

Personalized Emphasis Framing for Persuasive Message Generation

Tao Ding and Shimei Pan

Department of Information Systems
University of Maryland, Baltimore County
{taoding01, shimei}@umbc.edu

Abstract

In this paper, we present a study on personalized emphasis framing which can be used to tailor the content of a message to enhance its appeal to different individuals. With this framework, we directly model content selection decisions based on a set of psychologically-motivated domain-independent personal traits including personality (e.g., extraversion) and basic human values (e.g., self-transcendence). We also demonstrate how the analysis results can be used in automated personalized content selection for persuasive message generation.

1 Introduction

Persuasion is an integral part of our personal and professional lives. The topic of generating persuasive messages has been investigated in different fields with varied focuses. Psychologists focus on the cognitive, social and emotional processes of a persuader and a persuadee to understand what makes a communication persuasive (Hovland et al., 1953; Petty and Cacioppo, 1986; Smith and Petty, 1996). Marketing researchers are interested in applying theories of persuasion in promoting consumer products (Szybillo and Heslin, 1973; Han and Shavitt, 1994; Campbell and Kirmani, 2000; Kirmani and Campbell, 2004; Ford, 2005; Hirsh et al., 2012). Natural Language Generation (NLG) researchers are interested in studying the relation between language usage and persuasion in order to build automated systems that produce persuasive messages (Guerini et al., 2011; Reiter et al., 2003).

It is also generally believed that persuasion is more effective when it is custom-tailored to reflect the interests and concerns of the intended audience (Noar et al., 2007; Dijkstra, 2008; Hirsh et al., 2012). A proven tailoring tactic commonly used by politicians, marketing executives, as well as public health advocates is content framing (Meyerowitz and Chaiken, 1987; Maheswaran and Meyers-Levy, 1990; Grewal et al., 1994; Rothman and Salovey,

1997). Previous framing research has mainly focused on two types of framing strategies: *emphasis framing* and *equivalence framing*. To *emphasis frame* a message is to simplify reality by focusing on a subset of the aspects of a situation or an issue and make them more prominent in a communication to promote certain definition, causal interpretation and moral evaluation (Entman and Rojecki, 1993). For example, in political debating, the topic of *nuclear energy* can be framed as an *economic development* issue, a *safety* issue or an *environmental* issue. In marketing, the same car can be framed as a *low cost* car, a *performance* car, or a *green* car. With different framing strategies, the authors try to appeal to individuals with different beliefs and concerns. In contrast, *equivalence framing* focuses on presenting content as either loss-framed or gain-framed messages. For example, a smoking cessation message can employ a gain-frame like “You will live longer if you quit smoking”, or a loss-frame such as “You will die sooner if you do not quit smoking”. Even though the messages are equivalent factually, the frames can influence a receiver’s behavior either to encourage a desirable behavior or to avoid an unwanted outcome (Tversky and Kahneman, 1981). In this study, we focus on personalized *emphasis framing* which selectively emphasize the aspects of an entity (e.g., a car) to enhance its appeal to a given receiver.

Using emphasis framing as the framework for personalized content selection, we can take advantage of rich findings in prior framing research that link content selection decisions to a set of psychologically-motivated domain-independent personal characteristics. This has made our work more generalizable than those relying on application-specific user characteristics (e.g., use an individual’s smoking habit to tailor a smoke cessation message). Since content framing is a part of the content determination process, the model we propose is a part of the *content planner* in a Natural Language Generation (NLG) system (Reiter and

Dale, 1997).

There are three main contributions of this work.

1. To the best of our knowledge, this is the first effort in building an automated model of emphasis framing for personalized persuasive message generation.
2. We made content selection decisions based on a set of psychologically-motivated application-independent user traits, such as personality and basic human values, which makes our work more generalizable than those relying on domain-specific user characteristics and preferences.
3. We propose a cascade content selection model that integrates personalized content selection patterns in automated persuasive message generation.

2 Related Work

In the following, we summarize the research that is most relevant to our work including prior psychology and communication studies that link emphasis framing with personal traits. Since building computational models of emphasis framing was not the primary goal in these studies, we also include literature on personalized Natural Language Generation.

2.1 Emphasis framing and Personal Traits

There is a large body of social, marketing and communication theories on framing effects. (Zaller, 1992; Zaller and Feldman, 1992) point out that framing essentially reorganizes information to increase accessibility of an issue dimension by highlighting one cognitive path that had previously been in the dark. Others argue that the framing effect is due to a change in the rank order of the values associated with different aspects through the interaction with the content found within a message (Nelson et al., 1997; Chong and Druckman, 2007; Jacoby, 2000). The human decisions are controlled partly by the formulation of the problem and partly by the norms, habits, and personal characteristics of the decision-maker (Tversky and Kahneman, 1981).

