

Modeling Uncertainty in Personalized Emotion Prediction with Normalizing Flows

Piotr Miłkowski^{†1}, Konrad Karanowski^{†1}, Patryk Wielopolski¹, Jan Kocoń¹, Przemysław Kazienko¹, Maciej Zięba^{1,2}

¹ Department of Artificial Intelligence, Wrocław University of Science and Technology, Poland

² Tooploox, ul. Tęczowa 7, 53-601 Wrocław, Poland

{piotr.milkowski,konrad.karanowski,patryk.wielopolski,jan.kocoon,kazienko,maciej.zieba}@pwr.edu.pl

Abstract—Designing predictive models for subjective problems in natural language processing (NLP) remains challenging. This is mainly due to its non-deterministic nature and different perceptions of the content by different humans. It may be solved by Personalized Natural Language Processing (PNLP), where the model exploits additional information about the reader to make more accurate predictions. However, current approaches require complete information about the recipients to be straight embedded. Besides, the recent methods focus on deterministic inference or simple frequency-based estimations of the probabilities. In this work, we overcome this limitation by proposing a novel approach to capture the uncertainty of the forecast using conditional Normalizing Flows. This allows us to model complex multimodal distributions and to compare various models using negative log-likelihood (NLL). In addition, the new solution allows for various interpretations of possible reader perception thanks to the available sampling function. We validated our method on three challenging, subjective NLP tasks, including emotion recognition and hate speech. The comparative analysis of generalized and personalized approaches revealed that our personalized solutions significantly outperform the baseline and provide more precise uncertainty estimates. The impact on the text interpretability and uncertainty studies are presented as well. The information brought by the developed methods makes it possible to build hybrid models whose effectiveness surpasses classic solutions. In addition, an analysis and visualization of the probabilities of the given decisions for texts with high entropy of annotations and annotators with mixed views were carried out.

Index Terms—artificial neural networks, natural language processing, human profile modelling, probabilistic technique

I. INTRODUCTION

Human affective states, including emotions, strongly depend on the individual, the stimulant eliciting them, and the associated context [1]. Therefore, the reasoning of a person's perception based on machine learning bears a significant degree of uncertainty. It refers to the reaction to any content, including text reading. We can say that disagreements in human textual inferences are inherent [2]. Most solutions to subjective problems in natural language processing (NLP), like recognition of emotions, hate speech, sarcasm, sense of humor, sentiment, and many others, rely on generalized perspectives. They consider only text and its single generalized interpretation. Then, the commonly used solution is to simplify multiple distinct views, i.e., annotations provided by many annotators using majority voting or other methods to achieve

a sole perception. Overall, we can identify two sources of uncertainty: (1) humans, who are unsure and imprecise in their annotations (this is a hidden factor), and (2) a community of annotators. The latter refers to discrepancies between people in understanding the problem, and perception of a given text [3]–[5]. The standard measures for inter-rater agreement are Krippendorff's alpha [6] or Fleiss' kappa [7]. However, they provide only a single value characterizing the set of all annotations for all texts. Yet another (3) source of uncertainty: the trained model itself. It means that the model is not capable of precisely learning about concepts (what is *joy* or *hate speech*?) and relations from the available learning samples. This leads to errors and proximate reasoning. Simultaneously, emotions can be considered multidimensional objects, which requires multi-task learning [8] and further complicates the problem of uncertainty modeling. Most of the proposed approaches for subjective modeling in the NLP domain focus on deterministic predictions. In this work, we propose to enrich the family of emotional methods by introducing Emotional Normalizing Flow – an entirely probabilistic framework that utilizes conditional Normalizing Flows to model uncertainty. We postulate to represent the considered tasks as multivariate regression problems and represent the distribution of the outputs with conditional flows. This approach allows us to model complex multimodal distributions of multidimensional outputs. The experiments and validation were carried out on emotion detection (ten tasks) and hate speech (two tasks). We examine various choices of flow models and compare their performance with the mixture of Gaussians, showing the superiority of Emotional Normalizing Flow compared to the selected baseline. Moreover, we show that incorporating personalization into our model leads to better distribution adjustment measured with negative log-likelihood (NLL) value.

To summarize, the contributions of this work are as follows:

- We introduce a novel approach for probabilistic modeling in subjective NLP-based problems;
- We examine the impact of personalization on the quality of the model and show that in most of the considered experimental cases, additional information about the reader leads to better probability adjustment;
- We show that our approach outperforms the standard baseline that utilizes a mixture of Gaussians;

[†]These authors contributed equally to this work.

- We propose a hybrid approach utilizing Normalizing Flows and personalization that outperform previous models.

II. RELATED WORK

Initial work on emotion recognition in the text was based mainly on frequency analysis of words defined in lexicons of emotions [9], [10]. These lexicons contained words with assigned categories of basic emotions, e.g., joy, anger, sadness [11], [12]. Emotions occurring most frequently at the lexical level were then assigned to the entire text. With the development of text classification methods based on machine learning, datasets containing texts manually annotated with emotions began to emerge [13]–[20]. Due to annotators’ subjective perception of emotions, and thus low inter-annotator agreement, it was common to assign emotion labels to text based on majority voting [13], [18]. Based on such prepared data, text classification models were trained. Initially, such models as SVM [21], BiLSTM, and GRU [22] were used. Currently, transformer-based models such as BERT perform best in the task of emotion recognition [18], [23]. The aforementioned approaches require data for which the inter-annotator agreement is high. However, there are some data sets such as Wiki-Detox [24], Sentimenti [16] or Measuring Hate Speech [25], which contain an annotator identifier linked to their affective annotation. They also include multiple annotations for a given text from multiple annotators. For such data, new personalized approaches have recently been developed, in which the context of the annotator is taken into account in the model learning process [26]–[36]. This makes it possible, for example, to answer the question of what emotions a particular text evokes for a particular user. Recent method proposals also focus on neuro-symbolic approaches to explain decisions made [37], usage of large-scale pre-trained language model (PLM) for prompt-based classification tasks such as sentiment analysis and emotion detection [38], using recently popular large language models (LLMs) [39], or methods of complex persona attribute extraction [40]. However, the methods mentioned above do not model the uncertainty associated with the community’s subjective perception of emotions and the degree of indecision of the annotators themselves.

In this paper, to model uncertainty described in the Introduction we adapt the concept of Normalizing Flows. The best-known Normalizing Flow models such as NICE [41], RealNVP [42], MAF [43], and CNF [44] were originally used for density estimation and image generation tasks. These models were further extended and used as components for more sophisticated tasks or even for other domains of applications. In Computer Vision, there were proposed models such as RegFlow [45] for probabilistic future location prediction, Flow Plugin Network [46], PluGeN [47], and StyleFlow [48] models for conditional image generation. For the tabular data, recently, TreeFlow [49] was proposed that utilizes a combination of tree-based models with conditional Normalizing Flows to estimate uncertainty for uni- and multi-variate regression problems. In terms of Natural Language

Processing and Normalizing Flows, only Discrete Flow [50] was proposed to model character-level datasets using Normalizing Flows dedicated to the discrete data. To the best of our knowledge, no probabilistic approach has been proposed to model distributions of uncertainty in personalized natural language processing, and our Emotional Normalizing Flow is the first probabilistic model proposed for multi-task prediction of personalized emotions.

