Do Engineering Students Learn Ethics From an Ethics Course?

Roman Taraban, William M. Marcy, Mark S. LaCour, Dashiell Pashley, & Klara Keim

Texas Tech University
Lubbock, TX 79409 U.S.A.

E-mail: roman.taraban@ttu.edu

Abstract

The goal of the present research is to develop machine-assisted methods that can assist in the analysis of students' written compositions in ethics courses. As part of this research, we analyzed Social Impact Assessment (SIA) papers submitted by engineering undergraduates in a course on engineering ethics. The SIA papers required students to identify and discuss a contemporary engineering technology (e.g., autonomous tractor trailers) and to explicitly discuss the ethical issues involved in that technology. Here we describe the ability of three machine tools to discriminate differences in the technical compared to ethical portions of the SIA papers. First, using LIWC (Language Inquiry and Word Count) we quantified differences in analytical thinking, expertise and self-confidence, disclosure, and affect, in the technical and ethical portions of the papers. Next, we applied MEH (Meaning Extraction Helper) to examine differences in critical concepts in the technical and ethical portions of the papers. Finally, we used LDA (Latent Dirichlet Allocation) to examine differences in the topics in the technical and ethical portions of the papers. The results of these three tests demonstrate the ability of machine-based tools to discriminate conceptual, affective, and motivational differences in the texts that students compose that relate to engineering technology and to engineering ethics. We discuss the utility and future directions for this research.

1. Introduction

Advances in science and engineering inevitably raise ethical issues. Engineering educators are aware of this. Thus, ethics is a fundamental topic in engineering education and is codified in ABET goals for engineering students. Assessment of student performance in an ethics course demands a qualitatively different approach than that applied in courses like statics or thermodynamics. The latter courses often involve fixed constants, physical principles, and solving equations, and there is often an objective answer against which to judge students' work. In contrast, ethics is more verbal, involving discussion and essay forms of interaction. In ethics courses, students may be required to participate in online discussions, post to blogs, and submit research papers. Students are asked to critically analyze situations and events as well as exercise judgment regarding the attitudes, beliefs, and behaviors of those involved.

The goal of this paper is to present our current exploratory work in developing machine-assisted methods that could aid in the analysis of students' written compositions in ethics courses. We are developing these tools in a sophomore-level course that is offered to engineering majors at our university. This course promotes ethical reasoning through an introduction to ethical theories and contemporary ethical issues in engineering, technology and society. The course materials and assignments cover *intuitionism*, which equates a person's intuitive reaction to ethical issues with ethical values, three ethical theories – i.e., *utilitarianism, respect for persons*, and *virtue ethics* – and the National Society of Professional Engineers (NSPE) Code of Ethics. As part of the course requirements, students analyze and respond to ethical issues in contemporary social settings involving engineering dilemmas. A major course requirement is a Social Impact Analysis (SIA). The organization of the SIA papers is twofold. First, students freely identify and present a contemporary engineering technology (e.g., autonomous tractor trailers, fracking, drones, ethical hacking) in some detail. They then identify and analyze the positive and negative social consequences of that technology. They are required to incorporate knowledge from one or more of the ethical theories into their analyses.

In the present study we tested three machine-based tools (i.e., LIWC 2015, MEH, and LDA) for their ability to analyze differences in ethical versus non-ethical content in students' SIA papers. In the first test, we assessed the ability of machine analysis to identify differences in students' thinking and behavioral dispositions in the ethics versus non-ethics content in students' SIA papers, using LIWC 2015 software. In the second test, we assessed the ability of machine analysis to discriminate the key concepts in ethics versus non-ethics content in the SIA papers, using MEH software. In the third test, we assessed the ability of machine analysis to organize the concepts in the ethics versus non-ethics content in the SIA papers into coherent topics, using LDA software. In the next sections, we provide a brief overview of each of the software programs.

2. Background Literature

2.1 LIWC 2015 (Linguistic Inquiry and Word Count)

http://liwc.wpengine.com/

LIWC 2015 is a word-counting software program that compares word use against specialized dictionaries associated with 90 distinct variables. LIWC, and other related programs like those described below, are based on the assumption that the words that a person uses reveal information about the person's cognitions, motivations, attitudes, and emotions.

LIWC has been tested in a large number of studies and wide range of contexts. As one example, Robinson, Navea, and Ickes [1] found that the content of students' self-introduction essays at the start of the semester predicted their course performance. Specifically, relative word usage in specific categories in the students' brief essays was significantly correlated with course grades. Applying LIWC to students' college admissions essays, Pennebaker and colleagues [2] found that categorical differences in language use predicted students' college grades. Carroll [3] applied LIWC to students' essays in a sophomore-level critical thinking course and found differences in affect in students' essays at the beginning compared to end of semester.

