Learning to Generate Personalized Product Descriptions

Guy Elad
Technion
sguyelad@cs.technion.ac.il

Ido Guy
eBay Research
idoguy@acm.org

Slava Novgorodov
eBay Research
snovgorodov@ebay.com

Benny Kimelfeld
Technion
bennyk@cs.technion.ac.il

Kira Radinsky
Technion
kirar@cs.technion.ac.il

ABSTRACT

Personalization plays a key role in electronic commerce, adjusting the products presented to users through search and recommendations according to their personality and tastes. Current personalization efforts focus on the adaptation of product selections, while the description of a given product remains the same regardless of the user who views it. In this work, we propose an approach to *personalize product descriptions* according to the personality of an individual user. To the best of our knowledge, this is the first work to address the problem of generating personalized product descriptions. We first learn to predict a user's personality based on past activity on an e-commerce website. Then, given a user personality, we propose an extractive summarization-based algorithm that selects the sentences to be used as part of a product description in accordance with the given personality. Our evaluation shows that user personality can be effectively learned from past e-commerce activity, while personalized descriptions can lead to a higher interest in the product and increased purchase likelihood.

ACM Reference Format:
Guy Elad, Ido Guy, Slava Novgorodov, Benny Kimelfeld, and Kira Radinsky. 2019. Learning to Generate Personalized Product Descriptions. In *The 28th ACM International Conference on Information and Knowledge Management (CIKM'19), November 3–7, 2019, Beijing, China.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3357384.3357984

1 INTRODUCTION

High-quality product content, presented on an e-commerce website, has been empirically shown to positively influence user engagement, and in turn conversion rates and volumes of sales [31, 37]. Yet, the content of a product page is typically identical for all users who view it. In this work, we propose to personalize product content, descriptions, in particular, and adapt it according to the individual user who is viewing the product's page on an e-commerce website. We focus on product descriptions, as they pose the richest and most elaborate type of content, which can be changed in a more meaningful way for different users, compared to other content types, such as title [21], image [22], and attributes [45]. Personalization has been a key driving force in the e-commerce domain in the past two decades [23], but research has focused on surfacing personalized recommendation of products [1], rather than on the content of the products already presented to the user. To the best of our knowledge, this is the first work to address the task of personalizing product descriptions.

At the core of our suggested personalization technique is a personality model. We hypothesize that different personalities react differently to a given content. Although there is no single method for defining personality, the five-factor personality taxonomy [32, 39, 58], also known as the "Big Five" model of personality traits, is considered as one of the most well-studied and is widely accepted by psychologists as the current definitive model for the measure of personality. We therefore adopt this model to characterize an e-commerce user's personality. Among their many practices, the five personality traits have been shown relevant and influential over a broad set of consumer behaviors and characteristics in online shopping [16, 28, 59, 60].

In the most direct way, the five personality traits for a given individual are inferred from a rather long psychological questionnaire with dozens of questions. Requesting e-commerce users to complete such questionnaires is not likely to yield high response rates, due to the substantial effort required. Moreover, personality may considerably change over time [57]. Our first experiment therefore tackles the problem of automatically identifying an e-commerce user's personality. Previous work has shown that the personality can be predicted from social media behavior, such as tweets, Instagram photos, or Facebook activity [19, 20, 55]. In this work, we examine how a personality can be predicted based on the user's past e-commerce activity, including the number of purchases and their temporal patterns, the categories of the purchased products, their price distribution, and their titles. We conduct a user study wherein participants are asked to perform a psychological personality survey and provide access to their e-commerce purchase history. We develop classifiers, one per each of the five traits, and demonstrate that the personality can indeed be effectively predicted from historical user activity on an e-commerce website. We also perform a detailed analysis of the different features, providing insights on what historical purchase characteristics correlate with the different personality traits.

With high volumes of new items offered for sale every day, it is impractical to manually curate product descriptions per personality. Moreover, little research has been conducted to understand the required characteristics of such descriptions. Therefore, we first present a study of the language characteristics of the purchased product descriptions by different personality traits. Based on this

CIKM '19, November 3–7, 2019, Beijing, China

ACM ISBN 978-1-4503-6976-3/19/11…$15.00
https://doi.org/10.1145/3357384.3357984

analysis, we design an algorithm that uses extractive summarization to produce a product description adapted to a given personality. The description is generated by selecting a small set of sentences out of a larger set of candidates, ultimately yielding short descriptions of three sentences. Such descriptions are especially suitable for a quick consumption on a mobile device, as e-commerce mobile applications have demonstrated a tremendous growth and already account for a major portion of the overall e-commerce traffic [34].

An example is shown in Figure 1. The example presents two personalized descriptions generated by our algorithm—one adapted to a high Extraversion and Openness personality (left), and the other adapted to a low Extraversion and Openness personality (right). It is evident that the descriptions are substantially different, as the high Extraversion and Openness description highlights the external characteristics of the product, while the "opposite" description highlights the functionality.

To evaluate the personalized descriptions of our algorithm, we conduct two experiments, with 75 and 31 participants, respectively. In the first experiment, participants' personality is mapped using a questionnaire, while in the second it is predicted based on their historical e-commerce activity, reflecting an end-to-end application. After inferring the personality, participants of both experiments are asked to compare the personalized description against two alternative descriptions (also of three sentences): one is generated using a traditional (non-personalized) summarization technique and the other is created using personalization with the opposite personality. Results of both experiments show that the descriptions adapted to the user's personality are found more appealing and are also more likely to positively influence a hypothetical purchase decision.

The remainder of this paper is organized as follows. In the next section, we shortly review the Big Five personality model. In the following section, we discuss related work. In Section 4, we describe our first experiment for personality prediction from e-commerce user activity. The following section analyzes the language models characterizing different personality traits; these models serve as the basis for the personalized description generation algorithm, described in Section 6. In Section 7, we describe the experiments for evaluating the personalized descriptions. We conclude and suggest future directions in Section 8.

Our primary contributions can be summarized as follows:

- Our research is the first to suggest the personalization of product descriptions on e-commerce websites.
- We present the first study of user personality prediction using e-commerce activity.
- We present an algorithm for generating personalized product descriptions according to users' Big Five personality traits.
- We demonstrate that personalized descriptions can lead to higher user interest and purchase prospects.

2 THE OCEAN PERSONALITY MODEL

The Big Five personality traits, also known as the OCEAN model, classifies the human personality across a taxonomy of five broad dimensions: *Openness* to experience, *Conscientiousness*, *Extraversion*, *Agreeableness*, and *Neuroticism* (hence the acronym OCEAN). The model is universal and culture-independent: a study that inspected people from more than 50 different cultures found that the five traits could be globally used to describe personalities. Moreover, psychologists believe that the model is not only universal, but also has biological origins [26]. The five traits are commonly described as follows.

- **Openness** to experience: refers to a sense of curiosity about others and the world. This usually means that the person is creative, willing to consider new ideas, and has a good imagination.
- **Conscientiousness**: describes a detail-oriented nature, with good impulse control, who spends time preparing and enjoys finishing tasks on time.
- **Extraversion**: characterized by excitability and high amounts of emotional expressiveness. Extroverts are people who enjoy being active with others and make friends easily. They usually enjoy being the center of attention and feel energized when around other people.
- **Agreeableness**: usually relates to qualities such as trust, altruism, and kindness. Agreeable people enjoy helping and contributing to the happiness of others. They tend to be more cooperative.
- **Neuroticism**: describes a tendency to have unsettling thoughts and feelings. These individuals usually have mood swings and get easily stressed and upset.