III. BACKGROUND

a) Generalized and Personalized Approach to Subjective NLP Problems.: In the classic approach to the task of text classification or regression, we assume a training set of the form $\mathcal{D} = \{(\mathbf{t}_i, \mathbf{y}_i)\}_{i=1}^N$, where $\mathbf{t}_i \in \mathcal{T}$ is the i -th text document and \mathbf{y}_i is its annotation. However, many NLP tasks, such as recognizing emotions in a text or detecting hate speech, can be subjective because each person perceives these phenomena. This leads to a situation when we can have more than one annotation \mathbf{y} for the same text \mathbf{t} , as different people may annotate the same texts differently. Therefore, a training set is in the form of $\mathcal{D} = \{(\mathbf{t}_i, \mathbf{p}_i, \mathbf{y}_i)\}_{i=1}^N$, where \mathbf{y}_i is the annotation given by person $\mathbf{p}_i \in \mathcal{P}$ for text $\mathbf{t}_i \in \mathcal{T}$.


One approach to subjective tasks in NLP is the so-called *generalized*. It assumes that the model predicts the result based solely on the text and returns the same prediction for every user. Generalized models usually consist of two parts: **text encoder** (language model), which creates **text representation** e_t and **classifier** or **regressor** (usually fully-connected layer) that gives prediction \hat{y} . However, recent studies [28], [51], [52] show that this approach should not be considered correct, as adding information about the annotator significantly improves model quality and yields better results. The approach that combines information about the text and the human is so-called *personalized*. Compared to the generalized, personalized model adds another component called **profile extractor**, that creates **human representation** e_p . The comparison of generalized and personalized approaches is shown in Fig 1.

There are few existing architectures [28], [51] utilizing this fact. Still, all of them are deterministic, meaning none model uncertainty as a direct optimization of negative log-likelihood.


b) Normalizing Flows.: Normalizing Flows [53] are a class of generative models that enables estimation of the uncertainty of prediction thanks to the access to log probability function and thus enable direct optimization of negative log-likelihood (NLL). The goal of the model is to transform base distribution $p_U(\mathbf{u})$ (usually Gaussians with independent components) to the complex distribution of the data $p_Y(\mathbf{y})$ using a series of K invertible functions that can be written as $\mathbf{u} = \mathbf{f}_K \circ \dots \circ \mathbf{f}_1(\mathbf{y})$. For that purpose, Normalizing Flows utilize the change-of-variable formula and then the NLL \mathbf{y} is given by

$$\log p_Y(\mathbf{y}) = \log p_U(\mathbf{u}) - \sum_{k=1}^K \log \left| \det \frac{\partial \mathbf{f}_k}{\partial \mathbf{z}_{k-1}} \right|. \quad (1)$$

To specify the exact Normalizing Flow model, we need to define transformations $\mathbf{f}_1, \dots, \mathbf{f}_K$. Here, multiple models were



(a) Generalized deterministic model.



(b) Personalized deterministic model.

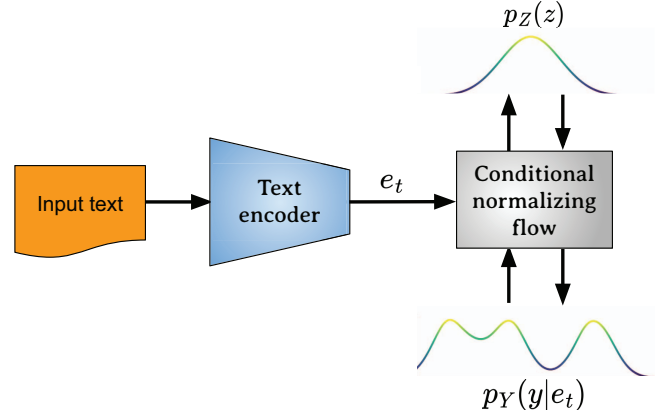
Fig. 1: Comparison of (a) generalized and (b) personalized deterministic models. (a) consists of two parts: a text encoder (language model) that creates text embedding e_t and a classifier or regressor (mostly fully-connected layer) that provides prediction \hat{y} . This approach is not considered suitable for subjective NLP tasks, like emotion recognition, because it does not respect the individual perception of the text. Model (b) fixes this problem by adding a profile extractor in the form of user representation e_p . It allows human individual characteristics to be included in the inference process. Both models are deterministic, giving us limited, spotlight information about subjective tasks.

proposed such as NICE [41], RealNVP [42], MAF [43] or Continuous Normalizing Flows [44].

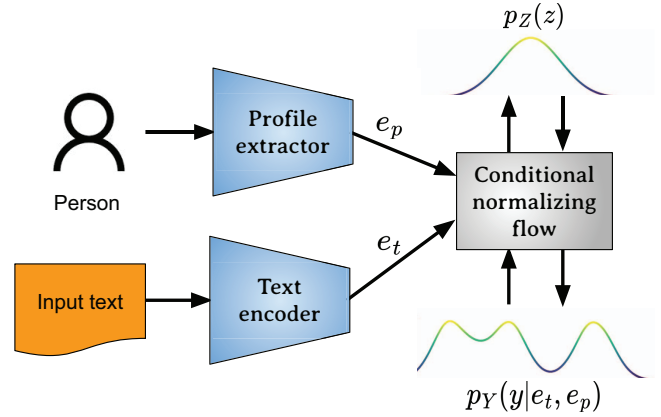
IV. OUR APPROACH

In this section, we introduce Emotional Normalizing Flow - the probabilistic model for subjective uncertainty modeling in the NLP domain. The general schema of the proposed approach is provided in Fig. 2. The model is composed of *Profile extractor* that is responsible for creating the representation of the person, e_p , and *Text encoder* that creates embedding e_t directly from the input text. Both components can be represented by various models (trainable and fixed), and we elaborate on this further in this section.

The extracted vectors e_p and e_t are further delivered to the conditional flow represented by the complex transformation function $f(\cdot)$. The role of the function is to transform multivariate regression outputs \mathbf{y} to \mathbf{z} that represents the variable in the base space, assuming given vectors, e_p and e_t . Formally, we have $\mathbf{z} = \mathbf{f}(\mathbf{y}, e_p, e_t)$, where \mathbf{f} is invertible with respect to



(a) Generalized flow-based probabilistic model.



(b) Personalized flow-based probabilistic model.

