Building a Personalized College Major Selection Web Page: Technical Report -- February 15, 2016

Ryan L. Boyd and James W. Pennebaker
Department of Psychology

Overview

Nearly 25% of incoming UT-Austin students are unable to get their first two choices for a college major. Historically, these students have been given an extensive list of all potential majors from which to choose. Many students simply lack awareness of the various majors and have no background knowledge that could be helpful in determining whether specific majors would suit their interests or skills.

The purpose of this project was to rely on students' admissions essays to provide a more coherent and tailored set of recommendations when students are selecting an alternative college major. The logic underlying this project is based on decades of empirical research demonstrating that psychological information can be extracted from the language of students' admissions essays via automated computer analyses. The results of these analyses can then be used to inform the "college major options" webpage so that potential majors most closely aligned with their interests and skills will be displayed first in a recommendation system.

Background to the Approach

Over the last several years, a revolution in computer text analysis has been taking place. It is now possible to analyze text files and determine the basic content of the files as well as the psychological properties of the authors. An example of the power of text analysis methods with admissions essays was recently undertaken at UT-Austin. Relying on admissions essays from over 25,000 students, it was possible to identify the thinking styles of students in a way that predicted their 4-year GPAs as well as college entrance scores and rates well above chance (Pennebaker, Chung, Frazee, LaVergne, and Beavers, 2014).

Very broadly, there are two common strategies to analyze open-ended text relevant to the Admissions Office project. The first is to explore the content, or "topics", that students write about. The second is to examine the writing style of the essays. Whereas writing topic reveals the themes of the essays by focusing on nouns, regular verbs, adjectives, and some adverbs, people's writing styles reveal their personality, social connections, and thinking styles by analyzing a group of small forgettable words such as pronouns and articles.

Linking Writing Content and Style to College Majors

The ultimate goal of the project is to link the content and style of the admissions essays to students' tendencies to excel within their concluding major. To build such a system, we have gone through the following steps that included:

1. Assembling a database of at least 10 years of college admission essays written by UT applicants who were admitted to UT and attended at least one semester. In addition to the essays themselves, we also collected standard demographic information (age, sex, social class, college entrance exam scores, high school rank), first choices for major, ultimate major selected, yearly GPA, years in college and graduation information, GPA rank for their final degree earned.
2. Analyzing the admission essays in multiple ways including identifying and classifying the dominant topics of the essays, as well as the use of LIWC to identify the rates of use of various function and emotion words, including summary variables linked to academic performance. For each student, then, we will have a complete breakdown of meaningful language dimensions for their admissions essays.
3. Linking language dimensions to success in college in general and in their major specifically. Particular attention will be paid to students who did not get into the majors of their choice. The initial purpose of the linking analyses will be to build models using machine learning methods to predict successful college performance using exclusively language variables (i.e., models will not include demographic information such as gender, socioeconomic background, ethnicity, etc.).
4. Building and testing specific models to determine how well we can predict successful college experiences. The models will be cross validated in multiple ways to evaluate their validity.
5. Delivering a set of algorithms that can be built into a web-based system so that college students can see what their college admissions essays reveal about their potential choices of a major.

Data and Methods

Student Admission Essays

During the admissions process, students ($N = 79,330$) composed between 1 and 4 essays ($M = 2.22$ essays per student) -- there was a starting total of 168,874 essays for the entire sample. Each student's essays were written in response to different prompts, such as describing a significant event in one's life or discussing the imagined qualities of an ideal classmate. Essays with fewer than 100 words ($N = 797$) were removed from the sample.

Linguistic analysis of essays. Each essay was analyzed using two broad approaches: a "top-down" and "bottom-up" method (Boyd & Pennebaker, 2015a).

Top-down language analyses. For top-down analyses, we used the LIWC2015 application and dictionary (Pennebaker, Booth, Boyd, & Francis, 2015) to score texts along ~90 psychologically relevant dimensions. Using this method, each essay was quantified by counting the percentage of words that belong to the various categories of language identified within the LIWC2015 dictionary. These dimensions included both function word and content word categories that have established psychometric properties and have been extensively validated in the past (e.g., Tausczik & Pennebaker, 2010). All essays were scored along several additional “descriptive” measures including word count, words per sentence, and punctuation use.

A primary advantage of this approach is that students’ use of function words (e.g., articles, pronouns, conjunctions, etc.) can be easily scored using top-down language analyses. Previous research has found that people who use higher rates of certain categories of function words (e.g., articles, prepositions) tend to perform better in school and have a thinking style similar to an “amateur scientist” -- these individuals often fare well in majors related to STEM fields. Other classes of function words (e.g., pronouns, negations, auxiliary verbs) are associated with “narrative” thinking styles and are socially oriented relative to the “amateur scientist” thinkers -- these individuals often fare well in liberal arts majors, for example.