The model has been broadly applied to different tasks, such as predicting success at work [5], academic achievements [35], usefulness within the professional context [12], and relationship partner preferences [3, 43]. In addition, there has been prior work on motivating online shopping via the Big Five model [10, 28].

3 RELATED WORK

Textual product descriptions received research attention in the recent years, indicating their importance to the e-commerce domain. Descriptions can be viewed as an extension of the information conveyed about the product in other textual fields, such as the title [21] and attributes [45], providing additional details about the product and the reason to its purchase. Several works studied the extraction of attribute-value pairs from product descriptions, in order to enrich the product's structured representation using

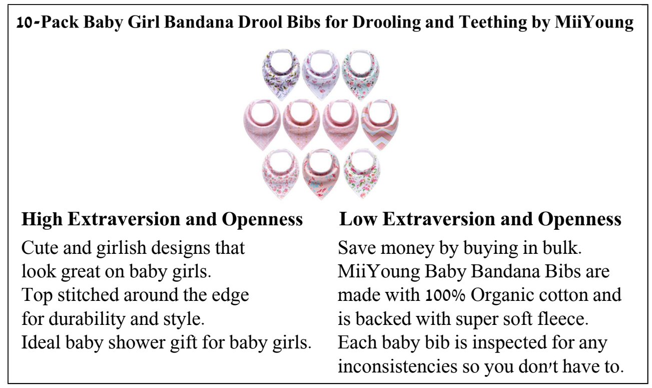

Figure 1: Example product with two personalized description versions: for high Extraversion and Openness personality (left), and for low Extraversion and Openness (right).

both supervised [49] and unsupervised [56] approaches. Dumitru et al. [14] applied text mining and clustering techniques to product descriptions in order to recommend product features for a given domain. Pryzant et al. [50] showed that product descriptions can be used for predicting the product's business outcome. Finally, a recent work by Novgorodov et al. [44] proposed a method for generating non-personalized product descriptions from user reviews.

In our work, we focus on personalized product descriptions. We split the process of generating the personalized descriptions into two separate steps: (1) predicting the user personality from online purchase activity, and (2) creating a personalized description for a given personality.

3.1 Personality Prediction

Numerous studies explored the connection between the language used and consumed by individuals and their personality traits [38, 47, 55]. With billions of users participating and sharing self-authored content, social media provides a tremendous opportunity for personality modeling. Many of the studies have focused on Twitter and Facebook, with various methods for predicting personality based on tweets and status updates [19, 20, 55]. For example, Golbeck et al. [20] collected Facebook activity information and found, among other things, that Extraversion and Neuroticism correlated with the number of friends. In addition, they used linguistic features, such as word counts in a user-authored posts, for personality prediction. A related work by the same authors [19] used Twitter data to predict the personality using both the MRC Psycholinguistic Database and the Linguistic Inquiry and Word Count (LIWC) tool, which is widely used to classify words into psychologically-meaningful categories. Schwartz et al. [55] analyzed over 700 million words and phrases originating from 75K Facebook posts and reported characterizing phrases for the different personality traits. For example, highly neurotic people used the terms "boring," "sick of," and "depressed," whereas "success," "beautiful day" and "blessing" were used by individuals with low Neuroticism. Instagram, a photo-sharing social network, also attracted research interest in the context of personality prediction. Ferwerda et al. [17] found that the filters on Instagram, reflecting the desired look of a user's photos, can serve as a useful feature for predicting personality. A later work explored personality prediction using both visual (e.g., hue, valence, saturation) and content features extracted from the photos [18]. In the music domain, a recent work by Qiu et al. [51] found that the lyrics of preferred songs is more predictive of personality than melody.

Very few works have been dedicated to the connection between user personality and e-commerce. Yang et al. [60] studied the tie between an individual's brand preferences and personal traits. The features were obtained from both psychometric surveys and automated social media analysis. Huang and Yang [28] investigated the relationship between personality traits and online shopping motivations by mapping the Big Five model to five motivation types: adventure, idea, sociality, lack of sociality, and convenience. To the best of our knowledge, no previous work has attempted to predict the user personality through online shopping patterns.

3.2 Personalized Content Generation

In recent years, there has been a considerable growth in the amount of text data and the need for automatic text summarization methods. Summarization techniques are generally divided into two main approaches: *extractive* [15] and *abstractive* [42]. Extractive summarization typically works at the sentence level and selects a subset of sentences from the original text that should be included in the summary. The abstractive approach selects words based on semantic understanding, and produces new sentences that capture the meaning of the original text. In this work, we take an extractive approach to generate the personalized descriptions.

Personalized textual content generation attracted research interest in various domains. Roy at al. [54] proposed a technique for the automatic generation of different marketing messages, targeted to different groups (students, designers, developers and managers). Ding et al. [13] studied persuasive message generation. They learned a mapping between personality traits and aspect importance to automatically customize a message content to enhance its appeal to receivers. Carenini and Moore [9] developed a model for automatic argument generation and showed that addressing user preferences increases its effectiveness. In a recent work, Krishna et al. [36] presented a framework for summary generation that takes into account the linguistic preferences of the specific audience who consume the generated summary. In the field of computer-based learning, Reichelt et al. [52] showed that personalization of learning materials can increase motivation and learning outcomes. Zander et al. [61] studied the effect of personalization on students' attention allocation using eye-tracking methods. They found that the personalized version of learning materials was more appealing and inviting. To the best of our knowledge, this work is the first to explore the generation of personalized product descriptions in the e-commerce environment.

4 PERSONALITY PREDICTION

Our first experiment examines whether and how a user's personality can be predicted from their historical purchase activity in an e-commerce website. Participants were recruited using the Figure Eight crowdsourcing platform (formerly known as CrowdFlower)[1]. Each participant was asked to answer the 44-question version of the Big Five Inventory personality test, commonly referred to as BFI-44 [32]. In addition, participants provided basic demographic information, including age, gender, nationality, and last but not least, their user name on a large global e-commerce website, which allowed us to retrieve their purchase history on the site. After filtering out participants with invalid user names, or without purchase history (fewer than 10 purchases on the site in total), we were left with 447 participants. The average participant age was 35.5 (std: 10.3, median: 33, min: 17, max: 70), with 54.1% males and 45.9% females. Participants originated from 31 different countries, including the United States (34.3%), Italy (11.2%), the United Kingdom (10.7%), Spain (10.1%), Canada (9%), and Russia (7.4%). On average, the number of purchases per participant on the e-commerce website was 100.7 (std: 197.5, median: 33, max: 2099). While this sample may not perfectly represent the buyer population on the platform, it includes active buyers, who are the initial target population for

[1] https://www.figure-eight.com

Table 1: Average, standard deviation, 30th percentile, median, and 70th percentile for each of the five personality traits as derived for 447 participants in the personality prediction experiment (score scale is 1–5).

Trait	Avg	Std	30%	Median	70%
Openness	3.56	0.58	3.28	3.50	3.90
Conscientiousness	3.66	0.66	3.25	3.62	4.00
Extraversion	2.95	0.79	2.50	3.00	3.37
Agreeableness	3.63	0.63	3.33	3.67	4.00
Neuroticism	2.86	0.82	2.37	2.87	3.25

our personalized descriptions. Table 1 presents the statistics of the personality trait scores derived from the questionnaire for our 447 participants.

