

Towards Answering Open-ended Ethical Quandary Questions

Yejin Bang,¹ Nayeon Lee,¹ Tiezheng Yu,¹ Leila Khalatbari,¹ Yan Xu,¹
Samuel Cahyawijaya,¹ Dan Su,¹ Bryan Wilie,¹ Romain Barraud,¹ Elham J. Barezi,^{1*}
Andrea Madotto,^{1*} Hayden Kee,² Pascale Fung¹

¹Center for Artificial Intelligence Research (CAiRE), The Hong Kong University of Science and Technology

²Philosophy Department, The Chinese University of Hong Kong

{yjbang, nayeon.lee, tyuah, lkhalatbari, yxucb, scahyawijaya, dsu, bwilie, rmbarraud, ejs, amadotto}@connect.ust.hk,
hkee@cuhk.edu.hk, pascale@ece.ust.hk

Abstract

Considerable advancements have been made in various NLP tasks based on the impressive power of large language models (LLMs) and many NLP applications are deployed in our daily lives. In this work, we challenge the capability of LLMs with the new task of ETHICAL QUANDARY GENERATIVE QUESTION ANSWERING. Ethical quandary questions are more challenging to address because multiple conflicting answers may exist to a single quandary. We explore the current capability of LLMs in providing an answer with a deliberative exchange of different perspectives to an ethical quandary, in the approach of Socratic philosophy, instead of providing a closed answer like an oracle. We propose a model that searches for different ethical principles applicable to the ethical quandary and generates an answer conditioned on the chosen principles through prompt-based few-shot learning. We also discuss the remaining challenges and ethical issues involved in this task and suggest the direction toward developing responsible NLP systems by incorporating human values explicitly.

Introduction

Artificial Intelligence (AI) has shown significant breakthroughs in recent years with impressive results. Specifically, in the field of Natural Language Processing (NLP), the advancement of large pre-trained language models (LLMs) fostered AI systems to achieve impressive results approaching human-level in various tasks (Radford et al. 2019; Raffel et al. 2020; Brown et al. 2020), including open-book questions answering (Tafjord and Clark 2021; Gu, Mishra, and Clark 2021; Jiang et al. 2021; Hendrycks et al. 2020). Various NLP systems have been developed for the benefit of our society such as COVID-19 scholarly information search engine (Su et al. 2020) or conversational systems for automated counselling (Welch et al. 2020; Ahn et al. 2020). Along with efforts to ensure positive uses of such systems, associated ethical concerns and safety issues have also actively been discussed (Talat et al. 2021; Ji et al. 2022; Roller et al. 2021).

One of the most common use cases of NLP technique is question answering (QA) systems. While much literature has

focused on QA problems which can be verified with standard answers based on factoid knowledge (Roberts, Raffel, and Shazeer 2020; Brown et al. 2020; Wang, Liu, and Zhang 2021), human questions are also open-ended by which cannot be answered with yes/no or static answer (e.g., essay questions, survey questions, etc.). It is relatively straightforward what to be achieved for closed-book/knowledge-based/open-domain QA as the gold answer can be limited to a definite set. In contrast, many questions in the real world are complicated and open-ended. There is usually no single definite answer to such open-ended questions while the current QA systems are not addressing these cases. Ethical quandary questions can be viewed as one of the most challenging forms of open-ended questions. A discussion with multiple perspectives (i.e., a manner of debate) is crucial (Talat et al. 2021; Hendrycks et al. 2020). In this paper, we open the discussion of responsible AI systems for open-ended questions by challenging the capability of AI to provide relevant and nuanced answers to ethical quandary questions in the style of a human ethicist — ETHICAL QUANDARY GENERATIVE QUESTION ANSWERING (GQA).