Bottom-up language analyses. The second approach that we used to quantify student essays is known as a “bottom-up” method, also commonly referred to as “topic modeling” or “topic discovery” methods. A bottom-up approach to language analysis is one that seeks to statistically determine units of meaning in a text based solely on how people naturally use words, particularly in how words occur together. For example, when analyzing student writing, we might find that a subset of articles commonly has the words “senate”, “vote” and “law” being used together at consistently high rates. A different subset of essays may have relatively low rates of those words, but instead consistently use the words “movie”, “actor”, and “tickets” together. Statistically, then, it is possible to sift out different “themes” based solely on what words commonly occur together across all texts. As with other content words measures, themes extracted in this way have been shown to reveal psychologically meaningful dimensions related to phenomena such as health outcomes (e.g., Lowe et al., 2013; Stanton et al., 2015; Wolf et al., 2010) and academic performance (LeFebvre et al., 2015). In short, these themes are often important in understanding *to what* people pay attention and where their thoughts tend to concentrate.

A standard, bottom-up method from the social sciences is known as the “Meaning Extraction Method” (MEM). Using this method, all essays were preprocessed to convert words to their most basic forms (e.g., “running”, “ran”, and “runs” would all be converted to the word “run”). Following this, a list of the most commonly-used words were identified -- words used in at least 0.5% of all essays -- all essays were then scored for either the presence (1) or absence (0) of each word. A data reduction technique known as Principal Components Analysis (PCA) was then used to extract themes from this data that were most prevalent throughout all essays. Examples of these themes are shown in Table 1.

Theme	Example Words
Perception	eye, voice, silence, soft, warm, smell, stomach, darkness
Culture	background, diversity, belief, cultural, religion, custom
Environment	pollution, planet, global, climate, oil, energy, water, natural
Medical	doctor, hospital, diagnose, medical, surgery, treatment, illness, cure
Family	mother, father, parent, child, sister, brother, marry, divorce, family
Sciences	engineering, physics, math, science, chemistry, solve, technology
Poverty	poverty, homeless, shelter, donate, poor, compassion, starve, needy
UT-Austin	university, UT, campus, longhorn, McComb, degree, education, graduate
Religion	christian, church, god, faith, bible, worship, prayer, spiritual, muslim

Table 1. Example themes extracted from student admission essays. Note that the label given to each theme (e.g., “medical”, “religion”, etc.) are simply subjective labels that can be used to refer to word clusters. These labels have no impact on the theme extraction process.

In total, 43 themes were extracted from student essays. These themes ranged from social topics (family, friendship) to leisure activities (hiking, literature) to academic concerns (coursework, education). All essays were then scored using an approach similar to that described above. Each essay was scored along each of the 43 themes by calculating the percentage of words belong to each theme. These scores were then used in a second PCA, resulting in 6 “meta” themes (e.g., themes about “obesity” and “alcohol abuse” commonly co-occurred, resulting in a broader, superordinate theme that reflects health).

Data aggregation and preparation. The above-mentioned linguistic analyses resulted in ~142 total measures of student language for each essay. All measures were merged for each essay, then averaged across essays for each student. This resulted in a single set of “average language use” measures for each student.

Student Performance and Outcome Data

We started with admissions and long-term performance data for 79,330 students ranging from the 2004 through 2014 academic years. This data tracked students from the time of admission through their first 6 years of enrollment at the University of Texas at Austin. Relevant

to the current analyses were admissions-related information, such as each student's preferred majors (e.g., "Biology" and "Archaeology"), SAT/ACT scores, and major assigned during the first year of enrollment. This data also tracked students through their first 6 years of enrollment and included performance measures such as GPA and graduation status at the end of each year.

Using this information, we calculated several "flags" for each student relevant to the current analyses. In this case, we are primarily interested in identifying those students who a) did not receive their desired major in their first year at UT-Austin, b) later obtained a degree from the UT-Austin, and c) performed particularly well among those students who obtained the same degree. Each student was marked to note if they did *not* receive their desired majors during their first year at UT-Austin (total $N = 19,147$). Degrees that were earned by less than 75 students in this subsample were omitted from further analyses for the sake of reliability (final $N = 50$ degrees). Within each type of degree issued by UT-Austin (e.g., Biology, Communication Studies, etc.) students were additionally marked as having "excelled" if their final GPA on record was in the top 25% for students earning that degree. Of these 19,147 students who did not receive their desired major at intake, 3,966 (20.7%) had a final GPA within the top 25% for their degree earned.