Although most research agrees that the characteristics of a receiver play an important role in framing effectiveness, there is a significant disagreement

in what characteristics of a receiver result in framing effects. For example, (Anderson, 2010) states that people with prior attitudes toward an issue can be influenced by frames, while Slothuus (Slothuus, 2008) and Tabor et al. (Taber et al., 2009) did not find a framing effect for those with strong values associated with an issue prior to exposure to the frame. The mixed results may be due to the fact that many of these studies did not take into account that people with different traits (e.g., different personality) may react to framing strategies differently.

Recently, personalized framing, especially personality-based framing research has become a hot topic. Among them, Hirsh (2012) investigates whether a persuasive appeal's effectiveness can be increased by aligning message framing with a recipient's personality profile. In this study, for a given mobile phone, they constructed five advertisements, each designed to target one of the five major personality traits. Their results indicate that advertisements were evaluated more positively when they cohered with participants' personality. In a separate study, (Conroy et al., 2012) found that certain personality traits, particularly openness, agreeableness, and conscientiousness mediate framing effects when participants were presented with different frames of political and health issues such as civil liberties, medical treatments, energy, affirmative action, and gun control.

Inspired by the above research, we also employ psychologically-motivated trait models to capture individual characteristics. In addition to personality, we also incorporate basic human values since framing effects were also shown to be related to personal beliefs and motivations. As a result, we have significantly increased the scope of our study over prior research. Moreover, unlike prior research where only messages hand-crafted by experts were used, we are interested in building computational models to automatically select a subset of the aspects to highlight based on personal traits.

2.2 Personalized NLG

There is also a large body of work on personalized Natural Language Generation (NLG). For example, STOP is a Natural Language Generation (NLG) system that generates tailored smoking cessation letters based on a user's responses to a four-page

smoking questionnaire (Reiter et al., 2003); PER-SIVAL customizes the content of search summaries based on its relevance to a given patient’s health record (McKeown et al., 2003); MATCH (Johnston et al., 2002) is a multimodal dialogue system that tailors the content of its responses based on a user’s restaurant preferences; M-PIRO (Burenhult, 2002) tailors the words and the complexity of museum object descriptions for different audiences (e.g. adults, children, and experts); PERSONAGE (Mairesse and Walker, 2011) and CRAG 2 (Gill et al., 2012) vary linguistic styles to project intended personality in spoken utterances. In addition, Carenini and Moore (Carenini and Moore, 2006) employed a multiattribute utility theory-based user preference model for personalized evaluative argument generation. Among them, STOP, PER-SIVAL and MATCH use domain-specific user models while M-PIRO, PERSONAGE and GRAG2 employ domain independent user properties, such as expertise and personality. For PERSONAGE and GRAG2, personality traits are mainly used to adapt linguistic styles. So far, there has not been much work focusing on using domain-independent user traits to automatically adapt message content to improve its persuasive appeal.

3 Acquiring Personal Traits

Since prior study often links framing effects to individual characteristics such as personality and individual motivations and beliefs, here we focus on two widely-accepted trait models in psychology: the Big5 personality model (Goldberg, 1993) and Schwartz’s basic human value model (Schwartz, 2003). Figure 1 shows the description of each of the Big5 personality traits along with each of the five basic human value traits.

To acquire the personality and value traits of a person, traditionally, psychometric tests, such as the IPIP test for Big 5 personality (Yarkoni, 2010a) and the PVQ survey for values (Schwartz, 2003), were used. Recent research in the field of psycholinguistics has shown that it is possible to automatically infer personal traits from one’s linguistic footprint, such as tweets, Facebook posts and blogs (Yarkoni, 2010b; Celli and Polonio, 2013; Chen et al., 2014). Unlike psychometric tests, automated trait analysis

Model	Factors	Description
Big5 Personality Model	Openness to Experience	The degree of intellectual curiosity, creativity and a preference for novelty and variety a person has.
	Conscientiousness	A tendency to be organized and dependable, show self-discipline, act dutifully, and prefer planned rather than spontaneous behavior
	Extraversion	Energy, positive emotions, assertiveness, sociability and the tendency to seek stimulation in the company of others, and talkativeness.
	Agreeableness	A tendency to be compassionate and cooperative. Also a measure of one’s trusting, helpful, well-tempered nature.
	Neuroticism	The tendency to experience unpleasant emotions easily, of low emotional stability and impulse control.
Schwartz’s basic human values	Openness to change	the desire for independence and new, novel, exciting and challenging experiences:
	Self-enhancement	The pursuit of self-interests including social status, prestige, resources and personal success through demonstrating competence
	Conservation	Self-restriction and conform to tradition, social expectations and norms; Pursue safety, harmony and stability; resist to change
	Hedonism	Pursuit of pleasure and sensuous gratification for oneself.
	Self-transcendence	Concern and pursue the enhancement and protection of the welfare and interests of others as well as the protection for nature.