Fig. 2: Comparison of (a) generalized and (b) personalized flow-based probabilistic models. Model (a), as in the case of generalized deterministic, uses only information about the text. However, unlike it, it models the conditional probability distribution $p_Y(\mathbf{y}|e_t)$ using Normalizing Flow. Model (b) extends the concept of the personalized deterministic model in a similar way to (a) but it exploits representations of both the text e_t and user e_p to model conditional probability distribution $p_Y(\mathbf{y}|e_p, e_t)$ for the subjective output predictions representing emotions.

\mathbf{y} , $\mathbf{y} = \mathbf{f}^{-1}(\mathbf{z}, e_p, e_t)$. Moreover, the complex transformation \mathbf{f} can be decomposed into a sequence of simple functions, $\mathbf{z} = \mathbf{f}_K \circ \dots \circ \mathbf{f}_1(\mathbf{y}, e_p, e_t)$, where the K is number discrete transformations. With such assumptions, the probability distribution for \mathbf{y} that represents the distribution over the regression outputs can be calculated using the formula:

$$p_Y(\mathbf{y}|e_p, e_t) = p_Z(\mathbf{z}) \cdot \prod_{k=1}^K \left| \det \frac{\partial \mathbf{f}_k}{\partial \mathbf{z}_{k-1}} \right|, \quad (2)$$

where $\mathbf{z}_0, \dots, \mathbf{z}_K$ are intermediate steps after discrete transformations, assuming $\mathbf{z}_0 := \mathbf{y}$, and $\mathbf{z}_K := \mathbf{z}$. $p_Z(\mathbf{z})$ is the assumed base distribution for \mathbf{z} with the known density

function, usually represented by Gaussian. Consequently, we have direct access to the density function for that conditional distribution. Therefore we can calculate the likelihood function for a set of input-output pairs to evaluate the quality of the model. We can sample an infinite number of output values assuming given inputs and interpret the results.

The proposed model can quickly adapt to the problems without personalization, simply skipping \mathbf{e}_p conditioning in the flow. Our approach is independent of the conditional Normalizing Flow type, and we experimentally compare the performances of the most popular models. We follow the methodology of incorporating conditional components described in [46].

a) Profile extractor: Vector \mathbf{e}_p contains information about the user specific to the personalization architecture used. This can include information such as the deviation of responses from the majority voice, metadata about the user, user identifier [28], the correlation of the text’s context with historical evaluations, or other features unique to the recipient of the text. It also can be randomly initialized and tuned during the learning process by backpropagation [51].

b) Text encoder: In the case of \mathbf{e}_t vector, text representation is implemented using Transformer language models. An attentional weight is assigned for a given text input, divided into individual tokens. The assigned values are then used to calculate the weighted sum of the resulting vectors [54]. It is possible to fine-tune the language model using the loss function of the final model.

c) Training the model: To trained Emotional Normalizing Flow we use the dataset $\mathcal{D} = \{(t_n, p_n, y_n)\}_{n=1}^N$, composed of t_n textual input, p_n features of the person, and corresponding subjective annotation y_n given by the person p_n for text t_n . We train our model directly by optimizing the negative log-likelihood function:

$$\mathcal{L} = - \sum_{n=1}^N \log p_Y(y_n | \mathbf{e}_{n,p}, \mathbf{e}_{n,t}), \quad (3)$$

where $\mathbf{e}_{n,p}$ is a vector, that represents profile of the person p_n , and $\mathbf{e}_{n,t}$ is an embedding of the text t_n . The model can be trained in a two-stage mode or end-to-end paradigm depending on the form of *Profile extractor* and *Text encoder*. In the first case, the embeddings \mathbf{e}_p and \mathbf{e}_t are extracted in the first stage, and parameters of the flow are trained while optimizing \mathcal{L} . Alternatively, suppose the *Profile extractor* or *Text encoder* are represented by differentiable architectures. In that case, the entire system can be optimized end-to-end, directly minimizing the negative log-likelihood function.

V. EXPERIMENTS

In this subsection, we evaluate our approach on a set of challenging datasets, investigating the impact of adding contextual information about the person in the model. Moreover, we compare flow-based probabilistic models to a simple Gaussian Mixture Model. Then, we compare our solution to the deterministic models using sampling from flow and

discretization. Finally, we mix deterministic and probabilistic approaches to create a hybrid model.

A. Datasets

a) Wikipedia-Detox: The Wikipedia Detox project has created a crowd-sourced dataset that contains one million annotations covering 100,000 discussions of page edits on Wikipedia [24]. These were often filled with toxic statements, verbal aggression, and even personal attacks. Each comment was annotated by about ten annotators provided by the Crowdflower service.

The collection containing toxic statements consists of 160,000 texts. It includes a binary determination of toxicity (where: 0 = non-toxic, 1 = toxic), as well as a rating from -2 to 2 (where: 2 = very healthy, 0 = neutral, and -2 = very toxic).

Sets for personal attacks and verbal aggression consist of 100,000 of the same comments. In addition to the binary marks for aggression (0 = neutral or friendly comment, 1 = aggressive or attacking), aggression is put on a scale analogous to toxicity from -2 to 2 (where: 2 = very friendly, 0 = neutral, and -2 = very aggressive). Personal attacks are divided into types: quoting, recipient, third party, or another type of attack. In addition to texts shared between these collections, the same applies to annotators. Thus, we can use knowledge from one collection to benefit from it in another or a collective approach. Those willing to participate in the study also completed questionnaires so that we have demographic information about them available.

b) Emotion Simple: This collection consists of 100 texts marked on 10 scales by 5,365 annotators [55]. Texts are opinions posted on websites. This gives 53.65 annotations per text and 1.69 markings from a single user. The texts were rated for eight basic emotions (sadness, anticipation, joy, fear, surprise, disgust, trust, and anger) and emotional arousal on a scale from 0 to 4 for each dimension. In addition, the tenth aspect rated is the valence expressed on a scale from -3 to 3 (where -3 = negative, 0 = neutral, and 3 = positive). In the set of individuals with two marks, those with three or more annotations also appear.

c) Emotion Meanings: In [56], a huge collection containing 6,000 assessed word collocations was prepared and published. It contains dimensions and scales analogous to the Emotion Simple collection – the basic emotions from Robert Plutchik’s Wheel of Emotions [57].

The scale of the collection makes it one of the most interesting and, simultaneously, the most difficult for personalizing emotion detection. It has 303,143 annotations from 16,101 people who participated in the study. Each collocation has been evaluated 50.67 times, and a single annotator has an average of 18.83 annotations.

The difficulty in working with these data is also because these are not full-fledged textual statements containing context but just two words. An example item from the collection: "colorful beads". Annotator data include information such as

gender, age, education, size of residence, relationship status, income, or political views.