For the personality prediction task, per each of the five traits, the positive class included the participants with a personality score above the 70th percentile (referred to as *High* polarity) and the negative class included the participants with a personality score below the 30th percentile (*Low* polarity; see Table 1). This task definition is aligned with previous studies of personality prediction in other domains [29, 38, 55].

For each user, we created a basic set of features, considering their past purchases on the e-commerce website during the years 2014-2018, as follows:

- Purchase history: number of purchases per year; dates of the earliest and most recent purchases.
- Product price statistics (min, 25th percentile, median, 75th percentile, max).
- Temporal patterns: percentage of purchases per day of the week (weekday, weekend); percentage of purchases per time of the day (daytime, evening, night).
- Percentage of purchases in each of the main categories. We considered a category (e.g., Fashion or Books & Magazines) as a *main category* if at least 5% of the participants purchased a product that belongs to this category. Overall, 35 categories were deemed as main categories based on this definition.
- Titles of the products purchased by the user. For the title representation, we used fastText word embeddings[2] pre-trained on Common Crawl and Wikipedia [25, 33], weighted based on each word's TF-IDF score [4].[3]

In addition, we included as features the participant's demographic attributes: age range, gender, and nationality.

We created five binary classifiers, one for each of the personality traits. Each classifier's task was to predict, given the participant's features as described above, whether the polarity of the personality trait would be *High* or *Low*. We experimented with both Logistic Regression [40] and XGBoost [11] , with 5-fold cross validation to tune the hyper-parameters and evaluate the classifiers. For Logistic Regression, we tuned the regularization strength and norm (L1 versus L2). For XGBoost, we also tuned the maximum depth of a tree, minimum split loss, subsample ratio of the training instances, subsample ratio of features per tree, and the learning rate and number of rounds (trees). As an evaluation metric, we used the area under the ROC curve (AUC). In addition, we used ANOVA F-value for feature selection to avoid overfitting [46].

[2]https://fasttext.cc/

[3]We experimented with other embedding and weighting techniques, but only report the combination that yielded the best performance. We conjecture that the fastText representation of each word as a bag of character n-grams, in addition to the word itself, helped capture spelling mistakes and amplifiers (e.g. "amazinggg"), which are common in product titles.

Table 2: AUC and Accuracy performance of the XGBoost and Logistic Regression classifiers for the personality prediction task.

Personality trait	AUC		Accuracy	
	XGBoost	LR	XGBoost	LR
Openness	0.84	0.80	78%	78%
Conscientiousness	0.86	0.85	79%	79%
Extraversion	0.89	0.86	85%	83%
Agreeableness	0.83	0.82	77%	75%
Neuroticism	0.85	0.84	81%	79%

Table 3: AUC performance of the XGBoost classifier for the personality prediction task when using ('Only') or disregarding ('Exclude') different feature families.

Feature family	Only					Exclude				
	O	C	E	A	N	O	C	E	A	N
Purchases	0.59	0.57	0.67	0.64	0.54	0.82	0.85	0.87	0.82	0.84
Price	0.60	0.57	0.62	0.64	0.60	0.83	0.85	0.86	0.81	0.84
Temporal	0.56	0.52	0.60	0.60	0.50	0.82	0.85	0.86	0.82	0.84
Categories	0.67	0.65	0.70	0.66	0.69	0.82	0.84	0.85	0.81	0.83
Titles	0.81	0.83	0.85	0.79	0.82	0.68	0.73	0.78	0.73	0.73
Demographic	0.60	0.73	0.63	0.67	0.68	0.82	0.83	0.86	0.81	0.84

Table 2 presents the AUC and the accuracy scores for both the Logistic Regression and XGBoost classifiers. It can be seen that the XGBoost classifier consistently achieved slightly better results than Logistic Regression, for each of the five personality traits. Performance is rather similar across the five traits, ranging, for XGBoost, from AUC of 0.83 for Agreeableness to 0.89 for Extraversion. Overall, these results indicate that personality traits can be effectively predicted from a user's history of online purchases.

We also set out to explore the importance of the different feature families, as listed above, to the personality prediction task. Table 3 shows the AUC of the XGBoost classifier when using only specific feature families, and when using all features except for these feature families (ablation tests). Evidently, the title features show the strongest performance of all feature families. Using title features alone yielded performance that was close to the overall performance as reported in Table 2 for each of the OCEAN traits. In addition, the title feature family was the only one whose removal substantially degraded the performance of the classifiers, indicating that the title captures discriminative signals that are not captured by any of the other features.

Other than titles, the category and demographic features showed relatively high performance. Table 4 presents, for each of the top categories, which account together for over 50% of the purchases, the portion of products purchased by users with *High* and *Low* personality traits (out of all products purchased by users with the respective traits). Differences can be observed across all categories

Table 4: Percentage of products purchased in top categories by each of the *High* ('H') and *Low* ('L') personality traits. Boldfaced portions are at least 25% higher than their *Low*/*High* counterparts.

Trait		Cell Phones& Accessories	Fashion	Home& Garden	Toys& Hobbies	Jewellery& Watches	Office	Books& Magazines	Baby
Openness	H	**14.96%**	13.16%	9.29%	4.06%	**6.33%**	3.75%	1.09%	**1.80%**
	L	10.24%	11.05%	8.10%	**6.61%**	3.82%	3.26%	**2.89%**	1.12%
Conscientiousness	H	11.78%	11.48%	9.44%	5.72%	3.50%	**2.80%**	2.22%	**1.61%**
	L	12.57%	11.95%	7.83%	6.68%	4.26%	1.47%	**3.34%**	1.05%
Extraversion	H	**16.97%**	13.57%	8.90%	4.00%	**5.92%**	1.66%	2.71%	1.28%
	L	10.00%	11.50%	8.61%	**7.81%**	3.37%	**3.24%**	2.67%	1.21%
Agreeableness	H	**14.91%**	13.18%	**10.49%**	5.96%	**4.90%**	1.88%	1.69%	**1.51%**
	L	9.76%	12.96%	7.39%	7.27%	3.77%	**3.74%**	1.55%	0.76%
Neuroticism	H	9.00%	**13.43%**	9.49%	5.40%	**5.29%**	2.25%	2.29%	1.58%
	L	**14.70%**	9.62%	10.82%	6.59%	3.00%	**3.41%**	2.74%	1.49%

and personality traits, which explains why the categories are a fairly effective discriminative feature family. For example, buyers with *Low* openness have substantially higher portions in books than buyers with *High* Openness. Buyers who are introverts, conscientious, and disagreeable have almost double portions in office products than their extroverts, unconscientious, and agreeable counterparts, respectively. Extroverts purchase higher portions of their products in cellphones and jewellery/watches, while neurotic individuals purchase more fashion and jewellery/watches products.

5 PERSONALITY TRAIT LANGUAGE MODEL

Having the BFI-44 questionnaire data (from the previous section, used for inferring the user personality) along with user identifiers on a large e-commerce website, we set out to explore the language that characterizes the content purchased by users with different personalities. For each trait and polarity (*Low* or *High*), we created a corpus as follows: we considered all descriptions of products purchased by users with the given trait and polarity. We filtered out short descriptions of 15 words or fewer (7.9%), then non-English descriptions (31.31%), and then limited the number of descriptions per user to 35 (sampled uniformly at random), to avoid a strong bias towards heavy users (46.4% of the descriptions were removed due to this criterion). In total, we were left with 4195 descriptions (from 385 users), which were mapped to the different trait-and-polarity pairs, with at least 778 descriptions associated with each pair. The average number of descriptions per user was 10.9 (std: 11.83, median: 5).

