Journal of International Technology and Information Management

Volume 29 | Issue 3 Article 3

2021

The Design, Development and Validation of a Persuasive Content Generator

Sam Khataei
Carleton University, contact@amirsamkhataei.com

Michael J. Hine
Carleton University, mike.hine@carleton.ca

Ali Arya
Carleton University, Ali.Arya@carleton.ca

The Design, Development and Validation of a Persuasive Content Generator

Sam Khataei (*Carleton University*)
Michael J. Hine (*Carleton University*)
Ali Arya (*Carleton University*)

ABSTRACT

This paper addresses the automatic generation of persuasive content to influence users' attitude and behaviour. Our research extends current approaches by leveraging individuals' social media profiles and activity to personalize the persuasive content. Unlike most other implemented persuasive technology, our system is generic and can be adapted to any domain where collections of electronic text are available. Using the Yale Attitude Change approach, we describe: the multi-layered Pyramid of Individualization model; the design, development, and validation of integrated software that can generate individualized persuasive content based on a user's social media profile and activity. Results indicate the proposed system can create personalized information that (a) matches readers' interests, (b) is tailored to their ability to understand the information, and (c) is supported by trustable sources.

INTRODUCTION

Persuasion has been defined as a purposeful attempt to change attitudes and/or behaviors (Fogg, 2003) and is an essential part of education, marketing, sales, and political communication where content consumers should be convinced that a certain product, skill, or approach is suitable and desirable for them. Given persuasive content is meant to influence opinions, beliefs, attitudes, or behaviours of intended audiences it can also critically affect individual, group and organizational decision-making.

The use of information and communications technology (ICT) in persuasion is advantageous as ICT is ubiquitous, can deal with huge amounts of data, and it can scale effectively (Fogg, 2003). It enables the collection and management of massive repositories of electronic data and information about global and cultural issues that can form the basis of content that can then be pushed out to potential recipients. In addition, persuasive opportunities can be amplified by using the Internet and its associated technologies as enormous numbers of potential content consumers can be easily reached (Oinas-Kukkonen & Harjumaa, 2009). This reality has resulted in leveraging the use of digital documents for persuasion purposes. However, since individuals possess different respective interests, characteristics, and abilities, these shared digital artefacts will not be equally persuasive to all users.

To provide persuasive communication, users are often categorized according to characteristics like age, marital status, education, and occupation and provided with content specific to their characteristics. For example, television/online advertisements will target a program/website's viewers, or financial institutions provide different promotional material targeting different demographics based on their perceived common needs. Although *categorization* can improve the persuasiveness of the content, it assumes individuals within a category are similar, which is not necessarily true. Individuals may have different points of view on a topic because of respective personal experiences, epistemic beliefs (Kardash & Scholes, 1996; Mason et al., 2006), or even dispositions toward knowledge (Crowson & DeBacker, 2008). More effective persuasion can be achieved through *personalization*, where the content is tailored for a specific person rather than to a category of people. For example, customized online ads based on browsing history and social media newsfeeds based on existing connections and recent activities target an individual user. Despite such initial attempts, there is a lack of a systematic approach to personalizing documents derived from electronic repositories for persuasion purposes.

Existing research on persuasive software systems has focused mainly on changing behaviour, leaving attitude a relatively under-addressed construct in systems research (Oinas-Kukkonen & Harjumaa, 2009; Torning & Oinas-Kukkonen, 2009). As attitude represents long term patterns of behaviours based on opinions and beliefs, it is a critical construct that needs attention to move research in persuasive systems forward. In addition, most existing persuasive systems have not been designed to work in multiple domains. For example, a system designed to assist in weight loss is not easily adaptable to persuade people about energy conservation.

To our knowledge, there is currently no systematic software solution for domain-independent Persuasive Content Generation (PCG) that provides personalized persuasive content based on electronic text repositories to influence attitude.

Past difficulties in designing such tools include the lack of personal data and associated computational models capable of manipulating content to facilitate persuasion. While it is acknowledged in the literature that ICT can amplify persuasion by allowing customized and individualized content (Slattery et al., 2020), the pervasive use of social media has created new opportunities to learn about individuals' preferences, connections and trustable sources through mining their social media profiles and content that can be used to personalize content

This paper addresses the problem of generating personalized persuasive content. The authors have previously proposed a general idea for personalized PCG (Khataei & Arya, 2015). Building on this initial work, we report on the full design and evaluation of our software framework and custom components that enable automatic generation of persuasive content based on: (1) the content consumer's social media profiles and activity; and (2) the original persuader's objectives. The persuader's involvement includes providing a database of source material and a set of rules that manage the content generation. We relied on Twitter data as a source of personal information, and the Yale Attitude Change (YAC) (Hovland et al., 1953) as the main theoretical model as its focus is on the under-addressed construct of attitude. We propose a software framework and implementation that is not domain specific with three layers corresponding to YAC's persuasive factors. These layers are the foundation for personalized and persuasive content that we believe will be retained.

This paper is organized as follows. First, a brief review of the most relevant persuasion and persuasive technologies literature is presented. This is followed by a description of the design of our PCG system with its reusable modular web-based structure. A mixed-method experiment used to validate the system is then reported. Finally, limitations and future research are discussed.

PERSUASION AND PERSUASIVE TECHNOLOGIES

One of the early works in the field of persuasion is the Yale Attitude Change (YAC) approach (Hovland et al., 1953). According to YAC, effective persuasion of an audience consists of four steps: gaining the audience's attention, adjusting the message's comprehension level to a level the intended user can consume, ensuring argument acceptance, and finally ensuring the message will be remembered. YAC has shaped marketing and advertising (Aronson et al., 2009) strategies and is influential because of its simplicity. Most well-known theories of persuasion (Oinas-Kukkonen & Harjumaa, 2008; Ajzen, 1991; Cacioppo et al., 1986) are either inspired by YAC or share common elements with it. Furthermore, the YAC approach and its dissection of how a system persuades are used in notable Persuasive Systems Design (PSD) models. Oinas-Kukkonen and Harjumaa (2009) have emphasized how PSD leveraged YAC on human persuasion through rational processing of arguments. One of the main PSD design principles is articulating how the credibility of a system makes it more persuasive. System credibility consists of seven principles and aligns with YAC's acceptance step. As highlighted in Oinas-

.

Kukkonen and Harjumaa's study, one-time or permanent behavior change is achieved more easily than attitude change (2009).

The study of persuasive technologies first emerged in the late 1990s (Fogg, 1998). Fogg defines *persuasive technologies* as technology designed to change attitudes or behaviors of users through persuasion and social influence but not through coercion (Fogg, 2002). How researchers individuate persuasive strategies varies. Cialdini developed six principles(Cialdini, 2004). Fogg described forty approaches under a general definition of persuasion (Fogg, 2002), and others have listed over 100 distinct tactics (Rhoads, 2007). Scholarly writing on these technologies began increasing in 2005 (Torning & Oinas-Kukkonen, 2009). More recently, Kaptein et al. studied how persuasive technologies can dynamically adapt to how users are persuaded (Kaptein et al., 2015). While these technologies all aim to directly change user behaviour, altering behaviour by first changing attitude has not received enough attention. According to Torning and Oinas Kukkonen (2009), five out of six persuasion-related studies address behavioural change and not attitude change. While behaviour involves the expression of feelings or action, attitude involves the mind's predisposition to ideas, values, or people. Attitude structures can be described in terms of three components:

- The affective component, involving a person's emotions about the attitude object (e.g., "I am scared of spiders").
- The behavioural (or conative) component, involving how attitude influences how individuals act or behave (e.g., "I will avoid spiders and scream if I see one");
- The cognitive component, involving a person's belief and knowledge about an attitude object (e.g., "I believe spiders are dangerous")

The ABC attitude model uses the above three components (Ellis, 1962) and assumes the link between attitudes and behaviour: that people are rational, that they always attempt to behave rationally, and that a person's behaviour should be consistent with their attitude(s). While this principle may be sound, people do not always follow it and sometimes behave in seemingly illogical ways, such as by smoking cigarettes while knowing that smoking causes lung cancer and heart disease. LaPiere shows that the cognitive and affective components of attitude do not always match behaviour(LaPiere, 1934). YAC, on the other hand, is more suitable in PCG systems as it is designed explicitly for influencing attitude which will then causes behaviour changes.