B. Setups

The dataset was divided into training, validation, and testing splits. Users and texts were not mixed between sets to bring the evaluation as close to the real-world scenario. Each experiment consisted of 10-fold cross-validation, and obtained results were averaged. Statistical significance tests were performed: t-test with Bonferroni correction to address the problem of multiple comparisons. In the tables within the rows, comparisons were made between models without and with personalization. Bold indicates the best result, and underline indicates the absence of statistically significant differences for each dataset. Within the “Type“ column, the best probabilistic model type or no significant difference between the two was similarly marked for each dataset separately.

a) *Baselines.*: We have three reference points. To check the impact of personalization, we compared personalized models with a baseline that uses only textual information (TXT-Baseline); it is a generalized approach. To investigate the impact of normalizing flows, we compared them with a Gaussian Mixture Model to have a reference point in the form of another, less complex probabilistic method. Finally, we compared our method with deterministic approaches.

b) *Models for conditional normalizing flows.*: In our experimental evaluation, we consider Emotional Normalizing Flow with various types of conditional normalizing flows. For single-dimensional datasets: *Wikipedia Detox: Toxicity*, *Wikipedia Detox: Aggression* and *Wikipedia Detox: Attack*, we used MAF (**maf**) and CNF (**cnf**). For multi-dimensional *Emotions Meanings* and *Emotions Simple*, we used two extra flows: RealNVP (**real_nvp**) and NICE (**nice**). We compared the results against the baseline that uses mixtures of Gaussians to model the probability (**gmm**).

c) *Models for personalization.*: We investigate three approaches to respect the personalization context: OneHot, HuBi-Formula, and HuBi-Medium [51]. They are confronted with TXT-Baseline (generalized, non-personalized) that does not contain any information about the annotator. We exploit LaBSE [58] as a language model in every experiment.

C. Experimental scenarios

a) *Experiment 1 - Comparison of generalized and personalized solutions in the probabilistic approach.*: The first approach verifies the performance of Emotional Normalizing Flow with fixed hyperparameters on multiple data sets and tasks: *Wikipedia Detox: Toxicity*, *Wikipedia Detox: Aggression*, *Wikipedia Detox: Attack*, *Emotion Meanings* and *Emotion Simple*. We also verified the ability of the proposed Emotional Normalizing Flow to transfer knowledge between thematically similar multidimensional text labels. For this purpose, the *Wikipedia Detox: Aggression* and *Wikipedia Detox: Attack* datasets were joined, as they contain annotations for the same texts performed by the same annotators. As a result, we obtained a dataset with multi-dimensional labels. This

experiment aimed to examine the effect of personalization models on the prediction of probability distributions, thus verifying whether the additional information provided to the model reduces its uncertainty and comparing Normalizing Flows to Gaussian Mixture Model.

b) *Experiment 2 - Investigating the effect of hyperparameters selected for personalization and Emotional Normalizing Flow methods on the most difficult dataset.*: The second approach was to verify the maximum possible reduction of model uncertainty by tuning the model hyperparameters to a given set and checking which normalizing-flow model obtained the best results. Due to limited resources, we decided to perform this experiment using only *Emotion Meanings* dataset. The parameters that we tuned were: the number of hidden features, number of layers, number of blocks per layer, dropout probability, batch normalization within layers, batch normalization between layers, learning rate, the size of hidden layers used to prepare user embeddings, and the size of the output of these embeddings.

c) *Experiment 3 - Comparison of probabilistic and deterministic approaches.*: To compare with classical methods [8], which are deterministic, it was necessary to prepare conversions of the Emotional Normalizing Flow output to the form of exact values. Included in the body of the paper is the application of two best normalizing flows (RealNVP and CNF) for multidimensional datasets (*Aggression & Attack* [classification task] and *Emotion Simple* [regression task]).

For the first type of task, each text or text-user pair was sampled using an iterative method. In the preparation step, we increased the number of samples in the test part of the dataset so that the value from 0 to 1 with a step of 0.1 for the class was tested as a possible context. Iterations were done twice for values of 0 and 1 in the opposite class. Next, an exponential was applied to the 44 probabilities of the resulting sample (22 per class for each text). Within the values for the opposite sampling, (e.g., [0.5, 0] and [0.5, 1]) of a given dimension were summed, and then for each stopper (0.0, 0.1, ..., 1.0) divided by the sum of all values for the dimension. If the probability mass prevailed on the side from 0.0 to 0.5, it was considered that the class was not assigned and vice versa for the other part of the axis.

It was impossible for a 10-dimensional set for the regression task to sample each possible dimension in all values separately because of the number of possible combinations. Each item from the test subset was replicated 100 times containing random real values from 0 to 1 in each class. Majority voting was then conducted to determine the most likely response for the scale of each dimension. In the collection, each dimension had a value analogous to the slider setting during annotation. For this reason, the task was treated as an ordinary regression, and the resulting values were rounded to the nearest possible position. This assumption was used for both values from the deterministic and probabilistic approaches.

d) *Experiment 4 - Hybrid approach (utilizing knowledge from the text and uncertainty modeling).*: The combined approach, hereafter referred to as hybrid, was done in two

steps. In the first, the learned Emotional Normalizing Flow models were sampled in the same way as in Experiment 3, but for all the texts in the collection. Then, the network input was extended to the deterministic model with an additional feature. A vector containing the resulting probabilities for each text was entered along with its embeddings and, in the case of approaches with personalization, the user profile. This vector contained all the values from the sampling, and no additional mathematical operations were performed on it.

D. Results

a) *Results of Experiment 1.:* The first experiment proved that adding personalization reduces the uncertainty of probabilistic models, Tab. I. For *Wikipedia Detox* datasets (*Aggression*, *Attack* and *Toxicity*), all personalized models received significantly lower negative log-likelihood values compared to the non-personalized TXT-Baseline. For all three tasks, the best architecture was HuBi-Medium combined with CNF. For *Aggression & Attack* dataset, personalization improved most cases' results. The best results were obtained by OneHot combined with RealNVP and HuBi-Formula combined with CNF. In the case of *Emotion Simple* and *Emotion Meaning*, personalization also reduced model uncertainty in most of the cases. For both datasets, the best results were obtained by the HuBi-Medium model combined with RealNVP. It is worth noting that compared to Gaussian Mixture Model, Normalizing Flows always obtain lower negative log-likelihood values. It suggests that target variables, i.e., emotions, have complex distributions, and using a simple probabilistic approach is not enough.

b) *Results of Experiment 2.:* In the second experiment, we carried out hyperparameter tuning on the most challenging dataset: *Emotion Meanings*, Tab. II. All possible combinations of hyperparameters were considered when performing the grid search. The results seem to confirm earlier speculations about MAF's better ability to deal with multidimensional problems compared to other approaches. Moreover, none of the variants indicated the benefit of using text alone as input.

c) *Results of Experiment 3.:* In the third experiment, we compared results obtained by deterministic models and Emotional Normalizing Flow, Tab. III. It was carried out on two datasets: combined *Wikipedia Detox: Aggression & Attack* and *Emotions Simple*. We also decided to use only two Normalizing Flow Models that performed the best on both of these datasets: RealNVP and CNF.