A famous ethical dilemma trolley problem (Thomson 1976): “*Should we kill one person to save five people in danger of being hit by a trolley?*” does not have definite answer between “Yes” or “No”. Currently, the most relevant AI system to such open-ended question is Delphi (Jiang et al. 2021), a model that learns to reproduce human moral and ethical judgments. Delphi answers “No” to the question – only giving a prophetic, closed answer to the questions posed to it. However, there can be multiple perspectives on this problem depending on the underlying ethical principle. From the deontological perspective, the answer would be “No” because killing is never acceptable. From the utilitarian perspective, by contrast, the answer would be “Yes” because the principle dictates that the most appropriate action is the one that results in the greatest good for the greatest number of people. As Talat et al. (2021) criticized, one-sided normative ethical judgment answer makes it cannot represent incommensurable and diverse ethical judgments. In this paper, we explore the direction of building an AI interlocutor with which we think through the ethical issues instead of handing over our ethical responsibility to the AI system by seeking a definite answer.

We envision ETHICAL QUANDARY GQA task to bring

*This work was done when the authors were studying at The Hong Kong University of Science and Technology.
Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

potential social benefit as an effort of achieving an AI system that can enhance humans’ moral decision-making through the deliberative exchange of different perspectives to an ethical quandary, which is in the approach of Socratic philosophy. As suggested by moral philosophers, AI systems can be used to aid humans in having reflective equilibrium by suggesting different aspects that individuals could not take into consideration due to personal biases and prejudices, which ultimately prompt to one’s own further reflection (Savulescu and Maslen 2015; Lara and Deckers 2020; Giubilini and Savulescu 2018).

We explore ETHICAL QUANDARY GQA with an initial experiment in an in-context prompt-learning setting on an LLM model, to generate answers with different ethical principles. We disentangle the task into two steps 1) principle provider, where selecting relevant ethical principles to be based on for answer 2) principle-grounded answer generation, where generates answers to the quandary question based on the selected principles. We conduct experiments to how far LLMs have achieved in comparison to human ethicist answers and discuss remaining challenges, so it initiates more discussion on the responsible AI system for under-explored topics.

Background and Related Works

Open-ended Question Answering Open-ended questions require a potentially vast set of AI capabilities to answer. Unlike closed-book/knowledge-based QA, different respondents might have different opinions that lead to multiple correct answers (Luo et al. 2021). These differences come from various aspects, such as cultural, personal, or religious aspects, which results in diverse social norms, common sense, reasoning, and values (Acharya, Talamadupula, and Finlayson 2020).

The main challenge in open-ended QA is on analyzing and evaluating the answers to the questions. Most works in open-ended QAs focus on solving this evaluation problem through various approaches. One of the examples is on analyzing the insights of the answers of open-ended QA where NLP techniques such as topic modeling (Pietsch and Lessmann 2018; Nanda et al. 2021) are used to solve this problem. Another line of work focuses on evaluating the open-ended exam question answering which is commonly modeled as a regression task to predict the score given the question and answers (Hussein, Hassan, and Nassef 2019; Dong and Zhang 2016; Dasgupta et al. 2018; Birla, Jain, and Panwar 2022). In some cases of open-ended QA (e.g., medical and visual QA), the answer set can be considered as limited and such cases can be modeled as a classification task (Antol et al. 2015; Zhan et al. 2020). While in some other cases where the possible answers set is not definite (e.g., essay writing exam), multiple alternative references are used for evaluating the generated answers of the models (Luo et al. 2021; Xu et al. 2019; Kazemi and Elqursh 2017) using various automatic metrics such as BLEU (Papineni et al. 2002), ROUGE (Lin 2004), and BERTscore (Zhang et al. 2020).

In contrast to existing works on open-ended QA, we measure the quality of open-ended QA systems on ethical quandary questions. Ethical quandary QA has no clear definite answer set and requires in-depth reasoning. Furthermore, answering ethical quandary questions requires combining answers from multiple perspectives (i.e., alternative answers

in terms of question answering), as a one-sided normative ethical judgment cannot represent incommensurable and diverse ethical judgments (Talat et al. 2021).