Figure 1: Description of Two Trait Models

allows us to infer personal traits for a large number of people, which makes it possible to scale up automated personal persuasion for a very large population (e.g., millions of social media users).

4 Acquiring Author Framing Strategy

Framing effects are often subtle and may be influenced by many factors, such as the credibility of the authors, the personality of the receivers and the context of the communication. In the first study, we investigate whether it is feasible to build a personalized content selection model based on a writer’s (a.k.a. an author’s) content framing strategies.

To investigate this, we first randomly generated ten cars, each include eight aspects: *safety, fuel economy, quality, style, price, luxury, performance and durability*. The value of each aspect was randomly generated on a 5-point Likert scale: “1 (very bad)”, “2 (bad)”, “3 (average)”, “4 (good)”, and “5 (excellent)”. We also conducted a large-scale personality and basic human value survey on Amazon Mechanical Turk (AMT). We used the 50-item IPIP survey (Goldberg, 1993) to obtain a Amazon Mechanical Turk worker (a.k.a. Turker)’s personality scores and the 21-item PVQ survey (Schwartz, 2003) to obtain his/her basic value scores. To en-

sure the quality of the data from AMT, we added two qualification criteria. A qualified Turker must (1) have submitted over 5000 tasks (2) with an acceptance rate over 95%. The survey also included several validation questions, which are pairs of questions that are paraphrases of each other. If the answers to a pair of validation questions are significantly different, the user data were excluded from our analysis. After removing invalid data, we collected the traits of 836 Turkers. The raw personality scores, ranging from 10 to 50, and raw value scores, ranging from 1 to 6, were computed directly from the survey answers. The normalized trait scores, ranging from 0 to 1, were computed using their rank percentiles in this population.

In addition, we designed two Human Intelligence Tasks (HITs) on AMT: a content customization task and a validation task. In the content customization task (a.k.a. Task 1), a Turker was asked to select one car aspect to emphasize in his message for a receiver. The validation task (a.k.a. Task 2) was used to validate whether a receiver prefers the message customized for her or not.

Specifically, in Task 1, the Turkers were asked to imagine that they work for a marketing firm on a campaign to promote a new car. Each Turker was given the specification of a car (randomly selected from the 10 randomly generated cars) and a receiver (randomly selected from the 836 Turkers whose trait scores were known to us). The Turker was asked to write a campaign message to persuade the receiver to buy the car. But the Turker can only select one of the eight car aspects to include in his message. Since customizing a message based on an interaction of all ten traits can be very challenging for a Turker, we used a simplified trait profile in our study. The simplified trait profile contains only two traits: the most prominent personality trait and the most prominent value trait. The prominence of a trait was defined based on the normalized trait score. The more different a trait score is from the median (.50), the more prominent the trait is. For comparison, for the same car, we also asked the same writer to select a car aspect for someone who has an opposite trait profile. The opposite trait profile is defined as the one that is most different from the given trait profile (with the lowest cosine similarity). After the writer selected a car aspect, he also wrote a campaign message us-

ing the selected aspect. Overall, after removing invalid data, we collected 490 pairs of messages for 131 pairs of receivers.

To validate the framing effect, in Task 2, we asked a new set of Turkers (receivers) to first complete an IPIP personality survey and a PVQ human value survey. Based on the survey results, we computed the trait profile for each of them. In addition, for each receiver in Task 2, we matched his/her trait profile with the 131 pairs of trait profiles collected in Task 1. The profile with the highest matching score (computed based on cosine similarity) was selected and its associated message pair was retrieved.

Then we presented the receiver with a pair of messages, one created for someone with matching trait profile, the other for someone with the opposite trait profile. We also randomized the order these messages were presented. Finally, we asked the receivers to rate which message they prefer more. If the framing strategies used by the Turkers (authors) in Task 1 were effective, then the Turkers (receivers) in Task 2 will prefer the messages tailored for them more than the ones tailored for someone with the opposite trait profile. Overall, after filtering out invalid data, we have collected the results from 145 receivers. Among them, 77 prefer the messages written for them, while 68 prefer the messages written for someone with the opposite trait profiles. We performed a *sign test* to determine whether the difference is statistically significant and the result was negative ($p < 0.2$).

Although moderate personalization effects were found in previous framing research, only expert-crafted messages were used (Hirsh et al., 2012). Here, when Turkers (mostly non-experts) were asked to customize the messages based on a receiver's traits, no significant effects were found. Since authors' emphasis framing strategies were not effective, we can not directly use authors' data to learn their emphasis framing strategies. Next, we present several experiments designed to automatically derive emphasis framing strategies based on a receiver's traits and his/her aspect selection decisions.

