Developing Machine-Assisted Analysis of Engineering Students' Ethics Course Assignments

Dr. Roman Taraban, Texas Tech University

Roman Taraban is Professor in the Department of Psychological Sciences at Texas Tech University. He received his Ph.D. in cognitive psychology from Carnegie Mellon University. His interests are in how undergraduate students learn, and especially, in critical thinking and how students draw meaningful connections in traditional college content materials.

Mr. Mark Stephen LaCour Jr., Texas Tech University

Mark is a doctoral student in the Department of Psychological Sciences at Texas Tech University. He completed his thesis on the role of calculator usage on mathematical cognition at the University of Louisiana at Lafayette.

Dr. William M. Marcy P.E., Texas Tech University

Professor and Director of the Murdough Center for Engineering Professionalism and The National Institute for Engineering Ethics Texas Tech University Lubbock, Texas

Mr. Richard A. Burgess II, Texas Tech University

Richard Burgess currently works as an Instructor in the Murdough Center for Engineering Professionalism (MCEP) and National Institute for Engineering Ethics (NIEE) in the Whitacre College of Engineering at Texas Tech University. He oversees the day to day operations of the Center's distance learning courses for both engineering students and practicing engineers. Additionally, he teaches an on-campus ethics course for undergraduate students. Burgess provides guest lectures on ethics throughout the Whitacre College of Engineering. Burgess has also worked to incorporate ethics into K-12 STEM education. The push to increase the number of students pursuing STEM careers needs to be accompanied by a sophisticated understanding of the complexity of technology. Ethics is a key part of this complexity and the next generation of STEM professionals will need the skills to effectively engage the ethical challenges they will face. Burgess is a regular presenter on incorporating ethics in a K-12 setting. A theme throughout these roles is the importance of teaching ethics and promoting ethical reflection in a way that is both accessible and substantive. This is a challenge that Burgess is keenly interested in. He holds bachelor's and master's degrees in Philosophy and is currently a PhD student in Systems and Engineering Management program in the Texas Tech Industrial Engineering Department.

Developing Machine-Assisted Analysis of Engineering Students' Ethics Course Assignments

Abstract

Our research concerns engineering ethics education. We were drawn to this topic by a recent paper titled "Do Ethics Classes Teach Ethics?", but more so by ABET criteria 3f and 3h regarding the development of ethical responsibility in engineering students. The purpose of the present project is to use the learning and analytical capabilities of IBM Watson Natural Language Classifier to analyze capstone papers submitted by undergraduates in a course on engineering ethics. The capstone papers that we analyzed required students to identify and discuss a contemporary engineering technology (e.g., autonomous tractor trailers) and to explicitly discuss the ethical issues involved. In the two tests described here we assessed whether Watson-NLC could classify sentences from students' papers as either related to ethics or not related to ethics. Additionally, we consider the utility of these simple machine-based classifications. Our longer-term goals are to use Watson-NLC to identify the ethical theory or theories from the course that students adopt to frame their ethical positions, to assess the effectiveness of students' ethical arguments, and to assess changes in ethical thinking across the semester.

Introduction

Advances in science and engineering inevitably raise ethical issues. For instance, should we build an offshore oil platform if it would reduce dependence on foreign oil but is a threat to the ecosystem? What new hazards for people do self-driving cars and trucks create, and who should assume liability if the technology fails? Should we promote computer hacking if it is subversive but increases national security? In our work with students in an undergraduate engineering ethics course, we have found that ethical issues and questions like these deeply engage students. It would be beneficial to students and consistent with the mission of engineering education to develop this enthusiasm into rigorous and informed ethical reflection. In this project we explore how an intelligent machine, specifically, IBM Watson Natural Language Classifier (Watson-NLC), can assist in this work.