Torning and Oinas-Kukkonen (2009) and, later, Hamari et al.(2014) provide overviews of the history of applications in persuasive-systems design. According to Hamari et al.'s comprehensive literature review, the two most prevalent foci of persuasion-related studies are health and/or exercise (47.4%) and ecology (21.1%),

.

which included technologies aimed at conserving energy. Most of the above studies, some of which have resulted in commercial products or services, base behaviour persuasion on control sensors. Fit4Life, for example, combines the input of a number of sensors to help users achieve their weight-loss goals (Purpura et al., 2011).

The model tracks individual progress and generates actionable items by tracking sensor data and comparing that data against user goals. While these applications are intended to persuade users by controlling the behaviour, they are not designed to influence user attitude. That is, in the absence of the application, it is not guaranteed that the user follows the regular exercise to achieve the weight-loss goals.

Busch and Patil used personalized content generated through surveys and gamification to train people to follow cyber-security best practices (Busch et al., 2016). Gamification uses game concepts and mechanisms, such as competition and leader boards, in nonentertainment applications (McGonigal, 2012; Hunter, 2011). In Busch and Patil's study, users started with reading personalized content, then they were presented with a quiz, a challenge, and a score. To encourage individuals to compete, the system posted scores on a board. While gamification can be used to increase engagement and potentially persuasion, the personalization in existing work is primarily in the form of categorization of users based on age, gender, or personality/gamer type rather than unique individual features(Busch et al. 2016; Orji et al., 2014).

Personalization, as an alternative to one-size-fits-all approaches, has been investigated in some recommender and persuasive applications(Orji & Moffatt, 2018), but they have been mostly focused on personalized experience such as in educational and persuasive games or recommended physical activities(Dharia et al., 2018; Hagen et al., 2016; Macvean & Robertson, 2013; Orji et al., 2013).

Less attention has been paid to personalization of content, particularly persuasive and educational text.

In our literature review, we did not encounter a model that translated persuasion theories to a software that automatically generate persuasive, custom-made, personalized content. As briefly reviewed in this section, the following characteristics are collectively or partially missing in the existing research on persuasive systems:

- Reliance on a comprehensive theoretical model Attention to attitude vs. behaviour
- Personalization based on individual features as opposed to simple categories
- Focus on educational text as content

.

While existing content generation methods may suffer from these shortcomings, there are existing components that we can use and/or improve in a new comprehensive and integrated system.

SYSTEM DESIGN AND DESCRIPTION

Our proposed system is based on a YAC-aligned layered model of individualization, social media data to create a user profile, and a flexible authoring structure as described in the following sections. While the approach can use any social network, on the initial system focuses on Twitter. We generally followed the rapid application development methodology (Martin, 1991) which includes requirements analysis; emphasizes reusable code and components; and the generation of prototypes in a build, demonstrate and refine cycle.

YAC Layers and Pyramid of Individualization

A large population can be divided into smaller population segments based on common characteristics, interests, and needs. While these factors allow people to be grouped, within each group, individuals can be differentiated by personal characteristics like experiences, friends and family, and opinions. We modeled these distinguishing factors as a multilayered pyramid that uses the general population for its base and narrows to a specific individual through layers of distinguishing characteristics.

Figure 1 shows how we divided an individual's basic characteristics into four main layers of the Pyramid of Individualization (PoI): segmentation (demographic characteristics such as age and gender), comprehension, personality traits, and person-specific individualization. The first three layers are usually shared amongst a large group of people. In contrast, the individualization layer includes personal features—like family history or personal opinions—that are shared within smaller crowds and can be truly individual (mutually exclusive). We represent these intimate features and final individualization at the pyramid's tip.

.

Figure 1. Pyramid of Individualization (PoI)

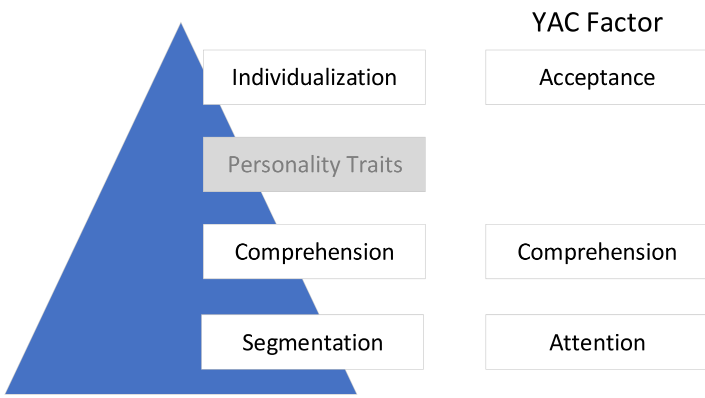

Although we identified personality traits as a PoI layer, it is outside the scope of the current version of our software framework and the implemented PCG system. Future phases of this research will aim at incorporating that layer.

Our multi-layer PoI model aligns with YAC and intentionally reflects its stages of persuasion. Segmentation makes sure that the topic is of general interest to the target group of people to attract their initial attention. The comprehension layer focuses on making the content understandable. Personality traits, once implemented, will help with these two YAC steps. The individualization layer is designed to make the content more acceptable. While none of PoI layers directly focuses on the YAC retention step, we hypothesize that together they will make the content more personalized and as such more memorable.

YAC identifies recipient-attention attainment as the first step towards persuasion success. "Interesting," "relevant," and "enjoyable" are parameters that increase the likelihood of gaining reader attention (Vesanen, 2007). With respect to this study, our approach to segmentation focuses on determining a persuadee's general areas of interest by categorization and mining social media feeds.

YAC suggests that audience acceptance and retention first require that an audience comprehends the presented message. For example, through three experiments, Eagly demonstrated how lowering comprehensibility lessened message acceptance (Eagly, 1974). This directly relates to classical (manually tuned) readability indices like the Flesch-Kincaid Index (Kincaid et al., 1975), the Gunning Fog Index (Gunning, 1952), and the Coleman-Liau Index (Coleman & Liau, 1975). Moreover, the criteria "easy to follow," "less complex," and "easy to understand" have been identified by researchers as a means to assess content comprehensibility(Collins-Thompson, 2014; Van der Sluis et al., 2014). Our approach involves assessing the comprehensibility level of online content that has been linked to through

.

persuadee's social media posts and then ensuring that the persuasive content aligns with that comprehensibility level.

Creating trust and credibility is a persuasiveness strategy that directly relates to the YAC acceptance stage. To be more granular, parameters like "trustworthy," "accurate," "authentic," and "believable" have been commonly accepted as a means to increase the likelihood of reader acceptance of content (Appelman & Sundar, 2016; Ravikumar et al., 2012). This study relies on the definition of 'trust' by Cho et al.(2015). They define trust as, "The willingness of the trustor (evaluator) to take risk based on a subjective belief that a trustee (evaluatee) will exhibit reliable behavior to maximize the trustor's interest under uncertainty (e.g., ambiguity due to conflicting evidence and/or ignorance caused by complete lack of evidence) of a given situation, based on the cognitive assessment of past experience with the trustee". This definition is particularly helpful because it provides clear reference to the relationship between two parties, the trustor's interest as the key goal, and the trustee's past actions as a key factor. In the context of SNS and Twitter, trustee's past actions are operationalized as tweets that are exchanged between trustor and trustee.

Our approach to individualization involves using our proposed User Trust Graph (UTG) (Arya et al., 2018) built from the persuadee's Twitter follower/followee structure to identify credible sources from which additional content can be appended to the persuasive content.

PCG System Overview

Our PCG system is shown in Figure 2. It is a collaboration between the following actors:

- System: the personalization system that assembles and prepares the content.
- Author (persuader): the content designer who creates the document templates and rule files that determine the main content topics and desired content sentiment. Content could convey a positive, neutral or negative sentiment.
- Reader (persuadee): the intended receiver of personalized content who logs into the system with proper authentication through their social-network account.
- Support: a trustable source (for example, a close friend, family member, public figure or organization) who has published material the system can use as supporting opinion

.