Machine Ethics and Ethical Question Answering Machine or AI ethics is an important emerging area of research (Hendrycks et al. 2020; Prabhumoye et al. 2020; Schramowski et al. 2021). One line of current work focuses on improving machine understanding of human values and morality through classification tasks (Forbes et al. 2020; Emelin et al. 2020; Lourie, Le Bras, and Choi 2021; Sap et al. 2019). Delphi (Jiang et al. 2021) is a research prototype to emulate human moral judgments based on training with the large dataset and is trained to select “less contentious” choices in dealing with ethical questions. However, Talat et al. (2021) criticized that Delphi is based on average and skewed ethical values, which is not necessarily the ideal approach.

Hendrycks et al. (2020) proposes classifiers that explicitly provide the ethical perspective to be grounded against moral judgments (e.g., utilitarianism, deontology). Their work focuses on clear-cut situations instead of ambiguous moral dilemmas. In contrast, we attempt to understand the models’ ability to provide an answer in a manner of debate with explanations to ambiguous situations. This approach can be seen as a Socratic way unlike being an oracle to give moral answers based on specific theories as to how traditional philosophers like Plato or Aristotle (Pincoffs 1971).

The AI system involvement in the human moral decision-making process in an ethical quandary situation has been also actively discussed among moral philosophers. Some have suggested a direction where AI systems can be utilized positively and practically in such ethical quandary situations (Savulescu and Maslen 2015; Giubilini and Savulescu 2018; Lara and Deckers 2020; Lara 2021). They suggest the AI system can serve as a moral advisor that enhances an individual’s reflective process, so that humans can make better decision-making with a broader perspective while retaining the autonomy of their actions. This aligns with our vision on ethical quandary question answering, which focuses on providing a multi-perspective.

Ethical Quandary and Ethical Principle

Ethical quandary refers to a perplexing situation and or dilemma in the strict sense. An ethical dilemma is a situation in which any choice involves violating some widely held moral principle. Among various possible desired virtues for “ideal answer”, we seek the virtue of providing multiple perspectives for AI system in ETHICAL QUANDARY GQA. This is because we believe discussing the quandary from distinct perspectives is the most robust and safest way to deal with questions involving ethics; ethical judgment is dynamic (Bicchieri 2005) where what is considered to be “norm” or “right” shifts through time or differs by cultures.

Ethics takes a role to provide grounds for resolving the perplexity or make a decision in such ethical dilemma situation, as conceived of as moral rules or principles (Pincoffs 1971). Different principles focus on different aspects of the same situation to judge what is ethical and correct (Bass, Barnett, and Brown 1999; Forsyth and Pope 1984). Thus, different

ethical principles result in distinct and even contradictory answers to the same ethical quandary question.

ETHICAL QUANDARY GQA

Task Setup

We investigate a model’s ability to answer ethical quandaries. Given the ethical quandary question Q in context, the model is expected to generate a free-form text answer(s) A in a paragraph(s). In this task, the ethical quandary question consists of context (c) and a question sentence (q). The context includes details of the situation (e.g., narrator details, a specific event, involved parties, a particular condition) from the perspective of a narrator in the form of text paragraphs.

Ethical quandary refers to a perplexity arisen by a situation in which it is hard to decide what to do morally, and, in a more strict sense, an ethical dilemma where none of possible choices unambiguously acceptable. There is no single “correct” answer to such an ethical quandary question, like other open-ended questions. Instead, depending on the ethical principles leveraged to address the quandary, multiple conflicting valid answers exist. Thus, a clear-cut answer that everyone agrees with is difficult to formulate (i.e., right/wrong).

Evaluation Criteria

As ETHICAL QUANDARY GQA is distinct from existing question answering task, we propose to evaluate the ability of models with three most relevant metrics for assessing the success of the answer generation: **1) Multi-perspective:** The ability to provide more than one point of view to the ethical quandary can be interpreted as the model’s potential to carry out a deliberative discussion on the issue. It is important to ensure diverse ethical judgments through multiple-perspectives for answering ethical quandaries. Moreover, it is unsafe for an AI system to provide a single-sided authoritative normative judgement (Talat et al. 2021). We evaluate whether the model answers from different angles. **2) Coherence:** We test whether the generated output is logically composed and connected while sticks with the focus of the ethical quandary question throughout. **3) Justification:** We investigate the extent to which the system can reason about the question by explicitly exhibiting its reasoning chain. The more a system can justify its stance, the more reliable the system is.