With a collection of descriptions representing users with *Low* and *High* personality traits who purchased the respective products, we aimed to examine the language difference between the *Low* and *High* collections for each of the five traits. Specifically, we sought for distinctive terms that characterize the language model of the *High* collection of descriptions compared to the *Low* collection, and vice versa. To this end, we used Kullback-Leibler (KL) divergence, which is a non-symmetric distance measure between two given distributions [7]. In particular, we calculated the unigrams that contributed the most to the KL divergence between the language model of the *Low* collection and the language model of the *High* collection, and vice versa, per each of the OCEAN traits. Formally, we computed the *contribution* of each word w across personality traits $T \in \{\mathsf{O}, \mathsf{C}, \mathsf{E}, \mathsf{A}, \mathsf{N}\}$ and polarity $p \in \{\mathsf{L}, \mathsf{H}\}$, as follows:

Table 5: Most distinctive unigrams for *High* vs. *Low* ('*High*') and *Low* vs. *High* ('*Low*') Extraversion and Agreeableness based on descriptions of products purchased by users with corresponding personality traits.

Extraversion		Agreeableness	
High	Low	High	Low
color	sale	perfect	kit
picture	adult	weekends	version
slim	old	kids	vehicle
gift	suitable	cute	heavy
pictures	easily	color	inside
new	includes	satisfaction	original
design	fast	great	classic
style	priority	human	white
brand	policy	comfortable	audio
galaxy	available	creative	memory
apple	need	beach	ready
greatly	grey	day	temperature
party	affordable	christmas	computer
decoration	strength	toy	used
popular	quickly	happy	indicator
premium	comfort	summer	instruction
creative	white	wedding	adjustment
friends	saving	environmental	combination
enjoy	book	good	sound
games	electrical	lovely	additional

$$\mathrm{wc}_T^p(w) = \mathrm{df}_T^p(w) \times \log\left(\frac{\mathrm{df}_T^p(w)}{\mathrm{df}_T^{-p}(w)}\right) \tag{1}$$

where $\mathrm{df}_T^p(w)$ is the fraction of descriptions containing the word w in the language model of T and p, and $-p$ denotes the opposite polarity (switching between L and H).

Inspecting the top unigrams in the list of most distinctive terms for *Low* and *High* personality traits, we can observe trends that coincide with different traits. These trends are demonstrated in Table 5, which lists the most distinctive unigrams for *Low* and *High* Extraversion and Agreeableness. It can be observed that for extrovert users (*High* Extraversion), visual and innovative product aspects are reflected through many of the terms, e.g., `style`, `new`, `design`, `brand`, `picture`, `slim`, `decoration`, and `color`. Well-known brands and models, such as `Apple` and `Galaxy` and expressions of fun end excitement, e.g., `party`, `friends`, `enjoy`, `games`, `popular`, `premium`, `feeling`, and `greatly`, are also among the unigrams characterizing extroverts. For introverts (*Low* Extraversion), on the other hand, terms that reflect practice and functionality are among the top unigrams, e.g., `suitable`, `affordable`, `comfort`, `strength`, `easily`, `fast`, and `priority`, as well as more descriptive details, such as `includes`, `electrical`, `old`, and `adult`.

For *High* Agreeableness, many of the distinctive terms are positive adjectives e.g., `perfect`, `cute`, `great`, `creative`, `happy`, `lovely`, `good`, and relate to people and family, e.g., `kids`, `human`, `wedding`, and leisure, e.g., `weekends`, `christmas`, `summer` and `beach`. On the other hand, *Low* Agreeableness is associated with more concrete and pragmatic terms, e.g., `ready`, `used`, `instruction`, `combination`, `version`, `kit`, `original`, `audio`, `temperature` and personal-use products, such as `vehicle` or `computer`.

Another demonstration of the differences between the language models tied with personality traits is the use of color words. Table 6 shows the top color unigrams according to their KL divergence

Table 6: Most distinctive color unigrams for *High* vs. *Low* ('*High*') and *Low* vs. *High* ('*Low*') Extraversion and Conscientiousness based on descriptions of products purchased by users with corresponding personality traits.

Extraversion		Conscientiousness	
High	Low	High	Low
brown	grey	green	black
red	white	brown	yellow
black	silver	cyan	purple
orange	mint	blue	grey
gold	blue	olive	silver

for *High* versus *Low* Extraversion and Conscientiousness, and vice versa. For Extraversion, *High* values are associated with warm colors and black, while for *Low* Extraversion, the distinctive colors include cold colors and white. These findings are in line with past literature, which found that extroverts prefer warm colors while introverts are attracted to cold hues [8, 53]. For *High* Conscientiousness, the list includes mostly cold colors of green and blue, while the *Low* list is topped by black and also includes grey and silver. Generally for Conscientiousness, colors were more prominent on the *Low* list: 17 different colors were at the top quartile of distinctive terms (75th percentile and above), while for *High* only 4 colors were within the top quartile. Color may play a smaller role for highly conscientious individuals, who are task-focused and dutiful [39].

The emerging differences between the description language of products purchased by users with *High* versus *Low* personality traits, as surfaced through the KL divergence, are used as a key building block of the algorithm that produces personalized descriptions. The next section describes this algorithm in detail.

6 PERSONALITY-BASED DESCRIPTION

In this section, we present our algorithm for producing personalized product descriptions. In this algorithm, we apply extractive summarization over a given product description. Similarly to the well-known LexRank algorithm for extractive summarization [15], our algorithm constructs a sentence graph, and applies random walks thereof to produce the summarization. In our case, the graph is built to adapt to the user personality, and it incorporates the personality language model of the user for creating the summary. Hence, we refer to our algorithm as *Personalized LexRank*.

The algorithm takes as input the user's personality-trait profile (identified by *High* or *Low* on each trait) and a list of candidate description sentences, and produces as output a sequence of k selected sentences. Formally, we associate with a user a *personality profile* (or just *profile* for short)

$$P : \{\mathsf{O}, \mathsf{C}, \mathsf{E}, \mathsf{A}, \mathsf{N}\} \rightarrow \{\mathsf{L}, \mathsf{H}\}$$

that maps every trait T to a polarity $P(T)$, which is either L (*Low*) or H (*High*). We are given as input a profile P and a collection S of candidate sentences.

We define the *average contribution of a term w with respect to a profile P*, denoted ac(w), as follows.

$$\mathrm{ac}_P(w) = \frac{1}{5}\sum_{i=1}^{5} \mathrm{wc}_T^{P(T)} \tag{2}$$

Recall that wc_T^p is defined in Equation (1). The average contribution of a *sentence* $s \in S$ to a given profile P is defined by:

$$\mathrm{ac}_P(s) = \frac{1}{|s|}\sum_{w \in s} \mathrm{ac}_P(w) \tag{3}$$

where $|s|$ is the number of words in the sentence.

