GlobalTrait: Personality Alignment of Multilingual Word Embeddings

Farhad Bin Siddique,^{1,2} Dario Bertero,^{1,2} Pascale Fung^{1,2,3}

¹Electronic and Computer Engineering Department

²Center for Artificial Intelligence Research (CAiRE)

³EMOS Technologies Inc.

The Hong Kong University of Science and Technology

Clear Water Bay, Hong Kong

[fsiddique, dbertero]@connect.ust.hk, pascale@ece.ust.hk

Abstract

We propose a multilingual model to recognize Big Five Personality traits from text data in four different languages: English, Spanish, Dutch and Italian. Our analysis shows that words having a similar semantic meaning in different languages do not necessarily correspond to the same personality traits. Therefore, we propose a personality alignment method, GlobalTrait, which has a mapping for each trait from the source language to the target language (English), such that words that correlate positively to each trait are close together in the multilingual vector space. Using these aligned embeddings for training, we can transfer personality related training features from high-resource languages such as English to other low-resource languages, and get better multilingual results, when compared to using simple monolingual and unaligned multilingual embeddings. We achieve an average F-score increase (across all three languages except English) from 65 to 73.4 (+8.4), when comparing our monolingual model to multilingual using CNN with personality aligned embeddings. We also show relatively good performance in the regression tasks, and better classification results when evaluating our model on a separate Chinese dataset.

Introduction

According to (Allport 1937), personality refers to the characteristic pattern in a person's thinking, feeling, and decision making. It is a quality of a person across a relatively long period and is different from emotions, which can be perceived in the moment. We can think of it as, personality is to emotion what climate is to weather. The Big Five model of personality (Goldberg 1993) is a common way of quantifying a person's personality, and is recognized by most psychologists around the world. It tries to represent the traits as scores across five dimensions:

- **Extraversion vs Introverted** (*Extr*) - sociable, assertive, playful vs aloof, reserved, shy;
- **Conscientiousness vs Unconscientious** (*Cons*) - self-disciplined, organised vs inefficient, careless;
- **Agreeableness vs Disagreeable** (*Agr*) - friendly, cooperative vs antagonistic, faultfinding;
- **Neuroticism vs Emotionally Stable** (*Emot*) - insecure, anxious vs calm, unemotional;

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

- **Openness to Experience vs Cautious** (*Openn*) - intellectual, insightful vs shallow, unimaginative.

Personality traits affect the usage of language in people (Mairesse et al. 2007), and it is an integral part of human-human interaction (Long and Martin 2000; Berry, Willingham, and Thayer 2000). As we develop smarter dialogue systems, future virtual agents need to detect and adapt to different user personalities in order to express empathy (Fung et al. 2016). Although traditionally the user personality can be identified by having the user fill out a self-assessment form such as the NEO Personality Inventory (Costa and McCrae 2008), this method is not feasible for many applications where we may wish to identify user personality, such as dialogue systems. Therefore, work on automatic personality assessment has become increasingly important recently with the rise in popularity of applications such as Human Resources (HR) screening, personalized marketing, and other social media related user-profiling.

Currently, personality labeled data is scarce, especially in the multilingual setting, which makes it essential to use the relatively larger size of data in English to help recognize traits in other languages. Previous work on multilingual personality recognition have tried to use the word-level or character-level similarity across languages to develop a multilingual model (Liu, Perez, and Nowson 2016; Siddique and Fung 2017). However, our experiments show that words used by people with different personality traits differ among languages or cultures, which is not captured by distributional semantics alone, making it necessary to learn a personality-based mapping of words to express each personality trait. Therefore, we propose GlobalTrait, which trains personality trait-based alignment of multilingual embeddings from the source language(s) to the target language, such that words that correlate positively to each trait are closer together in the multilingual (global) vector space. We show that taking such mapping or alignment of embeddings as input to our model gives us better multilingual results in the task of personality recognition.

Related Work

Automatic personality recognition has been done since as early as 2006 (Oberlander and Nowson 2006), where the personality of blog authors were identified using Naive

(a) Words corresponding positively to Extraversion trait

(b) Words corresponding positively to Conscientiousness trait

Figure 1: Multilingual embeddings of most significant words corresponding to the Extraversion and Conscientiousness trait for the four languages, plotted in 2-Dimension using t-SNE. We see that positively correlated words for each trait are clustered per language, which shows that words corresponding to each language has different semantic/contextual meaning compared to the other languages. Therefore, we learn a mapping from each source language to our target language English, so the words are aligned based on each of the personality traits.

