 INTRODUCTION

An increase in training data does not necessarily result in a solution for the learning problem. Nevertheless, the quantity of data remains decisive for the quality of a supervised classifier. Originating from the field of computer vision, many different methods to artificially create such data exist, which are referred to as *data augmentation*. For images, transformations such as rotations or changes of the RGB channel are useful, as the resulting model should be invariant for these.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

https://doi.org/10.1145/3544558

^{*}Science and Technology for Peace and Security | Pankratiusstraße 2, 64289 Darmstadt, Germany.

This work has been co-funded by the German Federal Ministry of Education and Research (BMBF) in the project CYWARN (13N15407), by the LOEWE initiative (Hesse, Germany) within the emergenCITY center, as well as the German Federal Ministry of Education and Research and the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity ATHENE. The calculations for this research were conducted on the Lichtenberg high performance computer of the TU Darmstadt.

Authors' address: M. Bayer, M.-A. Kaufhold, and C. Reuter, PEASEC, Technical University of Darmstadt, Pankratiusstraße 2, 64289 Darmstadt, Germany; emails: {bayer, kaufhold, reuter}@peasec.tu-darmstadt.de.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

^{0360-0300/2022/12-}ART146 \$15.00

Similar to computer vision, speech recognition uses procedures that change sound or speed. In contrast, research on data augmentation in Natural Language Processing (NLP) faces the difficult task of establishing such universal rules for textual data transformations which, when executed automatically, maintain labeling quality [1, 2]. Research in this area was therefore much more limited before 2019, despite existing extensive areas of application [3]. Nowadays, this challenge remains but is being addressed by many scientists from different research fields as more possibilities and complex mechanisms open up. Within these fields, researchers strive to meet various goals, e.g., generating more data for low-data regimes, balancing imbalanced dataset classes, or securing against adversarial examples. Thus, textual data augmentation comes in many contrasting forms that will be grouped and explained in this survey. We will provide in-depth analysis and also relate the methods to the state-of-the-art, as they now face another challenge due to the advent of transfer learning. For example, Longpre et al. [4] demonstrate that many data augmentation methods cannot achieve gains when using large pre-trained language models, as they already are invariant to various transformations. They hypothesize that data augmentation methods can only be beneficial, if they create new linguistic patterns that have not been seen before. Keeping this in mind, the survey is closed with a meta-perspective on the methods. This survey is therefore intended to contribute to data augmentation and general text classification by highlighting the following aspects:

- Goals and applications (C1). We highlight the goals and applications of data augmentation that we derive from the comprehensive review. These have only been presented to a limited and incomplete extent in previous research articles.
- Comprehensive survey on data augmentation in text classification (C2). Our survey provides a holistic overview of the data augmentation field in text classification. While methods for other NLP disciplines are mentioned, the listing is not complete, nor are the methods set in relation to each other as the text classification data augmentation methods are.
- Data-structure-driven taxonomy and method-oriented categorization (C3). The text classification data augmentation methods are clustered according to a data-structure-based, high-level taxonomy and then subdivided into more fine-grained method groups. This is also present in the surveys from Shorten and Khoshgoftaar [5] and Wen et al. [6] and is adapted for the text classification domain.
- Method-driven overview and in-depth details (C4). The textual data augmentation methods are explained clearly and concisely while including necessary details for delimitation and comparison. Contrasting to other works, our extensive study contains 12 groups with more than 100 different approaches.
- State-of-the-art review (C5): Within the literature survey we examine the latest state-of-theart considerations, for example, the limited benefit of textual data augmentation methods with large pre-trained models that are often neglected in current works.
- Relating methods (C6). Throughout this survey, the methods are set in relation to conception
 and performance comparisons, while taking the underlying models and application contexts
 into account.
- -Future research perspectives (C7). We identify future research opportunities that are either necessary for a state-of-the-art comparison or sensible to look into because of current challenges and promising directions for textual data augmentation.

The survey article is structured as follows: The article introduces the foundations of data augmentation in Section 2. This section is then broadened by the consideration of the goals and applications. Section 3 is subdivided into the various data augmentation groups and contains the explanations, as well as tabular overviews of the methods. In Section 4, an analysis of the data augmentation methods from a more global perspective is given and various future research directions are discussed. Section 5 outlines the limitations of data augmentation and provides a conclusion for this survey.

2 BACKGROUND: FOUNDATIONS, GOALS, AND APPLICATIONS OF DATA AUGMENTATION

In many machine learning scenarios, not enough data are available to train a high-quality classifier. To address this problem, data augmentation can be used. It artificially enlarges the amount of available training data by means of transformations [7]. In the well-known LeNet by LeCun et al. [8], early versions of data augmentation have already been observed. The notion of data augmentation comprises various research in different sub-areas of machine learning. Many scientific works merely relate data augmentation to deep learning, yet it is frequently applied in the entire context of machine learning. Therefore, this article adopts the notion of data augmentation as a broad concept, encompassing any method that enables the transformation of training data. However, following common understanding in research, semi-supervised learning is not regarded as a form of data augmentation and is only thematized if sensible in this survey.

An important term relating to data augmentation is label preservation, which describes transformations of training data that preserve class information [9]. For example, in sentiment analysis, an entity replacement within a sentence is often sufficient for label preservation, but randomly adding words may alter the sentiment (e.g., an additional "not" could invert the meaning of a sentence). In many research works, label preservation is adapted to also cover transformations changing the class information, if the label is adjusted correctly. Additionally, many transformations do not maintain the correct class in every case, but with a high probability. Shorten and Khoshgoftaar [5] define this probability as the safety of a data augmentation method. When this uncertainty is known, it could be directly integrated in the label. Otherwise, methods like label smoothing [10] can model a general uncertainty.

The goals of data augmentation are manifold and encompass different aspects. As mentioned above, training data are essential for the quality of a supervised machine learning process. Banko and Brill [11] show that only the creation of additional data can improve the quality of a solution in the confusion set disambiguation problem, while the choice of the classifier does not lead to a significant change. The model selection and development will remain a crucial aspect of machine learning. Yet, scholars suggest that in some situations, the choice for higher investments in algorithm-choice and -development instead of corpus-development should be carefully considered [11]. Closely connected to this is the big data wall problem, which Coulombe [9] mentions in his work on data augmentation. It describes that big companies benefit from the special advantage of having access to a large amount of training data. Consequently, the already large GAFAM-Companies (Google/Alphabet, Amazon, Facebook/Meta, Apple, and Microsoft) expand their predominance over smaller businesses due to their data superiority. An ideal data augmentation method could approach these points and decrease the dependency on training data even though full elimination is not likely.

Additionally, creating training data for various classification problems is accompanied by high labeling costs. In many instances, assessment and labeling by experts are necessary to prevent incorrect training examples. These aspects can, for example, be especially stressed concerning the field of crisis informatics [12, 13]. Creating relevant classifiers for emergency services and responders is only possible during crises and requires resources and time from personnel needed elsewhere to, e.g., act as first responders, therefore in the worst-case costing lives [14]. Similarly, training data for medical image processing is very valuable. Due to the rareness of certain diseases, the privacy of patients, and the requirement of medical experts, it is particularly challenging to

Original text	Altered text
South Africa's historic Soweto township	South Africa's historic Soweto township
marks its 100th birthday on Tuesday in a	marks its 100th birthday on Tuesday in a
mood of optimism.	moo P of optimism.
57% World	95% Sci/Tech
Chancellor Gordon Brown has sought to	Chancellor Gordon Brown has sought to
quell speculation over who should run the	quell speculation over who should run
Labour Party and turned the attack on the	the Labour Party and turned the attack
opposition Conservatives.	on the o B position Conservatives.
75% World	94% Business

Table 1. Examples for Adversarial Attacks Adapted from Ebrahimi et al. [21]

provide medical datasets [5]. In a related sense, many domains, such as cybersecurity, have a timecritical factor that requires training data to be collected as quickly as possible so that, in terms of the cybersecurity domain example, threats can be responded to quickly [15, 16]. With regard to such classification problems, data augmentation could help minimize the required amount of data needed to be labeled and to solve interlinked problems.

Data augmentation is particularly significant for the field of deep learning. Work such as that by Minaee et al. [17] has already extensively investigated the quality of deep learning algorithms in text classification, but there are many application scenarios where there is not enough data to produce high-quality classifiers. For example, Srivastava et al. [18] have also demonstrated that deep neural networks in general are particularly powerful but encompass a tendency to overfit; faced with unseen instances, they might generalize badly. This observation can be illustrated with help of the bias-variance dilemma. On the one hand, deep learning algorithms are, due to their deep and non-linear layers, very strong models with a lower biaserror. On the other hand, they show a high variance for different subsets of training data [19]. This problem can be solved by arranging the algorithm to prefer simple solutions or by providing a bigger amount of training data. The first option is aimed at methods of regularization, such as dropout or the addition of an L2 norm via the model's parameters in the loss-function. The second option is frequently realized by means of data augmentation and could, in this context, also be considered as a type of regularization. According to Hernandez-Garcia and König [20], data augmentation is a preferred regularization method, as it achieves generalization without degrading the models' representational capacity and without re-tuning other hyperparameters. While other methods reduce the bias error, data augmentation's objective is to keep it constant and is used to solve the problem at the root [5]. Nonetheless, data augmentation still depends on the underlying classification problem and can therefore not be effectively applied in all circumstances.

In the context of deep learning models, so-called adversarial examples/attacks are generated more and more frequently. These small changes in the input data, which are almost unrecognizable to humans, mislead the algorithms to make wrong predictions [21]. Table 1 shows two different genuine examples in which the smallest changes in the texts alter the classification prediction. Alzantot et al. [22] further present an algorithm that generates semantically and syntactically similar instances of training data, successfully outwitting sentiment analysis and entailment models. With the help of adversarial training, these automatic adversarial example generators can be used as data augmentation methods, as done, for example, in [23, 7, 24], or [25], in order for models to be less susceptible to such easy alterations. If the amount of data are taken into consideration, it stands out that certain classification problems are often heavily unbalanced, for instance, only a small amount is relevant (positive) while the irrelevant (negative) data are prevalent [26]. For



Fig. 1. Taxonomy and grouping for different data augmentation methods.

example, in an entire corpus for topic classification or crisis identification, only a few data actually relate to the topics or the crisis in question. Zhong et al. [27] term a dataset as unbalanced, if the distribution of classes within it is not approximately equally balanced. Data augmentation may help to enhance the amount of data for a certain class in order for balanced class distributions to be present and thus for a classifier to be able to be modeled more robustly [28, 29].

Data augmentation can also be helpful in sensitive domains. Dealing with confidential or privacy-related data, one can decrease the usage of real-world data by crafting artificial data. It is even possible to only train the algorithm on the newly created data, in order to prevent drawing any conclusions on non-artificial training data from a deployed model. For example, Carlini et al. [30] have demonstrated a method for extracting training data from large language models that could contain private information. For training such a privacy ensuring model, special data augmentation techniques that are able to anonymize the data have to be used.

Data augmentation exists in different types and areas of application. A taxonomy of the types in the textual domain can be seen in Figure 1. The augmentation methods can be divided into the transformation of raw data (data space) and processed representations of data (feature space) [5]. These representations are transformed types of data, for example, activation vectors of a neuronal network, the encoding of an Encoder-Decoder Network, or LSTM hidden states, respectively embeddings of data. Abstracting from the textual realm, in many cases, data augmentation depends on the underlying problem (text classification, image recognition, and so on.); and is therefore applied in different ways in different areas. Procedures generic enough to be used across different areas are for the most part limited to the feature space.

The most substantial research on data augmentation exists in the field of computer vision. This is due to the intuitive construction of simple label-preserving transformations. Data augmentation methods in computer vision are, among other, geometric transformations [7, 24], neural style transfers [31–33], interpolation of images [34], random partial deletions [35], and **generative adversarial network (GAN)** data generation [21]. Sophisticated techniques can additively improve the accuracy baseline for different problems by around 10 to 15 percent [35]. Another area of application for data augmentation is speech processing. Researchers have successfully used acoustic transformations of the input data. Ko et al. [36] have achieved up to 4.3 points better accuracy values by modifying speed. Furthermore, interfering with vocal tract length [35] or adding noise [34] may also enhance the quality of classifiers. The application of data augmentation in the textual

realm is considered a difficult task, since textual transformations preserving the label are difficult to define [1, 2]. Nevertheless, several simple and sophisticated methods have been developed in this and adjacent research areas.

3 TEXTUAL DATA AUGMENTATION METHODS

In the following, different data augmentation methods for textual data are summarized, explained, and subdivided in different groupings. Mainly methods focusing on the application of text classification are included, although augmentation methods for other tasks in the textual realm are also mentioned if they fit into the group. In this survey, text classification is considered a problem in the field of NLP, where units such as sentences, paragraphs, or documents are categorized into class labels [17]. For example, generative or sequence-tagging tasks, where either text has to be generated or the words of the units have to be tagged individually, are not regarded as tasks in this sense. This means that augmentation methods for tasks such as topic classification, sentiment analysis, or spam identification are focused, described, and analyzed in detail. Other tasks, on the other hand, like generative question answering, part of speech tagging, or machine translation are only mentioned in a non-comprehensive way. Therefore, in the context of text classification, our article provides a comprehensive overview containing the necessary details for researchers and practitioners. For a more general perspective on NLP augmentations (including sequence-tagging, parsing, text generation, and so on.), we recommend the reader to have a look at the work of Feng et al. [37], which is not as detailed in text classification as our work but presents a broader task view. In contrast to this task-driven view by Feng et al. [37], we are taking a method-oriented perspective while conducting a data-structure-driven, high-level categorization (see Figure 1). Contrary to other surveys in the field of data augmentation, we focus on setting the augmentation methods into context by comparing the conception and performance, with regard to the underlying models and application context. In this way, the listed augmentation groups contain an explanation with details on the differences within the group and a comprehensive overview of how the methods differ and which results they produce. This allows the reader to gain insights into which data augmentation technique might be most promising for the own usecase and what specifics need to be considered, while it is also possible to follow the data-structure-based taxonomy. In the end, we discuss important future research directions by setting all methods into context, which can help accelerate developments in this field.

In the next section, data augmentation methods relevant to textual contexts are summarized and grouped. Generally, the methods are described in a sensible order for the specific group. In groups with many similar approaches, we summarize the most important information in tabular form. We also extract information regarding improvements. The improvement indications are intended to give a quick overview of how well a method may perform, but are not in-depth informative or comparable on their own. For a more detailed perspective, the models and datasets are also displayed. This should provide a more holistic perspective, although in-depth information has to be extracted from the respective articles themselves.

3.1 Data Space

Augmentation in the data space deals with the transformation of the input data in its raw form, i.e., in the case of this survey, into the readable textual form of the data.

3.1.1 Character Level.

3.1.1.1 Noise Induction. The addition of noise to the input data are one of the data augmentation methods with the smallest alterations, especially when applied on a character level. As explained

in more detail further on, the induction of noise can also be used at the word level as well as in the feature space.

In this context, the basic idea of the method of Belinkov and Bisk [38] is to add artificial and natural noise to the training data so that, in their case, neural machine translation (NMT) models are less susceptible to adversarial examples. As artificial noise, Belinkov and Bisk [38] describe operations like the random switching of single letters ("cheese" \rightarrow "cehese"), the randomization of the mid part of a word ("cheese" \rightarrow "ceehse"), the complete randomization of a word ("cheese" \rightarrow "eseehc"), and the replacement of one letter with a neighboring letter on the keyboard ("cheeae"). Similarly, Feng et al. [39] randomly delete, swap, and insert characters of texts (the prompt portion) that are used for fine-tuning text generators. For this, they moreover ignore the first and last character of a word. To measure the suitability for text generators, they intrinsically measure the diversity, fluency, semantic context preservation, and sentiment consistency. The applied method is better than the baseline in every respect. These augmentations are also usable in the text classification domain. Ebrahimi et al. [21] used an existing model, trained with the initial dataset, to generate adversarial examples. They used the direct input data to flip a letter if the change increased the loss of the existing model. If a new model is trained with the additional data once again, the error rate is improved and the success of adversarial attacks is significantly mitigated. Furthermore, they compared their approach with the adversarial method from the previously mentioned work of Belinkov and Bisk [38] and the feature space method from Miyato et al. [40] (see Section 3.2.1). Based on a CharCNN-LSTM on the AG News dataset, they achieve the best improvement in accuracy by obtaining an additional 0.62%. While the method of Miyato et al. [40] improved the score by only 0.24 points, it is interesting to see that the method of Belinkov and Bisk [38] even decreased the accuracy by 0.33 points. Coulombe [9] describes the induction of weak textual sounds through the aforementioned change, deletion, and addition of letters in words and, in addition, the alteration of upper and lower case and the modification of punctuation. The highest absolute accuracy improvement by 2.5% can be seen in comparison to the best functioning baseline. However, the evaluation was performed with basic architectures and no embeddings, wherefore further studies are needed to validate the usefulness in a current setting.