The algorithm constructs a *sentence graph*, which is a graph that has the set S of candidate sentences as nodes, an edge between every two sentences (hence, it is a complete graph), a weight on every node, and a weight on every edge. The weight of a node (sentence) s is viewed as its prior *personalized stochastic probability vector*, denoted PV(s):

$$\mathrm{PV}_P(s) = \frac{\mathrm{ac}_P(s)}{\sum_{s' \in S} \mathrm{ac}_P(s')} \tag{4}$$

The weight of an edge (s, s') is define using the *personalized similarity* between the sentences s and s', denoted $\mathrm{PS}_P(s, s')$. It is defined by the cosine similarity between the bag-of-word representations of the two sentences, weighted by the user personality language model that, intuitively, assigns higher values to words that affect the personality. Formally, we have the following.

$$\mathrm{PS}_P(s, s') = \frac{\sum_{w \in s \cap s'} \mathrm{tf}(w, s) \cdot \mathrm{tf}(w, s') \cdot \mathrm{idf}(w)^2 \cdot \mathrm{ac}_P(w)^2}{\sqrt{\sum_{w \in s} F(w, s)} \times \sqrt{\sum_{w' \in s'} F(w', s')}} \tag{5}$$

where $\mathrm{tf}(w, s)$ is the number of occurrences of w in s, and $\mathrm{idf}(w)$ is the inverse document frequency (IDF) of w, and

$$F(w, s) = \left(\mathrm{tf}(w, s) \cdot \mathrm{idf}(w) \cdot \mathrm{ac}_P(w)\right)^2 .$$

The weight of an edge $\{s, s'\}$ is obtained by normalizing $\mathrm{PS}_p(s, s')$ so that every $\mathrm{PS}_p(s, \cdot)$ defines a probability distribution over the sentences. We refer to this value as the *transition probability* and denote it by $\mathrm{Tr}_p(s, s')$.

$$\mathrm{Tr}_P(s, s') = \frac{\mathrm{PS}_P(s, s')}{\sum_{s'' \in S} \mathrm{PS}_P(s, s'')} \tag{6}$$

We now define a personalized random walker over the graph we constructed. The random walk is parameterized by a *personalized dumping factor* d that determines the probability of jumping from a sentence s to another sentence, based on the following policy:

(1) With probability d, jump to s' with probability $\mathrm{Tr}_P(s, s')$;
(2) With probability $1 - d$, jump to s' with probability $\mathrm{PV}_P(s')$.

Next, we compute the stationary distribution of every node in this random walk. Formally, we assign to each sentence $s \in S$ an index $i \in \{1, \dots, |S|\}$. Consequently, we view Tr_P as an $|S| \times |S|$ matrix and PV_P as a vector of length $|S|$. The stationary distribution is a probability vector q of length $|S|$ (i.e., $q_i \geq 0$ and $\sum q_i = 1$) that satisfies the equation:

$$q = [d \cdot \mathrm{Tr}_P + (1 - d) \cdot \mathrm{MV}_P]^{\mathsf{T}} q$$

where MV_P is the $|S| \times |S|$ matrix obtained by replicating PV_P $|S|$ times. We refer to the probability $q(s)$ of the sentence s as the *personalized LexRank score* of s. Intuitively, the parameter d controls the trade-off between the two policies (Tr_P and MV_P). Lower values of d promote sentences that are adapted to the profile P, whereas high d values promote sentences that are central to the description. We estimate the personalized LexRank score using a simple iterative algorithm, namely the *power method* [15]. Since the above defined Markov chain is irreducible and aperiodic, the algorithm is guaranteed to terminate. To generate a description with k sentences, we simply select the k sentences with the highest personalized LexRank score. [4]

7 DESCRIPTION EVALUATION

In this section, we describe our second set of experiments, which aim to examine the effect of personalization, using the personalized LexRank algorithm, on product descriptions. Specifically, we set out to explore two main questions: do users of e-commerce websites find the personalized descriptions more appealing [27] and, hypothetically, would the personalized descriptions increase the likelihood of purchasing the product. We experimented with descriptions of k=3 sentences. These represent relatively concise content, which can be quickly consumed in its entirety, and is especially suitable for small-screen devices, such as mobile phones [34].

7.1 Experimental Setup

For our experiment, we collected a set of 20 best-selling products[5] from the large e-commerce website used to inspect purchase activity in our first experiment (Section 4). These products spanned the top categories on the site, as listed in Table 4. The average product price was 50.38 US Dollars (std: 6.78, median: 27.99, min: 3.99, max: 219.91). Each product included a professionally-written description, with 15.04 sentences on average (std: 3.48, median: 14, min: 10, max: 22) and an average of 18.8 words per sentence (std: 6.78, median: 18, min: 11, max: 34). For each product, we generated three *versions* of 3-sentence descriptions, with its original description sentences serving as the initial set of candidate sentences, using: (1) vanilla LexRank (*'neutral'*); (2) personalized LexRank using the participant's personality traits (*'personalized'*); and (3) personalized LexRank using the participant's opposite traits (*'opposite'*). Personalization was performed based on a pair of traits out of the Big Five. We experimented with two such pairs: (1) Openness and Extraversion ('O+E'), and (2) Conscientiousness and Agreeableness ('C+A'). This type of pairing has been commonly suggested by researchers, who associated 'O+E' with sensation seeking [2] and latent inhibition [48], and 'C+A' with prosocial behavior [24], effortful control [30], and compliance with authority [6].

For each trait, we determined the user's personality as *High* or *Low* (e.g., introvert vs. extrovert for Extraversion) based on the median value across our participants. This way, participants were evenly split between the two polarities of each trait. The personalized version was generated using the combination of the participant's derived polarities (*Low* or *High*) for each of the two traits, while the opposite version was created using the two opposite polarities. For example, if a participant had high Openness and high Extraversion, the personalized description was based on this combination, while the opposite description was based on low Openness and low Extraversion.

For the neutral description, we used a damping factor of 0.15 in the random walk performed by LexRank [15]. For the personalized and opposite descriptions, we experimented with different values of the personalized dumping factor d (0.01, 0.1, 0.2, and 0.3), in a round robin order across our participants.

Table 7: Distribution of the number of identical sentences between different description versions.

	0	1	2	3
Neutral–Personalized	38.75%	48.75%	11.25%	1.25%
Personalized–Opposite	90%	10%	0	0
Neutral–Personalized, $d = 0.01$	50%	43.75%	6.25%	0
Neutral–Personalized, $d = 0.1$	50%	37.50%	12.50%	0
Neutral–Personalized, $d = 0.2$	45.83%	41.67%	8.33%	4.17%
Neutral–Personalized, $d = 0.3$	16.67%	66.67%	16.67%	0

Table 7 (upper section) shows the distribution of overlapping sentences across our description versions. While for nearly half of the products and personality traits, there was one overlapping sentence between the neutral description and the personalized description, only rarely was there an overlap of two sentences or more. As could be expected, the overlap between the personalized and opposite versions was even lower, with 90% of all products and personality traits having no overlapping sentences at all. Overall, the overlap between the versions was low enough to make their comparison interesting.

The lower section of Table 7 presents the overlap between personalized and neutral descriptions as a function of the personalization dumping factor d. As explained in Section 6, lower values of d yield a more personalized description. Indeed, the overlap with the neutral description was generally smaller for lower values of d.

For our first of two experiments, we recruited 75 users of the e-commerce website. Of the participants, 58.7% were males, with the rest females, while the average age was 32.2 years (std:10.9, median:29, min:18, max:68). Each participant filled the BFI-44 personality test [32]; from there, the participant's personality scores for each of the 5 traits were derived. Then, each participant was presented with the 20 products, including their title, image, and the three description versions. The order of the three versions was randomized and so was the order of all 20 products. Out of the 20 products, for 10 the descriptions were personalized based on 'O+E' and for the other 10 based on 'C+A.' Participants were asked to select the description that was most appealing out of the three and then, for each version, they were asked to rate: (1) to what extent the description is appealing to them, on a 5-point Likert scale, ranging from 1 (not at all) to 5 (very much); and (2) assuming they were interested in purchasing the product, how the description would have influenced their purchase decision on a 5-point Likert scale from -2 (very negatively) to 2 (very positively), with 0 representing no effect. Figure 2 demonstrates the user interface used to collect participants' feedback. [6]

Our final experiment ties the personality prediction and description evaluation experiments together. It examines an end-to-end setting, wherein the personality traits are predicted from past user activity on the e-commerce website and the descriptions are personalized based on the predicted personality. For this experiment, we recruited 31 participants, with valid user names and at least 10 purchases on the site, as in the personality prediction experiment

[4]We publicly publish our algorithm's code: https://github.com/guyelad88/GPPD
[5]During the month of July 2018 in the United States.