Ethics is a fundamental topic in engineering education [1] that is consistent with ABET goals for engineering students: *3f. An understanding of professional and ethical responsibility; 3h. The broad education necessary to understand the impact of engineering solutions in a global, economic, environmental, and societal context.* Teaching ethics in engineering is obligatorily different than teaching a technical engineering topic, like statics. Whereas the latter involves fixed constants, physical principles, and solving equations, ethics is more verbal, involving discussion and essay forms of interaction. Students may be required to participate in online discussions, post to blogs, and submit research papers. With large class sizes, instructor workload could become substantial. The pedagogical goal of the present project is to develop the means to better prepare engineering students for professional careers in which they exercise ethical judgement in their technical work and are cognizant of the impact of their decisions on individuals, the community, and beyond. The research goal is to develop ways in which machine-assisted analysis of students' written compositions in ethics courses can serve as an

instructional resource.

The context of our research is a sophomore-level course that is offered to engineering majors at our university. This course develops ethical reasoning through an introduction to ethical theories and contemporary ethical issues in engineering, technology and society. Course materials and assignments consider *intuitionism*, which is a person's intuitive reaction to ethical issues, three ethical theories – i.e., *utilitarianism*, *respect for persons*, and *virtue ethics* – and the National Society of Professional Engineers Code of Ethics. Course activities require students to analyze and respond to ethical issues in contemporary social settings involving engineering dilemmas. A major course requirement is a capstone paper incorporating Social Impact Analysis (SIA). The general purpose of SIA is to identify and analyze the positive and negative social consequences of engineering plans and projects. In students' SIA papers, they identify and discuss a contemporary engineering technology (e.g., autonomous tractor trailers, fracking, drones, ethical hacking). They are required to incorporate knowledge from one or more of the ethical theories into their analyses.

The goal of the present study was to use machine-learning to identify the ethical content in the capstone papers submitted by students in the ethics course. In the two tests described in this paper, we assessed whether Watson-NLC could classify sentences from students' papers as either related to ethics or not related to ethics. Further, we consider the utility of these simple machine-based classifications. Our longer-term goals are to use Watson-NLC to identify the ethical theory or theories from the course that students adopt to frame their ethical positions, to assess the effectiveness of students' ethical arguments, and to assess changes in ethical thinking across the semester.

Machine Learning and Text Analysis

Machine analysis of text is not a new idea. An influential model related to the analysis of the semantic content in texts is Latent Semantic Analysis (LSA) http://lsa.colorado.edu/. LSA uses a sparse matrix of correlations between words and documents in order to represent the interrelationships of concepts within a high-dimensional semantic space. These correlations reliably correlate with human judgments of similarity, or relatedness, of meaning. One limitation of LSA is its approach to structured language – e.g., a sentence – as a "bag of words." Thus the model is constrained by its inability to capture the logical, syntactic structure of language, which relates to meaning. For instance, in the sentence, "The dog chased the cat," the syntactic structure of the sentence will determine who is doing the chasing and who is being chased.

Another software application, called CohMetrix http://cohmetrix.com/ , provides measures of the properties of a text. These properties include the coherence of text, its readability level, and its syntactic complexity. CohMetrix provides measures of the coherence of mental representations that are based on a text, but not the semantic content of those representations.

IBM Watson – Natural Language Classifier

The machine-learning system that we are testing is the IBM Watson Natural Language Classifier (Watson-NLC). IBM has developed a suite of intelligent tools within the IBM Watson program.

One of these, the Natural Language Classifier, is based on deep learning, which currently applies the most powerful learning algorithms for intelligent machines http://www.ibm.com/watson/developercloud/nl-classifier.html . The Natural Language Classifier can be trained to classify information in any domain. Examples from the IBM website include:

- Classify SMS texts as personal, work, or promotional
- Classify tweets into a set of classes, such as events, news, or opinions
- Tackle common questions from your users that are typically handled by a live agent
- Trigger actions in an application, such as start another application, respond with an answer, or begin a dialog

In our application, Watson-NLC learns how words and phrases from students' essays relate to specific classifications. With sufficient training Watson-NLC can readily classify new instances.

Training Watson-NLC is straightforward. The input file to Watson-NLC is a .csv (comma-separated values) file consisting of two columns and multiple rows. Each row constitutes a training instance. The first column in any given row contains an instance of a classification. The second column in a row contains the classification for the instance. The classifications for each training instance are determined by human raters. An example of a .csv file for training Watson-NLC is shown in Table 1.