Bayes algorithm and n-gram features. (Mairesse et al. 2007) used two sources of lexical features, Linguistic Inquiry and Word Count (LIWC) (Pennebaker, Francis, and Booth 2001) and Medical Research Council (MRC) Psycholinguistics Database (Coltheart 1981) features, to identify personality from written and spoken transcripts. More recently, for tasks such as the Workshop on Computational Personality Recognition (Celli et al. 2013), people have worked on identifying personality from social media texts (Facebook status updates and Youtube vlog transcriptions). (Verhoeven, Daelemans, and De Smedt 2013) used 2000 frequent trigrams as features and trained a SVM classifier, and (Farnadi et al. 2013) used LIWC features to train SVM, Naive Bayes and K Nearest Neighbor (KNN) algorithms.

Deep Learning models such as Convolutional Neural Networks (CNNs) have gained popularity in the task of text classification (Kalchbrenner, Grefenstette, and Blunsom 2014; Kim 2014). This is because CNNs are good at capturing text features via its convolution operation, which can be applied on the text by taking the distributed representation of the words, called word embeddings, as input. Learning such distributed representation comes from the hypothesis that words that appear in similar contexts have similar meaning (Harris 1954). Different works have been carried out in the past to learn such representations of words, such as (Mikolov et al. 2013; Pennington, Socher, and Manning 2014), and more recently (Bojanowski et al. 2016). Cross-lingual or multilingual word embeddings try to capture such semantic information of words across two or more languages, such that the words that have similar meaning in different languages are close together in the vector space (Faruqui and Dyer 2014; Upadhyay et al. 2016). For our task we use a more recent approach (Conneau et al. 2017), which does not require parallel data and learns a mapping from the source language embedding space to the target language in an unsupervised fashion.

Methodology

We propose GlobalTrait, a text-based model that uses multilingual embeddings across languages to train personality alignment per trait, such that words in the languages that correspond positively to one trait are closer together in the multilingual vector space: which are then fed to a CNN model for binary classification of each trait. We first train multilingual embeddings in the given languages, followed by identifying the most significant words in each language corresponding positively to each trait, which are used to learn a mapping from each source to the target language. The target language is the language in which we have the most labeled data available, in our case English. The mapping is essentially a personality trait based alignment that tries to bring the distributional representation of words correlating positively to each trait closer together. The initial multilingual embeddings along with the GlobalTrait aligned embeddings are then fed into a two channel CNN model to extract the relevant features for classification via a fully connected layer followed by softmax.

Multilingual Embeddings Training

Learning distributed representation of words (word embeddings) comes from the hypothesis that words that appear in similar contexts have similar meaning (Harris 1954). Different works have been done to learn such representations (Mikolov et al. 2013; Pennington, Socher, and Manning 2014; Bojanowski et al. 2016). Cross-lingual or multilingual word embeddings try to capture such semantic information of words across two or more languages (Faruqui and Dyer 2014; Upadhyay et al. 2016). We use the methodology of Multilingual Unsupervised and Supervised Embeddings (MUSE) (Conneau et al. 2017) to first train multilingual embeddings across the four languages. MUSE tries to learn a mapping W of dimension $d \times d$, where $d = 300$ is the embedding dimension we use, such that:

Table 1: Some examples of significant words corresponding positively to each trait per language.

	English (En)	Spanish (Es)	Italian (It)	Dutch (Nl)
Extr	people, music, stories, followers, social	enhorabuena, gente, sociales, programa, tenemos	piaciuto, notizia, pubblicato, università, generazione	mensen, kennen, schoten, benieuwd, helemaal
Cons	research, work, time, analysis, project	gobierno, interesante, cuenta, contra, tiempo	direttori, giornalismo, palestra, rivoluzione, ricerca	raadscommissie, bibliotheek, grootdebat, harlopen, ambtenaren
Agr	life, feel, thanks, amazing, heart, right	gracias, estoy, estudiar, vienes, dormir	anche, siamo, utilizzato, fatto, grazie	inderdaad, altijd, genieten, hopelijk, geslaagd
Emot	want, today, really, shit, know, cause	pero, porque, alguien, sólo, vamos, quieres	sempre, proprio, fumetto, altro, scoprire	alleen, vragen, eigenlijk, waarom, volgende
Openn	love, new, watching, youth, skills, travel	mundo, auditori, trabajo, corazon, nuevo	mostra, mondo, aiutarmi, informati, completando	andere, iedereen, leven, lekker, misschien, volgende