[6]The product descriptions with their personalized versions, can be found in the GitHub repository: https://github.com/guyelad88/GPPD

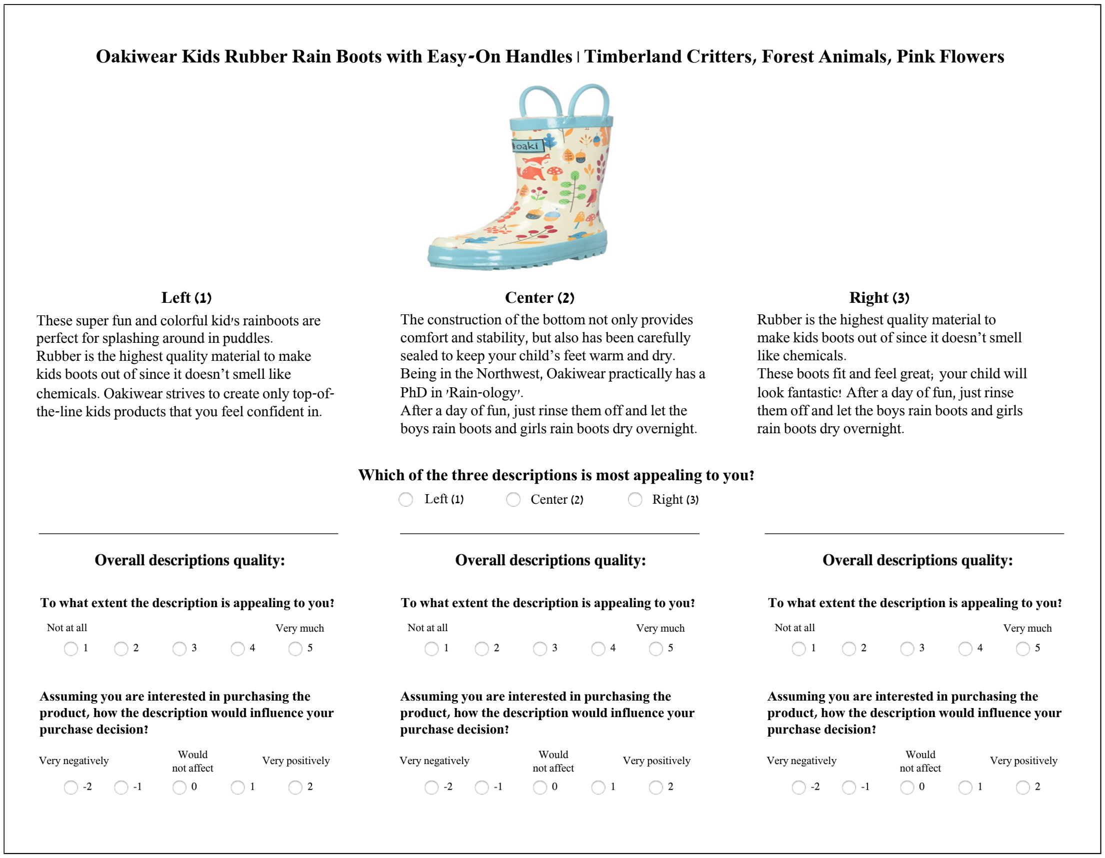

Figure 2: Product description evaluation interface.

Table 8: Percentage of "most appealing" selection for neutral ('Neu'), opposite ('Opp'), and personalized ('Per') descriptions.

Personality trait pair	Neu	Opp	Per
Openness + Extraversion	30.6%	31.9%	**37.5%**
Conscientiousness + Agreeableness	27.0%	34.9%	**38.1%**

described in Section 4. The average participant age was 30.3 years (std:5.3, median:29, min:25, max:55), with 54.8% males, and the rest females. The average number of purchases per participant on the site was 51.5 (std: 48.6, median: 29, max: 219). For each participant, the personality was predicted using the classifiers trained as described in Section 4. Then, using the predicted personality, each participant was presented with the 20 products (10 for 'O+E' and 10 for 'C+A'), including their title, image, and the three description versions, as shown in Figure 2.

7.2 Experimental Results

Table 8 presents the distribution of the "most appealing" selections in the first experiment, across the three description versions: neutral, opposite, and personalized. The personalized description received the highest portions, for both the 'O+E' and the 'C+A' trait pairs. The neutral and opposite versions consistently received lower portions, with the former, most noticeably for 'C+A', receiving the lowest portions. The average appeal ratings, depicted in Table 9, were significantly higher for the personalized descriptions of both 'O+E' and 'C+A' than the neutral and opposite descriptions, which received similar ratings to one another.[7] Overall, these results indicate that personalized descriptions have a stronger appeal than non-personalized descriptions, when personalized according to the user's appropriate traits.

Table 9: Average appeal rating for neutral ('Neu'), opposite ('Opp'), and personalized ('Per') descriptions.

Personality trait pair	Neu	Opp	Per
Openness + Extraversion	3.08	3.08	**3.21**
Conscientiousness + Agreeableness	3.21	3.20	**3.33**

Table 10 presents the average ratings for the hypothetical purchase influence. Here again, the ratings for the personalized descriptions are significantly higher than the ratings for both the neutral and opposite descriptions. The ratings for the last two are comparable, with a slight difference in favor of the neutral version.

[7] Statistical significance was measured using one-way ANOVA with Tukey post-hoc comparisons for $p<0.01$.

Table 10: Average rating of purchase influence for neutral ('Neu'), opposite ('Opp'), and personalized ('Per') descriptions.

Personality trait pair	Neu	Opp	Per
Openness + Extraversion	0.24	0.21	**0.35**
Conscientiousness + Agreeableness	0.34	0.33	**0.44**

Table 11: "Most appealing" selection portions, average appeal rating, and average purchase influence ratings for neutral ('Neu'), opposite ('Opp'), and personalized ('Per') descriptions across different segments.

Segment	Appeal selection			Appeal rating			Purchase rating		
	Neu	Opp	Per	Neu	Opp	Per	Neu	Opp	Per
Male	28.7%	33.8%	**37.5%**	3.11	3.14	**3.24**	0.24	0.25	**0.34**
Female	29.9%	32.2%	**37.9%**	3.18	3.12	**3.29**	0.35	0.28	**0.46**
Young	25.3%	36.4%	**38.3%**	3.17	3.27	**3.33**	0.36	0.40	**0.50**
Old	31.4%	31.3%	**37.3%**	3.02	2.96	**3.22**	0.13	0.04	**0.33**
d=0.01	24.6%	33.5%	**41.9%**	3.02	3.03	**3.23**	0.19	0.22	**0.41**
d=0.1	26.9%	33.2%	**39.9%**	3.21	3.17	**3.26**	0.32	0.26	**0.44**
d=0.2	30.2%	**35.2%**	34.6%	3.27	3.25	**3.30**	**0.40**	0.33	**0.40**
d=0.3	29.8%	31.8%	**38.4%**	3.11	3.12	**3.28**	0.26	0.25	**0.35**

Overall, we get an indication that the personalized descriptions are not only more appealing, but may also increase the likelihood of a transaction, which is the ultimate goal of the product's presentation on an e-commerce site.