Table 1. Portion of a CSV Input File to Watson-NLC

Example Input Sentences From Student Paper	Classifications
The fact that the United States was able to have large economic gains due to fracking made this a very viable route to phasing out coal power which is still very widely in use throughout the world.	Not Ethics
It is seen as a potential way to continue the growth in economy and electricity generation to eventually switch over to renewable resources, but for the time being has proven to not only be economically viable, but allowed the United States to gain influence over other countries who continue to develop and need new energy sources.	Not Ethics
Despite large benefits there have also been drawbacks.	Not Ethics
Possibly the largest concern about the use of fracking would be the effect it has on our fresh water supply; both contaminating it with chemicals creating more toxic water than water treatment plants can handle, and contaminating local water supplies.	Ethics
Since water is the main fluid being used in the process, this totals up to very large amounts of water being rendered unusable for drinking or farming.	Ethics

The .csv training file is submitted to Watson-NLC, which builds a classifier based on the instances provided to it. Figure 1 gives a sense of the Watson-NLC user interface.

After the classifier is built, it can be used to classify new instances. The output that Watson-NLC provides for old and new instances looks very much like a .csv input file. The difference is that the output includes an additional column, which shows Watson's classifications for each

Figure 1. The Watson-NLC User Interface.

instance and confidence ratings that the instances fit the classifications. When Watson-NLC is trained to classify instances into two possible categories, *Ethics* and *Not Ethics*, as in the present study, the classifier will report its confidence that any given instance fits the classification it assigns. The sum of confidence ratings for classifying an instance is always 1.00.

Figure 2 shows a sample of test outcomes. Column B shows the experimenters' original assignments and Column C shows the Watson-NLC assignments and confidence in those assignments.

Figure 2. Portion of a Watson-NLC Output File

	A	B	C
66	The chance of disasters are much less than risks people take daily	Ethics	Not Ethics (0.52)
67	When one feels as they are doing wrong, yet continues to do so, is when they should question their ethics, not just because potential risks	Ethics	Ethics (0.99)
68	With risk factor taken into account, companies go to extreme measures to avoid all potential incidents	Not Ethics	Not Ethics (0.97)
69	With the need for oil remaining obvious, deep sea drilling is our best alternative to shift away from drilling on the main land	Not Ethics	Not Ethics (0.86)
70	This frees up property once used for drilling sites and gives access/privacy back to original owners	Ethics	Ethics (0.94)
71	While the concern for aquatic animalâ□□s safety is serious, no people are directly harmed by this process, as they are with land drilling	Ethics	Ethics (0.99)

Tests of Learning in Watson-NLC

In this section we report two tests in which we trained Watson-NLC to classify ethical elements in students' SIA papers. In these tests, we trained Watson-NLC to simply classify the sentences in students' SIA papers into two categories: *Ethics and Not Ethics*, as suggested by Figure 2. The goal of these tests was to assess whether Watson-NLC could make these classifications, and whether Watson's agreement with human raters would improve with additional training.

For the first test, we chose ten SIA capstone papers at

random from the corpus of archived course papers. One of the researchers classified each sentence in the papers as *ethics* or *not-ethics*. A second researcher verified these classifications. In cases of disagreement between researchers, a final classification was determined through discussion and consensus. The final classifications were used to build a classifier with two categories: *ethics* and *not-ethics*. Ten new SIA papers were then used to test the classifier.

The classification results for ten new papers that Watson-NLC had not seen previously are shown in Table 2. Across the ten new papers, the average agreement of Watson-NLC with the human raters was 74.60%. Watson-NLC is clearly performing better than chance, but not in perfect agreement with the human classifiers. Given that the classifier was trained on only ten example papers, the results appeared quite promising.