Constructing Linguistic Algorithms for Students who Excelled

For this project, we were interested in identifying the language patterns that best differentiated students who fared particularly well within their earned degree after having not received their desired majors at intake. This was achieved using a 2-step approach. The first step consisted of a process known as "feature selection" whereby the combination of language variables most strongly and most consistently associated with excellence within each degree can be determined. Feature selection was performed using a recursive feature elimination approach using random forests paired with 5-fold cross-validation procedures. The use of random forests with cross-validation procedures served to prevent overfitting of the data, thereby providing more realistic estimates of the language variable combinations' ability to differentiate students who excelled from their peers. Feature selection procedures resulted in combinations of language measures that generally performed well in identifying students who excelled within their degree.

The second step of algorithm creation involved the creation of binomial models using the selected language measures within each degree. Binomial models were generated to provide a weighting scheme for the linguistic measures, as well as to further test the accuracy of language combinations using cross-validation procedures – all models were generated using 5-fold repeated cross-validation. The final results of this step were a series of weighted algorithms that can be used to compute the likelihood of a student's excellence in each degree (examples shown below; all $|\text{coefficients}| < .05$ removed; MEM themes are denoted by "Factor_"):

Chemistry Excellence Log Odds = $5.87 + (-0.52*\text{reward}) + (-0.78*\text{negate}) + (-0.26*\text{Factor_14}) + (0.32*\text{home}) + (-0.93*\text{Period}) + (0.21*\text{adverb}) + (-0.44*\text{Factor_39}) + (0.25*\text{Factor_1}) + (-1.37*\text{death}) + (0.27*\text{Factor_32}) + (0.18*\text{Factor_6})$

$$\begin{aligned} \text{Communication Studies Excellence Log Odds} = & -4.56 + (-0.17*\text{auxverb}) + \\ & (0.64*\text{Dash}) + (-0.09*\text{focuspresent}) + (0.23*\text{body}) + (0.95*\text{home}) + (-0.03*\text{Factor_7}) + \\ & (0.11*\text{AllPunc}) + (-0.27*\text{affiliation}) + (-0.17*\text{bio}) + (0.10*i) + (0.08*\text{prep}) + (- \\ & 0.85*\text{friend}) + (0.22*\text{Factor_41}) + (0.87*\text{anx}) + (-0.32*\text{family}) + (0.28*\text{Factor_31}) + \\ & (0.39*\text{power}) + (-0.41*\text{negate}) + (-0.13*\text{Comma}) + (-0.47*\text{Factor_8}) + (0.16*\text{Factor_25}) \\ & + (-0.14*\text{Factor_35}) + (0.29*\text{we}) + (0.08*\text{focusfuture}) + (0.26*\text{Factor_23}) \end{aligned}$$

These scores are then standardized into probabilities with the following formula:

$$\text{Degree Success Probability} = e^{\text{Degree Excellence Log Odds}} / (1 + e^{\text{Degree Excellence Log Odds}})$$

...and thus made to be comparable between degrees. Additionally, these scores should be standardized within-degree to ensure an unbiased ability to rank *relative* likelihoods from highest (most likely to excel in a given degree) to lowest (least likely to excel in a given degree). Cross-validated classification of students who achieved within-degree excellence using these algorithms demonstrated high accuracy (*M* accuracy = 80.0%, *SD* accuracy = 4.9%, range = 69.6-89.6%, *M* balanced accuracy = 71.0%, *SD* balanced accuracy = 11.2%).

Conclusions

We have developed a new, language-based system for identifying those degrees for which each student may be particularly well-suited in terms of both their interests and thinking styles. It is important to reiterate that these scores do not reflect who is *capable* of earning each type of degree issued by UT, and these scores do not account for other factors such as satisfaction with major or post-graduation success. Rather, these scores reflect the likelihood of each student to excel within each type of degree offered given based upon their admission essay language.

Nearly 1 out of 4 students who are accepted to UT do not enter their first year enrolled under their primary choices of major. Historically, these students have had to rely upon a decision-making process for subsequent major selection that is likely non-uniform and potentially biased in nature. The implementation of a new system for recommending potential major choices to students in a personalized, automatic, and statistically sound method will likely help guide incoming students to avenues with increased chances of success. By providing major recommendations to students who find themselves in the disparaging position of having to select an unintended major to pursue, we may improve student satisfaction at UT, retention/graduation rates, and goodness of fit between students and their goals.