$$W^* = \operatorname{argmin}_{W \in O_d(\mathbb{R})} \|WX - Y\|_F \quad (1)$$

where, $O_d(\mathbb{R})$ suggests that W is an orthogonal matrix consisting of real numbers, X and Y are the $d \times n$ matrices representing the word embeddings of n words in the source and target languages respectively. It is important that the mapping matrix W is orthogonal, so that we are performing a rotational mapping on the embedding space, which does not disrupt the monolingual semantic information of the original embeddings. Also, being orthogonal gives us the following Procrustes solution to equation (1):

$$W^* = UV^T, \text{ with } U\Sigma V^T = \text{SVD}(YX^T). \quad (2)$$

W mapping matrix is trained via a Generative Adversarial Network (GAN) training approach (Goodfellow et al. 2014; Ganin et al. 2016), where a generator and a discriminator network are both trained in parallel. It is a two player game where the discriminator tries to differentiate between a source embedding and a mapped embedding, and the generator tries to fool the discriminator by making WX as similar to Y as possible. For our monolingual embeddings used for each of the individual languages, we use the pre-trained word embeddings via fastText¹. The issue with multilingual embeddings is that it only captures the semantic information of words across languages, so words that share similar context will appear close together in the multilingual vector space. As we can see in Figure 1 (explained further in the *Experiments* section), words that correspond to the same trait do not always share a similar semantic meaning across languages. This motivates the need for a mapping of the embeddings from the source to the target language, hence the notion of personality trait-based alignment.

GlobalTrait - One Alignment Per Trait

We propose GlobalTrait, which does a personality trait-based alignment of the multilingual embeddings from the source to the target language. From the dataset, we use Term Frequency - Inverse Document Frequency (*tf-idf*) features to obtain the n most significant words that correspond positively to each trait per language, and get the multilingual

¹<https://github.com/facebookresearch/fastText>

Figure 2: Simplified architecture of our two channel CNN model, one channel taking the unaligned multilingual embeddings as input, while the other taking the GlobalTrait aligned embeddings, fed into a one layer CNN, and the extracted features are concatenated, followed by a max-pooling layer and a fully connected layer to softmax for binary classification.

embeddings corresponding to the words. Using these embeddings, we learn a second mapping from each source language to our target language, English, with the idea that the mapped or aligned embeddings will represent the trait, by being closer together in the vector space. There has to be one alignment per trait per language - a rotational mapping from the source to the target language space. We end up training 5 different mapping matrices for each source language, one for each of the Big Five traits.

Training Procedure We use the same training approach as MUSE to train our personality mapping. For each trait, we take the words having the highest significance in the source languages, and train a second mapping of their corresponding multilingual embeddings to the target language space. We can call such mapping of the Agreeableness trait, for example, as W_a , and therefore, we try to achieve the following equation, $W_a A = B$, where A and B are the trained multilingual embeddings of words in the source and target language respectively, which correlate positively to the Agreeableness trait. Let $X = \{x_1, \dots, x_n\}$ be the n multilingual embeddings in the source language of the words corresponding most positively to the Agreeableness trait. Like-

Figure 3: Words corresponding to the Agreeableness trait, left shows the unaligned multilingual embeddings, and right shows the GlobalTrait Personality aligned embeddings. We see that the alignment brings the multilingual embeddings to a common vector space (rather than being clustered per language), so we can get the personality trait based mapping between languages.

wise, $Y = \{y_1, \dots, y_m\}$ are the m multilingual embeddings of words in the target language, again for the Agreeableness trait. We try to train a discriminator such that it can differentiate between random samples taken from $W_a X = \{W_a x_1, \dots, W_a x_n\}$ and Y . We train W_a matrix such that the discriminator is unable to differentiate between the two, therefore making the mapping from source as close to the target as possible, bringing the notion of personality based alignment.