Table 2. Watson-NLC Classifications of Sentences in Ten New SIA Capstone Papers into *Ethics* and *Not-Ethics* Categories After Training on Ten SIA Papers

	Sentences Classified			
Student Paper	Sentences Agreed	Sentences Disagreed	Total Sentences	Percent Agreement
S11	97	24	121	80.17
S12	102	16	118	86.44
S13	42	10	52	80.77
S14	80	27	107	74.77
S15	43	23	66	65.15
S16	66	23	89	74.16
S17	64	35	99	64.65
S18	58	27	85	68.24
S19	59	19	78	75.64
S20	73	23	96	76.04
Average	68.40	22.70	91.10	74.60
SD	20.18	6.75	21.93	7.00

In order to test if the level of agreement between Watson-NLC and the human raters could be improved, an additional 30 papers were presented to Watson-NLC for training. After building a classifier from these 40 papers, Watson-NLC was tested on ten new papers. Average classification agreement of Watson-NLC with human raters increased to 79.30%. Detailed results are shown in Table 3.

Table 3. Watson-NLC Classifications of Sentences in Ten New SIA Capstone Papers into *Ethics* and *Not-Ethics* Categories After Training on Forty SIA Papers

	Sentences Classified			
Student Paper	Sentences Agreed	Sentences Disagreed	Total Sentences	Percent Agreement
s60	63	11	74	85.14
s61	103	20	123	83.74
s62	66	17	83	79.52

s63	96	16	112	85.71
s64	67	26	93	72.04
s65	75	34	109	68.81
s66	90	22	112	80.36
s67	60	22	82	73.17
s68	51	12	63	80.95
s69	98	18	116	84.48
s70	69	19	88	78.41
Average	76.18	19.73	95.91	79.30
SD	17.58	6.44	19.55	5.72

The trend of increased agreement with additional training is depicted in Figure 3. As Watson-NLC gains more experience with the classification task, agreement with human raters increases and disagreement decreases.

Figure 3. Watson-NLC Percent Agreement and Disagreement with Human Classifiers, by Training Input

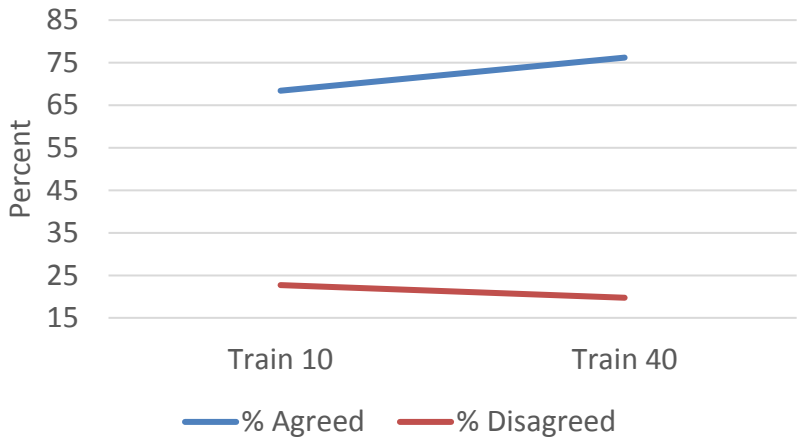

Potential Applications

These early results suggest several classroom applications. Figure 4 shows one student's complete paper that has been reduced in size to illustrate two properties. Sentences related to ethics are highlighted in green. A visual scan suggests only scant consideration of ethics in this paper, and a scattered distribution of ideas related to ethics.

If the output from Watson-NLC could be converted into a visual representation similar to that shown in the figure, it would have utility for the student and instructor as a rough indicator of whether the student fulfilled the SIA paper requirement to provide substantial discussion of ethical issues. Further, the proportion of the paper dealing with ethics could be easily calculated. Either form of output would provide the instructor with leverage in prompting students to further develop a position regarding the ethical issues associated with the engineering topic of the paper.

We do not want to attach more importance to this suggestion than it deserves, but rather to indicate that the machine resource has instructional utility even at this simple level of analysis.

Figure 4. Complete Student Paper with Ethics Text Highlighted in Green

Solar Panels, our clean and safe option for the future.

The sun is a natural and free source of energy. Also, it is silent and pollution free. Solar energy has a lot of potential and many advantages that with time, it might replace other resources of energy such as fossil fuels. Accordingly, to National Geographic, every hour of energy that the sun beams onto earth can give enough energy for the planet need for an entire year. Nowadays, solar energy technology produces less than 1% of the world energy, but it's implementation is growing exponentially.