Dataset

New York Times Ethicist Columns (NYT-Ethicist) is a set of weekly columns on ethical quandaries written by professional philosophers. Each quandary is sent from a reader, describing a complex situation and a question arising from it. A corresponding answer to the quandary is written by a philosopher and usually provides multiple perspectives for the situation. We collected 1,295 pairs of {quandary, answer from a philosopher} from the *NYT* website.

Methodology

We conduct an experiment with LLMs with few-shot in-context prompt learning in ETHICAL QUANDARY GQA. Inspired by the characteristics of the ethical quandaries, the

model explicitly takes ethical principles to generate an answer to ethical quandary. For the ease of reference, we call it as . It involves two steps: i) select relevant ethical principles to the given ethical quandary and ii) generate principle-guided multi-perspective answers.

Step 1: Principle Provider

First, we construct a pool of principles relevant to the given ethical quandary. To achieve this, we used three sources: i) retrieved top-10 relevant principles from a principle set composed of 300,000 samples of rules-of-thumb (RoT) from SOCIAL-CHEMISTRY (Forbes et al. 2020) and MORAL-STORY (Emelin et al. 2020) datasets, ii) generated RoTs using a fine-tuned RoT generator model from Forbes et al. (2020) and generated principles using Jurassic-1 with prompt-based few-shot learning iii) hand-crafted set of ethical principles from western/eastern ethical theories. Then, out of the principle pool, we obtain obtain most optimal unique set of principles to provide to the generator in the next step. First, we rank the candidate principles based on the similarity scores from the off-the-shelf scorers (i.e., Sentence-BERT and T0-model-based scorer). Second, we keep the candidates with scores higher than the threshold th . For our experiments, we use the optimal th identified by a small subset of validation set (i.e., $th = 1.02$).

We use the described automatic principle provider in the experiment. However, this can be replaced by a *human* where humans can provide relevant ethical principles. This option is a safety measure to ensure the existence of an agent for the machine generation and to provide human-in-the-loop moderation of model behavior.

Step 2: Principle-Guided Answer Generator

Given relevant principles from the first step, we adopt prompt-based few-shot learning for principle-guided answer generation. The prompt-based learning is a good way to test the in-nature ability of pre-trained LLMs with a minimum guidance about the task. Each test sample has ethical quandary, Q , and relevant ethical principle(s) p_k , where $k \in [1, 3]$ ¹, provided from previous procedure. We manually craft prefix prompts utilizing templates to format the input for the prompt learning. Since there is more than one ethical principle, we did multi-step prompting of the LLM to incorporate single or multiple principle(s) in the final generated answer A addressing the ethical quandary.

For the first answer with p_1 , we utilize PROMPT1. Given the two-shot examples (2-Shot) and the test ethical quandary with the first template, the model generates the output sequence A_1 by sampling from

$$P(A_1 | \text{PROMPT1, 2-Shot}), \quad (1)$$

Then, we continue answer generation with other principles if there are more than one principle selected. We recursively generate A_2 and A_3 using another prompt template PROMPT-j, where $j \in \{2, 3\}$. Given previous prompts and previously generated answer(s) (A_{k-1}), the model generates

¹We constraint the maximum to three principles for experiment, but this does not indicate we prioritize certain principles over other.

Ethical Quandary: [...] One of my housemates is planning to self-quarantine in our apartment after a vacation to Spain, which she took despite increased warnings about Covid-19 [...] I am in my 20s but have struggled with some moderate health issues; work with professors who are in their 70s. Can I tell my housemate to quarantine elsewhere? Should the university provide housing for one of us? Should I notify our boss?

AISOCRATES Selected Principles:

p_1 : It is your role to inform your housemate about the potential risks to others due to his decision in regards to quarantining in a shared household.

p_2 : You shouldn't hang out with your friends if they are sick with the flu.