5 Learning Emphasis Framing Strategies

To derive emphasis framing patterns based on a receiver’s traits and his/her aspect selection decisions, we designed another HIT (Task 3) on AMT to collect data. In Task 3, each Turker was asked to take the IPIP and PVQ surveys so that we can obtain his/her Big5 personality and value scores. In addition, we also asked him/her to rank all eight car aspects based on their importance to him/her. To control the influence of the value of a car aspect on a user’s aspect selection decision (e.g., if the value of “safety” is “poor” and the value of “fuel economy” is “good”, to promote the car, people almost always describe it as “a car with good fuel economy”, not “an unsafe car”, regardless of a receiver’s personality). In this study, we kept the values of all car aspects unspecified. After removing invalid data, our dataset has 594 responses, each contains a Turker’s personality and value scores as well as his/her rank of the eight car aspects. In the following, we describe how we analyze the relationship between aspect rank and personal traits.

5.1 Pattern Discovery with Regression

In our first study, we employed regression analysis to identify significant correlations between personal traits and aspect ranks. Specifically, we trained eight linear regression models, one for each of the eight car aspects. The dependent variable in each model is the rank of an aspect (from 1 to 8) and the independent variables are the ten user traits. In the regression analysis, we only focused on the main effects since a full interaction model with ten traits will require much more data to train. Since the raw scores of the personality and value traits use different scales, we normalized these scores so that they are all from 0 to 1. Table 1 shows the regression results.

Several interesting patterns were discovered in this analysis: (a) a positive correlation between the rank of “luxury” and “self-enhancement”, a trait often associated with people who pursue self-interests and value social status, prestige and personal success ($p < 0.0001$). This pattern suggests that to promote a car to someone who scores high on “self-enhancement”, we need to highlight the “luxury” aspect of a car. (b) the rank of “safety” is posi-

tively correlated with “conservation”, a trait associated with people who conform to tradition and pursue safety, harmony, and stability ($p < 0.005$). This result suggests that for someone values “conservation”, it is better to emphasize “car safety” in a personalized sales message. (c) “self-transcendence”, a trait often associated with people who pursue the protection of the welfare of others and the nature, is positively correlated with the rank of “fuel economy” ($p < 0.005$) but negatively correlated with the rank of “style” ($p < 0.005$). This suggests that for someone who values “self-transcendence”, it is better to emphasize “fuel economy”, but not so much on “style”. Other significant correlations uncovered in this analysis include a negative correlation between car “price” and “conservation” ($p < 0.005$), a negative correlation between car “safety” and “conscientiousness” ($p < 0.05$), and a positive correlation between “openness to change” and car “performance” ($p < 0.05$).

5.2 Pattern Discovery with Constrained Clustering

In the regression analysis, we only considered the main framing effects. In order to discover high-order interaction patterns with limited data, we want to use clustering to group people with similar traits together. In addition, we also want that the people in a cluster share similar aspect preferences. Otherwise, we won’t be able to link the trait patterns discovered in a cluster with specific aspect preferences. Thus, we employed constrained clustering in this analysis. With constrained clustering, we can ensure the homogeneity of the aspect preferences within each resulting cluster.

To facilitate this analysis, first we mapped the aspect ranks obtained in Task 3 into discrete categories. For a complete rank of eight car aspects, we mapped the top three ranked aspects to an “Important” class, bottom three to a “Not-Important” class, and the middle two to a “Neutral” class. In addition, we encoded the aspect homogeneity requirement as constraints. Typically, constrained clustering incorporates either a set of must-link constraints, a set of cannot-link constraints, or both. A must-link constraint is used to specify that the two data instances in the must-link relation should be placed in the same cluster. A cannot-link constraint is used

Table 1: Results of the Regression Analysis

	Safety	Fuel	Quality	Style	Price	Luxury	Perf	Durab
Agreeableness	0.39	-0.52	-0.53	0.54	0.81	0.004	-0.62	-0.27
Conscientiousness	-1.75 *	-0.31	0.80	0.29	-0.01	0.27	0.83	-0.12
Extroversion	0.69	-0.71	0.008	-0.25	-0.37	0.48	-0.07	0.224
Neurotisim	1.08	-0.01	-0.46	-0.11	-0.32	-0.07	0.18	-0.28
Openness	1.59	-0.05	0.01	-0.99	0.36	-0.53	-0.46	0.07
Conservation	1.99 **	-0.99	-0.66	0.84	-1.72 **	0.21	0.38	-0.03
Hedonism	1.47	-0.15	-0.69	0.16	0.51	-0.06	-0.82	-0.43
Openness to change	-2.15	0.08	0.58	0.48	-1.99 *	-0.38	2.29*	1.07
Self-enhancement	-1.39	-1.12	0.58	0.47	-0.31	2.41 ***	0.77	-1.41
Self-transcendence	1.33	2.37 **	1.36	-2.47 **	-0.91	-1.01	-0.33	-0.32