In the case of *Aggression & Attack* dataset, the results obtained by deterministic models were better for every architecture, including the non-personalized one. In the case of *Emotion Simple* dataset, the results obtained by probabilistic models significantly outperformed deterministic models. The best model was a combination of HuBi-Medium and CNF. This result seems to confirm that the probabilistic approach is especially effective on complex multi-dimensional tasks.

d) *Results of Experiment 4.:* In the fourth experiment, we mixed the deterministic approach with the probabilistic, to create a hybrid model, Tab. IV. In both *Aggression & Attack*

Dataset	Type	TXT-Baseline	OneHot	HuBi-Formula	HuBi-Medium
Toxicity	maf	0.0702	-0.0197	-0.0947	0.0053
	cnf	0.1231	-0.0707	-0.1202	-0.1378
	gmm	0.6422	0.6164	0.5829	0.5072
Aggression	maf	0.1705	0.1695	0.0526	0.0859
	cnf	0.1685	0.0978	0.0180	-0.0431
	gmm	0.8948	0.7783	0.7841	0.7446
Attack	maf	0.1669	0.1180	-0.0229	-0.0250
	cnf	0.1474	0.0811	-0.0427	-0.0950
	gmm	0.7512	0.7318	0.7105	0.6911
Aggression & Attack	maf	-1.3788	N/A	-0.3966	-0.3834
	nice	-0.8678	-1.1281	-1.0524	-1.0914
	real_nvp	-3.3482	-3.6181	-3.0235	-1.7208
	cnf	-3.5113	-2.7339	-3.7002	-2.1357
	gmm	3.1729	2.4028	2.5673	2.6858
Emotion Meanings	maf	0.5337	-0.0135	0.8525	0.1936
	nice	-1.9099	-1.8707	-2.0283	-1.4792
	real_nvp	-5.4775	-2.9377	-5.5189	-5.6377
	cnf	-3.7186	-1.9640	-3.4632	-4.8458
	gmm	5.9559	5.4858	4.8034	4.4719
Emotion Simple	maf	1.9130	2.6398	2.6393	1.9314
	nice	2.4254	1.8163	2.3502	1.9613
	real_nvp	2.5583	1.7845	2.3726	1.4355
	cnf	2.1220	1.8197	2.3347	1.9702
	gmm	4.2706	3.7496	4.2312	4.0910

TABLE I: Experiment 1 - negative log-likelihood values for all datasets, without hyperparameter tuning.

and *Emotion Simple* tasks, the results obtained by the hybrid approach outperformed previous methods by a large margin. For *Aggression & Attack* dataset, the best model turned out to be HuBi-Medium with CNF. For *Emotion Simple* dataset, HuBi-Medium with both RealNVP and CNF performed comparably well. The results of this experiment prove, that adding information about the model uncertainty makes big difference in the inference process, and helps to predict better for difficult and subjective tasks.

Dataset	Type	TXT-Baseline	OneHot	HuBi-Formula	HuBi-Medium
Emotion Meanings	maf	-11.5380	-14.0785	-10.6476	-12.1934
	nice	4.2167	-2.0243	-1.0978	-1.9208
	real_nvp	-2.3509	-4.5813	-5.2833	-7.0285
	cnf	-1.9623	-3.9381	-5.2247	-4.9381
	gmm	12.6564	9.6047	7.8908	8.4545

TABLE II: Experiment 2 - negative log-likelihood values for Emotion Meanings dataset, with hyperparameter tuning.

VI. CONCLUSIONS

In this paper, we proposed a novel Emotional Normalizing Flow approach to personalized NLP that opens up new perspectives on predicting reader behavior in a non-deterministic way. From the perspective of psychology and the variability of emotion sensation over time, the problem of emotion recognition is one of the most difficult and subjective tasks facing NLP. People do not perceive their emotions as zero-one, and most of the attempts so far classified their feelings in this way. The presented probabilistic approach based on normalizing flows provides more complex information about the uncertainty and diversity of possible emotional states. A comparative analysis of models for emotion recognition without and with personalization indicated that new methods are also effectively applicable in a non-deterministic setup. The generalized, non-personalized solution generates a completely different concentration of probability mass, directed toward a

Dataset	Method	TXT-Baseline	OneHot	HuBi-Formula	HuBi-Medium
Aggression & Attack [Macro F-1]	deterministic	0.5874	0.5954	0.7354	0.7847
	discrete(real_nvp)	0.4829	0.4682	0.5860	0.6788
	discrete(cnf)	0.5189	0.4738	0.6397	0.7374
Emotion Simple [R2]	deterministic	0.3936	0.5403	0.5574	0.5822
	discrete(real_nvp)	0.4472	0.6535	0.6582	0.6835
	discrete(cnf)	0.4428	0.6274	0.6685	0.7005

TABLE III: Experiment 3 - Comparison of classification and regression using the classical deterministic method and the result of sampling Emotional Normalizing Flow. Macro F-1 score for *Aggression & Attack* and R-squared for *Emotion Simple* datasets.

Dataset	Method	TXT-Baseline	OneHot	HuBi-Formula	HuBi-Medium
Aggression & Attack [Macro F-1]	deterministic	0.5874	0.5954	0.7354	0.7847
	hybrid(real_nvp)	0.8479	0.8254	0.8233	0.8743
	hybrid(cnf)	0.8693	0.9052	0.8400	0.9553
Emotion Simple [R2]	deterministic	0.3936	0.5403	0.5574	0.5822
	hybrid(real_nvp)	0.4722	0.7149	0.7237	0.7376
	hybrid(cnf)	0.4867	0.6899	0.69223	0.7388

TABLE IV: Experiment 4 - Comparison of the classical deterministic approach and hybrid models, which in addition use probabilistic knowledge. Macro F-1 score for *Aggression & Attack* and R-squared for *Emotion Simple* datasets.

quantitative approach. Personalization can shift the view of the problem in a contextual way by dedicating reasoning to a single user. Finally, we showed that adding information about model uncertainty significantly improves the ability to predict complex and subjective behaviors such as recognizing hate speech or emotions in a text. The hybrid model we created significantly outperformed the previous methods, becoming a new state-of-the-art on two very challenging tasks. Our future work will focus on applications of our approach to some other tasks such as active or reinforcement learning.

LIMITATIONS

One important issue related to the nature of normalizing flows is their ability to convert probabilities to disambiguate uncertain answers. At the moment, there are no reference datasets available that contain text and annotator information simultaneously with multiple markings of the same text by the same person. This is due to cost constraints in preparing such data. However, we have conducted experiments on datasets with different characteristics in which (1) one person marked several hundred texts [*Wikipedia Detox Datasets*] and (2) one text was evaluated dozens of times by different people [*Emotions Simple* and *Emotion Meanings* datasets]. In order to address the problem mentioned in the introduction, one text should have N annotations from the same person, e.g., a few days apart. If we gain access to or prepare such a dataset, we would be happy to conduct in-depth studies on it.