Figure 2. PCG detailed overview

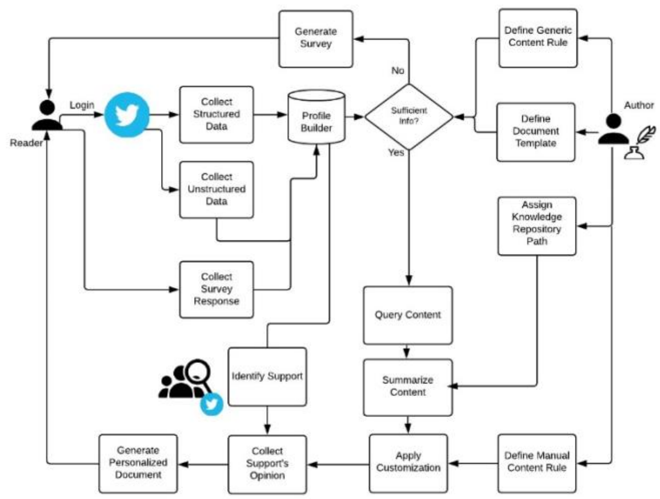

The PCG system has two major parts:

- User Interaction that includes components to process data related to the Author and Reader.
    - The author provides the inputs by storing the document template, content rules and knowledge repository path in a designated directory that is accessible by the system.
    - The reader signs in and authorizes the system to access his/her Twitter account and build the persuasive content through a web browser.

- PoI Layers that receive the users’ data and create the content

Author-related components are:

- Document template: The document contains the template for the persuasive content. The locations that require personalization by the system are identified with a unique ID in this template. It allows the aggregation of content from multiple sources.
- Content rules: the core logic prepared by the persuader and fed to the system to collect and summarize content.
- Knowledge repository: a collection of author-provided articles and documents.

The profile builder uses the collected data from the persuadee's social network account to assemble the user profile. The data from the user profile is later used to predict the user's comprehension and interest score.
The PoI layer components include the:

- Segmentation layer: the component where the system selects and assembles the initial content
- Comprehension layer: this ensures that persuasive content is worded at an appropriate level of complexity for the persuadee. The main engine that facilitates this result is the complexity assessment module for unstructured data in the user profiling subsystem.
- Individualization layer: the component where the system collects and appends related content (opinion) from users' credible sources as supporting information.

Rules Definition

Content rules are the essential core logic prepared by the persuader in XML format and fed to the personalization engine to collect, analyze, and personalize content. The rules are defined as XML elements that contain the two mandatory attributes called topic_id and polarity, as shown in Figure 3. The topic_id attribute uniquely identifies the rule, and the polarity attribute indicates the sentiment of the content that will be retrieved from the knowledge repository and ultimately displayed to the persuadee. The polarity attribute reflects author's (persuader's) expected attitude or emotion towards the content (text), i.e., whether it is positive, negative or neutral.
Within each rule, there can be many content elements. Each content element contains two attributes (segment, and weight), and a list of values. Segment indicates an area of interest within the given topic. For example, in Figure 4 , we see the overall topic being about solutions on fighting back climate change, and there are five content elements with cost management, cultural revolution, technology development, environmental education and resource contamination listed as segments. The weight attribute is used to boost the query clause. Knowledge repository contents matching this clause will have their score multiplied by the weight attribute.
For each segment, there is a list of values which reflect sub-areas of interest for the given segment. The values are separated with semicolons as demonstrated in a generic rule example figure. After user profiles are built from social media accounts, they are matched against the rule file. If a match occurs, a query clause is built, and content is pulled from the knowledge repository. If more than one

.

segment matches a user profile, then the weight attribute is used to determine the relative importance of each of the different segments.

Figure 3. Fighting climate change. A generic rule example

```xml
<rule topic="fight_climat_change" polarity="positive">
  <content fact="cost management" weight="15">renewable_energy;forest_management;land_management</content>
  <content fact="cultural revolution" weight="20">agricultural;consumption;production</content>
  <content fact="technology development" weight="25">renewable_energy;montreal_pact_ozone;efficiency_technology</content>
  <content fact="environmental education" weight="10">promoting_social_awareness;lifestyle_best_practices;human_factor</content>
  <content fact="resource contamination" weight="20">human_factor;nuclear_waste</content>
</rule>
```

As highlighted above, polarity (positive or negative sentiment) is a mandatory rule element attribute. When the system collects supporting opinion, it ensures that the sentiment associated to the opinion is aligned with the retrieved information from the knowledge repository. This way, the system avoids a contradiction between a supporting opinion and the original information.
The rule workflow is demonstrated in Figure 4. Besides creating normal (generic) rules, the author can add manual rules to overwrite the contents that have been retrieved by generic rules and identify specific content to be added. Doing so allows the author to prepare content and directly insert it into the final document.

Figure 4. Rule Workflow Overview

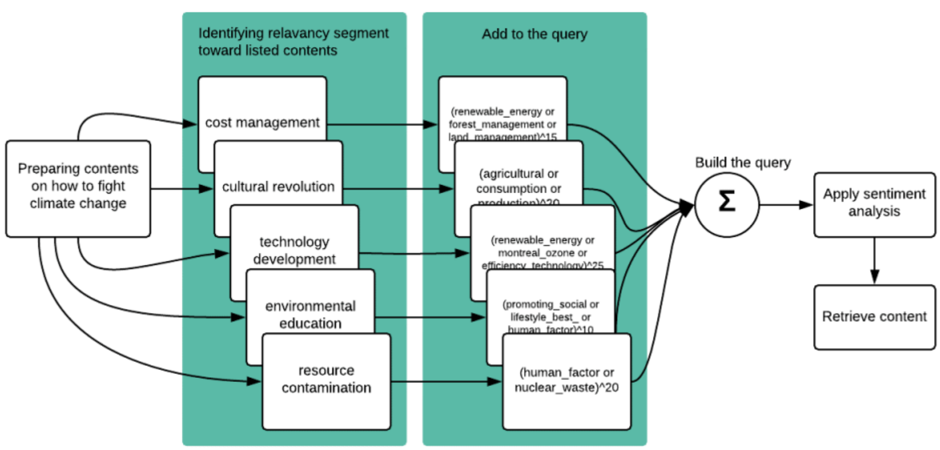

User Profiling

Before collecting content for the personalized document, the system needs to identify the target reader. The building of user profiles is shown in Figure 5 . The

.

primary data source for user-profile creation is the user's social media account. We used Twitter accounts because of their public-access flexibility using REST API. Our system creates user profiles from three information sources:

- Structured data: Twitter-profile information;
- Unstructured data: user tweets; and
- Persuadee direct input: additional information directly provided by the persuadee

Figure 5. User-profiling overview

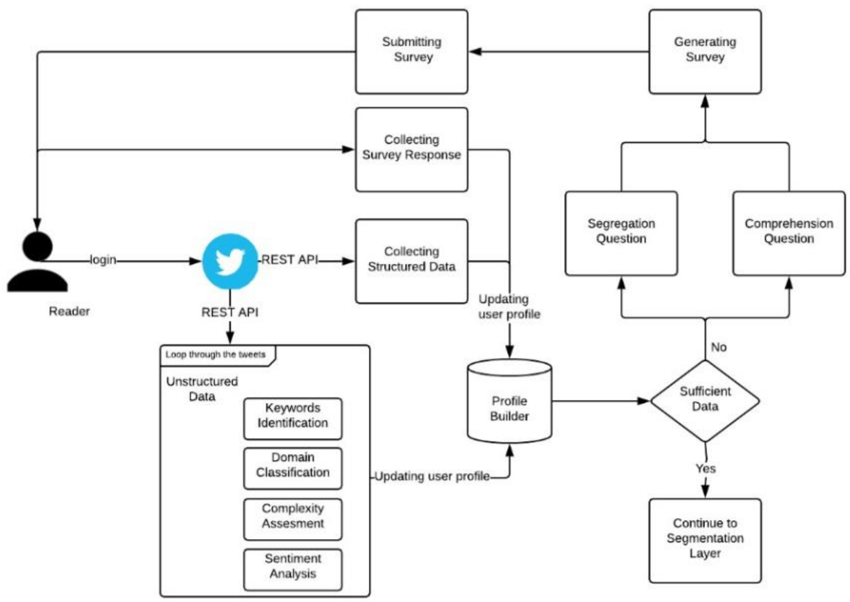

Structured Data

Personal user information, such as name, age, location, and a summary of interests, is available in most social network and microblog services for researchers.
Through open API, this information is easy to collect, and since the data is preorganized, it can be stored as is; however, data can be incomplete (a user may choose not to post bio details) or misleading. Furthermore, other relevant user attributes, such as explicit and implicit interests or political preferences, are usually missing. Cheng et al. (2010) estimated that only 26% of users report a specific

.

location, while the rest provide either general locations (e.g., states/provinces, countries) or nonexistent places. Pennacchiotti and Popescu (2011) conducted a pilot study of a similar nature to assess direct use of public profile information, such as gender and ethnicity, from Twitter. In a corpus of 14M active users in April 2010, they found 48% of users provided a short bio and 80% a location. Therefore, we decided not to rely solely on extracted data from user profiles.