- **Discriminator:** Let θ_D be the parameters of the discriminator, and $P_{\theta_D}(\text{source} = 1|z)$ is the probability that a vector z is the element of a source embedding according to the discriminator, and not the target (mapped) embedding. Therefore, for the Agreeableness trait, the discriminator loss function is as follows:

$$L_D(\theta_D|W_a) = -\frac{1}{n} \sum_{i=1}^n \log P_{\theta_D}(\text{source} = 1|W_a x_i) - \frac{1}{m} \sum_{i=1}^m \log P_{\theta_D}(\text{source} = 0|y_i) \quad (3)$$

- **Mapping Matrix:** We try to train W_a such that the discriminator is unable to differentiate between the source and the target (mapped) embeddings:

$$L_W(W_a|\theta_D) = -\frac{1}{n} \sum_{i=1}^n \log P_{\theta_D}(\text{source} = 0|W_a x_i) - \frac{1}{m} \sum_{i=1}^m \log P_{\theta_D}(\text{source} = 1|y_i) \quad (4)$$

Therefore, the discriminator and the weight matrix objective functions work alternative to each other, and they finally converge in a *min-max* solution. Similarly, we train a mapping matrix for each of the other four traits, which gives us the GlobalTrait personality aligned embeddings, where the embeddings closer in the multilingual vector space reflect the same trait.

Convolutional Neural Network

Deep learning models such as Convolutional Neural networks (CNNs) have gained popularity in the task of text classification (Kalchbrenner, Grefenstette, and Blunsom 2014; Kim 2014). Our CNN model is a two channel mode, where one channel takes the multilingual embeddings, while the other takes the GlobalTrait personality aligned embeddings as input. The first channel with multilingual embeddings is kept trainable, or we can call it a dynamic channel, which means the embeddings are also taken as training parameters, and can change as the training goes on. The other channel, where the personality alignment has been trained already, is kept static. For both channels, we choose window sizes to be 3, 4 and 5, which essentially extracts 3, 4 and 5-gram features from the text, and we have a max pooling operation that keeps the maximum features per window from both the channels. The total features are concatenated and passed to a fully connected layer with a single hidden layer and non-linear activation (tanh), ultimately mapping the features to a binary classification of each trait via softmax.

Experiments

Dataset

We used the 2015 Author Profiling challenge dataset (PAN 2015) (Rangel et al. 2015), which includes user tweets in four languages - English (en), Spanish (es), Italian (it) and Dutch (nl), where the personality labels were obtained via self-assessment using the BFI-10 item personality questionnaire (Rammstedt and John 2007). Only the training set was released to us from the PAN2015 website². Their test data was not available to us as we did not take part in the Author Profiling competition of 2015. The dataset consists of 152 English (14,166 tweets), 110 Spanish (9,879 tweets), 38 Italian (3,687 tweets), and 34 Dutch (3,350 tweets) users in total. Our task is to identify personality in user-level, so we concatenated all the tweets made by a single user to create one single training/test data point. As preprocessing, we tokenized each tweet using Twokenizer (Owoputi et al. 2013),

²<https://pan.webis.de/clef15/pan15-web/author-profiling.html>

Table 2: F-score results for the monolingual and multilingual performance of the CNN and logistic regression model in the binary classification task for each trait. Bold highlights the best performance per trait for each language.

Lang	Model	Extr	Agr	Cons	Emot	Openn	Average
en	Lgr-mono	64.0	53.3	64.5	72.5	60.5	63.0
	CNN-mono	74.4	48.2	72.8	74.9	67.7	67.6
es	Lgr-mono	73.2	72.1	70.0	56.9	69.0	68.2
	Lgr-multi	73.2	70.1	72.6	57.0	69.5	68.5
	Lgr-GlobalTrait	75.9	74.7	82.4	59.0	69.0	72.2
	CNN-mono	74.7	74.7	70.0	69.0	67.2	71.1
	CNN-GlobalTrait	79.4	76.0	83.3	67.3	67.0	74.6
it	Lgr-mono	53.3	60.3	52.5	60.4	63.3	58.0
	Lgr-multi	64.2	71.2	51.5	62.1	64.2	62.6
	Lgr-GlobalTrait	66.2	75.2	49.7	65.7	63.5	64.1
	CNN-mono	67.2	74.3	60.3	75.4	63.1	68.1
	CNN-GlobalTrait	64.0	77.5	58.3	78.0	63.2	68.2
nl	Lgr-mono	76.0	67.4	67.9	65.0	67.7	68.8
	Lgr-multi	74.2	58.2	66.8	66.2	67.0	66.5
	Lgr-GlobalTrait	76.9	52.9	62.0	68.2	66.4	65.3
	CNN-mono	76.4	60.6	61.8	78.6	64.6	68.4
	CNN-GlobalTrait	85.3	58.4	83.3	85.8	74.5	77.5