Natural noise, as defined by Belinkov and Bisk [38], covers spelling mistakes that are common in the respective language, using spelling mistake databases. Each word associated with a common mistake is replaced with the misspelled word, and if there is more than one, the mistake is randomly sampled. Belinkov and Bisk [38] receive varying BLEU scores with their artificial and natural noise methods; most noise operations made the model more robust against attacks with similar operations. Most importantly, natural noise almost consistently worsens a translation model regarding the baseline. Analogous to the natural noise defined by Belinkov and Bisk [38], Coulombe [9] also adds common spelling mistakes in the textual data and achieves good improvements when added to classifiers. The best baseline (XGBoost) was improved by an additional 1.5%. With such transformations, learners are better able to deal with spelling mistakes in prospective texts, even if mistakes are not present in the original training dataset. This variant of data augmentation can for example be of interest when dealing with texts originating from social networks.

3.1.1.2 Rule-based Transformations. Coulombe [9] implements rule-based transformations through the use of regular expressions. According to him, such rules are not easy to establish, since many surface-level transformations require deeper changes to preserve the grammar, and other transformations depend on the language. Valid transformations are, amongst others, the insertion of spelling mistakes, data alterations, entity names, and abbreviations. Coulombe concretely implements the transformation of verbal short forms to their long forms and vice versa ("I am" \leftrightarrow "I'm"). In the English language, this is semantically invariant if ambiguities are preserved [9]. With

this form of data augmentation, Coulombe achieves very good results [9]. The best baseline model (XGBoost) additionally gained 0.5% in terms of accuracy.

3.1.2 Word Level.

3.1.2.1 Noise Induction. Noise induction can also be applied on the word level. For example, the method of Xie et al. [35] encompasses two noise patterns. With "unigram noising", words in the input data are replaced by another word with a certain probability. By the method of "blank noising", words get replaced with "_". By the adoption of both patterns, the authors achieved improved results in their experiments.

Li et al. [41] are using syntactic and semantic methods as well as word dropout for the generation of noise. Syntactic noise is realized via the shortening of sentences and methods such as the alteration of adjectives or the relativization of modifiers, while semantic noise is generated by the lexical substitution of word synonyms (see Section 3.1.2.2). In contrast to these two methods, word dropout is more clearly comparable to noise. Random input neurons or rather words get masked out during the training of the network. According to the authors, their proposed methods achieve an improvement. Especially a combination of all methods promises an improvement of up to 1.7 points in terms of accuracy.

Two of the four sub-methods of the Easy Data Augmentation (EDA) method by Wei and Zou [2], i.e., random swap and deletion, should also be mentioned as methods of noise induction. In experiments, a combination of both sub-methods led to improved performance of the used classifier. EDA is very popular in the research field and was used as a method for comparisons in the works of Qiu et al. [42], Huong and Hoang [43], Anaby-Tavor et al. [44], Kumar et al. [45], Bayer et al. [46], Feng et al. [39], Luu et al. [47], and Kashefi and Hwa [48]. Wei and Zou [2] report that for a small dataset these two sub-methods gain higher improvements than the other two sub-methods that are based on synonym replacement and insertion (see Section 3.1.2). Nevertheless, Qiu et al. [42], Anaby-Tavor et al. [44], Bayer et al. [46], and Luu et al. [47] also report some cases in which EDA as a whole data augmentation method decreases the classification score. This result can be expected, as the methods random swap and deletion are not label preserving, for example, for sentiment classification: "I did not like the movie, but the popcorn was good" \rightarrow random_swap \rightarrow "I did like the movie, but the popcorn was not good". While Wu et al. [49] also use random swap and random deletion, they propose random span deletion, where consecutive words are deleted. This technique would lead to a worse label preservation, but it is only used for language modeling with contrastive learning. (see Section 3.4).

The training instances of one batch must have the same length when being fed into a neural network. For this purpose, the sequences are often zero-padded on one side. Rizos et al. [50] propose a specific noise induction method to augment the training data by shifting the instances within the confines of their padding so that the padding is not solely on one side. Evaluated by means of a hate speech detection dataset, the authors show that this method achieves additive performance gains of more than 5.8% (Macro-F1). Sun and He [51] also translate the instances by adding meaningless words either at the beginning or at the end. Unfortunately, they do not evaluate the impact of this method in isolation.

Xie et al. [52] propose a TF-IDF-based replacement method in which they are replacing uninformative words of an instance with other uninformative words. As the authors are combining this technique with round trip translation (see Section 3.1.4) and unsupervised data augmentation, it is not clear to which degree it benefits the task. Similarly, Choi et al. [53] replace casual features/words that are a determining factor for the label. In the contrastive learning scheme, they mask these words to generate counterfactual examples as well as other non-casual words to generate normal augmentations. More details and results of this procedure can be found in Section 3.4. More noise data augmentation methods related to other tasks can be found in the works of Cheng et al. [54], Li et al. [41], Wang et al. [55], Andreas [56], Guo et al. [57], Kashefi and Hwa [48], Sun and He [43], and Kurata et al. [58].

3.1.2.2 Synonym Replacement. This very popular form of data augmentation describes the paraphrasing transformation of text instances by replacing certain words with synonyms. One of the first applications of this substitution in the field of data augmentation was introduced by Kolomiyets et al. [59]. They substituted temporal expressions with potential synonyms from Word-Net [60]. As the authors argue, the replacement of one original token in a sentence will mostly preserve the semantics. Based on the time expression recognition task, the authors propose replacing the headword, since temporal trigger words are usually found there. While this application, however, showed no substantial improvements, the authors also proposed a language model replacement method that was more suited for the task at hand (see Section 3.1.2.4).

In later years, many researchers experimented with word replacements based on thesauri. The works of Li et al. [41], Mosolova et al. [61], Wang et al. [62], and much more partially or primarily execute synonym substitution in this way. Differences between the studies concern the specific words that are substituted, the synonyms that come into question, and the utilization of different databases. For example, X. Zhang et al. [63] and Marivate and Sefara [64] choose the synonyms for substitution on basis of the geometric distribution by which the insertion of a rather distant synonym becomes less probable. Furthermore, several approaches exclude stop words or words with certain POS-tags from the set of words considered for replacement. Interesting is also the second sub-method of EDA by Wei and Zou [2], where synonyms are not replacing specific words, but are randomly inserted into the instance. The replacement method, synonym-selection, database, and improvements of the various approaches are listed in Table 2.

Also to be emphasized is the more sophisticated integration into the learning process, as described by Jungiewicz and Pohl [66]. The authors replace words with synonyms only if the replacement with the chosen synonym maximizes the loss of the current state of the classifier model. Apart from this, there are approaches that adapt the general idea of thesauri-based replacements to perform augmentation on specific tasks, for example, in Kashefi and Hwa [48] and Feng et al. [39].

3.1.2.3 Embedding Replacement. Comparable to synonym substitution, embedding replacement methods search for words that fit as good as possible into the textual context and additionally do not alter the basic substance of the text. To achieve this, the words of the instances are translated into a latent representation space, where words of similar contexts are closer together. Accordingly, these latent spaces are based on the distributional hypothesis of distributional semantics [69, 70], which is currently mostly implemented in the form of embedding models. The selection of words that correspond to this hypothesis and are, thus, near in the representation space, implies that the newly created instances maintain a grammatical coherence, as displayed in Figure 2. Besides this advantage, Rizos et al. [50] argue that the "method encourages the downstream task to place a lower emphasis on associating single words with a label and instead place a higher emphasis on capturing similar sequential patterns, i.e., the context of hate speech". The benefits of this data augmentation technique in comparison to the synonym substitution method are that techniques based on the distributional hypothesis are more comprehensive and the context of texts is considered. This means that substitutions are not limited by a database, like WordNet, and that grammatically more correct sentences can be generated [71]. Furthermore, the general form of this approach can be beneficial for languages, which have no access to a large thesauri but a lot of general text resources, on the basis of which the self-supervised embedding models can be easily trained [9].

Synonym Database **Replacement Method** Synonym Selection Model Base Dataset Improvements [59] WordNet Headword replacement Not stated Logistic TempEval -1 (F1) Reuters (12) Regression -0.6Wikipedia (1) -0.1 [63, mytheas Randomly chosen number of Randomly based on Character [63]/[65] (Acc.) 65] (LibreOffice) words based on geometric geometric distribution. CNN AG News -0.38/-0.57 WordNetdistribution DBP +0.05/+0.13based Yelp P. -0.03/ Yelp F. +0.36Yahoo A. +0.22/0.65Amazon F. +0.1/0.1Amazon P. -0.17/-0.17[41] WordNet Substitutable words are nouns, +0.8 (Acc.) The remaining probability CNN MR verbs, adjectives, or adverbs of substitution is shared CR +1.2that are not part of a named among the synonyms Subj +0.5SST entity. based on a language model +0.1MR/CR 0.9 Each word is replaced with a score. certain probability. CR/MR 0.3 [9] XGBoost WordNet Only adverbs and adjectives, Most similar companion IMDB +0.5 (Acc.) information of the MLP (2 hidden +4.92sometimes nouns, more rarely synonym with the context verbs. layer) of the chosen word. WordNet CNN with -0.09/-0.21 [61] No pronouns, conjunctions, Uniform random Toxic prepositions, and articles for Comment (AUC) word embedding Classification replacement. Choosing uniform randomly. HIT IR-Lab [62] No time words, prepositions, Chi-square statistics ~+1 (Acc.) Character Hotel R. Tongyici Cilin and mimetic words. method CNN-SVM Laptop R. $\sim +1$ (Extended) Chi-square statistics method. Book R. ~+0.25 (Chinese) WordNet DNN [64] Verbs, nouns, and their Geometric distribution AG News ~+0.4 (Acc.) combination. Geometric Sentiment $\sim +0$ distribution Hate Speech ~-0.8 Kim CNN [66] WordNet & For Minibatch: Synonym that maximizes TREC +1.2 (Acc.) Augmentation with probability, Thesaurus. the loss. com POS-tag replacement, replacement of one word per sentence that maximizes loss. [2] WordNet No stop words. Uniform random CNN Classification SR/RI (Acc.) Choosing n random words to tasks (500) $\sim +1.9/\sim +2.0$ be replaced (SR) or from which (2000)~+1.2/~+0.9 the synonyms are inserted at a (5000)~+0.7/~+0.6 ~+1.0/~+0.9 random position (RI) (full) [1] WordNet Replacement of a word based CNN SST-5 -0.6 (Acc.) Temperature hyperparameter learned on a certain probability. SST-2 +0.5while training. Subi +0MPQA +0.2RT +0.1TREC -0.4[42] WordNet Replacement of a word based Temperature TextRCNN ICS -0.26 (Macro F1) on a certain probability. hyperparameter learned NEWS +1.63while training. [67] WordNet Not stated Not stated BERT SST-5 (40) -0.87 (Acc.) IMDB (40) -0.87 TREC (40) +0.01No stop words. 10% of [68] WordNet M-BERT CodiEsp-D +0.6 (F1) Not stated documents randomly selected. CodiEsp-P -0.7 (F1)

Table 2. Overview of Different Approaches of the Synonym Replacement Method

(Continued)

	Synonym Database	Replacement Method	Synonym Selection	Model Base	Dataset	Improvements
[39]	WordNet	Keywords replaced are ordered by their RAKE score (e.g., the probability of being a keyword).	Randomly selected. Replacement only with same POS-tag.	No model (intrinsic evaluation with different metrics)	Yelp-LR (small subset of Yelp Reviews)	+0.015 (SBLEU) -0.018 (UTR) -0.02 (TTR) -0.016 (RWords) 0 (SLOR) -0.007 (BPRO) +0.001 (SStd) 0 (SDiff)

Table 2. Continued

[..]*

*The full table (more methods) is available in the supplemental materials (online).



Fig. 2. Example of a sentence with predicted words that can be used to replace a word in the sentence [1].

Wang and Yang [72] use this kind of augmentation to better classify annoying tweets. They utilize k-nearest-neighbor to identify the most suitable embeddings as a substitution of the training data words. Compared to the baseline, they achieve an additive improvement of up to 2.4 points in the F1-Score with logistic regression. Marivate and Sefara [64], Rizos et al. [50], Huong and Hoang [43], and others utilize the embedding replacement in very similar ways. The greatest differences in terms of the method exist in the selection of words to be replaced (e.g., POS-tag based) and the selection of the replacing words based on the embeddings. An overview of the differences can be found in Table 3.

A major factor for poor results is that the embedding replacement does not necessarily guarantee the preservation of the contextual meaning of the instances. This, in turn, could lead to distortions of the label; e.g., "the movie was fantastic" and "the movie was horrible" are valid transformations but the sentiment is the opposite. A way to address this issue is the use of the counter-fitting method of Mrkšić et al. [74] for synonym embedding substitution, as carried out by Li et al. [41]. Counter-fitting is an approach that depicts word embeddings on the basis of a target function in a way that similarities between synonyms are rewarded and similarities between antonyms are sanctioned [74]. Li et al. [41] extend this approach by selecting the most fitting words with a higher possibility for the replacement. This is done by leveraging a language model that can give an indication of how well a given word fits into a sequence. However, the authors achieve rather mixed results with this method. The counter-fitting method offers considerably less replacement possibilities, since embeddings have to be trained on the downstream task, leading to a smaller coverage of their corpora words. Alzantot et al. [22] use this method in combination with a language model filtering in their adversarial example generator. They extend the approach by only incorporating the words that are maximizing the target label prediction probability (label

	Replacement Selection	Embedding Selection	Model Base	Dataset	Embedding Model	Improvements
[72]	Not stated	K-nearest-neighbor and cosine similarity	Logistic regression	Petpeeve dataset	UrbanDictionary W2V Twitter W2V GoogleNews W2V	+0.3 (F1) +1.7 +2.4
[64]	Random	Random with probability proportional to cosine similarity.	DNN	AG News Sentiment Hate Speech	Wikipedia W2V Wikipedia W2V GloVe Twitter	~0 (Acc.) ~+0.5 ~ -0.3
[50]	Every word	Cosine similarity threshold + POS-tag matching	CNN+LSTM/GRU	HON RSN-1 RSN-2	Word2Vec Hate Speech FastText Wikipedia GoogleNews W2V GloVe Common Crawl GloVe Common Crawl GloVe Common Crawl	-22.7 (Macro F1) +1.0 -3.3 +0.3 -0.2 0
[41]	Substitutable words are nouns, verbs, adjectives, or adverbs that are not part of a named entity. Each word is replaced with a certain probability.	Embeddings are found with counter-fitting. Candidates are replaced with a probability. The remaining probability is shared among the embeddings based on a language model score.	CNN	MR CR Subj SST MR/CR CR/MR	GoogleNews W2V	-0.6/-4.2 (Acc) +0.1/-3.7 +0.2/-1.4 -0.4/-4.2 +1.9/+0.4 +0.1/-3.0
[73]	No stop words or symbolic and numerical data	Cosine similarity threshold of 0.97	Manhattan LSTM model	Thai text similarity task	Thai2fit (Thai language)	+1.71
[]*						

Table 3. Overview of Different A	Approaches	of the Em	ibedding Rep	lacement Metl	100
----------------------------------	------------	-----------	--------------	---------------	-----

*The full table (more methods) is available in the supplemental materials (online).

preservation) of an already trained classifier. The authors report no improvements in terms of the task testing set, but they show that the model is safer regarding adversarial attacks. Embedding replacement methods are moreover used in specific task-dependent ways, such as by Kashefi and Hwa [48].

3.1.2.4 Replacement by Language Models. Language models represent language by predicting subsequent or missing words on the basis of the previous or surrounding context (classical and respectively masked language modeling). In this way, the models can, for example, be used to filter unfitting words, as already discussed in Section 3.1.2.3 in relation to the work of Alzantot et al. [22]. The authors generate similar words with GloVe embeddings and the counter-fitting method and utilize a language model to choose only words with a high probability of fit. In contrast to embedding replacements by word embeddings that take into account a global context, language models enable a more localized replacement [64]. Utilizing the subsequent word prediction, language models can also be used as the main augmentation method. Kobayashi [1] is, for example, using an LSTM language model to identify substitution words. However, language models do not only substitute words with similar meaning, but also with words that fit the context in principle [1]. This trait is encompassed with a greater risk of label distortion. To prevent the attachment of wrong labels to the new training data due to changed semantics, Kobayashi [1] modifies the language model so that it allows the integration of the label in the model for the word prediction (label-conditional language model). Inspired by this approach, Wu et al. [75] alter the architecture of the language model BERT [76] in a way that is labeled conditional (c-BERT). In an evaluation with different tasks, the authors showed that in a comparison to Kobayashi [1] and other

Publication	Method	Dataset	Improvements (Accuracy)
[75]	c-BERT	SST-2 MPQA TREC []*	+0.2 (CNN)/+0.5 (RNN) +0.5 (CNN)/+0.7 (RNN) +0.8 (CNN)/+0.2 (RNN) []*
[77]	c-BERT with consistency training	MLNI-m	+0.4 (RoBERTa-Base)
[44]	c-BERT	ATIS TREC WVA	-1.9 (BERT)/-0.8 (SVM)/-5.8 (LSTM) +1.1 (BERT)/+1.1 (SVM)/+6.5 (LSTM) +0.2 (BERT)/0.5 (SVM)/+2.4 (LSTM)
[67]	c-BERT integrated in reinforcement learning scheme	SST-5 (42) IMDB (45) TREC (45)	+1.17 (BERT)/+2.19 (normal c-BERT) +1.97 (BERT)/+1.97 (normal c-BERT) +0.73 (BERT)/+0.87 (normal c-BERT)
[78]	c-BERT and embedding substitution for multiple-pieces words	MRPC CoLA []*	+3.4 (TinyBERT) +21.0 (TinyBERT) []*

Table 4. Evaluation Results of the State-of-the-art Language Substitution Method c-BERT

*The full table (more results) is available in the supplemental materials (online).

approaches they were able to considerably increase the performance of a classifier (see Table 4). However, the c-BERT approach also has the disadvantage that the language model is fixed when applied, and in the case of low-data regimes, the augmentation might no longer be label preserving [67]. For this reason, Hu et al. [67] include the c-BERT method in a reinforcement learning scheme, which learns the task in a normal supervised fashion but is also able to simultaneously fine-tune the c-BERT LM. With this adaption, the authors significantly outperform the original c-BERT approach in a low-data regime setting. The results can be found in Table 4 together with the results of Anaby-Tavor et al. [44], who evaluated c-BERT as comparison, and Qu et al. [77], who employed the c-BERT model with supervised consistency training (see Section 3.4) on the MLNI-m task.