Unstructured Data

Users' tweets are a good example of unstructured data and cover insights on personal attitudes toward different topics (e.g., political orientation or ethnicity). To extract the information, the system relies on features like typical n-grams models, simple sociolinguistic features (e.g., presence of emotions), and communication behaviour (e.g., frequently retweeted content). After extracting a keyword from the tweet (first module), the system runs the domain-classification (categorization), and complexity-assessment modules.

Keyword Extraction

Well-known Term Frequency–Inverse Document Frequency (TF-IDF) (Sammut & Webb, 2010) and TextRank (Mihalcea & Tarau, 2004) are two of the most practical and conventional techniques for keyword extraction. Both techniques can extract user interests from a collection of tweets with great precision (Vu & Perez, 2013). They are not, however, appropriate for a single tweet since most terms in a tweet are used only once. We, therefore, used the default MAUI (Medelyan et al., 2010) toolkit as a baseline for automated keyword collection. MAUI enables the extraction of a list of potential keywords from a document and trains a decision tree using features like TF-IDF.

Domain Classification

Keywords on their own can be vague, ambiguous, and polymorphic. For example, a word like *bond* can have a variety of meanings in various contexts. Classifying keywords by domains helps identify keyword hyponymies and hypernyms: word X is a hyponym of word Y if X is a subtype or instance of Y. For example, *bond* in the finance domain is a hyponym of debt; at the same time, it is a hypernym of a broad range of asset types, such as government bonds, municipal bonds, and corporate bonds. Defining such a hierarchy for keywords can help predict user preferences. For instance, if a user prefers content about Java, C++, and Javascript, the person would likely be interested in other programming languages.

.

We used a Word-Class Lattice (WCL) supervised model to identify textual definitions and extract hypernyms from those definitions (Faralli & Navigli, 2013). After evaluating WCL with a dataset of 1000 first-paragraph sentences randomly sampled from Wikipedia, we found the tool adequately rates precision, recall, and accuracy of definition and hypernym extraction.

Complexity Assessment

Our personalization process assesses user comprehension via a user-document-keyword-complexity evaluation that ensures that the complexity of generated content matches user preference. The simplest way to perform a comprehensibility-based ranking for a given topic is to build a classifier that assigns a comprehensibility score. Inspired by Tan et al.'s approach (2012), we trained a logistic-regression classifier by using comprehensibility indices to extract pages from Simple English Wikipedia and English Wikipedia. This allowed us to have two versions of a text: normal and simplified. The maximum number of characters per tweet is 280. Since the above method relies on word length and sentence length, and since 280 characters are insufficient for calculating comprehension score, we only assessed the comprehension score for tweets that contained URLs to other documents. The PCG system then ensures that the persuasive content comprehensibility is consistent with the generated comprehensibility score.

Persuadee Direct Input

Insufficient tweets, or ambiguity and a lack of consistency in collected tweets is a major challenge for automated systems and can introduce misrepresentative data to user profiles. To address these issues, we built a survey-generator module. The system uses the survey-generator module when it fails to identify a user preference on a topic. We designed the survey to enable sufficient data collection to tune a predicted user-interest-and-comprehension score. Once the system accesses persuadee's tweets and determines the need for additional data collection, the persuadee's web browser is redirected to the system generated survey. The workflow is illustrated in Figure 6.

.

Figure 6. Collecting user data through PCG survey-generator module

The system generates two types of questions: segregation preference and comprehension preference:

- Segregation-preference questions are generated based on the segment attribute in the generic-rule file (e.g., in Figure 7, the questionnaire asks for the user's income level) and help the system segregate users and assemble persuasive content;
    - Comprehension-preference questions ensure users have the adequate background knowledge to comprehend the persuasive content. For instance, a well-known professor might possess advanced knowledge of computer science but be an intermediate mechanic; thus, he prefers easy-to-understand mechanical repair instructions. As Figure 7 shows, the system asked readers to rank their expertise level (from 1 to 5) of a given topic. This information is necessary for the system to adjust the level of presented background knowledge in the final personalized document.

.

Figure 7. RRSP Benefit document template, rule template, and survey layout

Document Template

```
<div id= "RRSP_benefits">
 Default text(content) given by the author. The text will be used as the
 fallback in case the system fails to identify a reliable personalized content.
</div>
```

Rule Template

```
<rule topic="RRSP_benefits" polarity="positive">
 <content segment="upper class income" weight="20">tax benefit</content>
 <content segment ="middle class income" weight="20">retirement fund</content>
 <content segment ="lower class income" weight="20">first time home buyer</content>
</rule>
```

Persuadee Direct Input

On scale of 1 to 5, where “1” means “Not Familiar” and “5” means “Very Familiar”, how familiar you are with the topic of “RRSP Benefits”? - Purpose: Determining comprehensibility

1 (Not comfortable) 2 3 4 5 (Very comfortable)

Please select any of below fields that match your personal interest. - Purpose: Assigning segregation

Lower class income Middle class income Upper class income

Information Assembler: Segmentation

Once an adequate user profile has been built the system will start to assemble content through the segmentation layer. It initiates by building a query to retrieve documents from a knowledge repository. The current implementation’s knowledge repository is a Solr Apache Lucene collection that contains articles for a specific topic (articles are indexed according to title and author-assigned keywords). The author gives the system the knowledge-repository URL. The Solr query to retrieve documents consists of two rule factors that are ordered according to priority:

.

- Rule content specialization within the content's domain (e.g., retirement fund in Figure 7)
- Rule segment User criteria required for content segregation (e.g., middle-class income in Figure 7)

When the system executes a given query by an author, it applies the sentiment-analysis module from the polarization of user-profile keywords compare the extracted content sentiment with the assigned sentiment value given in the rule. This ensures the contents sentiment polarity is consistent with the author's intent. In addition, building the query means a user does not necessarily fall into a single category. For example, because a user may relate to multiple domains (rule topic), the query will expand to include all categories.
Solr returns a list of relevant documents which are summarized using an unsupervised technique using TextRank algorithm. TextRank is easy to adapt and ranks all sentences in a text. It consists of similar tasks like building graphs for texts, where the graph vertices represent the units to be ranked (Mihalcea & Tarau, 2004).
Our system includes a modified TF-IDF weighting system (a common numerical statistic used in TextRank) (Salton & Mcgill, 1986) wherein TF-ISF (Allan et al., 2003) weights are computed instead. TF-ISF is a more suitable weighting system, since it ranks sentences instead of words. $TF_{i,j}$ is the term frequency of i^{th} index term in the j^{th} sentence, and ISF_i is the inverse-sentence frequency of i^{th} index term.

Like the TF-IDF model, the corresponding weight for a sentence is computed as:

$$ws_{ij}=tf_{ij}\,.\,isf_i \qquad (1)$$

Due to the effectiveness of cosine similarity, we have used it to measure the similarity weight between the two sentences (s_a and s_b).

$$w_{sim}(s_a, s_b) = \frac{\sum ws_{i,a}.ws_{i,b}}{\sqrt{\sum ws_{i,a}^2}\,.\sqrt{\sum ws_{i,b}^2}} \qquad (2)$$

To preserve content coherence, the system retrieves and reorders the top n sentences according to the source document's original order. If the system fails to retrieve

.

relevant content from the knowledge repository, it relies on author-given default content in the presentation template.

Opinion Mining: Individualization

To individualize the content, the system identifies trustworthy sources and pulls their relevant content to augment and individualize the overall content presented to the reader. Highly trustworthy communicators inspire positive attitudes toward the positions they advocate (Sternthal et al., 1978). Collecting supporting opinion from a trustable source requires multiple subprocesses (as Figure 8 illustrates).