Table 11 shows the "most appealing" portions, appeal ratings, and purchase influence ratings, for the three description versions across different population segments. The upper section focuses on gender, where it can be seen that the difference in favor of personalized descriptions is consistent and statistically significant for both males and females, across all metrics. It can be generally observed that the female ratings for both the appealing and the purchase influence questions are higher than the male ratings. The next section of Table 11 focuses on age: we compared the youngest 30% of the participants (age 26 and below) with the oldest 30% (33 and above). For both types of ratings, young participants had higher averages than the older participants, by a noticeable margin. The personalized descriptions received the best ratings from both age groups, with a more substantial gap from the other versions within the elderly group (in both groups the difference was statistically significant). Inspecting the results by the personalized dumping factor d (lower section of Table 11), which controls the personalization weight for personalized and opposite descriptions, indicates that the gap in favor of the personalized descriptions generally decreases as the value of d increases. As explained in Section 6, a smaller value of d assigns a higher weight to personalization. Therefore, these results indicate that a stronger personalization weight yields a more substantial advantage for the personalized descriptions.

Table 12 presents the results for neutral, opposite, and personalized descriptions in our end-to-end setting, as described in Section 7.1. It can be seen that the difference in favor of personalized

Table 12: Percentage of "most appealing" selection, average appeal rating, and average purchase influence ratings for neutral ('Neu'), opposite ('Opp'), and personalized ('Per') descriptions in an end-to-end setting.

Personality trait pair	Neu	Opp	Per
Most appealing selection	30.3%	32.0%	**37.7%**
Average appeal rating	3.14	3.19	**3.26**
Average purchase influence rating	0.33	0.31	**0.38**

descriptions is consistent across all three metrics. For both the "most appealing" portions and the appeal ratings, the difference was statistically significant, despite the relatively small number of participants.[8] The difference in favor of personalized descriptions slightly decreases across all metrics compared to the previous experiment, likely since personality prediction is not always accurate, as reported in Section 4. Overall, these results confirm that personalized description generation can be a feasible and productive application on e-commerce websites.

8 CONCLUDING REMARKS

Our experiments demonstrate that personalized descriptions, consistently across different metrics and different population segments, achieve a better outcome than non-personalized or oppositely-personalized descriptions, both in terms of user appeal and in terms of purchase likelihood. Assigning higher weight to the personalization factor leads to preferable results. Our personalization technique is based on user personality, which, as we show, can be predicted from past e-commerce activity, sparing the need of forcing users to fill traditional questionnaires in order to map their personality. Moreover, the personality inferred from e-commerce activity may often more reliably reflect the user's e-commerce persona [41]. Our results indicate that the embeddings of the titles of purchased products yield a particularly strong signal for personality prediction. Overall, a complete system can be built as part of e-commerce platforms, which predicts personality from user activity and in turn personalizes the product descriptions presented to a user based on their predicted personality.

Our results give a first indication that personalized descriptions can increase user satisfaction and conversion rates on e-commerce websites. They lay the foundation for further research on personalized product content. Future research may examine other personality combinations and taxonomies for the personalization task; additional methods for personalization that do not rely on personality; the personalization of other content types on the product page, such as its title, image, and attributes; and the in-vivo effect of personalized descriptions when deployed in e-commerce systems.

REFERENCES

[1] G. Adomavicius and A. Tuzhilin. 2005. Towards the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Transactions on Knowledge and Data Engineering* 17, 6 (June 2005), 734–749.

[2] Anton Aluja, Oscar Garcıa, and Luıs F Garcıa. 2003. Relationships among extraversion, openness to experience, and sensation seeking. *Personality and Individual Differences* 35, 3 (2003), 671–680.

[8] Statistical significance was measured using one-way ANOVA with Tukey post-hoc comparisons for $p<0.05$.

[3] Beth Emma Anderson. 2018. *Individual differences and romantic compatibility: the relationship between personality traits, eligibility and ideal partner preference.* Ph.D. Dissertation. UCL (University College London).

[4] Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *Proc. of ICLR.*

[5] Murray R Barrick and Michael K Mount. 1991. The big five personality dimensions and job performance: a meta-analysis. *Personnel psychology* 44, 1 (1991), 1–26.

[6] Laurent Bègue, Jean-Léon Beauvois, Didier Courbet, Dominique Oberlé, Johan Lepage, and Aaron A Duke. 2015. Personality Predicts Obedience in a Milgram Paradigm. *Journal of Personality* 83, 3 (2015), 299–306.

[7] Adam Berger and John Lafferty. 2017. Information retrieval as statistical translation. In *ACM SIGIR Forum*, Vol. 51. ACM, 219–226.

[8] Faber Birren. 1961. *Color psychology and color therapy; a factual study of the influence of color on human life.* Pickle Partners Publishing.

[9] Giuseppe Carenini and Johanna D Moore. 2006. Generating and evaluating evaluative arguments. *Artificial Intelligence* 170, 11 (2006), 925–952.

[10] Shu-Hui Chen and Kuan-Ping Lee. 2008. The role of personality traits and perceived values in persuasion: An elaboration likelihood model perspective on online shopping. *Social Behavior and Personality* 36, 10 (2008), 1379–1399.

[11] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proc. of KDD.* 785–794.

[12] Cary L Cooper and Ivan T Robertson. 2006. *International Review of Industrial and Organizational Psychology 2006.* Vol. 21. John Wiley & Sons.

[13] Tao Ding and Shimei Pan. 2016. Personalized emphasis framing for persuasive message generation. In *Proc. of EMNLP.* 1432–1441.

[14] Horatiu Dumitru, Marek Gibiec, Negar Hariri, Jane Cleland-Huang, Bamshad Mobasher, Carlos Castro-Herrera, and Mehdi Mirakhorli. 2011. On-demand feature recommendations derived from mining public product descriptions. In *Proc. of ICSE.* 181–190.

[15] Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research* 22 (2004), 457–479.

[16] Bruce Ferwerda, Markus Schedl, and Marko Tkalcic. 2015. Personality & emotional states: Understanding users' music listening needs. CEUR-WS. org.

[17] Bruce Ferwerda, Markus Schedl, and Marko Tkalcic. 2015. Predicting personality traits with instagram pictures. In *Proc. of EMPIRE.* 7–10.

[18] Bruce Ferwerda and Marko Tkalcic. 2018. Predicting Users' Personality from Instagram Pictures: Using Visual and/or Content Features? 157–161.

[19] Jennifer Golbeck, Cristina Robles, Michon Edmondson, and Karen Turner. 2011. Predicting personality from twitter. In *Proc. SocialCom.* IEEE, 149–156.

[20] Jennifer Golbeck, Cristina Robles, and Karen Turner. 2011. Predicting personality with social media. In *Proc. of CHI.* 253–262.

[21] Vishrawas Gopalakrishnan, Suresh Parthasarathy Iyengar, Amit Madaan, Rajeev Rastogi, and Srinivasan Sengamedu. 2012. Matching product titles using web-based enrichment. In *Proc. of CIKM.* 605–614.

[22] Anjan Goswami, Naren Chittar, and Chung H. Sung. 2011. A study on the impact of product images on user clicks for online shopping. In *Proc. of WWW.* 45–46.