Solar panels have been around us for quite a while and we might not even notice. For instance, many scientific calculators, stop lights, watches, spacecraft, rooftops, and more everyday electronic devices contain this technology already. However, these devices have very low efficient rates. Which means that they are not very effective in terms of absorbing solar energy. In addition, there are solar panels that are up to 48% efficient compared to the common solar panels which have an average of 18% efficiency. These efficient solar panels are expensive and not convenient to generate profit.

Accordingly, to energy matters, the person that discovered the photovoltaic effect is Alexandre Edmond Bacquerel, he claimed that electricity could be produced out of the sun energy that radiates into or planet in 1839. Moreover, it was too early for the implementation of photovoltaic cells at tat time and the levels of efficiency where incredibly low. Consequently, more than 100 years later, using the same principle that Bacquerel had discovered, but with the technology advancements at the time, he invented the solar cell. The basics of how a photovoltaic cell works consists of light striking certain element or compounds, which will cause the material and its surface to emit more electrons. The combination of these reactions make electrons flow possible and therefore allow electricity to happen.

Solar energy implemented in solar panels have advantages and disadvantages like any other form of energy as shown below.

Benefits of Solar energy:

Sustainable
Flexible
Environmentally friendly
Fast improvement
"Noise free"
Decrease of cost
Increase of demand
Low maintenance

Drawbacks of solar energy:
High cost to invest
Little pollution to manufacture
Needs large space (land)
Less efficient at high altitudes
Dependent on the weather

Based on the advantages and disadvantages list of solar panels, it seems that most of the problems associated with the manufacture of solar panels are of little importance compared to other forms of energy. Accordingly, to Professor Michael E. Wysession, Associate professor of Earth & Planetary Sciences at the Washington University in St. Louis mentions that the decrease in cost of solar panels decreases almost twice every year since 2010 and that it will continue within the next upcoming years. The cost to install solar panels in rooftops is going to be immediately profitable. On the other side, electric companies are not going to agree with their industry to be decreasing at that speed and the increase in the companies' prices for electricity consumption for those who do not have solar panels and still use conventional electricity. Nevertheless, this case will be applicable just for places in the world were direct sunlight is very prominent such as Texas, Arizona, Peru, Mexico, Australia, Cameroon, and Jordan. Jordan seem to be one of the most convenient places to start growing as a solar energy dependent government because Jordan has about 330 sunny days during the year. Also, the versatility of this energy source is broad. Places like the International Space Station (ISS), satellites use these high efficiently panels as well as on earth we use them for simple things like calculators, to save the usage of batteries. As technology advances, so rapidly, there is many unimaginable applications of these technologies being used today in the manner that backpacks have solar panels implemented to recharge phones, clear windows in buildings with the capacity to absorb and store energy for the rest of the day.

On a larger scale, solar energy is not only about Photovoltaic Panels (Solar PV), but there are solar energy thermal plants to collect the sun's energy as a heat source. United States, Japan, & Germany are some of the countries that implement this technology with more emphasis. However, Places like Peru and Mexico have been interested in acquiring these technologies from the Italian company Enel Green Power (EGP), which has offered them a price of approximately $48 MWh. Antonio Cammisecra, EGP's head of business, makes emphasis in the advances in solar energy that the prices are a clear ten of falling prices. Another interesting fact is when the summer comes. Normally the light bill price increases considerably because of the consumption of energy is larger. However, with solar panels, the summer makes you have more energy, which will be a perfect combination for the hot summer. Also, if that person has extra electricity they can sell their extra saved energy to someone who is going to be willing to do so.

Moreover, Tesla motors has associated with SolarCity to create a diversified energy conglomerate. Tesla motors is a company that sells electric cars, and now their plan is to start developing cars with efficient solar panels that will completely be absent of the usage of fossil fuels.