Figure 8. Overview of individualization layer

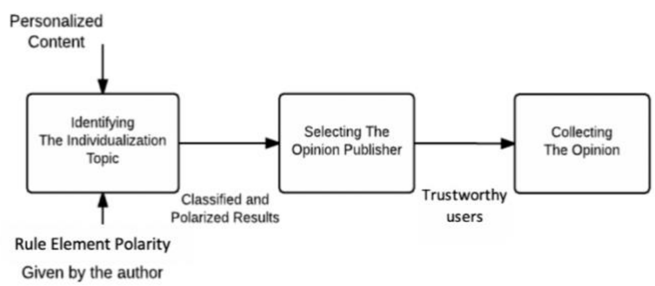

Identifying the Individualization Topic

Domain classification of content is the prerequisite for mining supporting arguments (opinions). As discussed earlier, the system collects content via user profiling and by applying author-driven content rules. Each rule element contains a polarity attribute that indicates an author's perspective toward a content element. To avoid contradiction between a supporting opinion and the content, the system first conducts a sentiment analysis on the content to identify content polarity. For instance, the following sentence implies a negative opinion about investing: *While stocks have historically performed well over the long term, there's no guarantee you'll make money on a stock at any given time, and you could lose all your principal*. Thus, the following would not make a good supporting opinion because it expresses the opposite perspective: *Just made a $3,000 investment in stock last year, and it turned into a $25,000 fortune within a couple years*.

.

Sentiment *Analysis*

Negative user preferences should be segregated from positive user preferences. The polarity of a sentiment is the point on the evaluation scale that corresponds to how positive or negative a sentiment is. Here we leverage from sentiment analysis on extracted keywords to determine whether user preference is negative, neutral or positive. Besides, by comparing the polarity of user sentiment toward the extracted keywords and the polarity of the presented content, we ensured user sentiment aligns with author-selected polarity. We performed sentiment analysis in the following cases:

- User keywords to determine user preferences
- Collecting opinionative information that is aligned with author selected rule polarity
- Collecting contents that is aligned with author selected rule polarity
- Evaluating trust between two users

To conduct the analysis, the sentiment-analysis module from the polarization of user-profile keywords relied on uClassify sentiment-analysis toolkit (https://www.uclassify.com/). uClassify treats sentiments as non-binary, and in this study, sentiments were classified using common categories like positive, negative, and neutral. Though, conducting sentiment analysis on user preference is not the only place where we leverage from uClassify. The next section further discusses how the system leverages from sentiment analysis on other related processes to create a persuasive content.

Searching for the Trustable Source

Online reputation and expertise are built through consistent and repeated contributions to online communities. We observe this in blogging communities and online review sites, among others where there isn't a strict limit on content creation. Twitter poses different challenges as a tweet has a length restriction, and we thus focus on the number of likes, re-tweets of posts, and our polarity analysis of the tweet content and responses. We use these interactions and analyses to build trust graphs that evaluate the trust between two people.

Neither a social graph nor structured data (such as total likes or retweets) are sufficient to measure user trustworthiness in a given social Twitter group. The sentiment associated with tweets shared between two parties is important. Existing approaches only use structured data to consider relationships among users.

To measure user trustworthiness in a given social Twitter group, we introduced the User-Trust Graph (UTG) (Arya et al., 2018) to estimate relationship-trust strength

.

from interactions (e.g., communication, tagging) and common interests. This graph is based on trust-score calculation as a function of time and tweet sentiment. As reported in (Arya et al., 2018), the UTG and has been pilot tested and initially validated.

The graph consists of nodes (corresponding to user accounts and tweets) and edges (corresponding to follow and retweet relationships). Unlike the Twitter social graph, which is relatively static, the user-trust graph is dynamic and reconstructed when: one user mentions the other in a tweet, both users share a retweet or hashtag, one user likes a tweet by the other.

We modelled interactions between two Twitter users through a trust graph with multiple links. In this user-trust graph, each node represents a user. There are two types of edges in this graph: directed and undirected. A directed edge between two nodes—A and B—exists if User A mentions User B in at least one tweet or likes a tweet from User B or retweets at least one of User B's tweets. When two users share a common hashtag or retweet, there is an undirected edge between two nodes (as Figure 9 illustrates).

Figure 9. User-trust graph

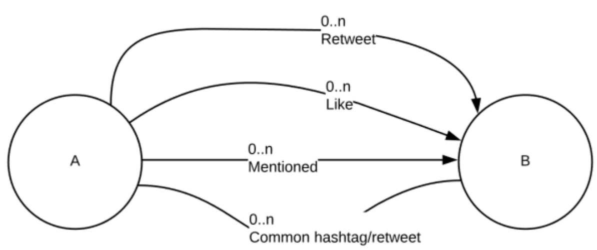

The system performs sentiment analysis on interactions between users and assigns respective positive, neutral, or negative scores. Despite the sentiment of a tweet's content, when a user retweets another person's tweet, the interaction is considered to have a positive score. But a retweet with a negative comment is considered to have a negative score. Intuitively, sum of interaction scores (sentiment) between two users is interpreted as trust strength between them.

Several PCG prototypes were built, demonstrated and refined. The following section describes the validation of the final system.

.

EVALUATION

Evaluation Overview

To evaluate the system, we employed a between-subjects experimental design where subjects were randomly assigned to a PCG condition (n=23)_ or a generic condition (n=25). Subjects in the PCG condition received fully personalized persuasive content generated by the PCG system while subjects in the generic condition received directly from a popular press source. Questionnaires related to persuasion were administered before the participation, right after the participation, and two weeks later. Climate change was chosen as a suitable domain to evaluate the system since many studies have attempted to persuade people to take action against climate change as it remains one of the largest problems we face globally. Changing people's attitudes about socio-scientific issues like climate change has unique challenges(Mason et al., 2006)(Lombardi & Sinatra, 2012): according to Sakamoto and Goldstone, "it is complex, multidimensional, and requires systems thinking, that is, the ability to think and reason abstractly about systems to appreciate their interactive nature" (Goldstone & Sakamoto, 2003). The choice of climate change is also consistent with Oinas-Kukkonen's stated requirement persuasive systems address global and cultural issues (Oinas-Kukkonen, 2013). Participants in the generic group were presented persuasive content from National Geographic[1].

The generic content was 598 words in length and had a Flesch-Kincaid readability score of 10.5 reflecting that the content is easy to read and easily understood by an average eleven-year-old student. The generic text was structured to demonstrate solid evidence for global climate change and to support that humans are contributing to the trend.

Hypotheses

As YAC was the model that informed the design of the PCG, we used the YAC factors *attention*, *comprehension*, *acceptance*, and *retention* as assessment mechanisms to evaluate the persuasiveness of the presented content. According to YAC (Hovland et al., 1953), to maximize the chance of persuading a user to take action or change opinion, first, their *attention* must be gained and the content's *comprehension* level must be adjusted so the user can read and understand the message.

As Vesanen and, later, Zhou et al. state, "Interesting," "relevant," and "enjoyable" are parameters that increase the likelihood of gaining reader attention (Zhou et al.,

[1]https://www.nationalgeographic.com/environment/global-warming/global-warming-solutions/

.

2012; Vesanen, 2007); we, therefore, proposed the following *attention* related hypotheses:

H1: The readers find the PCG content significantly more interesting than the generic content.

H2: The readers find the PCG content significantly more enjoyable than the generic content.

H3: The readers find the PCG content significantly more relevant than the generic content.

The criteria “easy to follow,” “less complex,” and “easy to understand” have been identified by researchers as a means to assess content comprehensibility (Collins-Thompson, 2014; Van der Sluis et al., 2014) therefore, we proposed the following *comprehension* related hypotheses:

H4: The readers find the PCG content significantly easier to follow than the generic content.

H5: The readers find the PCG content significantly less complex than the generic content.

H6: The readers find the PCG content significantly easier to understand than the generic content.

YAC defines the persuasion cycle as complete when the persuadee *accepts* and *retains* the presented information. Parameters like “trustworthy,” “accurate,” “authentic,” and “believable” have been commonly accepted as a means to increase the likelihood of reader acceptance of a content (Appelman & Sundar, 2016) we, therefore, defined the following *acceptance* related hypotheses:

H7: The readers find the PCG content significantly more trustworthy than the generic content.