Note: $p < 0.05$, ** $p < 0.005$, *** $p < 0.0001$

Table 2: Patterns Discovered in Clustering Analysis

Feature	Cluster	Accuracy	Label	Significant traits
Safety	1	0.7	Important	Extrave(+),Neuroti(+)
	2	0.64	Neutral	Conscie(+),Hedonis(+),Open(+),Self-en(+)
	3	0.71	Important	Conscie(+),Open(-)
Fuel	1	0.54	Neutral	Open(-),Self-en(-)
	2	0.54	Not-Important	Hedonis(+),Open(+),Self-en(+)
Quality	1	0.43	Important	Extrave(+),Neuroti(+)
	2	0.45	Non-Important	Hedonis(+),Open(+),Self-en(+)
	3	0.45	Not-Important	Conscie(+),Open(-)
Style	1	0.5	Not-Important	Hedonis(-),Open(-)
	2	0.55	Neutral	Conscie(+),Extrave(+),Neuroti(+)
	3	0.62	Neutral	Conscie(+),Hedonis(+),Open(+),Self-en(+)
	4	0.74	Not-Important	Conscie(+),Open(-)
Performance	1	0.73	Neutral	Extrave(+),Neuroti(+)
	2	0.5	Neutral	Conscie(+),Open(-)
	3	0.4	Not-Important	Hedonis(-),Open(-)
Durability	1	0.56	Not-Important	Extrave(+),Hedonis(+),Self-en(+)
	2	0.36	Important	Conscie(+),Hedonis(+),Open(+),Self-en(+)

Note: $CV < 0.12$ $P < 0.001$ $Diff > 0.2$

to specify that the two instances in the cannot-link relation should not be put in the same cluster. These constraints act as a guide for which a constrained clustering algorithm will use to find clusters that satisfy these requirements.

To encode the homogeneity constraint, for each car aspect (e.g. safety), we can simply add must-links between every pair of Turkers if they share the same aspect preference (e.g., both consider “safety” important) and add cannot-links for every pair of Turkers who do not share the same aspect preferences (e.g., one Turker considers “safety” “Important”, the other considers it “Not-Important”). But with both must-links and cannot-links, it is very likely we will get three big clusters, each is related to one of the three categories: Important, Neutral and Not-Important. Although the resulting clusters satisfy the aspect preference homogeneity requirement, they fail to group people with similar traits together. As a result, in this analysis, we only used cannot-links, which not only guarantees the homogeneity of aspect preferences, but also creates smaller clusters that group people with similar traits together.

We employed the Metric Pairwise Constrained KMeans algorithm (MPCK-MEANS) (Bilenko et al., 2004) to incorporate the aspect preference homogeneity requirement. The optimal cluster number K was determined empirically by running MPCK-MEANS with different K s, $K \in [3, 20]$ (3 is the minimum number of clusters since we have 3 different aspect preference categories).

To determine whether the resulting clusters capture any interesting patterns, we used two pattern selection criteria (a) a homogeneity criterion which requires that there is at least one trait whose values in the cluster is relatively homogeneous; (b) a distinctiveness criterion which requires that for the traits identified in (a), their cluster means need to be significantly different from the population means. For (a), we used the coefficient of variation (CV) as the homogeneity measure. CV, also known as relative standard deviation (RSD), is a standardized measure of dispersion of a probability or count distribution. It is often expressed as a percentage and is defined as the ratio of the standard deviation σ to the mean $|\mu|$. In the study, we required that all the CVs of

homogeneous traits to be lower than 0.12. For (b) we required that the differences of the means need to be significant based on an independent sample t-test with $p < 0.001$ and the difference of means is greater than 0.2.

Table 2 highlights some of the patterns discovered using this approach. In this table, we list the cluster id, cluster label (Important, Not-Important, Neutral), clustering accuracy, and significant traits in the cluster (“+” indicates that the cluster mean is higher than population mean, “-” means the opposite). For example, based on the Safety-1 pattern, people who are more extraverted (extrave(+)) and more neurotic (neurotic(+)) tend to consider “car safety” important. Similarly, based on pattern Safety-3, people who are more conscientious (conscie(+)) but less open (open(-)) tend to consider “safety” important. Other interesting patterns include: people who are less open (open(-)) and do not value hedonism (hedonism(-)) don’t consider performance very important (performance-3), and people who are more extraverted (extrave(+)), value hedonism and self-enhancement (hedonism(+), self-en(+)) do not think durability important (durability-1).

6 Apply Emphasis Framing in NLG

The patterns derived in the previous section can be used in personalized content selection for Natural Language Generation. In general, to promote a car, people tend to highlight the good aspects and avoid the bad aspects, regardless of a receiver’s personality. For example, people will likely to highlight the fuel economy aspect if a car is very fuel efficient while de-highlight the same aspect if a car is not fuel efficient. Thus, during content selection, to take the value of an aspect into consideration, we employ a cascade NLG model that integrates value-based content selection with trait-based personalization.