In the present analysis we were primarily interested in four variables that are defined as follows in the LIWC Manual [4]:

- **Analytical Thinking** - A high number reflects formal, logical, and hierarchical thinking; lower numbers reflect more informal, personal, here-and-now, and narrative thinking.
- **Clout -** A high number suggests that the author is speaking from the perspective of high expertise and is confident; low Clout numbers suggest a more tentative, humble, even anxious style.
- **Authentic -** A higher number is associated with a more honest, personal, and disclosing text; lower numbers suggest a more guarded, distanced form of discourse.
- **Tone** (emotion/sentiment) - A high number is associated with a more positive, upbeat style; a low number reveals greater anxiety, sadness, or hostility. A number around 50 suggests either a lack of emotionality or different levels of ambivalence.

Values for these variables are computed by LIWC 2015 in terms of percentiles, based on extensive prior research by Pennebaker and colleagues [4], making these variables especially attractive for small-sample analyses, as in the present study.

2.2 MEH (Meaning Extraction Helper)

https://meh.ryanb.cc/

MEH extracts word and phrase data from text data and calculates n-gram frequencies (e.g., frequencies of single words, 2-word sequences, 3-word sequences). In order to identify and extract the key concepts in a text, MEH deletes function words (e.g., *the, a, in, on*) and pronouns (e.g., *he, she, they*). In order to provide a more general description of the concepts in a text, MEH converts words to lemmas. This allows MEH to identify key concepts in the sample of texts.

There are several published studies in the research literature that apply MEH to qualitative, open-ended data. For example, MEH has been used to identify women's sexual self-schemas [5]. Participants completed open-ended essays regarding sex and sexuality. MEH was able to extract seven reliable themes from the essays: family and development, virginity, abuse, relationship, sexual activity, attraction, and existentialism. Identification of these themes was based on frequently used words across the participants' essays. In a study by Obschonka, Fisch, and Boyd [6], the authors extracted personality profiles by applying MEH to individuals' Twitter and Facebook posts. Both of these studies demonstrated the ability of MEH to carry out qualitative data analysis of open-ended texts.

2.3 LDA (Latent Dirichlet Allocation)

LDA is a software program for statistical text analysis. LDA is based on the assumption that a set of documents have a latent semantic structure that can be statistically inferred from correlations between words, across a sample of documents. LDA uses output from MEH in order to identify topics across a sample of written texts. Technically, a *topic* is a set of words that occurs consistently across a sample of texts in a particular context. LDA was originally developed by Blei, Ng, and Jordan [7] and has generated many computational variations. LDA has not been tested extensively on open-ended text data.

3. Method and Results

The materials for this research consisted of 80 archived papers. One paper was unusable, so the total corpus consisted of 79 papers. The papers were composed by students enrolled in a sophomore-junior level engineering ethics course at a public Research I university and were submitted anonymously for this research, which was approved by the Human Subjects Review Board at the university.

The format of the papers was a Social Impact Assessment (SIA). SIAs are generated in order to review the social effects of infrastructure projects and other development interventions. A rigorous SIA report describes and evaluates the consequences of the projects and suggests ways to mitigate these impacts. Students were instructed to explicitly consider the ethical implications of the engineering technology that they described and evaluated from a technological perspective in the SIA. Therefore, for the purpose of the analyses described next, we separated each paper into the texts in each student's SIA paper that were related to ethics and those that were not related to ethics. These two types of texts will be labeled Ethics and Non-Ethics texts in the remainder of this paper.

3.1 LIWC (Linguistic Inquiry and Word Count) Analysis

Summary results from the LIWC analyses for Ethics and Non-Ethics texts are shown in Figure 1. Each LIWC category was analyzed for differences using paired *t*-tests, which is an appropriate statistical test for related data. There was a small, marginally significant difference for the **Analytic Thinking** category [$t(77) = -1.92$, $p = .058$]. Non-Ethics texts (mean = 87.08) expressed slightly more formal, logical thinking than Ethics texts (mean = 85.61). There was a substantial and significant difference for the **Authentic** category [$t(77) = -6.43$, $p < .001$]. Non-Ethics texts (mean = 23.57) expressed an honest, disclosing narrative, whereas Ethics texts (mean = 15.22) expressed a more guarded narrative. There was a significant difference for the **Clout** category [$t(77) = 5.93$, $p < .001$]. Ethics texts (mean = 58.74) reflected more speaker confidence, whereas Non-Ethics texts (mean = 52.89) expressed a more anxious and tentative sense. There was a significant difference for the **Tone** category [$t(77) = 3.41$, $p = .001$]. Ethics texts (mean = 63.59) communicated a more upbeat positive affect, whereas Non-Ethics texts (mean = 54.01) expressed a more anxious affect.