H8: The readers find the PCG content significantly more accurate than the generic content.

H9: The readers find the PCG content significantly more authentic than the generic content.

.

H10: The readers find the PCG content significantly more believable than the generic content.

In the context of our study, retention refers to remembering the effect of the treatment. Some scholars define it as the persistence to perform (take action) a learned knowledge or behavior (Rubin & Wenzel, 1996; Ausubel, 1963). While individuals may remember the content due to factors other than the persuasiveness of content (such as their personality traits and memory), the overall statistical result can be a measure of persuasiveness as it includes a random set of participants. Prior research has used a two-week delay (Liebowitz & Frank, 2010)) to assess retention of a message.

H11: There is no significant change in PCG-group-participant intention to take action toward climate change between their post-treatment response to their two-weeks post-treatment response.

Finally, we assess the impact of the treatment on influencing participants' attitude and potential action towards climate change.

H12: There is a positive significant change in PCG-group-participant attitudes toward climate change from the pre-treatment to post-treatment responses.

H13: The readers of the PCG content are significantly more likely to have the intention to take action against climate change than the readers of the generic content.

Methodology

The content of the XML rule files followed Norwegian psychologist Espen Stoknes's advice on how to prepare persuasive contents (Stoknes, 2015) for climate change. Stoknes identified localized content with positive sentiment as an appropriate strategy to fight climate change. He argues against talking about global effects and says targeting messages to local perceptions will more effectively communicate the dangers of climate change. For instance, after Hurricane Sandy, New Yorkers understood sea-level rise; similarly, Californians now understand long-term drought. We downloaded 100 articles related to climate-change problems and solutions from Yale Environment 360, an online magazine published by *Yale School of Forestry & Environmental Studies* and Yale University. Each article contained 1000 to 2000 words and provided opinion, analysis, reporting, and debate on global environmental issues and features original articles by scientists, journalists, environmentalists, and academics.

.

In addition, we connected the PCG to *The Guardian* API so that the PCG could collect articles without having to download and host them separately.

Participants had to be 18+ years old, have a valid twitter account and be a fluent English speaker. They were recruited through a variety of postings on online channels at a major Canadian university. Once a potential participant visited the experiment website, they were randomly assigned to either the PCG or generic groups.

We designed three questionnaires to measure participants' opinions and feedback at three different points in time (Appendix I). Participants in both groups received all questionnaires. The pre-treatment questionnaire contained three questions answered before reading any content assessing participants' current views on climate change, and their ability to control environmental quality. The post-treatment questionnaire contained questions that are answered immediately after reading persuasive content. It includes a repeat of the three pre-treatment questions; an additional question about the subject's intention to take action against climate change; and a series of questions assessing attention, comprehension, and acceptance. The two-weeks post-treatment questionnaire was given two weeks after reading persuasive content, repeating the question assessing subjects' intention to take action against climate change (used in assessing retention). Two-week intervals have been used in other similar projects (Liebowitz & Frank, 2010) to assess retention of a message.

As we obtained results from Likert-scale data (ranges 1-5), with a limited number of responses, we could not assume our data were parametric, leading us to analyze them via nonparametric statistical methods. We used the Mann-Whitney test to assess whether two sets of data were significantly different from each other across and within treatments.

After submitting the questionnaire, participants were asked to log into the main system using their Twitter account. For those participants in the PCG treatment, the system-built user profiles by collecting the first 500 recent tweets from each Twitter account. The data processing task consists of extracting keywords and identifying the keywords' hyponymy (domain). Keywords given in user profiles demonstrated users' respective topic-relation factors. When the system was building the query, it looked for matching rule segments in the user profile`s keyword collection. The segment with the highest ratio match was picked for the query. If the segment failed to match any of the user's keywords, the system looked for matching rule segment in the user profile as a fallback. Once the query was completed, the system passed the query to Solr so that the most relevant documents from the knowledge repository could be retrieved.

To identity a quote on climate change from a trustable source, the system first generated a user-trust graph. We relied on users' Twitter interactions (likes, mentions, and retweets) to build trust graphs.

.

The system sorted graph results using a descending trust value; the system used the first fifty trustable sources to generate a list. As a fallback method, we created a second list from all the profiles that the user followed. The system passed both lists to Solr so Solr could retrieve the most relevant quotes from the knowledge repository. Based on the above, the PCG system presented the user with personalized content on climate change. Participants in the generic treatment received the National Geographic content as previously described.
After reading the content, the system asked the participants to complete the second questionnaire. The process of completing Questionnaire 1, reading the climate change content, and answering Questionnaire 2 was designed to take approximately fifteen minutes. However, we enforced no hard limit on how long participants spent on the experiment. Participants were also asked to respond to a follow-up questionnaire two weeks from the date of the study.

Results

The results for the attention questions are presented in Table 1. For the "Interesting" and the "Relevant" questions, the computed p-values are significant. In the "relevant" category, the PCG mean value is higher than the generic mean value. However, for the "interesting" question, the mean value for the Generic group is higher than that of the PCG group, which is the opposite of what was hypothesized.

Table 1. Attention-related results
****significant at .01**

	PCG		Generic		Mann-Whitney U	
Question	Mean	Variance	Mean	Variance	U	p-value
Interesting	4.045	0.426	4.167	0.667	297	< 0.0001**
Enjoyable	3.545	1.117	4.000	0.870	332	0.133
Relevant	4.364	0.433	2.917	1.036	68	< 0.0001**

Based on these results, we accept H3 and reject H1 and H2. Overall, there was partial support that the PCG content attracts more reader's attention than the generic persuasion content.
Results for the *comprehension* questions are shown in Table 2 below. As shown the computed p-value is significant for the "Easy to follow" question.

.

Other results were not significant. We therefore accept H4 and reject H5 and H6. Overall, there was partial support for the PCG generated content being more comprehensible than the generic persuasion content.

Table 2. Comprehension-related results
****significant at .01**

	PCG		Generic		Mann-Whitney U	
Question	Mean	Variance	Mean	Variance	U	p-value
Easy to Follow	4.545	0.450	3.792	1.129	152	0.002**
Complex	2.455	1.212	2.958	1.259	327	0.146
Understandability	4.227	0.755	3.792	1.303	208	0.124

Results for the 'acceptance' questions are show in Table 3 below. As shown the computed p-value is significant for two questions. Mean responses for both the "accurate" and "believable" questions are significantly higher for PCG than the generic group. Thus, by accepting H8 and H10 and rejecting H7 and H9, we found partial support for PCG content as being more acceptable than generic content.

Table 3. Acceptance-related results **significant at .01

	PCG		Generic		Mann-Whitney U	
Question	Mean	Variance	Mean	Variance	U	p-value
Trustworthy	4.136	0.790	4.208	0.694	273.5	0.075
Accurate	4.273	0.684	4.042	0.650	221	0.002**
Authentic	3.955	0.807	3.875	0.984	251	0.893
Believable	4.545	0.355	3.625	1.201	137.5	0.001**

To test retention of the PCG system we assessed the 'Intention to act' responses from immediately after reading the content with responses two weeks later. Results are shown in Table 4. As shown, there was a significant decrease in the responses to the 'intention to act' question indicating that the initial impact of the system's content was not retained over time. We thus reject H11.

.

Table 4. Retention Results for "Intention to Act against climate change" *significant at .05

	T2		T3		Mann-Whitney U	
Group	Mean	SD	Mean	SD	U	p-value
PCG	4.364	0.790	3.636	1.049	338	0.01*
Generic	3.958	0.908	3.708	0.908	336	0.259*

To test H12, we evaluated responses to the attitude towards climate change questions for the PCG subjects across T1 and T2. Results are shown in Table 5 below. As shown, the only significant change was in responses to 'environmental problems being exaggerated' which were reduced after reading the PCG content. There was thus partial support for H12.

Table 5. Results attitude towards climate change **significant at .01

	T1		T2		Mann-Whitney U	
Question	Mean	SD	Mean	SD	U	p-value
Evidence of Climate Change	4.500	0.598	4.636	0.492	216	0.717
Effect on Quality of Environment	2.318	1.211	2.182	1.006	249.5	0.871
Environmental Problems Exaggerated	3.182	0.733	2.591	0.590	346	0.009**

To test H13, we evaluated the responses to the 'Intention to act against climate change' question taken immediately after reading the persuasive content across the two treatments. As shown below in Table 6, responses from the PCG subject were significantly higher than those assigned to the generic treatment thus supporting H13.