The input to the cascade content selection model includes: (1) the values of all the car aspects; (2) the trait scores of a receiver; (3) the eight linear-regression models learned in Section 5.2, one for each aspect; (4) the interaction rules learned in Section 5.3; (5) n , the number of aspects needed in the output; (6) the value difference threshold δ_1 that determines whether the values of two or more aspects are significantly different; (7) the rank difference

threshold δ_2 that determines whether the ranks of two or more aspects predicted by the linear regression models are significantly different.

To select n aspects to emphasize, our system first ranks all the aspects based on their values. If the value of the n -th aspect v_n is significantly better than that of the $(n+1)$ -th aspect v_{n+1} (i.e., their difference is greater than δ_1), we output the top n aspects directly. Otherwise, for all the aspects whose values are either the same or not significantly worse than v_n , their ranks will be determined by the trait-based linear regression models. Moreover, after re-ranking relevant aspects based on the predicted ranks from the regression models, if the predicted rank of the n -th aspect r_n is significantly better than that of the $(n+1)$ -th aspect r_{n+1} (i.e., the rank difference is greater than δ_2), we just output the top n aspects in this list. Otherwise, for those aspects whose ranks predicted by the linear regression models are not significantly lower than r_n , we use the interaction rules discovered in Section 5.3 to further adjust their ranking scores (i.e., increase the rank by δ_2 if the cluster label is “Important”, or decrease by δ_2 if “Not-Important”). For each aspect, if more than one interaction rule applies, more accurate rules take precedence over less accurate rules. Finally, the system will output the top n aspects in the final list.

We use an example shown in Figure 2 to illustrate the cascade aspect selection process. In this example, we assume $n=3$, $\delta_1=1$ and $\delta_2=0.5$. We first sorted all the aspects based on their values. Since the values of “Fuel Economy” and “Luxury” are significantly better than the 3rd-ranked aspect “Price”, their ranks are not affected by personalized aspect selection. Similarly, since the values of “Performance” and “Style” are significantly lower than that of the 3rd-ranked aspect, their ranks are also not affected by personalization. Since the value differences among the rest 4 aspects, “Price”, “Durability”, “Quality” and “Safety” are all equal or not significant worse than v_3 , we used trait-based personalized ranks predicted by the regression models to re-rank them (the output ranks from the regression models are shown in the parentheses in the column “Regression-based Re-Ranking”). After re-ranking these aspects based on the predicted ranks, since the rank of the 3rd-ranked aspect “Price” (2.2) and that of “Safety” (2.5) is within δ_2 , we use the learned

Aspect	Value	Regression-based Re-Ranking	Interaction Rule-based Re-Ranking	Final Rank
Fuel Economy	5(Excellent)	1	1	1*
Luxury	5(Excellent)	2	2	2*
Price	3(Average)	3(2.2)	4(2.2)	4
Durability	3(Average)	6(6.7)	6(6.7)	6
Quality	3(Average)	5(4.5)	5(4.5)	5
Safety	3(Average)	4(2.5)	3(2.0)	3*
Performance	1(Very Bad)	7	7	7
Style	1(Very Bad)	8	8	8

Figure 2: A Cascade Content Selection Example

interaction rules to adjust their ranks. Since the predicted ranks of “Durability” and “Quality” are much worse than that of “Price”, their ranks are not affected by the interaction rules. To apply the interaction rules, assume for a given receiver, both his extraversion and neuroticism scores are much higher than the population average, the Safety-1, Quality-1 and Performance-1 rules are applicable. Since the Safety-1 rule predicts that “Safety” is “Important” to the receiver while none of the rules affects “Price”, the predicted rank for “Safety” is increased by δ_2 . After this adjustment, the ranks of all the aspects are shown in the “Final Rank” column. The top 3 aspects based on the final ranks are selected as the output (those marked with a *).

To evaluate the performance of the cascade content selection model, we conducted an additional AMT study. Given the specifications of the ten cars in Task1, we asked each AMT participant to select the top-n aspects to emphasize. Here $n=1$ and 3. In this task, aspect selection not only depends on the importance of an aspect to a receiver, but also the values of the aspects of a given car. We also acquired the personality and value scores of each Turkur based on the IPIP personality and PVQ value survey. Finally, we compared the output of our model with the aspects selected by the Turkers. We used top-n overlapping percentage as the evaluation metrics. Overall, we collected the aspect selection results from 38 Turkers, each on ten different cars. In total, we collected 380 data instances in our ground truth dataset. We have tested different δ_1 and δ_2 , the best results were obtained when $\delta_1 = 0$ and $\delta_2 = 0.5$. We compared our model with a baseline system which relies solely on the values of aspects to determine their ranks in the baseline system. If two or more aspects have the same value (e.g., the values of both “Price” and “Durability” are “3(Average)”, their ranks were determined randomly. Ta-

ble 3 shows the evaluation results. If only 1 aspect is needed in the output, the Top-1 agreement is 62% for the cascade model versus the baseline’s 54%. Similarly, if 3 aspects are needed in the output, the Top-3 agreement is 87% for the cascade model versus the baseline’s 46%. All the differences are statistically significant based on paired-t test ($p \leq 0.05$).