Countries developing Solar Energy

Germany was the leader producing solar panels until other countries like China, India, and the United States gain interest on developing these technologies. Moreover, in 2015, China became the largest producer of solar energy, and is currently making the biggest plant in the Gobi dessert. The amount of solar energy that China plans to produce in several yeas surpasses three times by what the United States has in mind.

Another country with ambition when it comes to producing solar enrgy is India. By 2022, India plans to produce 12.5% of their general energy needs by using solar energy. This will benefit considerably the economy of the country.

In addition, the state of Hawaii as of 2015, has a total of 12% of houses with houses with mounted solar panels.

Lastly, as mentioned earlier, Jordan is a country that has a large amount of solar activity throughout the entire year. Jordan operates a 160 Megawatt solar park, which plans to reduce its dependence in countries such as Egypt because of their dependence on their fossil fuel.

Types of Solar Energy

Most solar panels are made of silicon, which is the third most common element on earth, behind Iron and Oxygen. This panels made of silicon are divided into two categories, monocrystalline and polycrystalline.

Monocrystalline solar panels are traditionally more expensive and they are usually long silicon cylinders.

Polycrystalline panels have usually a blue color and are formed of molten silicon. This panels are more efficient and more common nowadays.

Micro Inverters are not a type of solar panels, but an implementation that made the solar panels more efficient. These inverters transform Direct current (DC) into Alternate Current (AC). Before micro inverters, if there was any disturbance in the solar panels, their efficiency will drop tremendously. But with the implementation of microinverters those problems are no longer existent since micro inverters make the solar panel work even if there are branches around it or dust on a part of the panels without affecting the whole area of the panel but the area being cover only.

Organic Photovoltaic solar panels are made of plastics or a different arrange of compositions. One of the main advantages of the organic solar photovoltaic panels is that they are also flexible and at a low price.

Another way to collect energy from the sun is called concentrated solar power (CSP) and it focuses the heat brought by the sun to boil the water in the equipment to run a steam powered generator. Solar troughs and solar towers are the two common Concentrated Solar power generators. Solar toughs are cylindrical throughs that used their energy at a tube in the middle that contains a liquid that warms up to heat water as a heat exchanger. The other technology used (Solar Towers) are large towers with mirrors around which heat the water around it to boil it. Unfortunately, solar towers are still very expensive. In addition, concentrated solar power (CSP) raise their temperatures considerably and in response to that, there has been found thousands of birds' dead within the area, and that makes a lot of people in the industry to disregard the usage of CSP and opt for photovoltaic panels. However, there is still hope for CSP once that those issues are being taking care of, since they are way more efficient than photovoltaic panels.

Another form of solar energy is the Stirling engine. And has an average efficiency of 30% to 35%. They run on a thermodynamic cycle of air, instead of water, it is made of a parabolic dish that causes the air inside to make work a piston inside the machine to produce heat and save energy.

Consequently, we can implement the idea of creating a solar panel that will contain water inside and that will desalinate the water with the solar energy. It would save the consumption of fresh water and to be prepared in case we run out of fresh water and have a backup plan.

In addition, some of the largest companies in the solar panel industry, offer financing, leasing, purchase options, savings, etc. Sunrun, Vivint solar, NRG Home Solar, XOOM Solar, and SunPower are just some of them. Another advantage is that if there is energy storage extra on your inverters, you would be able to sell it to neighbors that do not have solar panels but some electric companies are trying to charge customers for selling their solar energy, which in my perspective is unfair since these energies was obtained from the sun and the cost to make it was free. These is a principle of utilitarianism, since people is making their own energy by paying for solar panels, it is fair for them to the what they please with the energy that they store on their houses. However, electric companies try to find excuses to charge people for them to maintain a good business. In my perspective, light companies should get more involved in the development of solar panels as an option so that their business evolves just how the planet evolves also. There are many factors that must be taken in consideration when it comes to installing solar panels in rooftops such as roof size, placement of panels, sun exposure, energy needs, and cost vs incentives. In addition, when it comes to choosing a company to install solar panels, some of the things to take in consideration is the Leases, tax credits, warranties, monitoring, flexibility, life expectancy, and of course price of installation and equipment in general. Another thing to take in consideration is PPa's vs leases. PPa or Power purchase agreements, are affordable and easy to install. You would be paying only for the power generated not the equipment being used. However, many companies offer free installations. On the other side, leases are essentially equipment rentals, just like car rentals. Also, the customer will be paying a a fixed monthly pay rate.