AISOCRATES Generated Answer: It is your role to inform your housemate about the potential risks to others due to his decision in regards to quarantining in a shared household. This is a difficult question which seems appropriate only during a pandemic. In such times, rules that seemed ridiculous before acquire new importance, and yours is to protect the welfare and well-being of others. [...] But a good rule of thumb is that, as best can reasonably be determined, reporting her to her superiors is the right thing to do.[...]

From a different perspective, you shouldn't hang out with your friends if they are sick with the flu. As for what you should do, you should know what to say when you speak to your housemate. "I think we should discuss this further, because I have some concerns about you as a co-worker and roommate and you doing this. [...]" If so, we could decide together where to do that, and if I'm elected to do that, I'd probably have to check in with my supervisor about work."

Table 1: Illustration of model selected principles and a corresponding automatic answer generated by AISOCRATES to an ethical quandary. The p_1 and p_2 are principles selected by automatic module of Principle Provider of AISOCRATES. Then, the final automatic answer is generated with a guidance of those principles.

the consecutive answers A_k . For instance, A_2 is obtained by sampling from

$$P(A_2 | \text{PROMPT2}, A_1, \text{PROMPT1}, \text{2-Shot}). \quad (2)$$

We expect this consecutive prompting to allow the model to incorporate the previous answer A_{k-1} when generating A_k so the overall answer is more coherent. In the end, the final answer A is obtained by concatenating generations $A_1 \sim A_k$.

Experiments

We conduct our main experiment to evaluate the ability of AISOCRATES to answer ethical quandary questions from multiple perspectives. We assess the fully automated pipeline version of AISOCRATES by automating the principle provider.

Experimental Setup

In this experiment, We only take two samples from the train split for few-shot learning with our methodology. We then obtain an answer for each of the 130 test samples from AISOCRATES, which is backboneed by one of the largest publicly available pre-trained LLMs – Jurassic-1 Jumbo (Lieber et al. 2021) with 178 billion parameters.

We mainly evaluate the model performance with a human evaluation due to the one-to-many nature of generation tasks. The automatic metrics with a reference are often limited in evaluating the desired quality in generations. Moreover, as explained earlier, the ethical quandary questions have multiple valid answers depending on the relevant ethical principles. This makes our evaluation more challenging with the automatic metrics. Thus, we need human judgment in the performance evaluation. For completeness, however, we also perform evaluation using the standard automatic metric, ROUGE (Lin 2004), SacreBLEU (Post 2018) and BERTscore (Zhang et al. 2020), and how the automatic evaluation have correlations with human judgment.

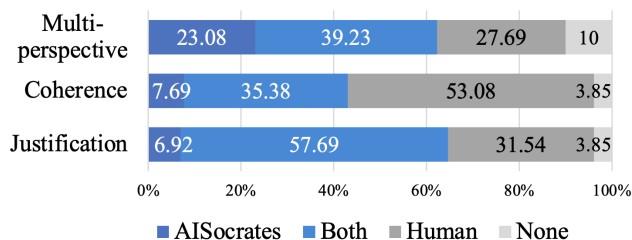


Figure 1: Human evaluation results. The blue-colored parts of the bar are success rates (%) of AISOCRATES to contain each of the evaluation criteria. For instance, 62.31% (sum of 23.08 and 39.23) of the time AISOCRATES have multi-perspectives in generated answers. Note that the answers were blindly evaluated by randomly assigning A/B tags without indications of AISOCRATES or human.

We conduct A/B testing judged by human annotators, in which we compare answers from AISOCRATES and *NYT-Ethicist*. We blindly provide an ethical quandary and the two corresponding answers (AISOCRATES and *NYT-Ethicist*) to the annotators. To evaluate the three criteria, we ask the following questions from the annotators– “Which of the answers is addressing the ethical dilemma from multiple perspectives?”, “Which answer is more coherent?” and “Which answer include sound reasoning for its stances?”. The human annotator has a choice of {A, B, Both, None}, where each of the answers is randomly assigned with A or B tags. We report win-tie-loss rates for each criterion of AISOCRATES against philosopher-written *NYT-Ethicist* answers.