3.2 MEH (Meaning Extraction Helper) Analysis

In order to identify the highest frequency concepts in the Ethics and Non-Ethics texts, we applied MEH. Table 1 summarizes the results. The Non-Ethics texts contained only a couple of concepts (i.e., *benefit, people*) that may have reflected issues of benefits of technology, and association of issues to people. However, this set of the 20 most frequent words in the Non-Ethics texts does not communicate a strong overall concern with ethics. The concepts from the Ethics portions of students' SIA papers is more strongly associated with ethics, as suggested by high-frequency terms like *people, ethical, benefit, utilitarianism, live, environment, ethics, safety, human, safe*.

3.3 LDA (Latent Dirichlet Allocation) Analysis

MEH identifies key concepts in texts. Of interest, is how concepts in students' work are organized around topics, where a *topic* is a set of words that occurs consistently across a sample of texts in a particular context. LDA was used to identity coherent topics in the students' SIA papers. When LDA was applied to the Non-Ethics content of the SIA papers, five prominent topics corresponded to the topics that students often chose to focus on in their papers: Topic 1: solar energy roadways; Topic 2: green concrete as a building material; Topic 3: artificial intelligence technology; Topic 4: hydraulic fracking; Topic 5: electric vehicles. The ten most frequent terms for each topic are shown in Table 2.

Five representative topics in the Ethics texts in students' SIA papers largely corresponded to the Non-Ethics topics. Topic 1: safe vehicles; Topic 2: benefits and consequences of technology; Topic 3: environmental concerns associated with oil fracking; Topic 4: combination of ethical concerns associated with solar highways and computer hacking; Topic 5: general ethical themes related to public health, safety, the environment, and engineering NSPE code.

4. Discussion

The present analyses confirmed several possibilities related to the utility of machine tools in the assessment of students' SIA papers in an ethics class. The LIWC 2015 analyses showed that ethics portions of students' papers differ significantly from non-ethics portions along several primary cognitive and affective dimensions. The MEH analyses showed that discussions of ethics draw on characteristically different key concepts than non-ethics discussions. Finally, the LDA analyses showed that the organization of topics differs in ethics and non-ethics discussions. Together, these analyses provide promising directions for future work to more closely connect the analyses from these several machine approaches in order to further decipher the content of SIA papers.

For instance, the MEH and LDA analyses that we have presented here show a clear preference among students for *utilitarian* theory or for choosing topics that are more amenable to a *utilitarian* analysis. Words associated with a utilitarian consequentialist approach such as "impact" and "cost" were common. Words associated with *respect for persons* theory (e.g., "privacy", "personal") or *virtue ethics* theory (e.g., "judgment", "vice") do not appear. In order for students to process and internalize the latter two theories, which is a fundamental goal of the course, instructors could urge them to write about topics that are more amenable to the principles and values in those theories. Students in the current sample chose topics amenable to discussion of economic cost-benefit analysis and safety issues—typical *utilitarian*-style analyses. If students were encouraged to discuss some of Apple's iPhone controversial privacy policies, for instance, they may process issues centering around the *respect for persons* theory, such as privacy and personal autonomy.

Analyses like those that we have attempted here could be informative to instructors, in part, by indicating the extent to which students internalized the ethics of the course. That is, the machine classifications could potentially provide useful data regarding the question of which ethical concepts were internalized by individual students, how these were organized into topics, and how these were related to analytical thinking, confidence in the material, and affective reactions, as indicated through an analysis of the words the students used to compose the SIA papers. Those data would give instructors feedback on where changes might be made in the course.

Successfully developing the means of using machine systems to assist in course assessments would allow instructors to provide more extensive and incisive feedback and guidance in ethics courses, by complementing instructors' assessments of students' work. This is a timely issue in any course, like engineering ethics, with high enrollments and that entails substantial student writing and requires considerable instructor time for scoring. We regard exploring these complementary assessment approaches as potentially having a high payoff in engineering ethics education and assessment.

A recent paper titled "Do Ethics Classes Teach Ethics?" [8] raises pertinent questions for engineering ethics education, questions about what is gained and how much change takes place through coursework. Machine-assisted assessments of the sort described here could provide some insight into this question by evaluating student work from early in the course and work produced later in the course [3].

Acknowledgement

This work is supported by the Texas Tech University Office of Global Communication.