Table 6. Results of Intention to Act Against Climate Change **significant at .01

	PCG		Generic		Mann-Whitney U	
Question	Mean	Variance	Mean	Variance	U	p-value
Intention to Act	4.364	0.790	3.958	0.908	197	<0.0001**

.

While the results are mixed, there is evidence that the PCG system can outperform generic persuasive content. Overall, there was partial support for each of the attention (H1, H2, and H3), comprehension (H4, H5, and H6), and acceptance (H7, H8, H9, and H10) evaluation hypotheses. Subjects in the PCG treatment found their content significantly more relevant, but significantly less interesting than those subjects in the generic treatment. Due to the PCG-segmentation components, PCG content had a higher relevancy mean value than the generic content. Subjects in the PCG treatment found their content significantly easier to follow, but not significantly less complex or easier to understand than the subjects in the generic treatment. Subjects in the PCG treatment found their content significantly more accurate and believable but not significantly more authentic or trustworthy than those subjects in the generic treatment. One of the PCG-content sections was dedicated to the latest personalized news related to climate change. Typically, during this phase of personalization, the PCG system looked for news articles within the participants' respective geolocations. Thus, participants were likely personally connected to the news-article content rather than scientific facts in the generic content. We also documented that enhancing the PCG content with news articles from reputable news agencies improved the content's believability. The PCG system was very effective in significantly increasing subjects' intention to take action against climate change (H13) but unfortunately said the impact was not retained when subjects were evaluated two weeks later (H11).

Finally, after exposure to the PCG content, subjects significantly reduced the degree to which they thought environmental problems are exaggerated. However, there was no significant change in subjects' thoughts about evidence of climate change or the degree to which they feel they can affect the quality of the environment. Results could be because climate change was already well-known among participants as subjects were like to belong to a moderate-to-highly educated group of people.

While the results are encouraging there are several limitations to the evaluation that may have contributed to the mixed results. First, the sample size is small and thus the results should be interpreted with caution. Also, the sampling is from a limited population as recruitment was done within a university environment. This resulted in a relatively highly educated set of participants who may already be knowledgeable about the impacts of climate change and thus persuading them may have limited possibilities. Finally, the measures used to assess *attention*, *comprehension*, *acceptance*, and *retention* were primarily indirect.

.

CONCLUSION AND FUTURE RESEARCH

This paper addresses the problem of delivering personalized persuasive content through a semi-automated system. The primary research goal was to provide a computational model, design and develop a system based on that model, and perform a preliminary evaluation to determine whether the system can improve the persuasiveness of content compared to generic persuasive text written by an expert. Unlike most other persuasion system efforts, the PCG can be used in multiple domains and is informed by a model (YAC) that is focused on attitude change.

The contributions of this paper are many. It introduces the PoI as a new model and basis for creating persuasive content. The new model extends the YAC factors into the context of persuasive text by defining different layers of personalization. In addition, a pluggable software framework that integrates newly designed re-usable software components with existing ones is presented.

These new components include: a rule-based component to create and control personalized content by the author (persuader); a User (persuadee) profile builder based on extracting information from social media data; and a user trust graph and related algorithms for selecting trustworthy supportive information. Overall, initial results indicate evidence of the PCG content being more persuasive than generic content.

Future work will include identifying partnerships with industry who typically rely on text for persuasive purposes. We will initially focus on public health institutions who have large portfolios on diverse topics where text is often used for educational purposes. This will require an investigation into privacy concerns. Additional work on the PCG system itself will include using other social media and aggregating different sources of data. This could include incorporating relevant personal tweets and opinions into the PCG content rather than relying on BrainyQuote or similar repositories. As the system evolves, we plan to validate an integrated persuasion scale incorporating the YAC factors of *attention*, *comprehension*, *acceptance*, and *retention*. Finally, we hope to expand the system to include the personality traits layer of the PoI model.

REFERENCES

Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*. https://doi.org/10.1016/0749-5978(91)90020-T

Allan, J., Wade, C., & Bolivar, A. (2003). *Retrieval and Novelty Detection at the Sentence Level*.

Appelman, A., & Sundar, S. S. (2016). Measuring Message Credibility. *Journalism & Mass Communication Quarterly*, *93*(1), 59–79.

https://doi.org/10.1177/1077699015606057
Aronson, E., Wilson, T. D., & Akert, R. M. (2009). *Social Psychology* (7th ed.). Pearson.
Arya, A., Hine, M., & Khataei, A. (2018). User trust graph: A model to measure trustworthiness. *CEUR Workshop Proceedings*, *2089*.
Ausubel, D. (1963). The psychology of meaningful verbal learning. In *Grune & Stratton*. http://psycnet.apa.org/record/1964-10399-000
Busch, M., Patil, S., Regal, G., Hochleitner, C., & Tscheligi, M. (2016). Persuasive Information Security: Techniques to Help Employees Protect Organizational Information Security. *Persuasive Technology*, *9638*, 339–351. https://doi.org/10.1007/978-3-319-31510-2_29
Cacioppo, J. T., Petty, R. E., Kao, C. F., & Rodriguez, R. (1986). Central and peripheral routes to persuasion: An individual difference perspective. *Journal of Personality and Social Psychology*, *51*. https://doi.org/10.1037/0022-3514.51.5.1032
Cheng, Z., Caverlee, J., & Lee, K. (2010). You are where you tweet. *Proceedings of the 19th ACM International Conference on Information and Knowledge Management - CIKM '10*, 759. https://doi.org/10.1145/1871437.1871535
Cho, J.-H., Chan, K., & Adali, S. (2015). A Survey on Trust Modeling. *ACM Computing Surveys*, *48*(2), 1–40. https://doi.org/10.1145/2815595
Cialdini, R. (2004). The Science of Persuasion. *Scientific American*. https://www.scientificamerican.com/article/the-science-of-persuasion/
Coleman, M., & Liau, T. L. (1975). A computer readability formula designed for machine scoring. *Journal of Applied Psychology*, *60*(2).
Collins-Thompson, K. (2014). Computational Assessment of Text Readability: A Survey of Past, Present, and Future Research. *Www-Personal.Umich.Edu*. http://www-personal.umich.edu/~kevynct/pubs/ITL-readability-invited-article-v6-4review.pdf
Crowson, H. M., & DeBacker, T. K. (2008). Belief, Motivational, and Ideological Correlates of Human Rights Attitudes. *The Journal of Social Psychology*, *148*. https://doi.org/10.3200/SOCP.148.3.293-310
Dharia, S., Eirinaki, M., Jain, V., Patel, J., Varlamis, I., Vora, J., & Yamauchi, R. (2018). Social recommendations for personalized fitness assistance. *Personal and Ubiquitous Computing*, *22*(2), 245–257. https://doi.org/10.1007/s00779-017-1039-8
Eagly, A. H. (1974). Comprehensibility of persuasive arguments as a determinant of opinion change. *Journal of Personality and Social Psychology*, *29*(6), 758–773. https://doi.org/10.1037/h0036202
Ellis, A. (1962). *Reason and emotion in psychotherapy.* Lyle Stuart.
Faralli, S., & Navigli, R. (2013). A Java Framework for Multilingual Definition and Hypernym Extraction. *51st Annual Meeting of the Association for*

.

Computational Linguistics: System Demonstrations. http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.386.4745

Fogg, B. (1998). Persuasive computers. *Human Factors in Computing Systems*. https://doi.org/10.1145/274644.274677

Fogg, B. J. (2002). Persuasive technology: Using Computers to Change What We Think and Do. In *Morgan Kaufmann*. https://doi.org/10.1145/764008.763957

Fogg, B. J. (2003). *Persuasive Computing: technologies designed to change attitudes and behaviors*. Morgan Kaufmann.

Goldstone, R. L., & Sakamoto, Y. (2003). The transfer of abstract principles governing complex adaptive systems. *Cognitive Psychology*, *46*, 414–466.

Gunning, R. (1952). *The Technique of Clear Writing*. McGraw-Hill.