In conclusion, I believe that solar energy is the future in humanity for energy consumption. The major disadvantage is that its cost is still up, but it has been decreasing considerably. Also, one of the main points of this type of energy is that it is environmentally friendly which is the contrary for other sources of energy. Moreover, as the efficiency of the solar panel increases, the required land needed to absorb energy is reduced. However, one of the few disadvantages that I see is that this technology is recommended in places with low latitudes and plenty of sunlight activity. Also, there needs to be plenty of land which is a problem in places like Europe, but not in places like Texas, Australia, Jordan, et

References

Wysession, Michael E. "Solar Power and Electricity | The Great Courses Plus." *The Great Courses Plus*. N.p., n.d. Web. 01 Nov. 2016.
"Follow the Sun." *The Economist*. The Economist Newspaper, 16 Apr. 2016. Web. 18 Nov. 2016.
Society, National Geographic. "Solar Power Energy Information, Solar Power Energy Facts - National Geographic." *National Geographic*. N.p., n.d. Web. 18 Nov. 2016.
"5 Best Solar Panel Companies for the Residential Homeowner." *Solar Power Authority*. N.p., n.d. Web. 18 Nov. 2016.

A second way in which Watson-NLC could be valuable to the student and instructor is in terms of the organization of ideas. Rather than scatter ideas related to ethical issues willy-nilly as they occur to the student, the instructor may prefer a more cohesive format in which the student

systematically first presents the technical aspects of the engineering topic and then brings in and develops the ethical concerns. Watson-NLC could be used to automatically provide a sense of organization and feedback to the student, based on the distribution of ethical and non-ethical content.

Future Directions

Our current classifier can be thought of as a Stage 1 classifier that separates ethics-related content from technical descriptions. Going forward, we intend to develop a Stage 2 classifier that will classify the ethics-related statements in the SIA papers into the several theories that are covered in the course – i.e., *utilitarianism, respect for persons,* and *virtue ethics* - as well as the NSPE code and *intuitionism*, which is the view that ethical issues and moral judgments are immediately apprehended. These classifications would be informative to instructors, in part, by indicating the extent to which students internalized the ethics of the course. Specifically, these machine classifications would provide useful data regarding the question of which ethical approaches were internalized. These data would give instructors feedback on where changes might be made in the course.

The current available data include only the capstone SIA papers. However, in future work we hope to also assess changes in ethical thinking from early to late in the course through the analysis of early and late student compositions. It is possible that the gains are small. Such an occurrence would be consistent with Haidt's [3] pessimistic view that so-called moral reasoning is really an epiphenomenon—they are simply one's *post hoc* rationalizations for what are ultimately moral intuitions—i.e. instances of ethical *intuitionism*. More optimistically, assessments of changes in ethical reasoning could help identify students who had not adequately internalized the course content and who could benefit from additional development.

Conclusions

The theoretical question of man vs machine has some merit and interest. More importantly, though, 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 their assessments of students' work. This is a timely issue in any course, like engineering ethics, with high enrollments, that entails substantial student writing, and that 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.

There is clearly a cost-benefit issue in attempting to train machine classifiers to carry out assessments in specialized courses. The difficulty of assessing ethical content in SIA papers is compounded by the fact that the course and papers involve multiple ethical theories and ethical codes and approaches. Clearly, building a classifier for a course that is offered only infrequently and that has low enrollments would produce a high cost to benefit ratio.

A recent paper titled "Do Ethics Classes Teach Ethics?" [4] 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.

References

1. Loui, M. C. (2006). Assessment of an engineering ethics video: Incident at Morales. *Journal of Engineering Education, 95*(1), 85-91.

2. Landauer, T. K. & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of the acquisition, induction, and representation of knowledge. *Psychological Review, 104,* 211-140.

3. Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review, 108*(4), 814-834.

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