References

[1] Robinson, R. L., Navea, R., & Ickes, W., Predicting final course performance from students' written self-introductions: A LIWC analysis. *Journal of Language and Social Psychology, 32*(4), pp. 469 – 479 (2013)

[2] Pennebaker, J. W., Chung, C. K., Frazee, J., Lavergne, G. M., & Beaver, D. I., When small words foretell academic success: The case of college admissions essays. *PLoS ONE, 9*(12), e115844 (2014)

[3] Carroll, D. W., Patterns of student writing in a critical thinking course: A quantitative analysis. *Assessing Writing, 12*, pp. 213–227 (2007)

[4] Pennebaker, J.W., Boyd, R.L., Jordan, K., & Blackburn, K., *The development and psychometric properties of LIWC 2015*. Austin, TX: University of Texas at Austin (2015)

[5] Stanton, A. M., Boyd, R. L., Pulverman, C. S., & Meston, C. M., Determining women's sexual self-schemas through advanced computerized text analysis. *Child Abuse & Neglect, 46*, pp. 78-88 (2015)

[6] Obschonka, M., Fisch, C., & Boyd, R., Using digital footprints in entrepreneurship research: A Twitter-based personality analysis of superstar entrepreneurs and managers. *Journal of Business Venturing Insights, 8*, pp 13-23 (2017)

[7] Blei, D. M., Ng, A. Y., Jordan, M. I., Latent Dirichlet Allocation. *Journal of Machine Learning Research, 3*(4-5), pp. 993–1022 (2003)

[8] Curzer, H., Sattler, S., DuPree, D. G., & Smith-Genthos, K., R., Do ethics classes teach ethics? *Theory and Research in Education, 12*(3), pp. 366-382 (2014)

Figure 1. LIWC Comparison of Non-Ethics and Ethics Texts

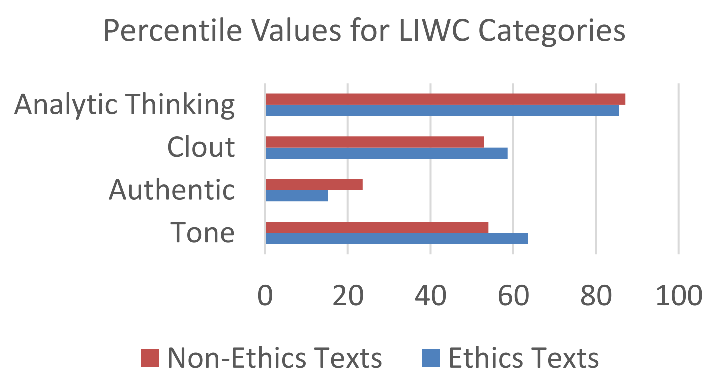

Table 1. Total Frequency-of-Mention of Most-Frequent Concepts in Ethics and Non-Ethics Portions of Students' SIA Papers

Ethics Concepts	Total Frequency		**Non-Ethics Concepts**	Total Frequency
people	412		company	487
ethical	369		water	400
water	270		technology	386
benefit	246		solar	375
utilitarianism	215		energy	334
technology	201		oil	283
live	181		cost	272
great	176		people	256
company	172		vehicle	246
engineer	172		benefit	244
public	169		time	221
environment	151		frack	201
ethic	139		car	197
safety	134		project	191
solar	128		concrete	189
human	124		panel	186
frack	123		system	181
amount	123		road	180
safe	121		produce	176
health	110		construction	173

Table 2. Ten Most Frequent Terms for Five Topics in Non-Ethics Content in Students' SIA Papers

Non-Ethics Topics				
Topic 1	Topic 2	Topic 3	Topic 4	Topic 5
solar	concrete	technology	water	vehicle
energy	construction	human	oil	car
panel	company	benefit	frack	engine
roadway	project	intelligence	gas	technology
road	building	artificial	drill	company
power	cost	artificial intelligence	company	electric
source	people	system	fracture	people
solar roadway	drone	robot	hydraulic	drive
solar panel	product	problem	hydraulic fracture	fuel
cost	material	government	industry	time

Table 3. Ten Most Frequent Terms for Five Topics in Ethics Content in Students' SIA Papers

Ethics Topics				
Topic 1	Topic 2	Topic 3	Topic 4	Topic 5
people	technology	water	ethical	engineer
live	ethical	frack	solar	people
human	consequence	oil	energy	ethic
utilitarianism	person	drill	roadway	public
car	future	industry	company	environment
amount	benefit	environmental	hack	health
vehicle	approach	gas	benefit	safety
life	cost	public	power	great
benefit	issue	impact	hacker	code
autonomous	live	environment	people	area