Jiao et al. [78] apply the already improved method by Wu et al. [75] and further adjust it in their work on TinyBERT. In doing so, the scholars reflect on the fact that the quality of the data generated with BERT is poor if many multiple-pieces words are included. To mitigate this problem, they propose to perform an embedding substitution on the base of GloVe embeddings [79] for such words. Further language model augmentations for different tasks are proposed by Gao et al. [80], Ratner et al. [81], Fadaee et al. [82], and Kashefi and Hwa [48].

3.1.3 Phrase and Sentence Level.

3.1.3.1 Structure-based Transformation. Structure-based approaches to data augmentation may utilize certain features or components of a structure to generate modified texts. Such structures can be based on grammatical formalities, for example, dependency and constituent grammars or POS-tags. Such approaches are therefore more limited to certain languages or tasks. Şahin and Steedman [83] are, for example, concerned with the augmentation of datasets from low resource languages for POS-tagging. By the method of "cropping", sentences are shortened by putting the focus on subjects and objects. With the "rotation" technique, flexible fragments are moved. The authors state that this method is dependent on certain grammatical sentence structures in different languages and probably only generates noise in the English language. Both methods are well suited for a multitude of low-resource languages. They were also tested by Vania et al. [84] for the augmentation of training data for dependency parsers for low-resource data.

Feng et al. [85] propose a method for changing the semantics of a text while trying to preserve fluency and sentiment. Given a set of phrases (replacement entities) to every instance, the so-called Semantic Text Exchange method first identifies phrases in the original text that can be replaced by a replacement entity based on the constituents. Then phrases similar to the identified phrases are replaced by a masked token. Subsequently, this is filled by an attention-based language model so that the similar words better suit the replacement entity. Feng et al. [39] adapt this approach by automatically selecting the 150 of the 200 most frequent nouns from the Semantic Text Exchange training set as replacement entity candidates and splitting their Yelp Review dataset into windows, as the method is only suitable for short texts. In an analysis with this dataset Feng et al. [39] reported that the Semantic Text Exchange method decreases fluency, diversity, and semantic content preservation.

An important work was proposed by Min et al. [86] who show that inversion (swapping the subject and object part) and passivation result in a higher generalization capability in natural language inference. In fact, considering their work in comparison with preliminary work [87]–[89] suggests that BERT is able to extract the relevant syntactic information from the instances but is unable to use this information in the task, as there are too few examples in the MNLI dataset demonstrating the necessity of syntax. Here, even a limited utilization of Min et al.'s [86] data augmentation methods already helps to mitigate this problem.

3.1.3.2 Interpolation. In numerical analysis, interpolation is a procedure to construct new data points from existing points [90]. While the formal interpolation versions are found in the feature space section, a sensible definition of interpolation in the data space of text is difficult to construct. However, the substructure substitution (SUB²) method by Shi et al. [91] is considered as such in this context due to its resemblance to the feature space methods. SUB² substitutes substructures (dependents, constituents, or POS-tag sequences) of the training examples if they have the same tagged label (for example, "a [DT] cake [NN]" in an instance can be replaced with "a [DT] dog [NN]" of another instance). The variant adapted for classification views all text spans of an instance as structures and is constrained by replacement rules that can be combined or completely left out. The replacement rules are only replacing (1) same lengths spans, (2) phrases with phrases, (3) phrases of the same constituency label, and (4) spans that come from instances with the same class label. The authors show that their methods outperform the baseline when applied to lowresource tasks. Their classification variant nearly doubles the accuracy on a subsample of the SST-2 and AG News datasets. Furthermore, they achieve better results than the language model augmentation c-BERT (Section 3.1.2.4). Similarly, Kim et al. [92] propose a data augmentation method based on lexicalized probabilistic context-free grammars that extract grammar trees from an input sentence and combines/substitutes them internally and with trees from other instances of the same class. Words are replaced with other words having the same POS-tag from the other sentences of the same class and WordNet synonyms. In this way, they can achieve a considerable performance improvement when applied in a few-shot, semi-supervised learning environment.

3.1.4 Document Level.

3.1.4.1 Round-trip Translation. Round-trip translation¹ is an approach to producing paraphrases with the help of translation models. A word, phrase, sentence, or document is translated into another language (forward translation) and afterward translated back into the source language (back translation) [93]. The rationale behind this is that translations of texts are often variable due to the complexity of natural language [9], which leads to various possibilities in the choice of terms or sentence structure. The process is depicted in Figure 3.

The approach is promising because of its good inherent label preserving and highly valuable paraphrasing capabilities. By the translation of the text, the content is preserved and only stylistic

¹Even though Coulombe [9], Yu et al. [94], Xie et al. [52], Qu et al. [68], and others use the term backtranslation for their data augmentation works as well, these approaches are assigned to the round-trip translation approaches because they execute forward and back translation.



Fig. 3. Round-trip translation process [94].

features based on the traits of the author are excluded or altered [95]. Some translation systems can propose several translation options; this is hinted at in Figure 2 ("k^2 paraphrases"). Yu et al. [94], Aroyehun and Gelbukh [71], Coulombe [9], Kruspe et al. [96], and others use this technique to generate artificial training data. Their works differ with regard to the used language and the subsequently applied filtering methods. These filtering methods are important, as the process of the twofold translation may be faulty [71]. Furthermore, Xie et al. [52] as well as Chen et al. [97] change the normal beam search generation strategy to random sampling with a temperature parameter to ensure a greater diversity. Details on the different round-trip translation applications are presented in Table 5.

3.1.4.2 Generative Methods. Generative methods are becoming increasingly interesting in recent data augmentation research. As the capabilities of language generation increased significantly, the current models are able to create very diverse texts and can thus incorporate new information. Here, Qiu et al. [42] introduce a variational autoencoder (VAE) based on a method that is used for text generation in their system. VAEs are probabilistic neural network structures that consist of an encoder network, which transforms input data into a latent representation, and of a decodernetwork, which transforms the latent representations back. The authors differentiate between unconditional and conditional VAEs. With unconditional VAEs, separate text generation models are trained for all classes, whereas with conditional VAEs, the label information is fed into the model as an additional input. Furthermore, they distinguish between sampling from the prior distribution, which leads to greatly diverse instances, and the posterior distribution, which produces text that is semantically closer to the training data. With the unconditional VAE and sampling from the prior distribution, they achieve the highest improvements of up to 2 F1-points (see Table 6). Malandrakis et al. [102] make similar efforts by evaluating VAEs for augmentation. While their objective is more narrowed, as they are interested in natural language understanding with limited resources, they analyze a broader variety of VAE augmentation variants. They also propose augmentation by conditional and unconditional VAEs with sampling from the posterior or prior distribution. Furthermore, they test two different learning objectives, where in the first the VAEs are used to reconstruct the input and in the second the VAEs take an instance of a particular class and try to construct another instance from that class. They also experiment with the addition of a discriminator network to the VAE that predicts the respective class from which an output appears to be. In intrinsic and extrinsic evaluations, the conditional VAEs with the reconstruction task are best performing. The discriminator variant achieves poor results, which stem from the little amount of available training data for the many different classes. Contrary to the improvements of Qiu et al. [42], the CVAEs outperform the VAE generation. An excerpt of the extrinsic evaluation is given in Table 6. However, it must be considered that the task at hand is very specific.

	Translation Model	Languages	Filtering	Model	Dataset	Improvements
[94]	Google's NMT [98]	$\mathrm{en} \to \mathrm{fr} \to \mathrm{en}$	No filtering	Convolution and self-attention	SQuAD	+1.5 (EM)/+1.1 (F1)
[9]	Google Translate API	Not stated	Excluding identical instances. Similarity threshold based on lengths.	XGBoost MLP 2 hidden layer	IMDB	+0 (Acc.) +5.8
[71]	Google Translate API	$en \rightarrow fr, es,$ de, hi $\rightarrow en$	No filtering	NBSVM LSTM BiLSTM LSTM-CNN CNN-BiLSTM []*	Aggression Detection	+0.19 (Macro F1) +7.39 +5.6 +19.45 +14.33 []*
[96]	Google Translate	Randomly selected	No filtering	Fusion CNN	TREC Incident Streams track	~ -1.2 (F1)
[64]	Google Translate API & Amazon translate	$en \rightarrow fr,$ $de \rightarrow en$	"We ensured that the [] texts carry the same meaning as the source text"	DNN	AG News Hate Speech	~+0.33 (Acc.) ~-2.3
[52]	WMT'14 English-French translation model	$\mathrm{en} \to \mathrm{fr} \to \mathrm{en}$	No filtering	Randomly initialized transformer	Yelp-5	+1.65 (Acc.)
[77]	WMT19 and released in FairSeq	$en \rightarrow de \rightarrow en$	No filtering	RoBERTa	MLNI-m	+0.9 (Acc.)
[99]***	Translation models from Britz et al. [100]	$en \rightarrow de, zh \rightarrow en$	No filtering	BERT	QQP RTE SST-2 CoLA []*	+0.4 (Acc.) +3.6 (Acc.) +0.7 (Acc.) +2.3 (Mcc) []*
[101]**	Not stated	Not stated	No filtering	Transformer base with consistency training	QQP RTE SST-2 []*	-0.2 (Acc.) +5.1 (Acc.) +0.7 (Acc.) []*

[..]*

*The full table (more methods and results) is available in the supplemental materials (online).

**Trained with consistency training.

***Trained with contrastive learning.

VEAs are also a main component of the NeuralEditor proposed by Guu et al. [103] that generates new texts based on edition vectors. For the training of the generative model, they take pairs of instances x' and x in the training data that are lexically similar, encode the differences of them and noise into an edition vector z, and try to generate x based on x' and z. It should be noted that the lexical similarity is just a rough approximation of semantic similarity. This represents a potential source of error, as, e. g., instances could be negated which in turn weakens the label preservation capabilities. However, this suffices for the purposes of the authors, as they only use the method for language modeling. Specifically, in this domain, they report improvements in terms of generating quality and perplexity. Raille et al. [104] propose an Edit-transformer, which is an adaptation of the Neural-Editor with the additional ability to function cross-domain so that the learned edits of a large dataset can be transferred to a smaller dataset. Besides the improvements in speed and language modeling, they also apply their method on three different classification tasks. The results are shown in Table 6.

Rizos et al. [50] create an RNN that, depending on a specific class, learns language modeling to generate training data thereafter. The class-specific RNN for augmentation is primed with a random start word from the class-specific training data. However, the authors state that this method produces the poorest results compared to embedding substitution and noise generation. In a similar

Publication Method Model Dataset Improvements [42] VAE Ensemble of BiLSTM, TextCNN, ICS (Zh) +0.04 (F1) CVAE + prior sampling TextRCNN, and FastText with News Category Dataset (EN) +2.02XGBoost as top-level classifier ICS (Zh) CVAE + posterior sampling -0.13News Category Dataset (EN) +1.55ICS (Zh) -0.06News Category Dataset (EN) +1.88 [102] VAE BiLSTM Movie +4.0 (Macro F1) CVAE + prior sampling Movie + Live Entertainment -0.5CVAE + posterior sampling Movie +5.9 Movie + Live Entertainment +1.7Movie +5.6 Movie + Live Entertainment +0.6 [118] CVAE BERT SNIPS (few shot) +8.00 SNIPS +0.06 (Acc.) FBDialog (few shot) +7.42FBDialog +0.0[104] Transformer-based sentence editor CNN SST-2 (20%) +0.87 (Acc.) CNN SST-2 (100%) -0.84 LSTM Amazon Reviews (1%) +1.12LSTM Amazon Reviews (4%) [..] +0.41 [..] RNN LM with random start word CNN+LSTM + GloVe++ HON [50] -1.8 (Micro-F1) RSN-1 priming +8.2RSN-2 -7.4 [68] CNN-LSTM LM priming with 30% of CNN-LSTM CodiEsp-P +3.1 (F1) a sentence [51] seqGAN +1.06 (F1) LSTM + pretrained embeddings Tan's task CNN + pretrained embeddings +0.9 LSCNN + pretrained embeddings +0.8[106] CS-GAN (GAN, RNN and CNN Amazon-5000 +1.6 (Acc.) reinforcement learning) Amazon-30000 -0.21 Emotion-15000 +0.77 NEWS-15000 +2.25[107] GPT-2 for rarer instances without Logistic Alerting No comparative regression/biLSTM/Bi-attentive filtering Information Feed results classification + ELMo + GloVe Prioritization BERT [44] CVAE ATIS (5) +7.3 (Acc.) TREC (5) +0.8WVA (5) -1.8 BERT [44] GPT-2 generation and classifier ATIS (5) +22.4 (Acc.) filtering ATIS (100) $\sim +0.5$ TREC (5) [..]* +4.0 [..] DistilGPT2 generation and classifier BERT AG-NEWS [108] +0.61 (Acc.) filtering CNN CyberTrolls +0.45BERT SST-2 +1.63GPT-2 with conditional fine-tuning, SST-2 (100) +15.53 (Acc.) [46] **ULMFit** -0.19 (Acc.) special prompting, and embedding SST-2 (700) filtering Layoff +4.84 (F1) Management Change +3.42 (F1) West Texas Explosions +3.81 (F1) Dublin [..]* -2.54 (F1) [..] [109] GPT-2 with conditional fine-tuning, RoBERTa MediaEval +0.55 (micro-F1) special prompting, and classifier FlauBERT CLS-FR +0.57filtering [110] GPT-2 with a reinforcement learning XLNet Offense Detection (20%) +1.3(F1)Offense Detection (40%) component for class conditional +4.3generation. Sentiment Analysis (20%) +1.2Sentiment Analysis (40%) +1.4Irony Classification (20%) +1.0Irony Classification (40%) +2.3[113] GPT-3 with prompt-based BERT (base) COLA (0.1%, 0.3%, 1.0%) +7.9, 3.2, -2.4 BERT (large) generation and pseudo-labeling SUBI (0.1%, 0.3%, 1.0%) +1.3, -1.8, -1.2SST-2 (0.1%, 0.3%, 1%, full) +20.9, 19.3, 5.7, 2.9 SST-2 (0.1%, 0.3%, 1.0%) [..] +23.7, 14.6, 3.0 [..]*

Table 6. Overview of Text Generation Methods

*The full table (more results) is available in the supplemental materials (online).

sense, Ollagnier and Williams [68] perform text generation using a language model (LSTM-CNN). In contrast, they split each document in a minibatch into sentences, then generate new sentences for 30% of them and utilize 30% of the beginning of a given sentence as a prompt.

Sun and He [51] use the seqGAN architecture [105] to generate artificial data on basis of a GAN. Comparable to computer vision, seqGAN consists of a generator network creating texts and a discriminator network examining the authenticity of the generated texts next to the real instances. As the discriminator network can only prove the authenticity after a sequence of words and thus gives delayed feedback to the generator, the generator network is trained as a reinforcement learning agent. Utilizing the method as a data augmentation technique, the authors only receive minor improvements of classification quality. Partially inspired by SeqGAN, Li et al. [106] propose CS-GAN, which consists of a GAN, RNN, and reinforcement learning component for sentence generation. The model receives the information about the label as a prior for the generator, which is implemented with by the RNN and RL components, which are then required by the discriminator to generate meaningful sentences. Subsequently, a classifier forces the output of sentences to fit the label. The results are listed in Table 6.