and replaced all usernames and URL mentions with generic words (*@username* and *@url*), so the model is not affected by mentions in the tweet that are not influenced by the user’s personality. Since we are interested in the binary classification of each big five trait, we carried out a median split of the scores, to obtain positive and negative samples (users) for each of the five traits. For our results shown, we carried out a stratified k-fold cross validation, by making k=5 splits of the training set into training/validation, and then show the average result across the 5 different validation sets.

Experimental Setup

For our MUSE training, we used a discriminator with 2 hidden layers, each having a dimension of 2048, and we ran our training for 5 epochs with 100,000 iterations in each epoch. When training the personality alignment, we took the top 3000 significant words corresponding positively to each trait per language. For our evaluation, we built a source to target language dictionary and used mean cosine distance as the validation metric. For our CNN model, we used 64 filters per filter size, and for our fully connected layer, we set the hidden layer dimension to 100, and we ran our training for 100 epochs for each model with batch size = 10. We used binary cross entropy as our loss function, and used Adam optimizer with learning rate of $1e^{-4}$.

Text Personality Analysis - Visualization of Embeddings

To check if words having similar semantic information across languages contribute to similar Big Five traits, we carried out some text-based analysis of our data. For each of the Big Five traits, we first obtained the most significant words in each language using term frequency-inverse document frequency (*tf-idf*) features. We then took the top 750 words for each language having the highest *tf-idf* and plot

their trained multilingual embeddings (trained using MUSE from monolingual embeddings in each language) on a 2-D space by performing t-distributed Stochastic Neighbor Embedding (*t-SNE*) on the 300-dimensional vectors. As an example, plots for the Extraversion and Conscientiousness traits are shown in Figure 1.

As we can see in the figure, for both traits there is very little overlap in the embedding space between the four languages, and most of the words are clustered per language. This shows that words corresponding to each trait might not have the same semantic meaning across the languages. Similar to the Extraversion and Conscientiousness trait, in all five traits, we see a similar trend, some overlap between English and Spanish, some overlap between Spanish and Italian, but very little or no overlap between Dutch and any of the languages. Therefore, this gives rise to the need for a certain mapping from each language to our target language, English. Some examples of words corresponding positively to the traits per language is shown in table 1.

Binary Classification

We first carried out binary classification of the users based on the median split of scores in each trait. Classification is of more importance to us as dialogue systems and other similar applications require us to classify each person into positive/negative for each trait, which can then be used to make decisions such as adapting to the given personality.

Baseline We implemented a simple logistic regression classifier, which takes in the average embeddings of the words as input features, in order to compare our aligned embeddings result with the monolingual counterpart. We present the results of the following experiments for comparison:

- **Lgr-mono**: logistic regression using monolingual embeddings

Table 3: Regression results of our multilingual CNN model using the GlobalTrait aligned embeddings, compared to two of the previous papers’ work on the same dataset. Bold indicates best performance (lowest RMSE) per language per trait.

Lang	Model	Extr	Agr	Cons	Emot	Openn	Average
es	Char Bi-RNN (Liu, Perez, and Nowson 2016)	0.148	0.143	0.157	0.177	0.136	0.152
	tf-idf linear regression (Sulea and Dichiu 2015)	0.152	0.148	0.114	0.181	0.142	0.147
	<i>CNN-GlobalTrait</i>	0.142	0.150	0.135	0.169	0.151	0.149
it	Char Bi-RNN (Liu, Perez, and Nowson 2016)	0.124	0.130	0.095	0.144	0.131	0.125
	tf-idf linear regression (Sulea and Dichiu 2015)	0.119	0.122	0.101	0.150	0.130	0.124
	<i>CNN-GlobalTrait</i>	0.107	0.128	0.120	0.147	0.134	0.127

- **Lgr-multi**: logistic regression using unaligned multilingual embeddings
- **Lgr-GlobalTrait**: logistic regression using our personality aligned multilingual embeddings
- **CNN-mono**: CNN using monolingual embeddings as input
- **CNN-GlobalTrait**: two channel CNN using multilingual embeddings plus the GlobalTrait aligned embeddings

In both ‘-multi’ and ‘-GlobalTrait’, models, the training set includes both the English and the source language’s training data, and is tested on the source language’s validation sets, while ‘-mono’ is just the monolingual model for the respective source language. Results are reported in table 2.