Due to the language of one of the datasets being different from English, a multilingual model was used to embed the text. This decision was made in order to allow for direct comparisons and cross-referencing. This would have added an unnecessary layer to the already relatively complex problem

that was addressed. It is possible to experiment with other language models as well using the source codes provided[†].

ACKNOWLEDGMENTS

This work was financed by (1) the National Science Centre, Poland, project no. 2021/41/B/ST6/04471; (2) Contribution to the European Research Infrastructure 'CLARIN ERIC - European Research Infrastructure Consortium: Common Language Resources and Technology Infrastructure', 2022-23 (CLARIN Q); (3) the Polish Ministry of Education and Science, CLARIN-PL; (4) the European Regional Development Fund, as a part of the 2014-2020 Smart Growth Operational Programme, projects no. POIR.04.02.00-00C002/19, POIR.01.01.01-00-0923/20, POIR.01.01.01-00-0615/21, and POIR.01.01.01-00-0288/22; (5) the statutory funds of the Department of Artificial Intelligence, Wrocław University of Science and Technology; (6) the Polish Ministry of Education and Science within the programme "International Projects Co-Funded"; (7) the European Union under the Horizon Europe, grant no. 101086321 (OMINO). However, the views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor European Research Executive Agency can be held responsible for them. The work conducted by Maciej Zieba and Patryk Wielopolski was supported by the National Centre of Science (Poland) Grant No. 2021/43/B/ST6/02853.

REFERENCES

- [1] L. F. Barrett, *How emotions are made: The secret life of the brain*. Pan Macmillan, 2017.
- [2] E. Pavlick and T. Kwiatkowski, "Inherent Disagreements in Human Textual Inferences," *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 677–694, 11 2019.
- [3] C. Beck, H. Booth, M. El-Assady, and M. Butt, "Representation problems in linguistic annotations: Ambiguity, variation, uncertainty, error and bias," in *Proceedings of the 14th Linguistic Annotation Workshop*, (Barcelona, Spain), pp. 60–73, Association for Computational Linguistics, Dec. 2020.
- [4] A. M. Davani, M. Díaz, and V. Prabhakaran, "Dealing with disagreements: Looking beyond the majority vote in subjective annotations," *Transactions of the Association for Computational Linguistics*, vol. 10, pp. 92–110, 2022.
- [5] E. Troiano, S. Padó, and R. Klinger, "Emotion ratings: How intensity, annotation confidence and agreements are entangled," in *Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, (Online), pp. 40–49, Association for Computational Linguistics, Apr. 2021.
- [6] K. Krippendorff, "Computing krippendorff's alpha-reliability," *Annenberg School for Communication Departmental Papers: Philadelphia*, 2011.
- [7] J. L. Fleiss, "Measuring nominal scale agreement among many raters.," *Psychological bulletin*, vol. 76, no. 5, p. 378, 1971.
- [8] P. Miłkowski, S. Saganowski, M. Gruza, P. Kazienko, M. Piasecki, and J. Kocoń, "Multitask personalized recognition of emotions evoked by textual content," in *EmotionAware 2022: Sixth International Workshop on Emotion Awareness for Pervasive Computing Beyond Traditional Approaches at PerCom 2022*, (online), pp. 347–352, mar 2022.
- [9] C. Strapparava and R. Mihalcea, "Learning to identify emotions in text," in *Proceedings of the 2008 ACM symposium on Applied computing*, pp. 1556–1560, 2008.

[†]<https://github.com/CLARIN-PL/personalized-emotion-prediction-with-normalizing-flows>