Wang and Lillis [107], Anaby-Tavor et al. [44], Abonizio and Junior [108], Bayer et al. [46], Claveau et al. [109], and Liu et al. [110] use the GPT-2 model of Radford et al. [111], which achieves very good results in text generation, to create new complete instances. Concerning the adoption of the method, Wang and Lillis [107] only describe that they use rare instances as dependent examples for the generation. Anaby-Tavor et al. [44], on the other hand, develop a method that increases safety with regard to label preservation. In a first step, they further train the GPT-2 model with training data of a certain task (fine-tuning). In the process, they concatenate the respective label to every instance in order to facilitate the generation of new data for the respective class. Finally, a classifier determines which generated instances can actually be assigned to the class stated. The authors manage to achieve significant improvements in the classification of sentences. They show that their method outperforms conditional VAEs (unfortunately no sampling technique is described) and even EDA (Section 3.1.1.1) and c-BERT (Section 3.1.2.4) when applied to a severe low data regime. The results of their LAMBADA approach and CVEA implementation are given in Table 6. Abonizio and Junior [108] try to improve this approach by concatenating three random samples as a prompt for the generation. Furthermore, they are using DistilGPT2 by Sanh et al. [112], which is substantially faster and smaller. As can be seen in Table 6, the method consistently outperforms the baseline. While LAMBADA and PREDATOR are only applicable to short texts as instances, Bayer et al. [46] designed a GPT-2 based approach to augment short as well as long text tasks. In this way, very high label preservation and diversity is to be achieved by fine-tuning the language model on the class-specific data, generating data prompted with specialized training data tokens, and a filtering method based on document embeddings. They can achieve high improvements for constructed and real-world low data regimes. However, they also discuss the limitations of their method and useful applications in terms of specific datasets and tasks. The results can also be seen in Table 6. Similarly, Claveau et al. [109] fine-tune the GPT-2 model using the class-specific data and input a random word from the original texts for a generation. Afterward a classifier is applied to filter the generated data instances. They evaluate their approach using English and French datasets (see Table 6). Liu et al. [110] use a reinforcement learning component after the softmax prediction of the GPT-2 model to predict the tokens depending on the class for which the instance is to be generated. The authors tested their method with various model architectures. It consistently improved all of them in all tasks, especially the larger pre-trained models, like BERT and XLNet. The results for XLNet are shown in Table 6. Yoo et al. [113] are among the first authors using the much larger GPT-3 model by Brown et al. [114] for data augmentation, which has considerably better generation capabilities. Such large language models are expensive and hard to be fine-tuned on the training data, which is why their augmentation method GPT3Mix selects some examples from the dataset and incorporates them with the label into sensible prompts for the model to be conditioned on. The newly created instances are then extracted from the generated text and a pseudo-label probability is calculated with the GPT model. As shown in Table 6, the method achieves outstanding performance increases on scarce and one full dataset. The authors further demonstrate that their method is superior to other augmentation methods, such as EDA [2], round-trip-translation, and Tmix [97], and that the performance increases if larger models for the classification are used. Nevertheless, given the size of the GPT-3 network and the correspondingly large training dataset, it might even be able to replicate some of the training (or even test) instances that were left out in the creation of a scarce dataset.

In the generative method, proposed by Lee et al. [115], the first step is to subdivide the data into slices (informed by or based on the labels). Then, a generative model is trained on these slices to predict an instance in the slice based on a subsample of instances in that slice. This model is subsequently used to generate new data for underrepresented slices by priming it with instances from it. This way, the authors gain several improvements in text classification, intent classification, and relation extraction tasks with state-of-the-art results for the latter two. Furthermore, Ding et al. [116] and Chang et al. [117] propose methods using generative models for tasks other than text classification.

3.2 Feature Space

Data augmentation in the feature space is concerned with the transformation of the feature representations of the input.

3.2.1 Noise Induction. As in the data space, noise can also be introduced in several variants in the feature space. For example, Kumar et al. [118] employ four such techniques for the ultimate goal of intent classification. One of those methods applies random multiplicative and additive noise to the feature representations, as shown in [58]. However, in contrast, they are not transforming the created representations back into the data space. Another method called Linear Delta calculates the difference between two instances and adds it to a third (all from the same class). The third method, which interpolates instances, is further elaborated in Section 3.2.2.2 (see Table 8). For their fourth method, the authors are adapting the Delta-Encoder by Schwartz et al. [119] for textual data. There, an autoencoder model learns the deltas between instance pairs of the same class, which is then utilized to generate instances of a new unseen class. In a normal testing setting, the methods only slightly improve the classification results, while in a few-shot setting all methods are highly beneficial.

Several feature space data augmentation methods stem from the adversarial training research field. As explained in the background section, the models are trained with adversarial examples, i.e., little perturbed training data instances that would change the prediction or maximize the loss. This can be formally written as follows [120]:

$$\min_{\theta} \mathbb{E}_{(\mathbf{Z}, \mathbf{y}) \sim \mathbf{D}} \left[\max_{||\boldsymbol{\delta}|| \leq \epsilon} \mathbf{L} \left(f_{\theta} \left(\mathbf{X} + \boldsymbol{\delta} \right), \mathbf{y} \right) \right],$$

where Θ are the model parameters and δ describes the perturbation noise added to the original instances (within a norm ball). Further, *D* is the data distribution, *y* the label, and *L* a loss function. The training of the network (outer minimization) can still be solved by **stochastic gradient descent (SGD)**, while the search for the right perturbations (i.e., inner maximization) is non-concave [120]. As described by Zhu et al. [120], **projected gradient descent (PGD)** [121, 122] can be used to solve this. Unfortunately, several convergence steps (K) to get a good result make it computationally expensive [120]. Shafahi et al. [123] and Zhang et al. [124] propose two methods

	Model	SST-2 Acc	STS-B P/S Corr	MNLI- m/mm-Acc	QQP Acc	RTE Acc	QNLI Acc	MRPC F1	CoLA Mcc
Baseline	RoBERTa-L	96.4	92.4	90.2	92.2	86.6	94.7	90.9	68.0
Adversarial Training	PGD	96.4	92.4	90.5	92.5	87.4	94.9	90.9	69.7
	FreeAT	96.1	92.4	90.0	92.5	86.7	94.7	90.7	68.8
	FreeLB	96.7	92.7	90.6	92.6	88.1	95.0	91.4	71.1
	ALUM*	96.6	92.1	90.9	92.2	87.3	95.1	91.1	68.2
	ALUM	-	-	91.4	-	-	-	-	-
	SMART	96.9	92.8	91.1	92.4	92.0	95.6	92.1	70.6
Cutoff**	Token	96.9	92.5	91.0	92.3	90.6	95.3	93.2	70.0
	Feature	97.1	92.4	90.9	92.4	90.9	95.2	93.4	71.1
	Span	96.9	92.8	91.1	92.4	91.0	95.3	93.8	71.5

Table 7. Comparison of Different Noise Inducing Methods on the GLUE task

*only adversarial pre-training.

**supervised consistency training.

The values in bold mark the method that performs best in the respective task.

that calculate the gradient with respect to the input (for PGD) on the same backward pass as the gradient calculations regarding the network parameters during a training step. This mitigates the additional calculation overhead of PGD. In detail, Free adversarial training (FreeAT) by Shafahi et al. [123] trains the same batch of training examples K times so that several adversarial updates can be performed. You Only Propagate Once (YOPO) by Zhang et al. [124] accumulates the gradients with respect to the parameters from the K steps and updates the parameters accordingly. Zhu et al. [120] also propose a method called Free Large-Batch (FreeLB), which is similar to YOPO, as it also accumulates the parameter gradients. On several tasks, this method consistently exceeds the results of the baseline and the other two methods. The results of the GLUE dataset are given in Table 7. Miyato et al. and Miyato et al. [40, 125] change the normal adversarial training rule so that no label information is needed and call it virtual adversarial training. Without going into exact details, virtual adversarial training regularizes the standard training loss with a KL divergence loss of the distribution of the predictions with and without perturbations, where the perturbations are chosen to maximize the KL divergence. While the virtual adversarial training method is suitable for semi-supervised learning, we are particularly interested in the supervised setting. Their method improves the supervised DBpedia topic classification task baseline classifier by 0.11 points of accuracy, leading to a +0.03 increase in accuracy in comparison to the conventional adversarial training method. Jiang et al. [126] propose the adversarial method SMART, which relies on the virtual adversarial training regularization. They introduce the utilization of the Bregman proximal point optimization with momentum to solve the virtual adversarial training loss, which prevents the model from aggressive updates [40]. The authors show in their experiments that the method significantly improves the baseline and is also able to achieve better results than the other methods discussed in this paragraph (for an overview, see Table 7). Furthermore, they demonstrate robustness enhancement and domain adaption capabilities in several evaluation applications.

Wang et al. [127] and Liu et al. [128] developed methods for enhancing the pre-training of language models with adversarial training. Wang et al. [127] simply generate adversarial examples on the output embeddings in the softmax function of the language models. Thereby they manage to reduce the perplexity of the AWD-LSTM and QRNN models on different datasets, which leads, for example, to a reduction of 2.29 points with respect to the Penn Treebank dataset with the AWD-LSTM model. However, it is not clear how the training of bigger pre-trained language models like BERT and RoBERTa would have been influenced by this method. This is addressed in the work



Fig. 4. Visualization of the different cutoff methods [101].

of Liu et al. [128] with their method called Adversarial training for large neural Language Models (ALUM), which introduces noise to the input embeddings. The authors build their system based on the virtual adversarial training by Miyato et al. [40], as they noticed that it is superior to conventional adversarial training for self-supervision. Furthermore, they found out that they can omit the Bregman proximate point method and the adversarial training proposed by Jiang et al. [126] and Shafahi et al. [123] when they are using curriculum learning, where the model is first trained with the standard objective and then with virtual adversarial training. They report promising generalization and robustness improvements with the largest transformer models. For example, RoBERTa models can be improved with the ALUM continual pretraining by +0.7 on the MNLI task, while standard continual pretraining does not introduce further gains. The results on the GLUE dataset are given in Table 7. The authors also tested the robustness of the models with three different adversarial datasets, where ALUM achieves significant improvements in all tasks. In another evaluation setting, they combine adversarial pretraining with adversarial fine-tuning. ALUM improves all evaluation scores of the standard pre-trained models. This model reaches the best performances and significantly outperforms the other models in all tested tasks, e.g., with an increased accuracy of +0.4 more than without tuning the SNLI dataset. The improvement on the MNLI task is given in Table 7.

With regard to the generative adversarial training methods of the feature space, it is also of interest to investigate how the newly created examples can be transformed into the data space to enable their inspection. This is done in the works of Liu et al. and Wan et al. [129, 130]. Wan et al. attempt to improve the classification behavior of a grammatical error correction system by training with adversarial examples. Such an example, extracted from the application of loss-increasing noise in the hidden representation of a transformer encoder, is mapped to the data space by a transformer encoder that was trained autoregressively. Then they use a similarity discriminator based on the model to filter instances that are not similar to their initial counterparts. Liu et al. [129] also use a transformer autoencoder architecture to generate data space instances. In contrast to the work of Wan et al. [130], they generate the noisy instances from the input embeddings, subsequently filter instances based on unigram word overlap, and try to improve machine question generation and question answering tasks. Both methods significantly improve the baselines and other methods.

Given the constraint that adversarial training can be computationally expensive, Shen et al. [101] propose three simple and efficient data augmentation methods of the feature space (see Figure 4). Token cutoff sets the entire embedding of a single word to 0, while the feature cutoff sets one embedding dimension of each word in the input to 0. The third method, span cutoff, employs token cutoff across a coherent span of words. With each method, several different, slightly modified instances can be created, which the authors see as different perspectives/views that can be integrated in a multi-view learning fashion through consistency training. This means that the model should predict similar outputs across different views (details can be found in Section 3.4). The authors evaluate their model on the GLUE task and compare it with adversarial training algorithms as



Fig. 5. Illustration of the interpolation method SMOTE.

well as round-trip translation. In three out of eight tasks, an improvement over all other methods could be achieved (see Table 7). They extend the cutoff strategies to work with language generation, and thereby significantly outperform the baseline as well as the adversarial training method of Wang et al. [127].

3.2.2 Interpolation Methods. For textual data, interpolation methods are mostly limited to the feature space since there is no intuitive way for combining two different text instances. Nevertheless, the application in the feature space seems reasonable, as the interpolation of hidden states of two sentences creates a new one containing the meaning of both original sentences [97, 131]. Besides this, from a learning-based perspective, interpolation methods have a high value for machine learning models. Possible explanations for the success of interpolation methods may stem from the balancing of classes, the smoothening of the decision border (regularization) [132], and the improvement of the representations [133].

For example, the **Synthetic Minority Over-sampling Technique (SMOTE**) approach in its original context was developed for the purpose of oversampling the minority class, which, as described in the background section, inherently leads to better classification performances. In fact, a balancing of a class can easily be achieved by simply copying the minority class. However, Chawla et al. [132] show that simple oversampling leads to more specific decision boundaries than applying SMOTE in the balancing of classes. Interpolation methods can smoothen the boundary, as shown in Figure 5. Smoothened and more general decision borders signify that an algorithm can generalize better and, in relation to training data, is accompanied by less overfitting. In this context, when applying interpolation methods to representations of the input data, Verma et al. [133] empirically and theoretically prove that representations are flattened with regard to the number of directions with significant variance. This is desirable since data representations capture less space, meaning that a classifier is more uncertain for randomly sampled representations and a form of compression is achieved which leads to generalization [133–135].

3.2.2.1 SMOTE Interpolation. SMOTE is an interpolation method of feature space representations of input data. With SMOTE, various neighbors close to a specific instance are searched within the feature space in order to be interpolated with the following formula:

$$\tilde{x} = x_i + \lambda * dist(x_i, x_j),$$

where (x_i, y_i) is the source instance and (x_j, y_i) is a close neighbor with the same class label. dist(a, b) is a distance measure and $\lambda \in [0, 1]$. Unlike mixup, only instances of the same class get interpolated. The rationale behind the calculation of neighbors with the same class labels is that the interpolations tend to be class preserving, leading to a higher safety of the technique. However, this leads to a limited novelty and diversity of the created instances.

SMOTE is rudimentarily illustrated in Figure 5. In the illustration, a binary classification problem is shown, in which a learning algorithm has learned the specific decision border. To achieve a balanced class distribution, a new instance is added to the blue class by utilizing SMOTE. This addition achieves, apart from a balancing of the dataset, an adjustment of the decision boundary. The

new boundary is less specific and thus contributes to more general decisions. SMOTE in combination with textual data augmentation is applied, for instance, in the work of Wang and Lillis [107]. Unfortunately, the authors do not describe how and at which point of the network the method is applied.

3.2.2.2 Mixup Interpolation. Mixup by Zhang et al. [136] is an interpolation method similar to SMOTE. In the simplest adoption, the convex interpolation is implemented with the following formulas:

$$\tilde{x} = \lambda x_i + (1 - \lambda) x_j$$
, whereas x_i , x_j are input vectors

 $\tilde{y} = \lambda y_i + (1 - \lambda) y_i$, whereas y_i , y_j are one – hot – coded labels

 (x_i, y_i) and (x_j, y_j) are sampled from the training data and λ is either fixed in [0, 1] or $\lambda \sim \text{Beta}(\alpha, \alpha)$, for $\alpha \in (0, \infty)$.

Mixup is a general technique that can be applied to all kinds of equal dimensional data. However, text cannot trivially be represented in equal dimensions [137]. As a very general method, Verma et al. [133] propose the idea of applying mixup within a randomly selected hidden layer of a neural network. Despite the fact that the authors only perform the tests on image datasets, this approach paves the way for the application of mixup in many textual-related tasks. The results are very promising, and for textual evaluations we advise the reader to look at the methods described below (Table 8), which oftentimes can be seen as specializations of the approach by Verma et al. [133] for textual data. Marivate and Sefara [64] state that they use mixup on representations of bag of word models, TF.IDF models, word embeddings, and language models. Unfortunately, the authors do not explicitly describe how the interpolation is performed. This raises questions about how to interpolate instances of different sizes, when, for example, word embedding vectors are used. Marivate and Sefara [64] report about 0.2, 0.4, and 0 points gain for the AG News, Sentiment 140, and Hate Speech detection task. In contrast, Qu et al. [77] describe the internal implementation of their interpolation. For the interpolation, they draw two instances from a mini-batch and linearly combine their input embedding matrices in the way described above. They improve the baseline on the MNLI-m task by an additional 0.6% in terms of accuracy. Guo et al. [138] propose two variants, wordMixup and senMixup, where the interpolation is applied in the word embedding space and on the final hidden layer of the neural network before it is passed to a softmax layer. For wordMixup the sequences have to be zero padded so that the dimensions are the same. For senMixup, this is not necessary, as the hidden embeddings generated are of the same length each. The improvement results of both methods with regard to the CNN model with pretrained GloVe embeddings (trainable), which is the best baseline, are presented in Table 8. Guo [139] further advances the wordMixup approach by using a nonlinear interpolation policy. The policy is constructed to mix each dimension of the individual word embeddings in a given sentence separately. Furthermore, the labels are also interpolated nonlinearly, while they are learned adaptively based on the mixed embeddings. This way, a much larger variety of generated examples can be created. While this procedure outperforms the other two variants in most tasks, it can also have a negative effect on the classification quality, as shown in Table 8. Similar to the senMixup method, Sun et al. [137] apply mixup to the output of transformer models. Furthermore, they only activate mixup in the last half of the training epochs to learn good representations first. The improvements on the GLUE benchmark are listed in Table 8. In a very similar way, Chen et al. [97] propose TMix, which is also able to interpolate the hidden representations of an encoder. Indeed, TMix is able to interpolate at every layer of the encoder, similar to Verma et al. [133]. Based on the work of Jawahar et al. [140], who analyzed the types of information learned in different layers of BERT, the authors narrowed down their approach and opted for 7, 9, and 12 as interpolation layers as they contain the syntactic and semantic information. The improvements of TMix are also shown in Table 8.