Results We achieve an average F-score of 74.6 in Spanish, 68.2 in Italian, and 77.5 in Dutch when using our multilingual CNN model with the GlobalTrait aligned embeddings, which are the highest performance achieved in each of the three languages. As we can see in table 2, CNN-GlobalTrait performs the best except for two traits in Spanish and three traits in Italian. The discrepancies can be due to the imbalanced nature of the dataset, and the logistic regression being a simple classifier, can converge better than CNN, especially for smaller datasets. For logistic regression, using multilingual and then GlobalTrait aligned embeddings improves on the monolingual results. In general, our multilingual results perform better than monolingual, except for one trait in Spanish, two traits in Italian, and one trait in Dutch. This shows that we can use the features retrieved from English to help us recognize personality in the other languages, and our personality alignment makes it easier for such kind of transfer learning.

Regression

We also carried out regression experiments to compare our model with a recent paper (Liu, Perez, and Nowson 2016) that tries to perform multilingual personality recognition on the same dataset. They use a character to word to sentence for personality traits (C2W2S4PT) model, which uses a two layer bi-directional RNN model with Gated Recurrent Units (GRU) followed by a fully connected layer to achieve the results.

For our regression task, we used the scores given in the dataset, and did not carry out the median split anymore. We kept our same training procedure as our classification task

to train the multilingual and the personality aligned embeddings using MUSE. Our CNN model was also the same except for our last fully connected layer, where we did not have the softmax layer and instead of cross entropy, we used mean-squared error as our objective (loss) function:

$$L(\theta) = \frac{1}{n} \sum_{i=1}^n (y_{t_i} - \hat{y}_{t_i})^2 \quad (5)$$

where y_{t_i} is the ground truth personality score of the t_i tweet, and $\hat{y}_{t_i}$ is the predicted score, θ being the collection of all parameters being trained. As our evaluation metric, we use Root Mean Squared Error (RMSE), which tries to measure the performance via the average error of the model across all users:

$$RMSE_{user} = \sqrt{\frac{1}{n} \sum_{i=1}^U (y_{user_i} - \hat{y}_{user_i})^2} \quad (6)$$

where y_{user_i} and $\hat{y}_{user_i}$ are the true and predicted personality trait score of the i^{th} user, and U is the total number of users. Table 3 shows our results compared to the (Liu, Perez, and Nowson 2016) model (we only show our results for Spanish and Italian, since they did not use the Dutch data in their paper). We also compare our results to the PAN 2015 participants (Sulea and Dichiu 2015) who used character n-gram based *tf-idf* features to train a regression model, and achieved one of the highest results in the competition.

As we can see from our results in table 3, our model gets the best performance in Extraversion trait for both languages, and it performs comparably in the other traits. This could mean that our GlobalTrait personality alignment works better for Extraversion, such that the words in different languages that correspond positively to Extraversion are indeed closer together in the multilingual vector space. However, it does not perform as well in the Openness trait, for example. It is important to take into consideration that (Sulea and Dichiu 2015) uses a monolingual model for each language, and therefore is not expandable to multiple languages. Also, unlike (Liu, Perez, and Nowson 2016) our model does not use character based RNN, which enables us to train on languages that do not share the same characters as English. To show this, we carried out separate classification experiments on a Chinese dataset.

Table 4: Binary classification F-score results on Chinese dataset, showing the comparison of GlobalTrait to other monolingual and multilingual models. All the multilingual models are trained using both Chinese and English training set, and tested on the Chinese validation sets (results shown are average of 5 different train/validation splits).