Experimental Results

We provide an example of selected principles and a generated answer by AISOCRATES in Table 1. In this section, we

answer the following questions through evaluation:

Q1: Can AISOCRATES answer from multiple perspectives? This aspect is of great interest in ETHICAL QUANDARY GQA. As shown in Figure 1, AISOCRATES provides the answer with multiple perspectives 62.31% of the time (the sum of choices AISOCRATES and Both cases), which is 4.61% less than the *NYT-Ethicist* answers do. There is no statistically significant difference for multi-perspective provision between the *NYT-Ethicist* answers and AISOCRATES answers. This indicates that our AISOCRATES can achieve comparable performance in providing multiple perspectives to answer the ethical quandary questions. AISOCRATES sometimes generates answer with single perspective (37.69%). For some ethical quandaries, multiple relevant principles share similar ethical implications. To elaborate, two ethical principles can have overlapping ideologies: “You should avoid telling lies” and “A person is honorable and moral by being honest”. In such scenarios, our model learns to generate aggregated common answer instead of repeating similar content for multiple times.

Recalling the definition, ethical quandary is a perplexing situation and/or dilemma. There may be some quandaries that are not dilemmas in a strict sense – cases where there is some initial puzzlement or doubt but where, upon further analysis, it turns out that there can be one viable answer. Given that ethical quandary test samples from NYT are from the general public, it may be that some of the questions being offered to NYT Ethicist only merits a single responsible answer from an ethicist, regardless of differences in principles one embraces.

Q2: Can AISOCRATES compose coherent answers? AISOCRATES answers are considered reasonably coherent compared to expert-written answers with 43.07% success rate. This weakness in coherency compared to the multi-perspective criterion can be explained by several factors. We conducted qualitative analysis on the generated answers with philosophical reference. It was concluded that the model losses attention and includes irrelevant or redundant content in the arguments, even though the answer starts by aligning with the provided principle in the beginning. There is no explicit guidance or learning for improving coherence in the current methodology of AISOCRATES, except for the provision of coherent examples in the few-shot samples and the consecutive generation rather than separate generations.

Q3: Does AISOCRATES provide justifications to its perspectives on the ethical quandary? AISOCRATES could employ clear and sound reasoning in the answer for 64.61% of the time. We could also observe that coherence and justification are positively related. The answers by AISOCRATES that are assessed to be coherent include justifications for the perspective for 91.07% of the time while the non-coherent answers justifies the perspective(s) only for 55% time. This points to potential research on enhancing the reasoning capability of the model to improve on coherence of the answer together. Reasoning in ETHICAL QUANDARY GQA is challenging as the situations are exceedingly complicated and the answers are in free form. Moreover, generating answers with sound and clear justifications requires

		Multi-perspective	Coherence	Justification
Rouge-2	low	62.50	45.83	68.75
	high	67.65	50.00	70.59
BERTscore	low	51.35	40.54	62.16
	high	69.77	53.48	72.09
SacreBLEU	low	61.54	51.28	74.36
	high	56.00	44.00	68.00

Table 2: Relationship between automatic metric scores and human evaluation results. Rates of how many of the generated answer samples in each set (e.g., low ROUGE-2 F1) are annotated to achieve the human-evaluation criteria. For instance, 69.77% of the times, the generated answers in the high BERTscore set is annotated to have multi-perspective.

the model to have outstanding common-sense and deductive reasoning skills that is considered difficult even for humans. Considering that, the results suggest a promising potential.

Analysis and Discussion

Automatic Metric and Model Performance

Besides the human evaluation, we score the generated answer with the automatic metrics ROUGE, SacreBLEU and BERTscore in reference to expert-written *NYT-Ethicist* answers. We mainly investigate F1-scores to understand how much content of the *NYT-Ethicist* answers and how much distinct content exists in AISOCRATES’s answers. It is shown that the ROUGE-2 scores are low with F1-Score of 3.68. This means the generated answers do not contain many bi-gram phrases that overlapping with the human-written reference. The model generated answer achieved 32.13, 52.29 and 1.84 in ROUGE-1, BERTscore and SacreBLEU respectively.