Method	Technique for textual application	Model	Datasets	Improvements
mixup by Marivate and Sefara [64]	Not stated	DNN	AG News Sentiment 140 Hate Speech	+0.2 (Acc.) +0.4 +0
[118]	Interpolation of the BERT CLS output	BERT-base- english-uncased	SNIPS (few shot) SNIPS FBDialog (few shot) FBDialog	+8.36 (Acc.) +0.0 +7.92 +0.08
[77]	Interpolation of the embedding matrices	RoBERTa-base	MNLI-m	+0.6 (Acc.)
wordMixup by Guo et al. [138]	Interpolation of zero-padded word embeddings	CNN	Trec SST-2 MR []*	+1.6 (Acc.) +0.2 +1.5 []*
senMixup by Guo et al. [138]	Interpolation on the final hidden layer	CNN	Trec SST-2 MR []*	+1.2 (Acc.) +0.3 +0.8 []*
Nonlinear Mixup by Guo [139]	Nonlinear interpolation of padded word embeddings	CNN	Trec Subj MR []*	+2.6 (Acc.) -0.5 +3.6 []*
Mixup-Transformer by Sun et al. [137]	Interpolation after last layer of the transformer	BERT-large	CoLA SST-2 MNLI-mm []*	+2.68 (Corr.) +0.81 (Acc.) -0.01 (Acc.) []*
TMix by Chen et al. [97]	Interpolation of the <i>m</i> th BERT layer (7, 9, and 12 randomly chosen per batch)	BERT-base + average pooling + two-layer MLP	AG News (10) AG News (2500) IMDB (10) IMDB (2500) []*	+4.6 (Acc.) +0.2 +1.8 +0.5 []*
TMix evaluated by [113]	Interpolation of the <i>m</i> th BERT layer (7, 9, and 12 randomly chosen per batch)	BERT-base	SST-2 (0.1, 0.3, 1.0%) MPQA (0.1, 0.3, 1.0%) RT20 (0.1, 0.3, 1.0%) []*	-0.2, -1.5, -2.1 +0.2, 2.9, 0.0 +2.3, 0.6, -1.9 []*
Intra-LADA [141]	Interpolation of an instance with a randomly reordered version of itself	BERT-base- multilingual-cased + linear layer	CoNLL (5%) CoNLL (100%) GermEval (5%) GermEval (100%)	+0.24 (F1) +0.03 (**) +0.29 +0.04 (**)
Inter-LADA [141]	Interpolation of the nearest neighbors and sometimes randomly selected instances	BERT-base- multilingual-cased + linear layer	CoNLL (5%) CoNLL (100%) GermEval (5%) GermEval (100%)	+1.32 (F1) +0.64 +0.49 +0.33
Intra-Inter-LADA [141]	Combination of Intra- and Inter-LADA	BERT-base- multilingual-cased + linear layer	CoNLL (5%) CoNLL (30%) GermEval (5%) GermEval (30%)	+1.57 (F1) +0.59 +0.53 +0.78

Table 8. Overview of Different Approaches of the Replacement Method "Mixup Interpolation"

*The full table (more results) is available in the supplemental materials (online).

**Included in the pretraining.

Similarly, Chen et al. [141] propose an interpolation augmentation method in which the hidden layer representations of two samples are interpolated. However, they noticed that this method is not suitable for sequence tagging tasks. For this reason, they propose Intra- and Inter-LADA. Intra-LADA aims at reducing noise from interpolating unrelated sentences by only interpolating an instance with a randomly reordered version of itself. This way, they can increase the performance in every tested task (see Table 8). However, Chen et al. [141] also hypothesize that their Intra-LADA algorithm is limited in producing diverse examples. This limitation leads to Inter-LADA, which sets a tradeoff between noise and regularization by interpolating instances that are close together. The closeness is estimated through kNN based on sentence-BERT [142] embeddings and extended by an occasional sampling of two completely random instances. As it can be seen in Table 8, Inter-LADA oftentimes performs better than Intra-LADA. The combination of both can further improve the results.

ACM Computing Surveys, Vol. 55, No. 7, Article 146. Publication date: December 2022.

3.3 Combination of Augmentation Methods

In augmentation research, a common technique is to combine several data augmentation methods to achieve more diversified instances [143]. Here, the combination can mean either the application of multiple separate or stacked methods. For the first kind, Sun and He [51] propose word-level and phrase-level methods that they apply separately. While the results of the word-level and phraselevel methods differ insignificantly, the combination of both groups of methods produced very good results. Similarly, Li et al. [41] combined their proposed methods, which led to better results for the in-domain evaluations. In the work of Bonthu et al. [144] round-trip translation, random swap, random deletion, and random synonym insertion are separately combined, which leads to the best improvement of an LSMT classifier. Furthermore, in contrastive learning, it makes sense to use more than one data augmentation strategy since the goal is to learn meaningful representations that can be fostered by many different views. For example, Yan et al. [145] and Wu et al. [49] use several simple methods of data augmentation for the contrastive learning objective. Details on contrastive learning and the results of the works can be found in the next section. The method of stacking data augmentation techniques, on the other hand, is not always feasible. It is, for example, in most cases not possible to first apply a feature space method and then a data space method. Qu et al. [77] tested this with round-trip translation, cutoff, and adversarial examples. Round-trip translation and the training with adversarial examples produced the best results. Marivate and Sefara [64] stack round-trip translation, synonym, and embedding replacement with mixup. In two out of three evaluation settings, this procedure reduces the minimal error.

For the combination of augmentation methods, the meta-learning augmentation approach by Ratner et al. [81] is also of interest. It describes the utilization of a neural network to learn data augmentation transformations [5]. Specifically, Ratner et al. use a GAN to generate sensible sequences of transformations that were defined beforehand. This approach is usable for the image as well as text datasets and the authors show that it can achieve a significant improvement when applied to a relation extraction task with augmentations based on language model replacements.

3.4 Training Strategies

While semi-supervision is not considered as data augmentation in this work, it can still be sensibly combined through consistency training. In its origin, consistency training is used to make predictions of classifiers invariant to noise [52]. This can be enforced by minimizing the divergences between the output distributions of real and noised instances. Additionally, as only output distributions are included in the process, this consistency can be trained with unlabeled data. Several authors analyze how consistency training behaves when data augmentation methods are used for noise. This process can be illustrated by taking an instance whose label is unknown, applying a label-preserving data augmentation method, and then learning the model to predict the same label for both instances. In this way, the model can learn the invariances and is able to generalize better. Xie et al. [52] show that they achieve very good results by employing consistency training with round-trip translation and a TF-IDF-based replacement method, with an absolute improvement of 22.79% in accuracy on an artificially created low-data regime based on the Amazon-2 dataset with BERT base. They are also able to outperform the state-of-the-art in the IMDb dataset with only 20 supervised instances. Chen et al. [97] even extend this approach within their MixText (TMix) system. First, they generate new instances with round-trip translation. Then, they guess the label of the original and augmented instances by taking a weighted average of the predictions of all of them. In the training, they randomly sample two instances and mix them together with TMix. If one of the two instances is from the original data, they are using the normal supervised loss, but if both instances are from the unlabeled or augmented data, they use the consistency loss, like Xie et al. [52]. Consistency training can also be applied in a supervised fashion as an additional

term in the training objective to enforce identical predictions. This is, for example, used in the cutoff method by Shen et al. [101]. They show in their ablation studies that this consistency term improves the accuracy results additively by 0.15%.

Ou et al. [77] combine supervised consistency training with contrastive learning. The contrastive learning scheme should bring the original and augmented instances closer together and the other instances further apart in the representation space. Contrastive learning can be applied in the pretraining phase of a language model so that meaningful representations are learned directly. Wu et al. [49] show that training a language model from scratch with this objective can lead to increased performances for downstream tasks. As augmentation methods, the authors use word deletion, span deletion, random reordering and synonym substitution, as well as combinations in sets of two. The evaluation of several tasks shows that there is no clear best augmentation method. Fang et al. [99] and Yan et al. [145] show that contrastive learning can also result in better sentence representations when using an already pre-trained model and further training the masked language modeling task with contrastive learning. While the work of Fang et al. [99] uses round-trip translation, Yan et al. [145] experiment with adversarial training, token mixing, cutoff and dropout. Qu et al. [77] and Choi et al. [53] even include contrastive learning into the supervised setting. As augmentation strategies, Qu et al. [77] use adversarial training combined with round-trip translation, and Choi et al. [53] use counterfactuals based on language model substitution. Combined with consistency training, Qu et al. [77] achieve even further improvements. A comparing overview can be found in Table 9 of the supplementary material (online).

Other training strategies in which the order of how the training examples are presented to the learning algorithm is altered are for example employed by Liu et al. [128], Yang et al. [146], and Claveau et al. [109]. Liu et al. [128] adopt a curriculum learning approach, where the algorithm first learns less difficult instances. Transferred to the data augmentation topic, the model is first trained with the original data and then with the augmented data. Yang et al. [146] reverse this step and first train the model with the augmented data and then with the original data. This way, the model can correct unfavorable behavior that is learned through noisy augmented data. They also tried an importance-weight loss in which the weights of the synthetic instances are lower but find that the other training method performs better.

3.5 Filtering Mechanisms

Mechanisms that filter the generated instances are especially important for methods that are not perfectly label-preserving. A simple mechanism is, for example, employed by Liu et al. [129], who remove generated instances based on the overlap of unigram words with their original equivalents. Similarly, other metrics could also be used, e.g., Levenshtein distance, Jaccard similarity coefficient, or Hamming distance. Wan et al. [130] use a similarity discriminator, initially proposed by Parikh et al. [147], which also measures the similarity of two sentences.

The generative methods by Anaby-Tavor et al. [44], Abonizio and Junior [108], and Claveau et al. [109] filter instances based on a classifier that was trained on the class data. This significantly reduces the diversity of samples, and the classifier cannot really be improved as it is already familiar with these instances. Bayer et al. [46] improve this by using embeddings to measure the quality of the generated instances with regard to the class and more importantly by incorporating the human expert in the loop who needs to determine the correct threshold. However, Yang et al. [146] consider another filtering mechanism in their work, which does not require human assistance and is very sophisticated due to the inclusion of two perspectives. Generally, Yang et al. [146] propose a generative method that is suitable for increasing the dataset size for question-answering tasks. While they propose to utilize language models for fine-tuning and generation of questions and answers, their filtering methods can be adapted for other data augmentation methods as well.

Space	Group	Work	Method description	Improvement
	Character Level Noise	[21]	Flip a letter, if it maximizes the loss	+0.62 Acc. (LSTM)
	Synonym Replacement	[66]	Only replace words with a synonym, if it maximizes the loss	+1.2 Acc. (Kim CNN)
	Embedding	[22]	Choosing embeddings based on the counter-fitting method	-0.6 - +1.9 Acc. (CNN)
	Replacement	[41]	Counter-fitting, language model selection, and maximizing the prediction probability	Safer model (LSTM)
Data Space	Language Model Replacement	[67]	c-BERT integrated in reinforcement learning scheme	+0.73 - +1.97 Acc. (BERT)
		[78]	c-BERT and embedding substitution for compound words	+1.9 - +21.0 Acc. (TinyBERT)
	Phrase Level Interpolation	[91]	Substitutes substructures	+20.6 - +46.2 Acc. (XLM-R)*
	Round-trip Translation	[52]	Random sampling with a temperature parameter	+1.65 Acc.
	Generative	[46]	Conditional GPT-2 with human assisted filtering	-2.54 F1 - +15.53 Acc. (ULMFit)*
	Methods	[110]	GPT-2 with a reinforcement learning component	+1.0 - +4.3 F1 (XLNet)*
		[126]	Virtual adversarial training with special optimization	+0.5 – +5.4 Acc. (RoBERTa-l)
P. d.	Noise	[128]	Virtual adversarial training with curriculum learning	-0.3 Corr +1.2 Acc. (RoBERTa-l)
Space		[101]	Embedding noising	+0.0 Corr +4.4 Acc. (RoBERTa-l)
		[137]	Interpolation after last layer of the transformer	-0.01 Acc +2.68 Corr. (BERT-l)
	Interpolation	[97]	Interpolation of a random BERT layer	+0.0 - +4.6 Acc. (BERT-b)*

Table 9. Collection of Some of the Most Advanced Data Augmentation Techniques for Text Classification

*Results contain tests on low data regime datasets.

[141]

A first filtering mechanism determines whether a generated instance is detrimental by measuring whether the validation loss increases when including the artificial instance. As this would require retraining the model with each generated example, the authors propose to use influence functions [148, 149] to approximate the validation loss change. Furthermore, the first train on the augmented instances and then on the original training data so that the model can adjust itself when unfavorable noise is included in the augmented instances. The other filtering mechanism tries to favor diversity by selecting examples that maximize the number of unique unigrams.

Interpolating neighbors and reordered versions

4 DISCUSSION: A RESEARCH AGENDA FOR TEXTUAL DATA AUGMENTATION

In the previous section, different data augmentation methods were grouped, explained, and compared in terms of performance and put into context with each other. One has to keep in mind that the results reported by the authors of the approaches linked in this survey article are restricted in their expressiveness and only show one perspective. Many results are limited to special kinds of models and datasets. Based on our findings, we identified an agenda for future research on data augmentation as follows:

4.1 Researching the Merits of Data Augmentation in the Light of Large Pre-trained Language Models

Generally, it is not possible to determine which augmentation method works best for a given dataset, nor predict which research direction will be the most appealing in the future.

+0.53 - +1.57 F1 (BERT-b)*

Nevertheless, some patterns in current approaches hint at the directions research can follow in order to overcome current obstacles and challenges. One of the most significant challenges, as formulated by Longpre et al. [4], concerns the usage of large pre-trained language models, which makes the utilization of several data augmentation methods obsolete. Large pre-trained models are currently state-of-the-art, nevertheless, we advise taking further advancements and findings in the research landscape into account, as for example deep belief networks [150], capsule networks [151], or task-specialized networks e.g., for sentiment analysis [152, 153]. Experiments with BERT or other bigger language models are therefore of particular interest. Similarly, several studies [44, 46, 67, 110, 129, 113] have shown that methods only slightly transforming instances with random behavior, such as with synonym replacement (Section 3.1.2.2), EDA (synonym replacement, random swap, deletion, and insertion in one) (Section 3.1.2.1), or by inserting spelling errors (Section 3.1.1.1), tend to be less beneficial in this setting than more elaborate ones. Particularly adversarial training (Section 3.2.1), cutoff (Section 3.2.1), interpolation (Sections 3.1.3.2 and 3.2.2), and some generative methods (Section 3.1.4.2) have shown significant improvements with large pre-trained language models. While replacement methods based on embeddings (Section 3.1.2.3) and especially language models (c-BERT) (Section 3.1.2.4) can also gain improvements in combination with those pre-trained models, several studies [44, 64, 77, 110] have shown that the previously mentioned methods can, in most cases, achieve improved results.

The described performance differences become apparent when approaching the challenge highlighted by Longpre et al. [4] from an intuitive perspective. Large language models map data to a latent space with representations nearly invariant to some transformations. For example, synonym replacement methods only replace words that are by definition very close to the representation space, leading to instances that are almost identical [61]. As Longpre et al. [4] hypothesize, data augmentation methods can only be helpful, if they are able to introduce new linguistic patterns. In such instances, using the mentioned methods and generative methods, in particular, might be sensible, as they are based on other large language models that can introduce a high novelty. However, the challenge proposed by Longpre et al. [4] does not have to be universally true. For example, the SUB² method by Shi et al. [91] only interpolates phrases from the training data and thus does not include unseen linguistic patterns but achieves high gains with a pre-trained model. Another interesting aspect concerns the experiments conducted by Yoo et al. [104], with which they demonstrate that their GPT-3-based generative augmentation method actually improves as the size of the pre-trained classifier increases. The authors hypothesize that larger classifiers have more capacity to better incorporate the GPT3Mix samples.

4.2 Improving Existing Data Augmentation Approaches

In general, the most promising data augmentation methods have limits and challenges that may be overcome with further research. **Generative models** or their output needs to be conditional on the specific class. Otherwise, the created instances might not preserve the label. This conditioning is oftentimes reached by training a model, which in turn requires enough data to be consistent. Bayer et al. [46] have shown that the conditional model can best replicate the data class if the problem definition and task data are relatively narrow. Tasks with a broad variety of topics in the data seem less suitable. This problem might be mitigated by adopting new conditioning methods. Currently, most approaches are extended by filter mechanisms. Existing mechanisms, as detailed in Section 3.5, have some drawbacks which might be reduced in the future. Another obstacle concerns generative models themselves, which can require many resources and time to create new instances [46]. Therefore, lightweight alternatives need to be tested in this setting, thus potentially preventing a high dependency on resources, which is referred to by Bayer et al. [46] as the high resource wall problem. Similarly, methods like **round-trip translation** are limited by the

underlying model. For example, Marivate and Sefara [64] hypothesize that round-trip translation might not be appropriate for social media data, where translation errors increase. This problem will be addressed in the future, as machine translation models improve their translation capabilities for such difficult instances.