Hagen, K., Chorianopoulos, K., Wang, A. I., Jaccheri, L., & Weie, S. (2016). Gameplay as exercise. *Conference on Human Factors in Computing Systems - Proceedings*, *07-12-May-2016*, 1872–1878. https://doi.org/10.1145/2851581.2892515

Hamari, J., Koivisto, J., & Pakkanen, T. (2014). *Do Persuasive Technologies Persuade? - A Review of Empirical Studies* (pp. 118–136). Springer International Publishing. https://doi.org/10.1007/978-3-319-07127-5_11

Hovland, C. I., Janis, I. L., & Kelley, H. H. (1953). *Communication and persuasion; psychological studies of opinion change.* Yale University Press.

Hunter, R. (2011). *The gamification handbook : everything you need to know about gamification.* Lightning Source. https://books.google.ca/books/about/The_Gamification_Handbook.html?id=fhxRngEACAAJ&redir_esc=y

Kaptein, M., Markopoulos, P., de Ruyter, B., & Aarts, E. (2015). Personalizing persuasive technologies: Explicit and implicit personalization using persuasion profiles. *International Journal of Human-Computer Studies*, *77*, 38–51. https://doi.org/10.1016/j.ijhcs.2015.01.004

Kardash, C. M., & Scholes, R. J. (1996). Effects of preexisiting beliefs, epistemological beliefs, and need for cognition on interpretation of controversial issues. *Journal of Educational Psychology*, *88*(2), 260–271. https://doi.org/10.1037/0022-0663.88.2.260

Khataei, A., & Arya, A. (2015). Personalized Presentation Builder for Persuasive Communication. *Communication Design Quarterly Review*, *3*(3), 25–32. http://dl.acm.org/citation.cfm?id=2792993

Kincaid, J. P., Fishburne, J. R. P., Rogers, R. L., & Chissom, B. S. (1975). *Derivation of New Readability Formulas (Automated Readability Index, Fog Count and Flesch Reading Ease Formula) for Navy Enlisted Personnel*. http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA006655

LaPiere, R. T. (1934). Attitudes vs. Actions. *Social Forces*, *13*(2).

.

https://doi.org/10.2307/2570339
Liebowitz, J., & Frank, M. (2010). *Knowledge Management and E-Learning*. CRC Press.
Lombardi, D., & Sinatra, G. M. (2012). College Students' Perceptions About the Plausibility of Human-Induced Climate Change. *Research in Science Education*, *42*(2), 201–217. https://doi.org/10.1007/s11165-010-9196-z
Macvean, A., & Robertson, J. (2013). Understanding exergame users' physical activity, motivation and behavior over time. *Conference on Human Factors in Computing Systems - Proceedings*, 1251–1260. https://doi.org/10.1145/2470654.2466163
Martin, J. (1991). *Rapid application development. Macmillan College Division*.
Mason, L., Scirica, F., & Salvi, L. (2006). Effects of beliefs about meaning construction and task instructions on interpretation of narrative text. *Contemporary Educational Psychology*, *31*. https://doi.org/10.1016/j.cedpsych.2005.12.001
McGonigal, J. (2012). *Reality is broken : why games make us better and how they can change the world*. Random House. https://books.google.ca/books?id=Aa4j6YC1Vu8C&dq=Why+Games+Make+Us+Better+and+How+They+Can+Change+the+World&hl=en&sa=X&ved=0ahUKEwjCj9DYvZ_dAhUFmVkKHTq-Bi8Q6AEILTAB
Medelyan, O., Perrone, V., & Witten, I. H. (2010). Subject metadata support powered by Maui. *Proceedings of the 10th Annual Joint Conference on Digital Libraries - JCDL '10*. https://doi.org/10.1145/1816123.1816204
Mihalcea, R., & Tarau, P. (2004). TextRank: Bringing order into texts. *Proceedings of EMNLP*, *4*(4), 404–411. https://doi.org/10.3115/1219044.1219064
Oinas-Kukkonen, H. (2013). A foundation for the study of behavior change suppor systems. *Personal and Ubiquitous Computing*, *17*(6), 1223–1235.
Oinas-Kukkonen, Harri, & Harjumaa, M. (2008). A systematic framework for designing and evaluating persuasive systems. *Persuasive Technology: Third International Conference, PERSUASIVE 2008*. https://doi.org/10.1007/978-3-540-68504-3-15
Oinas-Kukkonen, Harri, & Harjumaa, M. (2009). Persuasive Systems Design: Key Issues, Process Model, and System Features. *Communications of the Association for Information System*, *24*.
Orji, R., Mandryk, R. L., Vassileva, J., & Gerling, K. M. (2013). Tailoring persuasive health games to gamer type. *Conference on Human Factors in Computing Systems - Proceedings*, 2467–2476. https://doi.org/10.1145/2470654.2481341
Orji, R., & Moffatt, K. (2018). Persuasive technology for health and wellness: State-of-the-art and emerging trends. *Health Informatics Journal*, *24*(1), 66–

.

91. https://doi.org/10.1177/1460458216650979
Orji, R., Vassileva, J., & Mandryk, R. L. (2014). Modeling the efficacy of persuasive strategies for different gamer types in serious games for health. *User Modeling and User-Adapted Interaction*, *24*. https://doi.org/10.1007/s11257-014-9149-8
Pennacchiotti, M., & Popescu, A.-M. (2011). A Machine Learning Approach to Twitter User Classification. *Icwsm*, 281–288. http://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/viewFile/2886/3262
Purpura, S., Schwanda, V., Williams, K., Stubler, W., & Sengers, P. (2011). Fit4Life: The Design of a Persuasive Technology Promoting Healthy Behavior and Ideal Weight. *SIGCHI Conference on Human Factors in Computing Systems*, 423–432.
Ravikumar, S., Balakrishnan, R., & Kambhampati, S. (2012). Ranking tweets considering trust and relevance. *Proceedings of the Ninth International Workshop on Information Integration on the Web - IIWeb '12*, 1–4. https://doi.org/10.1145/2331801.2331805
Rhoads, K. (2007). *How Many Influence, Persuasion, Compliance Tactics Strategies Are There?* www.workingpsychology.com
Rubin, D. C., & Wenzel, A. E. (1996). One hundred years of forgetting: A quantitative description of retention. *Psychological Review*, *103*(4), 734–760. https://doi.org/10.1037/0033-295X.103.4.734
Salton, G., & Mcgill, M. (1986). *Introduction to Modern Information Retrieval*. McGraw-Hill.
Sammut, C., & Webb, G. I. (Eds.). (2010). TF--IDF. In *Encyclopedia of Machine Learning* (pp. 986–987). Springer US. https://doi.org/10.1007/978-0-387-30164-8_832
Slattery, P., Finnegan, P., & Vidgen, R. (2020). Persuasion: an analysis and common frame of reference for IS research. *Communications of the Association for Information Systems*, *46*.
Sternthal, B., Ruby, D., & Leavit, C. (1978). The Persuasive Effect of Source Credibility: Tests of Cognitive Response. *Journal of Consumer Research*, *4*(4), 252–260.
Stoknes, P. E. (2015). *What we think about when we try not to think about global warming : toward a new psychology of climate action*. Chelsea Green Publishing.
Tan, C., Gabrilovich, E., & Pang, B. (2012). To Each His Own : Personalized Content Selection based on Text Comprehensibility. *WSDM*.
Torning, K., & Oinas-Kukkonen, H. (2009). *Persuasive System Design: State of the Art and Future Directions*.
Van der Sluis, F., Van den Broek, E. L., Glassey, R. J., Van Dijk, E. M. A. G., &

.

De Jong, F. M. G. (2014). When complexity becomes interesting. *Journal of the Association for Information Science and Technology*. https://doi.org/10.1002/asi.23095

Vesanen, J. (2007). What is personalization? A conceptual framework. *European Journal of Marketing*. https://doi.org/10.1108/03090560710737534

Vu, T., & Perez, V. (2013). Interest mining from user tweets. *ACM International Conference on Conference on Information & Knowledge Management*. https://doi.org/10.1145/2505515.2507883

Zhou, X., Xu, Y., Li, Y., Josang, A., & Cox, C. (2012). The state-of-the-art in personalized recommender systems for social networking. *Artificial Intelligence Review, 37*. https://doi.org/10.1007/s10462-011-9222-1

.

Declarations

Funding: Not applicable

Conflicts of Interest: On behalf of all authors, the corresponding author states that there is no conflict of interest

.