Table 3: Cascade Content Selection Evaluation

	Cascade	Baseline
Top-1 agreement	0.62	0.54
Top-3 agreement	0.87	0.46

7 Discussion

In general, there are two main challenges in adapting a personalized content selection model trained in one domain to another domain: (1) adapting the *data model* from one domain (e.g., restaurant data) to another (e.g., movie data); (2) adapting a domain-specific *user model* (e.g., a user’s preferences of restaurant features such as “cuisine type”) to a different domain (e.g., a user’s preferences of movie features such as “movie genre”). Although our data model is in the automobile domain, we adopted a domain-independent user model motivated by psychological theories(e.g., personality and basic human values), instead of a domain-dependent user preference models (e.g, a user’s preferences of “fuel economy”). This allows us to more easily apply typical domain adaptation methods such as instance-based (Zadrozny, 2004) or feature-based transfer learning (Blitzer et al., 2006) to further adapt the system and generalize the current results.

8 Conclusions

In this study, we analyzed the relationship between an individual’s traits and his/her aspect framing decisions. Our analysis has uncovered interesting patterns that can be used to automatically customize a message’s content to enhance its appeal to its receivers. We also proposed a cascade content selection model to automatically incorporate the analysis results in automated persuasive message generation. Our evaluation results have demonstrated the effectiveness of this approach.

References

- Kristen D Anderson. 2010. Framing traits: The role of personality in framing effects.
- Mikhail Bilenko, Sugato Basu, and Raymond J Mooney. 2004. Integrating constraints and metric learning in semi-supervised clustering. In *Proceedings of the twenty-first international conference on Machine learning*, page 11. ACM.
- John Blitzer, Ryan McDonald, and Fernando Pereira. 2006. Domain adaptation with structural correspondence learning. In *Proceedings of the 2006 conference on empirical methods in natural language processing*, pages 120–128. Association for Computational Linguistics.
- G Burenhult. 2002. Generating multilingual personalized descriptions of museum exhibits-the m-piro project.
- Margaret C Campbell and Amna Kirmani. 2000. Consumers' use of persuasion knowledge: The effects of accessibility and cognitive capacity on perceptions of an influence agent. *Journal of Consumer Research*, 27(1):69–83.
- Sandra Carberry, Jennifer Chu-Carroll, and Stephanie Elzer. 1999. Constructing and utilizing a model of user preferences in collaborative consultation dialogues. *Computational Intelligence*, 15(3):185–217.
- Giuseppe Carenini and Johanna D Moore. 2006. Generating and evaluating evaluative arguments. *Artificial Intelligence*, 170(11):925–952.
- Fabio Celli and Luca Polonio. 2013. Relationships between personality and interactions in facebook. *Social Networking: Recent Trends, Emerging Issues and Future Outlook*, pages 41–54.
- Jilin Chen, Gary Hsieh, Jalal U Mahmud, and Jeffrey Nichols. 2014. Understanding individuals' personal values from social media word use. In *Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing*, pages 405–414. ACM.
- Dennis Chong and James N Druckman. 2007. Framing theory. *Annu. Rev. Polit. Sci.*, 10:103–126.
- Susan Conroy, Carmine M Pariante, Maureen N Marks, Helen A Davies, Simone Farrelly, Robin Schacht, and Paul Moran. 2012. Maternal psychopathology and infant development at 18 months: the impact of maternal personality disorder and depression. *Journal of the American Academy of Child & Adolescent Psychiatry*, 51(1):51–61.
- Arie Dijkstra. 2008. The psychology of tailoring-ingredients in computer-tailored persuasion. *Social and personality psychology compass*, 2(2):765–784.
- Robert M Entman and Andrew Rojecki. 1993. Freezing out the public: Elite and media framing of the us anti-nuclear movement.
- Christopher M Ford. 2005. Speak no evil: targeting a population's neutrality to defeat an insurgency. *Parameters: The US Army War College Quarterly, Summer*.
- Alastair J Gill, Carsten Brockmann, and Jon Oberlander. 2012. Perceptions of alignment and personality in generated dialogue. In *Proceedings of the Seventh International Natural Language Generation Conference*, pages 40–48. Association for Computational Linguistics.
- Lewis R. Goldberg. 1993. The structure of phenotypic personality traits. *American Psychologist*, 48(1):26.
- Dhruv Grewal, Jerry Gotlieb, and Howard Marmorstein. 1994. The moderating effects of message framing and source credibility on the price-perceived risk relationship. *Journal of consumer research*, pages 145–153.
- Marco Guerini, Oliviero Stock, Massimo Zancanaro, Daniel J OKeefe, Irene Mazzotta, Fiorella de Rosi, Isabella Poggi, Meiyi Y Lim, and Ruth Aylett. 2011. Approaches to verbal persuasion in intelligent user interfaces. In *Emotion-Oriented Systems*, pages 559–584. Springer.
- Sang-Pil Han and Sharon Shavitt. 1994. Persuasion and culture: Advertising appeals in individualistic and collectivistic societies. *Journal of experimental social psychology*, 30(4):326–350.
- Jacob B Hirsh, Sonia K Kang, and Galen V Bodenhausen. 2012. Personalized persuasion tailoring persuasive appeals to recipients personality traits. *Psychological science*, 23(6):578–581.
- Carl I Hovland, Irving L Janis, and Harold H Kelley. 1953. Communication and persuasion; psychological studies of opinion change.
- William G Jacoby. 2000. Issue framing and public opinion on government spending. *American Journal of Political Science*, pages 750–767.
- Michael Johnston, Srinivas Bangalore, Gunaranjan Vasireddy, Amanda Stent, Patrick Ehlen, Marilyn Walker, Steve Whittaker, and Preetam Maloor. 2002. Match: An architecture for multimodal dialogue systems. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics*, pages 376–383. Association for Computational Linguistics.
- Amna Kirmani and Margaret C Campbell. 2004. Goal seeker and persuasion sentry: How consumer targets respond to interpersonal marketing persuasion. *Journal of Consumer Research*, 31(3):573–582.
- Durairaj Maheswaran and Joan Meyers-Levy. 1990. The influence of message framing and issue involvement. *Journal of Marketing research*, pages 361–367.
- François Mairesse and Marilyn A. Walker. 2011. Controlling user perceptions of linguistic style: Trainable generation of personality traits. *Comput. Linguist.*, 37(3):455–488, September.