Lang	Model	Extr	Agr	Cons	Emot	Openn	Average
ch	Lgr-mono	58.2	59.0	57.4	56.9	56.5	57.6
	Lgr-multi	62.1	60.0	61.4	60.5	58.2	60.4
	Lgr-GlobalTrait	64.1	62.1	61.3	62.9	59.1	61.9
	CNN-mono	60.6	58.4	59.3	58.2	57.5	58.8
	CNN-GlobalTrait	64.2	61.9	63.0	62.5	60.1	62.3

Personality Classification on Chinese

Personality labeled data is currently rare in languages such as Chinese, which necessitates a model like GlobalTrait, enabling us to use English as additional training data to help us recognize personality in the Chinese test set. We use a Chinese personality labeled dataset called the BIT Speaker Personality Corpus (Zhang et al. 2017), collected and released to us by the Beijing Institute of Technology. It consists of 498 Chinese speech clips, each around 9-13 seconds and labeled with Big Five Personality scores given by five judges. We take a mean of the five scores for each clip, and carry out a median split for our binary classification task on each trait. We use an automatic speech recognition (ASR) system to get the speech transcriptions, and use Jieba segmenter³ to tokenize the Chinese text into words, since words in Chinese are not separated by a blank space.

We get pre-trained monolingual embeddings of Chinese from fastText and use MUSE to train multilingual embeddings in English and Chinese. We then use our GlobalTrait alignment method to map the positively correlated Chinese words in each trait to our target language space of English. The average of our 5 fold cross-validation experimental results are shown in table 4 and we train the same five models that were defined earlier. The results show us that, using our GlobalTrait aligned embeddings undoubtedly improves performance on the Chinese evaluation, which indicates a connection between the English and Chinese data captured via the personality alignment.

Final Discussion and Future Work

We have seen from our results that using the larger data available in English, we are able to improve our multilingual results, when applied to other languages such as Spanish, Italian, Dutch and Chinese. The personality alignment is particularly interesting, as it shows us how the words used to express different personality traits compare and contrast between multiple languages. Since we train a mapping for each trait per language, one word can have five different embeddings, based on the five different trait mappings. For example, our analysis show that the mapped embedding of the word ‘mundo’ (world) in Spanish is closest to the words ‘travel, flights, fresh’, etc. in English, for the Openness trait, while the same word ‘mundo’ for the Extraversion trait gets closest to ‘parties, love, life’, etc.

³<https://github.com/fxsjy/jieba>

We plan to explore more in the future to get more insight into such mappings for our GlobalTrait alignment, and also apply our model to other datasets in different languages. It will also be interesting to apply our GlobalTrait aligned embeddings to other models such as the Bi-directional RNN model we saw implemented by (Liu, Perez, and Nowson 2016), and other hierarchical attention networks, where our sequential data will be the tweets of a single user, based on the chronological order of the user tweets. Another work would be to include a much larger English personality labeled database, possibly the Facebook data released by myPersonality project⁴, where they collected data of around 154,000 users, with a total of 22 million status updates. Such a large database will help us find a better relation from other languages to English, thereby giving us a more meaningful personality alignment.

Conclusion

We propose the use of personality aligned embeddings, GlobalTrait, which maps the embedding space from the source language to our high-resource target language (English), thereby enabling us to get better multilingual results. We have shown in our paper that conventional methods that try to use monolingual or even multilingual word similarity for personality recognition may not always give better results, as words corresponding to personality traits might not have similar semantic meaning across multiple languages. Such a method like GlobalTrait can give us a better understanding of how people express personality across different cultures and languages, and therefore enable us to train better language-independent models for multilingual personality recognition.

Acknowledgements

This work was partially funded by CERG #16214415 of the Hong Kong Research Grants Council and RDC #1718050-0 of EMOS.AI.