[23] Anna Goy, Liliana Ardissono, and Giovanna Petrone. 2007. Personalization in e-commerce applications. In *The adaptive web.* 485–520.

[24] AM Grant and JM Berg. 2011. Prosocial Motivation at Work: When, Why, and How Making a Difference Makes a Difference.

[25] Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomas Mikolov. 2018. Learning Word Vectors for 157 Languages. In *Proc. of LREC.*

[26] Michael Gurven, Christopher Von Rueden, Maxim Massenkoff, Hillard Kaplan, and Marino Lero Vie. 2013. How universal is the Big Five? Testing the five-factor model of personality variation among forager–farmers in the Bolivian Amazon. *Journal of personality and social psychology* 104, 2 (2013), 354.

[27] Marc Hassenzahl. 2001. The effect of perceived hedonic quality on product appealingness. *International Journal of Human-Computer Interaction* 13, 4 (2001), 481–499.

[28] Jen-Hung Huang and Yi-Chun Yang. 2010. The relationship between personality traits and online shopping motivations. *Social Behavior and Personality: an international journal* 38, 5 (2010), 673–679.

[29] Francisco Iacobelli, Alastair J Gill, Scott Nowson, and Jon Oberlander. 2011. Large scale personality classification of bloggers. In *Proc. of ICACII.* 568–577.

[30] Lauri A Jensen-Campbell, Monica Rosselli, Katie A Workman, Mirtha Santisi, Jasmin D Rios, and Debbie Bojan. 2002. Agreeableness, conscientiousness, and effortful control processes. *Journal of Research in Personality* 36, 5 (2002), 476–489.

[31] Alice Jiang, Zhilin Yang, and Minjoon Jun. 2013. Measuring consumer perceptions of online shopping convenience. 24 (04 2013), 191–214.

[32] Oliver P John and Sanjay Srivastava. 1999. The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research* 2, 1999 (1999), 102–138.

[33] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of Tricks for Efficient Text Classification. In *Proc. of EACL.* 427–431.

[34] Gagandeep Kaur and Gagandeep Kaur. 2016. Mobile applications are major players in the world of e-commerce. *International Journal of Advanced Research in IT and Engineering* 5, 2 (2016), 13–21.

[35] Meera Komarraju, Steven J Karau, Ronald R Schmeck, and Alen Avdic. 2011. The Big Five personality traits, learning styles, and academic achievement. *Personality and individual differences* 51, 4 (2011), 472–477.

[36] Kundan, Aniket Murhekar, Saumitra Sharma, and Balaji Vasan Srinivasan. 2018. Vocabulary Tailored Summary Generation. In *Proc. of COLING.* 795–805.

[37] Moez Limayem, Mohamed Khalifa, and A. Frini. 2000. What makes consumers buy from Internet? A longitudinal study of online shopping. *IEEE Trans. Systems, Man, and Cybernetics, Part A* 30 (2000), 421–432.

[38] François Mairesse, Marilyn A Walker, Matthias R Mehl, and Roger K Moore. 2007. Using linguistic cues for the automatic recognition of personality in conversation and text. *Journal of artificial intelligence research* 30 (2007), 457–500.

[39] Robert R McCrae and Oliver P John. 1992. An introduction to the five-factor model and its applications. *Journal of personality* 60, 2 (1992), 175–215.

[40] Peter McCullagh and John A Nelder. 1989. *Generalized linear models.* Vol. 37.

[41] Filippo Menczer, W Nick Street, and Alvaro E Monge. 2002. Adaptive assistants for customized e-shopping. *IEEE Intelligent Systems* 17, 6 (2002), 12–19.

[42] Ramesh Nallapati, Bowen Zhou, Caglar Gulcehre, Bing Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint* abs/1602.06023 (2016).

[43] George A Neuman, Stephen H Wagner, and Neil D Christiansen. 1999. The relationship between work-team personality composition and the job performance of teams. *Group & Organization Management* 24, 1 (1999), 28–45.

[44] Slava Novgorodov, Ido Guy, Guy Elad, and Kira Radinsky. 2019. Generating product descriptions from user reviews. In *Proc. of WWW.* 1354–1364.

[45] Eun Joo Park, Eun Young Kim, Venessa Martin Funches, and William Foxx. 2012. Apparel product attributes, web browsing, and e-impulse buying on shopping websites. *Journal of Business Research* 65, 11 (2012), 1583–1589.

[46] Fabian Pedregosa et al. 2011. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* 12 (Nov. 2011), 2825–2830.

[47] James W Pennebaker and Laura A King. 1999. Linguistic styles: Language use as an individual difference. *Journal of personality and social psychology* 77, 6 (1999), 1296.

[48] Jordan B Peterson, Kathleen W Smith, and Shelley Carson. 2002. Openness and extraversion are associated with reduced latent inhibition: Replication and commentary. *Personality and Individual Differences* 33, 7 (2002), 1137–1147.

[49] Katharina Probst, Rayid Ghani, Marko Krema, Andrew Fano, and Yan Liu. 2007. Semi-supervised learning of attribute-value pairs from product descriptions. In *Proc. of IJCAI.* 2838–2843.

[50] Reid Pryzant, Young-Joo Chung, and Dan Jurafsky. 2017. Predicting sales from the language of product descriptions. In *Proc. of ECOM (SIGIR Workshops).*

[51] Lin Qiu, Jiayu Chen, Jonathan Ramsay, and Jiahui Lu. 2019. Personality Predicts Words in Favorite Songs. *Journal of Research in Personality* 78 (02 2019), 25–35.

[52] Maria Reichelt, Frauke Kämmerer, Helmut M Niegemann, and Steffi Zander. 2014. Talk to me personally: Personalization of language style in computer-based learning. *Computers in Human behavior* 35 (2014), 199–210.

[53] Carl Robinson. 1975. Color preference as a function of introversion and extraversion. *Perceptual and Motor Skills* (1975).

[54] Rishiraj Saha Roy, Aishwarya Padmakumar, Guna Prasaad Jeganathan, and Ponnurangam Kumaraguru. 2015. Automated linguistic personalization of targeted marketing messages mining user-generated text on social media. In *Proc. of CICLing.* 203–224.

[55] H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin EP Seligman, et al. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one* 8, 9 (2013).

[56] Keiji Shinzato and Satoshi Sekine. 2013. Unsupervised extraction of attributes and their values from product description. In *Proc. of ACL.* 1339–1347.

[57] Jule Specht, Boris Egloff, and Stefan C Schmukle. 2011. Stability and change of personality across the life course: The impact of age and major life events on mean-level and rank-order stability of the Big Five. *Journal of personality and social psychology* 101, 4 (2011), 862.

[58] Ernest C Tupes and Raymond E Christal. 1992. Recurrent personality factors based on trait ratings. *Journal of personality* 60, 2 (1992), 225–251.

[59] Wann-Yih Wu and Ching-Ching Ke. 2015. An online shopping behavior model integrating personality traits, perceived risk, and technology acceptance. *Social Behavior and Personality: an international journal* 43, 1 (2015), 85–97.

[60] Chao Yang, Shimei Pan, Jalal Mahmud, Huahai Yang, and Padmini Srinivasan. 2015. Using personal traits for brand preference prediction. In *Proc. of EMNLP.* 86–96.

[61] Steffi Zander, Maria Reichelt, Stefanie Wetzel, et al. 2015. Does Personalisation Promote Learners' Attention? An Eye-Tracking Study. *Frontline Learning Research* 3, 4 (2015), 1–13.