- Kathleen R McKeown, Noemie Elhadad, and Vasileios Hatzivassiloglou. 2003. Leveraging a common representation for personalized search and summarization in a medical digital library. In *Proceedings of the 3rd ACM/IEEE-CS joint conference on Digital libraries*, pages 159–170. IEEE Computer Society.
- Beth E Meyerowitz and Shelly Chaiken. 1987. The effect of message framing on breast self-examination attitudes, intentions, and behavior. *Journal of personality and social psychology*, 52(3):500.
- Thomas E Nelson, Zoe M Oxley, and Rosalee A Clawson. 1997. Toward a psychology of framing effects. *Political behavior*, 19(3):221–246.
- Seth M Noar, Christina N Benac, and Melissa S Harris. 2007. Does tailoring matter? meta-analytic review of tailored print health behavior change interventions. *Psychological bulletin*, 133(4):673.
- Richard E Petty and John T Cacioppo. 1986. *The elaboration likelihood model of persuasion*. Springer.
- Ehud Reiter and Robert Dale. 1997. Building applied natural language generation systems. *Natural Language Engineering*, 3(01):57–87.
- Ehud Reiter, Roma Robertson, and Liesl M Osman. 2003. Lessons from a failure: Generating tailored smoking cessation letters. *Artificial Intelligence*, 144(1):41–58.
- Alexander J Rothman and Peter Salovey. 1997. Shaping perceptions to motivate healthy behavior: the role of message framing. *Psychological bulletin*, 121(1):3.
- Shalom H Schwartz. 2003. A proposal for measuring value orientations across nations. *Questionnaire Package of the European Social Survey*, pages 259–290.
- Rune Slothuus. 2008. More than weighting cognitive importance: A dual-process model of issue framing effects. *Political Psychology*, 29(1):1–28.
- Stephen M Smith and Richard E Petty. 1996. Message framing and persuasion: A message processing analysis. *Personality and Social Psychology Bulletin*, 22:257–268.
- George J Szybillo and Richard Heslin. 1973. Resistance to persuasion: Inoculation theory in a marketing context. *Journal of Marketing Research*, pages 396–403.
- Charles S Taber, Damon Cann, and Simona Kucsova. 2009. The motivated processing of political arguments. *Political Behavior*, 31(2):137–155.
- Amos Tversky and Daniel Kahneman. 1981. The framing of decisions and the psychology of choice. *Science*, 211(4481):453–458.
- Tal Yarkoni. 2010a. The abbreviation of personality, or how to measure 200 personality scales with 200 items. *Journal of Research in Personality*, 44(2):180–198.
- Tal Yarkoni. 2010b. Personality in 100,000 words: A large-scale analysis of personality and word use among bloggers. *Journal of research in personality*, 44(3):363–373.
- Bianca Zadrozny. 2004. Learning and evaluating classifiers under sample selection bias. In *ICML*.
- John Zaller and Stanley Feldman. 1992. A simple theory of the survey response: Answering questions versus revealing preferences. *American journal of political science*, pages 579–616.
- John Zaller. 1992. *The nature and origins of mass opinion*. Cambridge university press.