References

- Allport, G. W. 1937. Personality: A psychological interpretation.
- Berry, D. S.; Willingham, J. K.; and Thayer, C. A. 2000. Affect and personality as predictors of conflict and close-

⁴<http://mypersonality.org/>

- ness in young adults' friendships. *Journal of Research in Personality* 34(1):84–107.
- Bojanowski, P.; Grave, E.; Joulin, A.; and Mikolov, T. 2016. Enriching word vectors with subword information. *arXiv preprint arXiv:1607.04606*.
- Celli, F.; Pianesi, F.; Stillwell, D.; and Kosinski, M. 2013. Workshop on computational personality recognition (shared task). In *Proceedings of the Workshop on Computational Personality Recognition*.
- Coltheart, M. 1981. The mrc psycholinguistic database. *The Quarterly Journal of Experimental Psychology* 33(4):497–505.
- Conneau, A.; Lample, G.; Ranzato, M.; Denoyer, L.; and Jégou, H. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.
- Costa, P. T., and McCrae, R. R. 2008. The revised neo personality inventory (neo-pi-r). *The SAGE handbook of personality theory and assessment* 2:179–198.
- Farnadi, G.; Zoghbi, S.; Moens, M.-F.; and De Cock, M. 2013. Recognising personality traits using facebook status updates. In *Proceedings of the workshop on computational personality recognition (WCPRI3) at the 7th international AAAI conference on weblogs and social media (ICWSM13)*. AAAI.
- Faruqui, M., and Dyer, C. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 462–471.
- Fung, P.; Bertero, D.; Wan, Y.; Dey, A.; Chan, R. H. Y.; Siddique, F. B.; Yang, Y.; Wu, C.-S.; and Lin, R. 2016. Towards empathetic human-robot interactions. *arXiv preprint arXiv:1605.04072*.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* 17(1):2096–2030.
- Goldberg, L. R. 1993. The structure of phenotypic personality traits. *American psychologist* 48(1):26.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in neural information processing systems*, 2672–2680.
- Harris, Z. S. 1954. Distributional structure. *Word* 10(2-3):146–162.
- Kalchbrenner, N.; Grefenstette, E.; and Blunsom, P. 2014. A convolutional neural network for modelling sentences. *arXiv preprint arXiv:1404.2188*.
- Kim, Y. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882*.
- Liu, F.; Perez, J.; and Nowson, S. 2016. A language-independent and compositional model for personality trait recognition from short texts. *arXiv preprint arXiv:1610.04345*.
- Long, M. V., and Martin, P. 2000. Personality, relationship closeness, and loneliness of oldest old adults and their children. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences* 55(5):P311–P319.
- Mairesse, F.; Walker, M. A.; Mehl, M. R.; and Moore, R. K. 2007. Using linguistic cues for the automatic recognition of personality in conversation and text. *Journal of artificial intelligence research* 30:457–500.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, 3111–3119.
- Oberlander, J., and Nowson, S. 2006. Whose thumb is it anyway?: classifying author personality from weblog text. In *Proceedings of the COLING/ACL on Main conference poster sessions*, 627–634. Association for Computational Linguistics.
- Owoputi, O.; O'Connor, B.; Dyer, C.; Gimpel, K.; Schneider, N.; and Smith, N. A. 2013. Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of the 2013 conference of the North American chapter of the association for computational linguistics: human language technologies*, 380–390.
- Pennebaker, J. W.; Francis, M. E.; and Booth, R. J. 2001. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates* 71(2001):2001.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Rammstedt, B., and John, O. P. 2007. Measuring personality in one minute or less: A 10-item short version of the big five inventory in english and german. *Journal of research in Personality* 41(1):203–212.
- Rangel, F.; Rosso, P.; Potthast, M.; Stein, B.; and Daelemans, W. 2015. Overview of the 3rd author profiling task at pan 2015. In *CLEF*. sn.
- Siddique, F. B., and Fung, P. 2017. Bilingual word embeddings for cross-lingual personality recognition using convolutional neural nets. *Learning* 21:22.
- Sulea, O.-M., and Dichiu, D. 2015. Automatic profiling of twitter users based on their tweets: Notebook for pan at clef 2015. In *CLEF (Working Notes)*.
- Upadhyay, S.; Faruqui, M.; Dyer, C.; and Roth, D. 2016. Cross-lingual models of word embeddings: An empirical comparison. *arXiv preprint arXiv:1604.00425*.
- Verhoeven, B.; Daelemans, W.; and De Smedt, T. 2013. Ensemble methods for personality recognition. In *Proceedings of the Workshop on Computational Personality Recognition*, 35–38.
- Zhang, Y.; Liu, J.; Hu, J.; Xie, X.; and Huang, S. 2017. Social personality evaluation based on prosodic and acoustic features. In *Proceedings of the 2017 International Conference on Machine Learning and Soft Computing*, 214–218. ACM.