For **adversarial examples**, Liu et al. [128], hypothesize that good generalizability performance stems from the perturbation of the embedding space, rather than the input space. However, data space adversarial training methods should not be disregarded too quickly, as Ebrahimi et al. [21] show that their data space method achieves better results than the virtual adversarial training by Miyato et al. [40]. A general challenge for adversarial training is that it can disturb the true label space in the training data. For example, adversarial example generators often rely on the belief that close input data points tend to have the same labels [120]. Concerning the data space methods, this is often not true for natural language tasks, where few words or even characters determine the class affiliation (e.g., sentiment classification: "I can't believe I like the movie" \rightarrow small transformation \rightarrow "I can't believe I like the movie"). Whether this applies to the adversarial example generators in the feature space needs to be evaluated. If so, research needs to find a way to exclude cases where small transformations disturb labels and at best include cases where stronger transformations still preserve the labels. For this purpose, inspecting feature space methods would be helpful. However, such an inspection is difficult to conduct due to their high-dimensional numerical representation. The same applies to the feature space's **interpolation methods**, where a back transformation to the data space is not trivial. Though, certain approaches, such as those from Liu et al. and Wan et al. [129, 130], use techniques such as encoder-decoder architectures capable of transforming the newly created instances to the data space. An inspection of interpolated instances could lead to interesting insights. This opens another research direction where the interpolation of instances in the data space could be further investigated. A method that initially implements this behavior is SUB² (Section 3.1.3.2), which interpolates instances of the data space through sub-phrase substitutions. This, however, does not result in a high diversity, which is particularly interesting. In this regard, further analysis of the GPT-3 language model by Brown et al. [114] could be valuable, as it shows very interesting interpolation capabilities in the data space.

However, even **avoidably inferior methods** can achieve better results, if they are integrated sensibly. The work of Jungiewicz and Pohl [66] can serve as an example. They perform synonym substitution only, if it increases the loss of the model. This demonstrates that some data augmentation techniques proposed in the different groups are advanced, sometimes adapting existing methods and refining them.

We highlight some advanced works of the different groups in Table 9 to show which research directions can be considered in the future. It must be emphasized that these methods are not necessarily the best in their groups. The selection is made by the author team on the basis of the information gathered while writing this survey.

4.3 Establishing more Comprehensive Evaluation Criteria and Standards for Method Comparison

A general problem in data augmentation research concerns that mostly only improvements with regard to the prediction performance on specific datasets are presented. While this metric is likely the most important one, other metrics, such as time and resource usage, language variety, or configurability, are also important for practitioners as well as for researchers. For example, the generative approaches based on GPT-2 seem very promising when considering prediction performance gain. Nevertheless, language variety is narrowed down, as the model is primarily trained on English data. Furthermore, only a few authors discuss the time required for the application of their data augmentation methods. The GPT-2 based method of Bayer et al. [46] takes up to 30 seconds for 146:30

generating one example, leading to several computing days for a 10-times augmentation of a small dataset. For instance, in the context of crisis informatics, this might take too long, as classifiers have to be created quickly for immediate incident management [154]. We, therefore, urge scientists developing data augmentation techniques to consistently describe the limitations of their approaches. For further data augmentation research, flexible standards should be established in order for methods to be compared more reliably, similar to other machine learning research fields, e.g. few-shot learning [155] or natural language generation [156]. It seems unrealistic that one or few general datasets can capture all peculiarities of data augmentation methods, especially not of those that one tailored to a specific problem. Nevertheless, a small benchmark that can be included in evaluations of upcoming data augmentation methods would be desirable. In the bestcase, such a benchmark would address different data augmentation goals, consisting of two or more datasets, from which one replicates a few-shot learning setting and the other a normal learning setting. With the growing usage of generative models, it might also be sensible to consider using datasets that are not part of current training datasets for language models, as an incorporation of testing data would lead to wrong conclusions. The benchmark should not be too large, in order to ensure specific evaluations can still be carried out. Researchers that try to develop such a benchmark, could also consider to specify how much data augmentation should be performed and what models should be used. When determining which model should be used, it might be useful to create an updatable benchmark, as proposed by Gehrmann et al. [156], which can be modified according to more recent state-of-the-art models.

4.4 Enhancing the Understanding of Text Data Augmentation

Shorten and Khoshgoftaar [5] highlight that while for some image data augmentation techniques it is easy to understand how they might improve the dataset and derived classifiers, however, for other techniques this improvement has not been explainable yet. This also applies to the text regime, where for example, data augmentation methods that paraphrase text without changing the meaning are naturally sensible, while methods applied in the feature space are much more complex to capture. Already the visualization of the data of feature space augmentations created by, for example, adversarial examples or interpolation methods, is much more complicated than in the image domain. As previously elaborated, existing approaches try to convert representations back into the data space by using encoder-decoder architectures [129, 130]. The resulting data space representations could then be investigated and used to better understand underlying data augmentation methods. Furthermore, a more in-depth understanding of why and when data augmentation works needs to be established. With the rise of large language models, the question emerges whether data augmentation methods paraphrasing input instances without incorporating new patterns may be obsolete [4]. Certain works have challenged this perspective, by demonstrating that even existing patterns can be beneficial for performance [91]. In this context, it is interesting to note that the augmentation method of Yoo et al. [113] provides better results when the size of the pre-trained language model increases.

4.5 Fostering the Usability of Data Augmentation Application

Most data augmentation methods are still research-based in their incremental development progress and therefore not suitable for every practitioner. A simple way to improve usability is to publish code and, in the best-case, develop libraries that can be used out of the box for augmenting a text dataset. Dhole et al. [157] propose the first large framework to include many text data augmentation methods and filtering mechanisms. The library by Papakipos et al. [158] is not as big for textual data augmentation methods but can be used for multiple modalities (audio, image, text, and video). While these are very useful libraries, the amalgamation of many procedures

comes with abstraction problems. For example, only individual data instances can be transformed and the augmentation procedure does not have access to the entire dataset, so that, for example, no interpolation procedures are implemented. In addition to creating libraries, it might be useful to explore augmentations with a good learning process integration. This can be considered as a criterion to simplify embedding the procedure in the general learning process. Resource utilization, speed, and general continuity in the learning process are crucial for this process. The first two criteria are becoming increasingly relevant as they are related to the current trend of data augmentation, i.e., the use of large underlying models that create a high resource and time execution overhead. As described above, this might be countered by utilizing more lightweight models. A low continuity in the learning process refers to the circumstance that a text data augmentation method is detached from the actual training process; or in the worst case, the learning procedure needs to be split into two halves. The former, also described as offline data augmentation by [37], means that the original data are augmented independently from the model training. A data augmentation technique is called online, if it is embedded into the learning process so that the artificial instances are stochastically included by the learning algorithm, which is, e.g., implemented in the work of Bonthu et al. [144]. The second form occurs, for example, when a feature space method needs to separate the normal network structure, in order to detach the encoder or embedding layer from the rest of the network. This results in a continuity problem of learning, so that, e.g., the encoder or embedding level cannot be trained further.

5 CONCLUSION

This survey provides an overview over data augmentation approaches suited for the textual domain. Data augmentation is helpful to reach many goals, including regularization, minimizing label effort, lowering the usage of real-world data, particularly in privacy-sensitive domains, balancing unbalanced datasets, and increasing robustness against adversarial attacks (see Section 2). On a high level, data augmentation methods are differentiated into methods applied in the feature and in the data space. These methods are then subdivided into more fine-grained groups, from noise induction to the generation of completely new instances. In addition, we propose several promising research directions that are relevant for future work. Especially in this regard, a holistic view on the current state-of-the-art is necessary. For example, the increasing usage of transfer learning methods makes some data augmentation methods obsolete, as they follow similar goals. Hence, there is a need for more sophisticated approaches that are capable of introducing new linguistic patterns not seen during pre-training, as suggested by Longpre et al. [4].

While data augmentation is increasingly being researched and seems very promising, it also has several limitations. For instance, many data augmentation methods can only create high-quality augmented data, if the original amount of data are large enough. Furthermore, as Shorten and Khoshgoftaar [5] describe, data augmentation is not capable of covering all transformation possibilities and eliminating all kinds of biases in the original data. Adopting the example of Shorten and Khoshgoftaar [5], in a news classification task, in which there are no articles containing sports, the standard data augmentation methods will most certainly also not create sport articles, even though this would be necessary. In contrast, data augmentation might induce new undesirable biases. For instance, language models like GPT can contain biases that are then propagated into the dataset [159]. The wide variety of techniques and some very sophisticated methods also bring another layer of complexity that needs to be understood. Moreover, data augmentation can be time-consuming, meaning that not all methods are feasible for time-critical machine learning development domains, e.g., in some areas of crisis informatics. An increased demand for resources, especially concerning training generative models, is inherent to data augmentation.

In order to mitigate some of the limitations and amplify the strengths of data augmentation, however, we proposed our research agenda, which comprises (1) researching the merits of data augmentation in the light of large pre-trained language models, (2) improving existing data augmentation approaches, (3) establishing more comprehensive evaluation criteria and standards for method comparison, (4) enhancing the understanding of text data augmentation, as well as (5) fostering the usability of data augmentation application.

REFERENCES

- S. Kobayashi. 2018. Contextual augmentation: Data augmentation bywords with paradigmatic relations. In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, vol. 2. 452–457. DOI: 10.18653/v1/n18-2072
- [2] J. Wei and K. Zou. 2020. EDA: Easy data augmentation techniques for boosting performance on text classification tasks. In Proceedings of the 9th International Joint Conference on Natural Language Processing. 6382–6388. DOI:10. 18653/v1/d19-1670
- [3] K. Kafle, M. Yousefhussien, and C. Kanan. 2017. Data augmentation for visual question answering. In Proceedings of the 10th International Natural Language Generation Conference. 198–202. DOI: 10.18653/v1/w17-3529
- [4] S. Longpre, Y. Wang, and C. DuBois. 2020. How effective is task-Agnostic data augmentation for pretrained transformers? In Proceedings of the Findings of the Association for Computational Linguistics Findings of ACL. 4401–4411. DOI:10.18653/v1/2020.findings-emnlp.394
- [5] C. Shorten and T. M. Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. J. Big Data 6, 1 (2019), 1–48. DOI: 10.1186/s40537-019-0197-0
- [6] Q. Wen, L. Sun, F. Yang, X. Song, J. Gao, X. Wang, and H. Xu. 2020. Time series data augmentation for deep learning: A survey. In IJCAI International Joint Conference on Artificial Intelligence. 4653–4660. DOI: 10.24963/ijcai.2021/631
- [7] L. Taylor and G. Nitschke. 2019. Improving deep learning with generic data augmentation. In Proceedings of the 2018 IEEE Symposium Series on Computational Intelligence. 1542–1547. DOI: 10.1109/SSCI.2018.8628742
- [8] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. 1998. Gradient-based learning applied to document recognition. Proceedings of the IEEE 86, 11 (1998), 2278–2323. DOI: 10.1109/5.726791
- C. Coulombe. 2018. Text Data Augmentation Made Simple By Leveraging NLP Cloud APIs. [Online]. Available: http://arxiv.org/abs/1812.04718
- [10] R. Müller, S. Kornblith, and G. Hinton. 2019. When does label smoothing help? Adv. Neural Inf. Process. Syst. 32, NeurIPS (2019). http://arxiv.org/abs/1906.02629
- M. Banko and E. Brill. 2001. Scaling to very very large corpora for natural language disambiguation. In Proceedings of the ACL. 26–33. DOI: 10.3115/1073012.1073017
- [12] L. Palen and K. M. Anderson. 2016. Crisis informatics-new data for extraordinary times: Focus on behaviors, not on fetishizing social media tools. *Science* (80-.) 353, 6296 (2016), 224–225. DOI:10.1126/science.aag2579
- [13] M. Bayer, M.-A. Kaufhold, and C. Reuter. 2021. Information overload in crisis management: Bilingual evaluation of embedding models for clustering social media posts in emergencies. *Proc. Eur. Conf. Inf. Syst.* 1–18. http://www. peasec.de/paper/2021/2021_BayerKaufholdReuter_InformationOverloadInCrisisManagementBilingualEvaluation_ ECIS.pdf.
- [14] M. A. Kaufhold, M. Bayer, and C. Reuter. 2020. Rapid relevance classification of social media posts in disasters and emergencies: A system and evaluation featuring active, incremental and online learning. *Inf. Process. Manag.* 57, 1 (2020), 102132. DOI:10.1016/j.ipm.2019.102132
- [15] M.-A. Kaufhold, A. S. Basyurt, K. Eyilmez, M. Stöttinger, and C. Reuter. 2022. Cyber threat observatory: Design and evaluation of an interactive dashboard for computer emergency response teams. In *Proceedings of the European Conference on Information Systems*. 6–18. Retrieved from https://aisel.aisnet.org/ecis2022_rp/99.
- [16] M.-A. Kaufhold, J. Fromm, T. Riebe, M. Mirbabaie, P. Kuehn, A. S. Basyurt, M. Bayer, M. Stöttinger, K. Eyilmez, R. Möller, C. Fuchß, S. Stieglitz, and C. Reuter. 2021. CYWARN: Strategy and technology development for cross-platform cyber situational awareness and actor-specific cyber threat communication. In *Workshop-Proceedings Mensch and Computer*. 1–9. Retrieved from https://dl.gi.de/handle/20.500.12116/37412.
- [17] S. Minaee, N. Kalchbrenner, E. Cambria, N. Nikzad, M. Chenaghlu, and J. Gao. 2021. Deep learning-based text classification. ACM Comput. Surv. 54, 3 (2021), 1–40. DOI: 10.1145/3439726
- [18] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 15, 1 (2014), 1929–1958.
- [19] S. Geman, E. Bienenstock, and R. Doursat. 1992. Neural networks and the Bias/Variance dilemma. Neural Comput. 4, 1 (1992), 1–58. DOI: 10.1162/neco.1992.4.1.1

ACM Computing Surveys, Vol. 55, No. 7, Article 146. Publication date: December 2022.

- [20] A. Hernández-García and P. König. 2018. Data augmentation instead of explicit regularization. arXiv:1806.03852. Retrieved from http://arxiv.org/abs/1806.03852.
- [21] J. Ebrahimi, A. Rao, D. Lowd, and D. Dou. 2018. Hotflip: White-box adversarial examples for text classification. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, Vol. 2. 31–36. DOI: 10.18653/ v1/p18-2006
- [22] M. Alzantot, Y. Sharma, A. Elgohary, B. J. Ho, M. B. Srivastava, and K. W. Chang. 2020. Generating natural language adversarial examples. In *Proceedings of the Empirical Methods in Natural Language Processing*. 2890–2896. DOI:10. 18653/v1/d18-1316
- [23] S. C. Wong, A. Gatt, V. Stamatescu, and M. D. McDonnell. 2016. Understanding data augmentation for classification: When to Warp? In Proceedings of the International Conference on Digital Image Computing: Techniques and Applications. DOI: 10.1109/DICTA.2016.7797091
- [24] P. Y. Simard, D. Steinkraus, and J. C. Platt. 2003. Best practices for convolutional neural networks applied to visual document analysis. In *Proceedings of the International Conference on Document Analysis and Recognition*. 958–963. DOI:10.1109/ICDAR.2003.1227801
- [25] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel. 2017. Domain randomization for transferring deep neural networks from simulation to the real world. In *Proceedings of the International Conference on Intelligent Robots and Systems, vol. 2017-Septe.* 23–30. DOI: 10.1109/IROS.2017.8202133
- [26] H. Inoue. 2018. Data augmentation by pairing samples for images classification. http://arxiv.org/abs/1801.02929.
- [27] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. 2020. Random erasing data augmentation. In 34th AAAI Conf. Artif. Intell. 13001–13008. DOI: 10.1609/aaai.v34i07.7000
- [28] M. Frid-Adar, E. Klang, M. Amitai, J. Goldberger, and H. Greenspan. 2018. Synthetic data augmentation using GAN for improved liver lesion classification. In *Proceedings of the Int. Symp. Biomed. Imaging*, vol. 2018-April. 289–293. DOI:10.1109/ISBI.2018.8363576
- [29] X. Zhu, Y. Liu, J. Li, T. Wan, and Z. Qin. 2018. Emotion classification with data augmentation using generative adversarial networks. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics*). vol. 10939. 349–360. DOI: 10.1007/978-3-319-93040-4_28
- [30] N. Carlini, F. Tramèr, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, Ú. Erlingsson, A. Oprea, and C. Raffel. 2020. Extracting training data from large language models. In *Proceedings of the 30th USENIX* Security Symposium. 2633–2650. [Online]. Available: http://arxiv.org/abs/2012.07805
- [31] X. Cui, V. Goel, and B. Kingsbury. 2014. Data Augmentation for deep neural network acoustic modeling. In Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing. 5582–5586. DOI: 10.1109/ICASSP. 2014.6854671
- [32] N. Jaitly and G. E. Hinton. 2013. Vocal tract length perturbation (VTLP) improves speech recognition. In Proceedings of the 30 th Int. Conf. Mach. Learn. 90 (2013), 42–51. http://www.iarpa.gov/.
- [33] A. Ragni, K. M. Knill, S. P. Rath, and M. J. F. Gales. 2014. Data augmentation for low resource languages. In Proceedings of the Proceedings of the Annual Conference of the International Speech Communication Association. 810–814.
- [34] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Satheesh, S. Sengupta, A. Coates, and A. Y. Ng. 2014. Deep speech: Scaling up end-to-end speech recognition. http://arxiv.org/abs/1412.5567.
- [35] Z. Xie, S. I. Wang, J. Li, D. Lévy, A. Nie, D. Jurafsky, and A. Y. Ng. 2017. Data noising as smoothing in neural network language models. In *Proceedings of the 5th International Conference on Learning Representations*, vol. 2015-Jan. 3586– 3589. DOI: 10.21437/interspeech.2015-711
- [36] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur. 2015. Audio augmentation for speech recognition. In Proceedings of the Proceedings of the Annual Conference of the International Speech Communication Association, vol. 2015-Jan. 3586–3589. DOI: 10.21437/interspeech.2015-711
- [37] S. Y. Feng, V. Gangal, J. Wei, S. Chandar, S. Vosoughi, T. Mitamura, and E. Hovy. 2021. A survey of data augmentation approaches for NLP. In *Findings of the Association for Computational Linguistics*. 968–988. DOI:10.18653/v1/2021. findings-acl.84
- [38] Y. Belinkov and Y. Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *Proceedings of* 6th International Conference on Learning Representations.
- [39] S. Y. Feng, V. Gangal, D. Kang, T. Mitamura, and E. Hovy. 2020. GenAug: Data augmentation for finetuning text generators. In Deep Learning Inside Out (DeeLIO): The First Workshop on Knowledge Extraction and Integration for Deep Learning Architectures, 29–42. DOI: 10.18653/v1/2020.deelio-1.4
- [40] T. Miyato, A. M. Dai, and I. Goodfellow. 2017. Adversarial training methods for semi-supervised text classification. In Proceedings of the 5th International Conference on Learning Representations.
- [41] Y. Li, T. Cohn, and T. Baldwin. 2017. Robust training under linguistic adversity. In Proceedings of the European Chapter of the Association for Computational Linguistics, vol. 2. 21–27. DOI: 10.18653/v1/e17-2004