- [10] F. S. Tabak and V. Evrim, "Comparison of emotion lexicons," in *2016 HONET-ICT*, pp. 154–158, IEEE, 2016.
- [11] R. Plutchik, "A general psychoevolutionary theory of emotion," in *Theories of emotion*, pp. 3–33, Elsevier, 1980.
- [12] P. Ekman, "An argument for basic emotions," *Cognition & emotion*, vol. 6, no. 3-4, pp. 169–200, 1992.
- [13] L. A. M. Oberländer and R. Klinger, "An analysis of annotated corpora for emotion classification in text," in *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 2104–2119, 2018.
- [14] J. Kocoń, A. Janz, and M. Piasecki, "Context-sensitive sentiment propagation in wordnet," in *Proceedings of the 9th global wordnet conference*, pp. 329–334, 2018.
- [15] J. Kocoń and A. Janz, "Propagation of emotions, arousal and polarity in wordnet using heterogeneous structured synset embeddings," in *Proceedings of the 10th Global Wordnet Conference*, pp. 336–341, 2019.
- [16] J. Kocoń, A. Janz, P. Miłkowski, M. Riegel, M. Wierzba, A. Marchewka, A. Czoska, D. Grimling, B. Konat, K. Juszczak, et al., "Recognition of emotions, valence and arousal in large-scale multi-domain text reviews," in *9th Language & Technology Conference*, 2019.
- [17] J. Kocoń, P. Miłkowski, M. Wierzba, B. Konat, K. Klessa, A. Janz, M. Riegel, K. Juszczak, A. Marchewka, et al., "Multilingual and language-agnostic recognition of emotions, valence and arousal in large-scale multi-domain text reviews," in *Language and Technology Conference*, pp. 214–231, Springer, 2019.
- [18] D. Demszky, D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, and S. Ravi, "Goemotions: A dataset of fine-grained emotions," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4040–4054, 2020.
- [19] J. Kocoń, J. Radom, E. Kaczmarz-Wawryk, K. Wabnic, A. Zajączkowska, and M. Zaśko-Zielińska, "Aspectemo: multi-domain corpus of consumer reviews for aspect-based sentiment analysis," in *2021 International Conference on Data Mining Workshops (ICDMW)*, pp. 166–173, IEEE, 2021.
- [20] A. Srivastava, A. Rastogi, A. Rao, A. A. M. Shueb, A. Abid, A. Fisch, A. R. Brown, A. Santoro, A. Gupta, A. Garriga-Alonso, et al., "Beyond the imitation game: Quantifying and extrapolating the capabilities of language models," *Transactions on Machine Learning Research*, 2023.
- [21] S. Mohammad, "# emotional tweets," in * *SEM 2012: The First Joint Conference on Lexical and Computational Semantics—Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pp. 246–255, 2012.
- [22] M. Abdul-Mageed and L. Ungar, "Emonet: Fine-grained emotion detection with gated recurrent neural networks," in *Proceedings of the 55th annual meeting of the association for computational linguistics (volume 1: Long papers)*, pp. 718–728, 2017.
- [23] C.-C. Hsu and L.-W. Ku, "Socialnlp 2018 emotionx challenge overview: Recognizing emotions in dialogues," in *Proceedings of the sixth international workshop on natural language processing for social media*, pp. 27–31, 2018.
- [24] E. Wulczyn, N. Thain, and L. Dixon, "Ex machina: Personal attacks seen at scale," in *Proceedings of the 26th international conference on world wide web*, pp. 1391–1399, 2017.
- [25] C. J. Kennedy, G. Bacon, A. Sahn, and C. von Vacano, "Constructing interval variables via faceted rasch measurement and multitask deep learning: a hate speech application," *arXiv preprint arXiv:2009.10277*, 2020.
- [26] S. Akhtar, V. Basile, and V. Patti, "Modeling annotator perspective and polarized opinions to improve hate speech detection," in *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, vol. 8, pp. 151–154, 2020.
- [27] J. Kocoń, A. Figas, M. Gruza, D. Puchalska, T. Kajdanowicz, and P. Kazienko, "Offensive, aggressive, and hate speech analysis: From data-centric to human-centered approach," *Information Processing & Management*, vol. 58, no. 5, p. 102643, 2021.
- [28] F. Miresghallah, V. Shrivastava, M. Shokouhi, T. Berg-Kirkpatrick, R. Sim, and D. Dimitriadis, "Useridentifier: implicit user representations for simple and effective personalized sentiment analysis," *arXiv preprint arXiv:2110.00135*, 2021.
- [29] K. Kanclerz, A. Figas, M. Gruza, T. Kajdanowicz, J. Kocoń, D. Puchalska, and P. Kazienko, "Controversy and conformity: from generalized to personalized aggressiveness detection," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 5915–5926, 2021.
- [30] P. Miłkowski, M. Gruza, K. Kanclerz, P. Kazienko, D. Grimling, and J. Kocoń, "Personal bias in prediction of emotions elicited by textual opinions," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021): Student Research Workshop*, (Online), pp. 248–259, Association for Computational Linguistics, Aug. 2021.
- [31] A. Ngo, A. Candri, T. Ferdinan, J. Kocoń, and W. Korczynski, "Studemo: A non-aggregated review dataset for personalized emotion recognition," in *Proceedings of the 1st Workshop on Perspectivist Approaches to NLP@ LREC2022*, pp. 46–55, 2022.
- [32] K. Kanclerz, M. Gruza, K. Karanowski, J. Bielaniec, P. Miłkowski, J. Kocoń, and P. Kazienko, "What if ground truth is subjective? personalized deep neural hate speech detection," in *Proceedings of the 1st Workshop on Perspectivist Approaches to NLP@ LREC2022*, pp. 37–45, 2022.
- [33] J. Bielaniec, K. Kanclerz, P. Miłkowski, M. Gruza, K. Karanowski, P. Kazienko, and J. Kocoń, "Deep-sheep: Sense of humor extraction from embeddings in the personalized context," in *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 967–974, IEEE, 2022.
- [34] T. Ferdinan and J. Kocoń, "Personalized models resistant to malicious attacks for human-centered trusted ai," in *The AAAI-23 Workshop on Artificial Intelligence Safety (SafeAI 2023)*, CEUR Workshop Proceedings, 2023.
- [35] W. Mieszczewski-Kowszewicz, K. Kanclerz, J. Bielaniec, M. Oleksy, M. Gruza, S. Woźniak, E. Dzięcioł, P. Kazienko, and J. Kocoń, "Capturing human perspectives in nlp: Questionnaires, annotations, and biases," in *The ECAI 2023 2nd Workshop on Perspectivist Approaches to NLP*, CEUR Workshop Proceedings, 2023.
- [36] J. Kocoń, J. Baran, K. Kanclerz, M. Kajstura, and P. Kazienko, "Differential dataset cartography: Explainable artificial intelligence in comparative personalized sentiment analysis," in *International Conference on Computational Science*, pp. 148–162, Springer, 2023.
- [37] E. Cambria, Q. Liu, S. Decherchi, F. Xing, and K. Kwok, "Senticnet 7: A commonsense-based neurosymbolic ai framework for explainable sentiment analysis," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pp. 3829–3839, 2022.
- [38] R. Mao, Q. Liu, K. He, W. Li, and E. Cambria, "The biases of pre-trained language models: An empirical study on prompt-based sentiment analysis and emotion detection," *IEEE Transactions on Affective Computing*, 2022.
- [39] M. M. Amin, R. Mao, E. Cambria, and B. W. Schuller, "A wide evaluation of chatgpt on affective computing tasks," *arXiv preprint arXiv:2308.13911*, 2023.
- [40] L. Zhu, W. Li, R. Mao, V. Pandelea, and E. Cambria, "Paed: zero-shot persona attribute extraction in dialogues," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 9771–9787, 2023.
- [41] L. Dinh, D. Krueger, and Y. Bengio, "Nice: Non-linear independent components estimation," *arXiv*, 2014.
- [42] L. Dinh, J. Sohl-Dickstein, and S. Bengio, "Density estimation using Real NVP," in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, OpenReview.net, 2017.
- [43] G. Papamakarios, T. Pavlakou, and I. Murray, "Masked autoregressive flow for density estimation," 2018.
- [44] W. Grathwohl, R. T. Chen, J. Bettencourt, I. Sutskever, and D. Duvenaud, "Fjord: Free-form continuous dynamics for scalable reversible generative models," *arXiv preprint arXiv:1810.01367*, 2018.
- [45] M. Zieba, M. Przewieźlikowski, M. Śmieja, J. Tabor, T. Trzcinski, and P. Spurek, "RegFlow: Probabilistic Flow-based Regression for Future Prediction," *CoRR*, vol. abs/2011.14620, 2020.
- [46] P. Wielopolski, M. Koperski, and M. Zieba, "Flow plugin network for conditional generation," *arXiv preprint arXiv:2110.04081*, 2021.
- [47] M. Wolczyk, M. Proszewska, L. Maziarzka, M. Zieba, P. Wielopolski, R. Kurczab, and M. Śmieja, "Plugen: Multi-label conditional generation from pre-trained models," in *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pp. 8647–8656, AAAI Press, 2022.

- [48] R. Abdal, P. Zhu, N. J. Mitra, and P. Wonka, "Styleflow: Attribute-conditioned exploration of stylegan-generated images using conditional continuous normalizing flows," *ACM Trans. Graph.*, vol. 40, no. 3, pp. 21:1–21:21, 2021.
- [49] P. Wielopolski and M. Zieba, "Treeflow: Going beyond tree-based gaussian probabilistic regression," *CoRR*, vol. abs/2206.04140, 2022.
- [50] D. Tran, K. Vafa, K. K. Agrawal, L. Dinh, and B. Poole, "Discrete flows: Invertible generative models of discrete data," in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada* (H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. B. Fox, and R. Garnett, eds.), pp. 14692–14701, 2019.
- [51] J. Kocoń, M. Gruza, J. Bielaniec, D. Grimling, K. Kanclerz, P. Miłkowski, and P. Kazienko, "Learning personal human biases and representations for subjective tasks in natural language processing," in *2021 IEEE International Conference on Data Mining (ICDM)*, pp. 1168–1173, IEEE, 2021.
- [52] P. Kazienko, J. Bielaniec, M. Gruza, K. Kanclerz, K. Karanowski, P. Miłkowski, and J. Kocoń, "Human-centred neural reasoning for subjective content processing: Hate speech, emotions, and humor," *Information Fusion*, 2023.
- [53] D. J. Rezende and S. Mohamed, "Variational Inference with Normalizing Flows," in *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, vol. 37 of *JMLR Workshop and Conference Proceedings*, pp. 1530–1538, JMLR.org, 2015.
- [54] J. Vig and Y. Belinkov, "Analyzing the structure of attention in a transformer language model," *arXiv preprint arXiv:1906.04284*, 2019.
- [55] P. Miłkowski, M. Gruza, K. Kanclerz, P. Kazienko, D. Grimling, and J. Kocoń, "Personal bias in prediction of emotions elicited by textual opinions," in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: Student Research Workshop*, pp. 248–259, 2021.
- [56] M. Wierzba, M. Riegel, J. Kocoń, P. Miłkowski, A. Janz, K. Klessa, K. Juszczyk, B. Konat, D. Grimling, M. Piasecki, *et al.*, "Emotion norms for 6000 polish word meanings with a direct mapping to the polish wordnet," *Behavior Research Methods*, pp. 1–16, 2021.
- [57] R. Plutchik, *The emotions*. University Press of America, 1991.
- [58] F. Feng, Y. Yang, D. Cer, N. Arivazhagan, and W. Wang, "Language-agnostic BERT sentence embedding," in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, (Dublin, Ireland), pp. 878–891, Association for Computational Linguistics, May 2022.

APPENDIX

A. Personalization architectures

Architectures for personalization have been modified to return a probability distribution instead of the probability of a class. Input features are passed to the normalizing flow as context. We are dealing with four architectures:

- Baseline (Fig. 3) - the input is just an embedding of text
- OneHot (Fig. 4) - the user represented as a one-hot vector is concatenated to the text embedding
- HuBi-Formula (Fig. 5) - the deviation of the user's response is taken as its representative feature
- HuBi-Medium (Fig. 6) - annotation-based learned user embedding combined with text embeddings provides the context

B. Implementation details

Experiments were performed using the code provided in the "Anonymous".

Grid search for hyperparameters of Normalizing Flows in experiment 2 was performed with the values specified in Tab. V and Tab. VI.




Fig. 3: TXT-Baseline architecture utilizing normalizing flows.

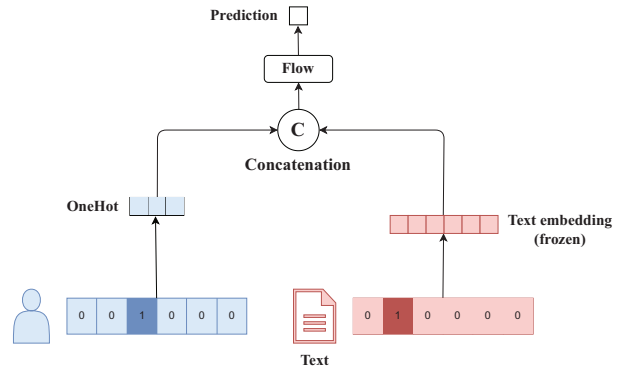


Fig. 4: OneHot architecture utilizing normalizing flows.

For all experiment purposes, we used a machine with AMD Ryzen Threadripper PRO 3955WX 16-Core Processor CPU, 2 x NVIDIA GeForce RTX 3090 GPUs, and 256 GB RAM.

C. Visualization of probabilities

For the combined set of Aggression & Attack, visualizations of the waveform of the probability function were prepared as a result of Experiment 3. described in the publication, are presented in Fig. 7 and Fig. 8.

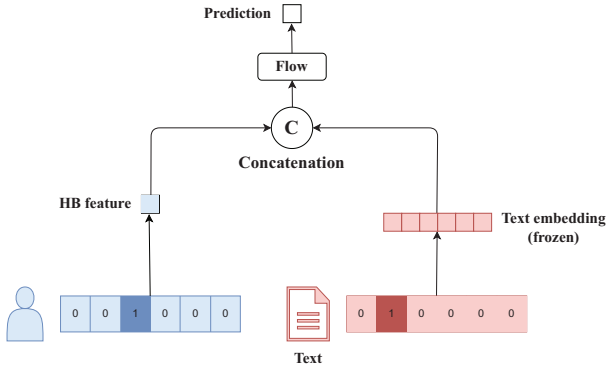


Fig. 5: HuBi-Formula architecture utilizing normalizing flows.

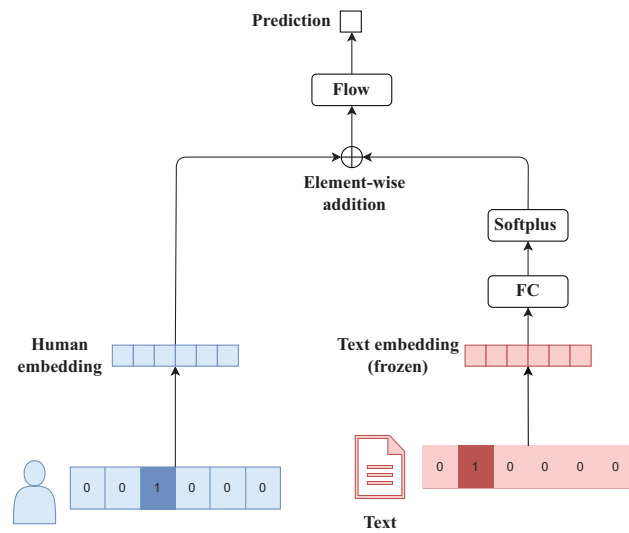


Fig. 6: HuBi-Medium architecture utilizing normalizing flows.

Hyperparameter	Values
hidden features	[2, 4, 6, 8]
num layers	[1, 2, 3, 4, 5]
num. blocks per layer	[1, 2, 3, 4]
dropout probability	[0.0, 0.1, 0.2, 0.4]
batch norm within layers	[True, False]
batch norm between layers	[True, False]

TABLE V: Hyperparameters for Normalizing Flows and their possible values during experiment 2. Note that for MAF, the *dropout probability* hyperparameter was not used at all.

Hyperparameter	Values for TXT-Baseline & OneHot & HuBi-Formula	Values for HuBi-Medium
embedding dim.	50	50
hidden dim.	50	[128, 256, 512, 786]
output dim.	-	[128, 256, 512, 768]
learning rate	[1e-5, 1e-4]	[1e-5, 1e-4]

TABLE VI: Hyperparameters for training and architectures, and their possible values during experiment 2.




Fig. 7: Visualizations of the waveform of the probability function for RealNVP with different architectures and for Attack and Aggression.




Fig. 8: Visualizations of the waveform of the probability function for CNF with different architectures and for Attack and Aggression.