- [42] S. Qiu, B. Xu, J. Zhang, Y. Wang, X. Shen, G. De Melo, C. Long, and X. Li. 2020. EasyAug: An automatic textual data augmentation platform for classification tasks. In *Proceedings of the Companion of the World Wide Web Conference*. 249–252. DOI: 10.1145/3366424.3383552
- [43] T. H. Huong and V. T. Hoang. 2020. A data augmentation technique based on text for Vietnamese sentiment analysis. ACM Int. Conf. Proceeding Ser. 1–5. DOI: 10.1145/3406601.3406618.
- [44] A. A. Tavor, B. Carmeli, E. Goldbraich, A. Kantor, G. Kour, S. Shlomov, N. Tepper, and N. Zwerdling. 2020. Do not have enough data? Deep learning to the rescue!. 34th AAAI Conf. Artif. Intell. 34, 06 (2020), 7383–7390. DOI: 10.1609/ aaai.v34i05.6233
- [45] V. Kumar, A. Choudhary, and E. Cho. 2020. Data augmentation using pre-trained transformer models. http://arxiv. org/abs/2003.02245.
- [46] M. Bayer, M.-A. Kaufhold, B. Buchhold, M. Keller, J. Dallmeyer, and C. Reuter. 2022. Data augmentation in natural language processing: A novel text generation approach for long and short text classifiers. *Int. J. Mach. Learn. Cybern.* DOI:10.1007/s13042-022-01553-3
- [47] S. T. Luu, K. Van Nguyen, and N. L.-T. Nguyen. 2020. Empirical study of text augmentation on social media text in vietnamese. http://arxiv.org/abs/2009.12319.
- [48] O. Kashefi and R. Hwa. 2020. Quantifying the evaluation of heuristic methods for textual data augmentation. In Proceedings of the 6th Workshop on Noisy User-generated Text (W-NUT'20). 200–208. DOI:10.18653/v1/2020.wnut-1.26
- [49] Z. Wu, S. Wang, J. Gu, M. Khabsa, F. Sun, and H. Ma. 2020. CLEAR: Contrastive learning for sentence representation. http://arxiv.org/abs/2012.15466.
- [50] G. Rizos, K. Hemker, and B. Schuller. 2019. Augment to prevent: Short-text data augmentation in deep learning for hate-speech classification. In Proceedings of the International Conference on Information and Knowledge Management. 991–1000. DOI: 10.1145/3357384.3358040
- [51] X. Sun and J. He. 2020. A novel approach to generate a large scale of supervised data for short text sentiment analysis. Multimed. Tools Appl. 79, 9–10 (2020), 5439–5459. DOI:10.1007/s11042-018-5748-4
- [52] Q. Xie, Z. Dai, E. Hovy, M.-T. Luong, and Q. V. Le. 2019. Unsupervised data augmentation for consistency training. In Advances in Neural Information Processing Systems, vol. 2020-Dec. http://arxiv.org/abs/1904.12848.
- [53] S. Choi, M. Jeong, H. Han, and S.-W. Hwang. 2022. C 2 L: Causally contrastive learning for robust text classification. 2022. Available: www.aaai.org.
- [54] Y. Cheng, L. Jiang, and W. Macherey. 2020. Robust neural machine translation with doubly adversarial inputs. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. 4324–4333. DOI: 10.18653/ v1/p19-1425
- [55] X. Wang, H. Pham, Z. Dai, and G. Neubig. 2020. Switchout: An efficient data augmentation algorithm for neural machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 856–861. DOI:10.18653/v1/d18-1100
- [56] J. Andreas. 2020. Good-Enough compositional data augmentation. In Proceedings of the ACL. 7556–7566. DOI:10. 18653/v1/2020.acl-main.676
- [57] D. Guo, Y. Kim, and A. Rush. 2020. Sequence-Level mixed sample data augmentation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing. 5547–5552. DOI: 10.18653/v1/2020.emnlp-main.447
- [58] G. Kurata, B. Xiang, and B. Zhou. 2016. Labeled data generation with encoder-decoder LSTM for semantic slot filling. *Proceedings of the Annual Conference of the International Speech Communication Association*, vol. 08-12-Sept. 725–729. DOI: 10.21437/Interspeech.2016-727
- [59] O. Kolomiyets, S. Bethard, and M. F. Moens. 2011. Model-portability experiments for textual temporal analysis. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, vol. 2. 271–276.
- [60] G. A. Miller, R. Beckwith, C. Fellbaum, D. Gross, and K. J. Miller. 1990. Introduction to wordnet: An on-line lexical database. Int. J. Lexicogr. 3, 4 (1990), 235–244. DOI: 10.1093/ijl/3.4.235
- [61] A. V. Mosolova, V. V. Fomin, and I. Y. Bondarenko. 2018. Text augmentation for neural networks. In CEUR Workshop Proc. vol. 2268. 104–109.
- [62] X. Wang, Y. Sheng, H. Deng, and Z. Zhao. 2019. Charcnn-svm for chinese text datasets sentiment classification with data augmentation. Int. J. Innov. Comput. Inf. Control 15, 1 (2019), 227–246. DOI: 10.24507/ijicic.15.01.227
- [63] X. Zhang, J. Zhao, and Y. Lecun. 2015. Character-level convolutional networks for text classification. In Proceedings of the Advances in Neural Information Processing Systems, vol. 2015-Jan. 649–657.
- [64] V. Marivate and T. Sefara. 2020. Improving short text classification through global augmentation methods. In *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 12279. 385–399. DOI:10. 1007/978-3-030-57321-8_21
- [65] X. Zhang and Y. LeCun. 2015. Text understanding from scratch. http://arxiv.org/abs/1502.01710.

ACM Computing Surveys, Vol. 55, No. 7, Article 146. Publication date: December 2022.

- [66] M. Jungiewicz and A. Smywiński-Pohl. 2019. Towards textual data augmentation for neural networks: Synonyms and maximum loss. Comput. Sci. 20, 1 (2019), 57–84. DOI: 10.7494/csci.2019.20.1.3023
- [67] Z. Hu, B. Tan, R. Salakhutdinov, T. Mitchell, and E. P. Xing. 2019. Learning data manipulation for augmentation and weighting. In Advances in Neural Information Processing Systems, vol. 32. 15764–15775.
- [68] A. Ollagnier and H. Williams. 2020. Text augmentation techniques for clinical case classification. CEUR Workshop Proceedings 2696, June (2020), 22–25. https://temu.bsc.es/codiesp/index.php/2019/09/19/resources/.
- [69] Z. S. Harris. 1954. Distributional Structure. WORD 10, 2-3 (1954), 146-162. DOI: 10.1080/00437956.1954.11659520
- [70] J. Firth. 1962. A synopsis of linguistic theory. In Proceedings of the Studies in Linguistic Analysis. 1–32.
- [71] S. T. Aroyehun and A. Gelbukh. 2018. Aggression detection in social media: Using deep neural networks, data augmentation, and pseudo labeling. In Proceedings of the 1st Work. Trolling, Aggress. Cyberbullying. 90–97.
- [72] W. Y. Wang and D. Yang. 2015. That's so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 2557–2563. DOI: 10.18653/v1/d15-1306
- [73] T. Phreeraphattanakarn and B. Kijsirikul. 2021. Text data-augmentation using text similarity with manhattan siamese long short-term memory for thai language. J. Phys. Conf. Ser. 1780, 1 (2021), 12018. DOI: 10.1088/1742-6596/ 1780/1/012018
- [74] N. Mrkšić, D. Séaghdha, B. Thomson, M. Gašić, L. Rojas-Barahona, P. H. Su, D. Vandyke, T. H. Wen, and S. Young. 2016. Counter-fitting word vectors to linguistic constraints. In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. 142–148. DOI:10.18653/v1/n16-1018
- [75] X. Wu, S. Lv, L. Zang, J. Han, and S. Hu. 2019. Conditional BERT contextual augmentation. In Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 11539. 84–95. DOI: 10.1007/978-3-030-22747-0_7
- [76] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. North Am. Chapter Assoc. Comput. Linguist. Hum. Lang. Technol. - Proc. Conf. 1, no. Mlm (2019), 4171–4186. http://arxiv.org/abs/1810.04805.
- [77] Y. Qu, D. Shen, Y. Shen, S. Sajeev, J. Han, and W. Chen. 2020. CoDA: Contrast-enhanced and diversity-promoting data augmentation for natural language understanding. http://arxiv.org/abs/1810.04805.
- [78] X. Jiao, Y. Yin, L. Shang, X. Jiang, X. Chen, L. Li, F. Wang, and Q. Liu. 2020. TinyBERT: Distilling BERT for natural language understanding. In *Proceedings of the Findings of the Association for Computational Linguistics Findings of* ACL. 4163–4174. DOI: 10.18653/v1/2020.findings-emnlp.372
- [79] J. Pennington, R. Socher, and C. D. Manning. 2014. GloVe: Global vectors for word representation. In Proceedings of the Conference on Empirical Methods in Natural Language Processing. 1532–1543. DOI: 10.3115/v1/d14-1162
- [80] F. Gao, J. Zhu, L. Wu, Y. Xia, T. Qin, X. Cheng, W. Zhou, and T. Y. Liu. 2020. Soft contextual data augmentation for neural machine translation. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. 5539–5544. DOI: 10.18653/v1/p19-1555
- [81] A. J. Ratner, H. R. Ehrenberg, Z. Hussain, J. Dunnmon, and C. Ré. 2017. Learning to compose domain-specific transformations for data augmentation. In *Proceedings of the Advances in Neural Information Processing Systems*, vol. 2017-Dec. 3237–3247. Retrieved from https://github.com/HazyResearch/tanda.
- [82] M. Fadaee, A. Bisazza, and C. Monz. 2017. Data augmentation for low-Resource neural machine translation. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, Vol. 2. 567–573. DOI:10. 18653/v1/P17-2090
- [83] G. G. Şahin and M. Steedman. 2018. Data augmentation via dependency tree morphing for low-resource languages. In Proceedings of the Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing. 5004– 5009. DOI: 10.18653/v1/d18-1545.
- [84] C. Vania, Y. Kementchedjhieva, A. Søgaard, and A. Lopez. 2020. A systematic comparison of methods for low-resource dependency parsing on genuinely low-resource languages. In Proceedings of the Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing. 1105–1116. DOI:10.18653/v1/d19-1102
- [85] S. Y. Feng, A. W. Li, and J. Hoey. 2020. Keep calm and switch on! Preserving sentiment and fluency in semantic text exchange. In Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing. 2701–2711. DOI: 10.18653/v1/d19-1272
- [86] J. Min, R. T. McCoy, D. Das, E. Pitler, and T. Linzen. 2020. Syntactic data augmentation increases robustness to inference heuristics. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. 2339– 2352. DOI: 10.18653/v1/2020.acl-main.212
- [87] R. Thomas McCoy, E. Pavlick, and T. Linzen. 2020. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the ACL*. 3428–3448. DOI:10.18653/v1/p19-1334

- [88] Y. Goldberg. 2019. Assessing BERT's syntactic abilities. 2019. http://arxiv.org/abs/1901.05287.
- [89] I. Tenney, P. Xia, B. Chen, A. Wang, A. Poliak, R. Thomas McCoy, N. Kim, B. Van Durme, S. R. Bowman, D. Das, and E. Pavlick. 2019. What do you learn from context? Probing for sentence structure in contextualized word representations. In 7th Int. Conf. Learn. Represent. Retrieved from https://github.com/jsalt18-sentence-repl/jiant.
- [90] J. F. Steffensen. 2006. Interpolation. Courier Corporation. https://www.amazon.de/Interpolation-Second-Mathematics-Steffensen-2006-03-17/dp/B01K3HR83G.
- [91] H. Shi, K. Livescu, and K. Gimpel. 2021. Substructure substitution: Structured data augmentation for NLP. In Findings of the Association for Computational Linguistics. 3494–3508. DOI: 10.18653/v1/2021.findings-acl.307
- [92] H. Kim, D. Woo, S. J. Oh, J.-W. Cha, and Y.-S. Han. 2021. ALP: Data augmentation using lexicalized PCFGs for fewshot text classification. http://arxiv.org/abs/2112.11916.
- [93] M. Aiken and M. Park. 2010. The efficacy of round-trip translation for MT evaluation. Transl. J. 14, 1 (2010), 1–10.
- [94] A. W. Yu, D. Dohan, M. T. Luong, R. Zhao, K. Chen, M. Norouzi, and Q. V. Le. 2018. QaNet: Combining local convolution with global self-attention for reading comprehension. In *Proceedings of the 6th International Conference on Learning Representations*.
- [95] E. Rabinovich, S. Mirkin, R. N. Patel, L. Specia, and S. Wintner. 2017. Personalized machine translation: Preserving original author traits. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics. 1074–1084. DOI: 10.18653/v1/e17-1101
- [96] A. Kruspe, J. Kersten, M. Wiegmann, B. Stein, and F. Klan. 2018. Classification of Incident-related Tweets: Tackling imbalanced training data using hybrid CNNs and translation-based data augmentation. In *Proceedings of the TREC*.
- [97] J. Chen, Z. Yang, and D. Yang. 2020. MixText: Linguistically-Informed interpolation of hidden space for semisupervised text classification. In Proceedings of the 58th Annu. Meet. Assoc. Comput. Linguist. 2147-2157. DOI:10. 18653/v1/2020.acl-main.194
- [98] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, J. Klingner, A. Shah, M. Johnson, X. Liu, Ł. Kaiser, S. Gouws, Y. Kato, T. Kudo, H. Kazawa, K. Stevens, G. Kurian, N. Patil, W. Wang, C. Young, J. Smith, J. Riesa, A. Rudnick, O. Vinyals, G. Corrado, M. Hughes, and J. Dean. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. 1–23. http://arxiv.org/abs/1609.08144.
- [99] H. Fang, S. Wang, M. Zhou, J. Ding, and P. Xie. 2020. CERT: Contrastive self-supervised learning for language understanding. Retrieved Jul. 14, 2022 from https://github.com/UCSD-AI4H/.
- [100] D. Britz, A. Goldie, M. T. Luong, and Q. V. Le. 2017. Massive exploration of neural machine translation architectures. In Proceedings of the Conference on Empirical Methods in Natural Language Processing. 1442–1451. DOI: 10.18653/v1/ d17-1151
- [101] D. Shen, M. Zheng, Y. Shen, Y. Qu, and W. Chen. 2020. A simple but tough-to-beat data augmentation approach for natural language understanding and generation. http://arxiv.org/abs/2009.13818.
- [102] N. Malandrakis, M. Shen, A. Goyal, S. Gao, A. Sethi, and A. Metallinou. 2019. Controlled text generation for data augmentation in intelligent artificial agents. In Proceedings of the 3rd Workshop on Neural Generation and Translation. 90–98. DOI:10.18653/v1/d19-5609
- [103] K. Guu, T. B. Hashimoto, Y. Oren, and P. Liang. 2018. Generating sentences by editing prototypes. In Proceedings of the Transactions of the Association for Computational Linguistics, Vol. 6. 437–450. DOI: 10.1162/tacl_a_00030
- [104] G. Raille, S. Djambazovska, and C. Musat. 2020. Fast cross-domain data augmentation through neural sentence editing. 2020. http://arxiv.org/abs/2003.10254.
- [105] L. Yu, W. Zhang, J. Wang, and Y. Yu. 2017. SeqGAN: Sequence generative adversarial nets with policy gradient. In Proceedings of the Conference on Artificial Intelligence. 2852–2858.
- [106] Y. Li, Q. Pan, S. Wang, T. Yang, and E. Cambria. 2018. A generative model for category text generation. *Inf. Sci. (Ny)*. 450, C (2018), 301–315. DOI: 10.1016/j.ins.2018.03.050
- [107] C. Wang and D. Lillis. 2020. Classification for crisis-related tweets leveraging word embeddings and data augmentation. In *Proceedings of the TREC*. 8. Retrieved from https://trec.nist.gov/.
- [108] H. Queiroz Abonizio and S. Barbon Junior. 2020. Pre-trained data augmentation for text classification. In Proceedings of the Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 551–565. DOI: 10.1007/978-3-030-61377-8_38
- [109] V. Claveau, A. Chaffin, and E. Kijak. 2021. La génération de textes artificiels en substitution ou en complément de données d'apprentissage (Generating artificial texts as substitution or complement of training data). In Proceedings of the Actes de la 28e Conférence sur le Traitement Automatique des Langues Naturelles Lille. 37–49. http://arxiv.org/ abs/2110.13016.
- [110] R. Liu, G. Xu, C. Jia, W. Ma, L. Wang, and S. Vosoughi. 2020. Data Boost: Text data augmentation through reinforcement learning guided conditional generation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 9031–9041. DOI:10.18653/v1/2020.emnlp-main.726

- [111] I. S. Alec Radford, Jeffrey Wu, Rewon Child, David Luan, and Dario Amodei. 2020. Language models are unsupervised multitask learners. OpenAI Blog 1, May (2020), 1–7. Retrieved from https://github.com/codelucas/newspaper.
- [112] V. Sanh, L. Debut, J. Chaumond, and T. Wolf. 2019. DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. http://arxiv.org/abs/1910.01108.
- [113] K. M. Yoo, D. Park, J. Kang, S. W. Lee, and W. Park. 2021. GPT3Mix: Leveraging large-scale language models for text augmentation. In Proceedings of the Findings of Findings of the Association for Computational Linguistics, Findings of ACL. 2225–2239. DOI: 10.18653/v1/2021.findings-emnlp.192
- [114] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. 2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, vol. 2020-Dec. [Online]. Available: http://arxiv.org/abs/2005.14165.
- [115] K. Lee, K. Guu, L. He, T. Dozat, and H. W. Chung. 2021. Neural data augmentation via example extrapolation. http://arxiv.org/abs/2102.01335.
- [116] B. Ding, L. Liu, L. Bing, C. Kruengkrai, T. H. Nguyen, S. Joty, L. Si, and C. Miao. 2020. DAGA: Data augmentation with a generation approach for low-resource tagging tasks. In *Proceeding of the Conference on Empirical Methods in Natural Language Processing*. 6045–6057. DOI: 10.18653/v1/2020.emplp-main.488
- [117] E. Chang, X. Shen, D. Zhu, V. Demberg, and H. Su. 2021. Neural data-to-text generation with LM-based text augmentation. In Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics. 758–768. DOI: 10.18653/v1/2021.eacl-main.64
- [118] V. Kumar, H. Glaude, C. de Lichy, and W. Campbell. 2019. A closer look at feature space data augmentation for fewshot intent classification. In Proceedings of the 2nd Workshop on Deep Learning Approaches for Low-Resource Natural Language Processing. 1–10. DOI: 10.18653/v1/d19-6101
- [119] E. Schwartz, L. Karlinsky, J. Shtok, S. Harary, M. Marder, R. Feris, A. Kumar, R. Giryes, and A. M. Bronstein. 2018. Delta-encoder: An effective sample synthesis method for few-shot object recognition. http://arxiv.org/abs/1806. 04734.
- [120] C. Zhu, Y. Cheng, Z. Gan, S. Sun, T. Goldstein, and J. Liu. 2019. FreeLB: Enhanced adversarial training for natural language understanding. http://arxiv.org/abs/1909.11764.
- [121] P. L. Combettes and J. C. Pesquet. 2011. Proximal splitting methods in signal processing. Springer Optim. Its Appl. 49 (2011), 185–212. DOI: 10.1007/978-1-4419-9569-8_10
- [122] T. Goldstein, C. Studer, and R. Baraniuk. 2014. A field guide to forward-backward splitting with a FASTA implementation. http://arxiv.org/abs/1411.3406.
- [123] A. Shafahi et al. 2019. Adversarial training for free!. In *Proceedings of the Advances in Neural Information Processing Systems*, vol. 32.
- [124] D. Zhang, T. Zhang, Y. Lu, Z. Zhu, and B. Dong. 2019. You only propagate once: Accelerating adversarial training via maximal principle. In *Proceedings of the Advances in Neural Information Processing Systems*, vol. 32. Retrieved from https://github.com/a1600012888/YOPO-You-Only-Propagate-Once.
- [125] T. Miyato, S. I. Maeda, M. Koyama, K. Nakae, and S. Ishii. 2016. Distributional smoothing with virtual adversarial training. In Proceedings of the 4th International Conference on Learning Representations, ICLR 2016.
- [126] H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and T. Zhao. 2020. SMART: Robust and efficient fine-tuning for pre-trained natural language models through principled regularized optimization. In *Proceedings of the ACL*. 2177–2190. DOI:10. 18653/v1/2020.acl-main.197
- [127] D. Wang, C. Gong, and Q. Liu. 2019. Improving neural language modeling via adversarial training. In Proceedings of the 36th International Conference on Machine Learning, ICML 2019. 11387–11397.
- [128] X. Liu, H. Cheng, P. He, W. Chen, Y. Wang, H. Poon, and J. Gao. 2020. Adversarial training for large neural language models. 2020. http://arxiv.org/abs/2004.08994.
- [129] D. Liu, Y. Gong, J. Fu, Y. Yan, J. Chen, J. Lv, N. Duan, and M. Zhou. 2020. Tell me how to ask again: Question data augmentation with controllable rewriting in continuous space. In *Proceedings of the Conference on Empirical Methods* in Natural Language Processing. 5798–5810. DOI: 10.18653/v1/2020.emnlp-main.467
- [130] Z. Wan, X. Wan, and W. Wang. 2021. Improving grammatical error correction with data augmentation by editing latent representation. In *Proceedings of the COLING*. 2202–2212. DOI: 10.18653/v1/2020.coling-main.200
- [131] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio. 2016. Generating sentences from a continuous space. In Proceedings of the Computational Natural Language Learning. 10–21. DOI: 10.18653/v1/k16-1002
- [132] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. 2002. SMOTE: Synthetic minority over-sampling technique. J. Artif. Intell. Res. 16 (2002), 321–357. DOI:10.1613/jair.953
- [133] V. Verma, A. Lamb, C. Beckham, A. Najafi, I. Mitliagkas, D. Lopez-Paz, and Y. Bengio. 2019. Manifold mixup: Better representations by interpolating hidden states. In *Proceedings of the 36th International Conference on Machine Learning*. 11196–11205.

- [134] N. Tishby and N. Zaslavsky. 2015. Deep learning and the information bottleneck principle. In Proceedings of the IEEE Information Theory Workshop. DOI: 10.1109/ITW.2015.7133169
- [135] R. Shwartz-Ziv and N. Tishby. 2017. Opening the black box of deep neural networks via information. http://arxiv. org/abs/1703.00810.
- [136] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. 2018. MixUp: Beyond empirical risk minimization. In Proceedings of the 6th International Conference on Learning Representations.
- [137] L. Sun, C. Xia, W. Yin, T. Liang, P. Yu, and L. He. 2021. Mixup-Transformer: Dynamic data augmentation for NLP tasks. 3436–3440. DOI: 10.18653/v1/2020.coling-main.305
- [138] H. Guo, Y. Mao, and R. Zhang. 2019. Augmenting data with mixup for sentence classification: An empirical study. http://arxiv.org/abs/1905.08941.
- [139] H. Guo. 2020. Nonlinear mixup: Out-of-manifold data augmentation for text classification. 34th AAAI Conf. Artif. Intell. 34, 04 (2020), 4044–4051. DOI:10.1609/aaai.v34i04.5822
- [140] G. Jawahar, B. Sagot, and D. Seddah. 2020. What does BERT learn about the structure of language? In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. 3651–3657. DOI: 10.18653/v1/p19-1356
- [141] J. Chen, Z. Wang, R. Tian, Z. Yang, and D. Yang. 2020. Local additivity based data augmentation for semi-supervised NER. In Proceedings of the Conference on Empirical Methods in Natural Language Processing. 1241–1251. DOI:10. 18653/v1/2020.emnlp-main.95
- [142] N. Reimers and I. Gurevych. 2020. Sentence-BERT: Sentence embeddings using siamese BERT-networks. In Proceedings of the Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing. 3982–3992. DOI: 10.18653/v1/d19-1410
- [143] D. Hendrycks, N. Mu, E. D. Cubuk, B. Zoph, J. Gilmer, and B. Lakshminarayanan. 2019. AugMix: A simple data processing method to improve robustness and uncertainty. http://arxiv.org/abs/1912.02781.
- [144] S. Bonthu, A. Dayal, M. S. Lakshmi, and S. Rama Sree. 2022. Effective Text Augmentation Strategy for NLP Models. Springer, 521–531. DOI: 10.1007/978-981-16-4538-9_51
- [145] Y. Yan, R. Li, S. Wang, F. Zhang, W. Wu, and W. Xu. 2021. ConSERT: A contrastive framework for self-supervised sentence representation transfer. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing. 5065–5075. DOI: 10.18653/v1/ 2021.acl-long.393
- [146] Y. Yang, C. Malaviya, J. Fernandez, S. Swayamdipta, R. Le Bras, J. P. Wang, C. Bhagavatula, Y. Choi, and D. Downey. 2020. Generative data augmentation for commonsense reasoning. In *Proceedings of the Findings of the Association* for Computational Linguistics Findings of ACL: EMNLP. 1008–1025. DOI:10.18653/v1/2020.findings-emnlp.90
- [147] A. P. Parikh, O. Täckström, D. Das, and J. Uszkoreit. 2016. A decomposable attention model for natural language inference. In Conference on Empirical Methods in Natural Language Processing. 2249–2255. DOI: 10.18653/v1/d16-1244
- [148] A. Robinson, R. D. Cook, and S. Weisberg. 1984. Residuals and influence in regression. J. R. Stat. Soc. Ser. A, 147, 1 (1984), 108. DOI: 10.2307/2981746
- [149] P. W. Koh and P. Liang. 2017. Understanding black-box predictions via influence functions. In Proceedings of the 34th International Conference on Machine Learning. 2976–2987.
- [150] I. Chaturvedi, Y.-S. Ong, I. W. Tsang, R. E. Welsch, and E. Cambria. 2016. Learning word dependencies in text by means of a deep recurrent belief network. *Knowledge-Based Syst.* 108, C (2016), 144–154. DOI:10.1016/j.knosys.2016. 07.019
- [151] W. Zhao, H. Peng, S. Eger, E. Cambria, and M. Yang. 2020. Towards scalable and reliable capsule networks for challenging NLP applications. In Proceedings of the ACL57th Annual Meeting of the Association for Computational Linguistics. 1549–1559. DOI:10.18653/v1/p19-1150
- [152] Y. Ma, H. Peng, T. Khan, E. Cambria, and A. Hussain. 2018. Sentic LSTM: A hybrid network for targeted aspect-based sentiment analysis. *Cognit. Comput.* 10, 4 (2018). DOI: 10.1007/s12559-018-9549-x.
- [153] P. K. Jain, W. Quamer, V. Saravanan, and R. Pamula. 2022. Employing BERT-DCNN with sentic knowledge base for social media sentiment analysis. J. Ambient Intell. Humaniz. Comput. 1, 3 (2022) 3. DOI:10.1007/s12652-022-03698-z
- [154] M.-A. Kaufhold. 2021. Information Refinement Technologies for Crisis Informatics: User Expectations and Design Principles for Social Media and Mobile Apps. Springer Fachmedien Wiesbaden, 2021.
- [155] J. Bragg, A. Cohan, K. Lo, and I. Beltagy. 2021. FLEX: Unifying evaluation for few-Shot NLP. 2021. http://arxiv.org/ abs/2107.07170.
- [156] S. Gehrmann, T. Adewumi, K. Aggarwal, P. S. Ammanamanchi, A. Anuoluwapo, A. Bosselut, K. R. Chandu, M. Clinciu, D. Das, K. D. Dhole, W. Du, E. Durmus, O. Dušek, C. Emezue, V. Gangal, C. Garbacea, T. Hashimoto, Y. Hou, Y. Jernite, H. Jhamtani, Y. Ji, S. Jolly, M. Kale, D. Kumar, F. Ladhak, A. Madaan, M. Maddela, K. Mahajan, S. Mahamood, B. P. Majumder, P. H. Martins, A. McMillan-Major, S. Mille, E. van Miltenburg, M. Nadeem, S. Narayan, V. Nikolaev, R. A. Niyongabo, S. Osei, A. Parikh, L. Perez-Beltrachini, N. R. Rao, V. Raunak, J. D. Rodriguez, S. Santhanam, J. Sedoc, T. Sellam, S. Shaikh, A. Shimorina, M. A. S. Cabezudo, H. Strobelt, N. Subramani, W. Xu, D. Yang, A. Yerukola, and

ACM Computing Surveys, Vol. 55, No. 7, Article 146. Publication date: December 2022.

J. Zhou. 2021. The GEM Benchmark: Natural language generation, its evaluation and metrics. In *Proceedings of the* 1st Workshop on Natural Language Generation, Evaluation, and Metrics. 96–120. DOI:10.18653/v1/2021.gem-1.10

- [157] K. D. Dhole, V. Gangal, S. Gehrmann, A. Gupta, Z. Li, S. Mahamood, A. Mahendiran, S. Mille, A. Srivastava, S. Tan, T. Wu, J. Sohl-Dickstein, J. D. Choi, E. Hovy, O. Dusek, S. Ruder, S. Anand, N. Aneja, R. Banjade, L. Barthe, H. Behnke, I. Berlot-Attwell, C. Boyle, C. Brun, M. A. S. Cabezudo, S. Cahyawijaya, E. Chapuis, W. Che, M. Choudhary, C. Clauss, P. Colombo, F. Cornell, G. Dagan, M. Das, T. Dixit, T. Dopierre, P.-A. Dray, S. Dubey, T. Ekeinhor, M. Di Giovanni, R. Gupta, R. Gupta, L. Hamla, S. Han, F. Harel-Canada, A. Honore, I. Jindal, P. K. Joniak, D. Kleyko, V. Kovatchev, K. Krishna, A. Kumar, S. Langer, S. R. Lee, C. J. Levinson, H. Liang, K. Liang, Z. Liu, A. Lukyanenko, V. Marivate, G. de Melo, S. Meoni, M. Meyer, A. Mir, N. S. Moosavi, N. Muennighoff, T. S. H. Mun, K. Murray, M. Namysl, M. Obedkova, P. Oli, N. Pasricha, J. Pfister, R. Plant, V. Prabhu, V. Pais, L. Qin, S. Raji, P. K. Rajpoot, V. Raunak, R. Rinberg, N. Roberts, J. D. Rodriguez, C. Roux, V. P. H. S., A. B. Sai, R. M. Schmidt, T. Scialom, T. Sefara, S. N. Shamsi, X. Shen, H. Shi, Y. Shi, A. Shvets, N. Siegel, D. Sileo, J. Simon, C. Singh, R. Sitelew, P. Soni, T. Sorensen, W. Soto, A. Srivastava, K. A. Srivatsa, T. Sun, M. V. T, A. Tabassum, F. A. Tan, R. Teehan, M. Tiwari, M. Tolkiehn, A. Wang, Z. Wang, G. Wang, Z. J. Wang, F. Wei, B. Wilie, G. I. Winata, X. Wu, W. Wydmański, T. Xie, U. Yaseen, M. Yee, J. Zhang, and Y. Zhang. 2021. NL-Augmenter: A framework for task-sensitive natural language augmentation. *Nivranshu Pasricha* 20 (2021), 75. http://arxiv.org/abs/2112.02721.
- [158] Z. Papakipos and J. Bitton. 2022. AugLy: Data augmentations for robustness. arXiv:2201.06494. Retrieved from http: //arxiv.org/abs/2201.06494.
- [159] I. Solaiman, M. Brundage, J. Clark, A. Askell, A. Herbert-Voss, J. Wu, A. Radford, and J. Wang. 2019. Release strategies and the social impacts of language models. http://arxiv.org/abs/1908.09203.

Received 9 July 2021; revised 9 May 2022; accepted 3 June 2022