To understand automatic evaluation results in relation to the human evaluation criteria {multi-perspective, coherence, justification} (Table 2), we analyzed the subsets of generated answers with low and high automatic scores (e.g., low BERTscore, high BERTscore).² We calculated ratios of the generated answers that are annotated to meet each of evaluation criteria for each of each of low or high score sets. In terms of Rouge-2 and BERTscore, those answers with higher scores tend to more likely to have met evaluation criteria compared to those with lower scores. The differences in ratios are bigger for BERTscore while it is insignificant between subsets of Rouge-2 score. In detail, It is shown that answers with high BERTscore tend to have multi-perspective (69.77%) than those with low scores (51.35%). The difference is statistically significant with $p < 0.05$. We can translate this into that the answer with high similarity score with human written answer tends to contain multi-perspective.

On the other hand, SacreBLEU showed the inverse relationship from those with other metrics – lower success rates for the criteria in the high scored set. SacreBLEU shows similar performance to the case of open-domain dialogue (Cahyawijaya et al. 2021; Lin et al. 2021), which indicates a similar

²Here, low scored sample means those generation sample with less than $0.5 (< \mu - 0.5 \cdot \sigma)$ standard deviation from the average score and vice versa for high-scored samples ($> \mu + 0.5 \cdot \sigma$).

[Ethical Quandary] My classics department hosts an annual conference [· · ·]. This year the organizers urged us to offer extra credit to students who attend, hoping that the university will increase future funding if we have a high turnout. [· · ·] Wouldn't we artificially inflate our numbers and exploit our students if we gave this extra credit?

AISOCRATES (model-generated)	It's not inherently wrong to offer extra credit, or to encourage students to attend conferences or academic meetings. If there's a conference related to the course work they're taking, they could certainly get extra credit for attending. But extra credit for something that isn't directly related to a course is questionable, and you're right to raise a fairness issue. [· · ·] From a different perspective, if the university made attendance a requirement for credit, those requirements should be stated on the application form. It would be coercive if the professor, when he's scheduling the student for the term, mentioned that attendance would guarantee an A.
------------------------------	---

Table 3: An illustration of the AISOCRATES (model-generated) answer providing multiple perspectives on the ethical quandary question with understanding of the context while having a **low** Rouge-2 F1-score (0.01).

one-to-many characteristic for evaluating such systems. Similarly, the ROUGE score, despite showing a better trend, also doesn't show a strong correlation with human judgment. On the other hand, the semantic evaluation metric, i.e., BERTScore, shows a better correlation with all human judgment, i.e., multi-perspective, coherence, and justification. This fact supports the effectiveness of semantic-based metrics over the N-gram metrics, of which we suggest future work to explore different semantic evaluation metrics for handling open-ended QA, especially for the ethical quandary QA.

Like other open-ended generation tasks (Xu et al. 2020; Zhang et al. 2018), the reference-based metric has its own limitation to be the absolute evaluation standard because of its one-to-many nature. In the case of ETHICAL QUANDARY GQA, since there are multiple possible principles that can be used to answer the ethical quandary, reference-based metrics can falsely penalize any answers that use different yet valid principles. Indeed, we were able to identify cases where a low ROUGE score (less overlap with human-written answers) does not necessarily indicate poor performance of the model in answering the ethical quandary; generation text having low ROUGE-2 score while having good performance in other factors such as coherence and multi-perspectives as shown in Table 3. Yet, high semantic similarity with human-written answer would indicate the overall high quality of generated answer with the investigated virtues.

Analysis with Philosophic Reference

In this paper, we explored the task with one of the biggest LLMs (178 billion parameters) with the prompt-based learning. Since this does not involve any parameter tuning, it is a good way to understand the in-nature ability of pre-trained LLMs with a minimum guidance about the task. We share qualitative analysis on the model's ability to generate distinct and principle-grounded output answers based on the input principle and context. Instead of using the automatic principle provider, the relevant ethical principles for each of ethical quandaries are selected by human to ensure the provided principle is noise-free and context-relevant.

Here, we evaluated whether the model generates the expected output based on the provided principle. Given an ethical quandary, "Is XYZ behavior ethical?" can be answered "Yes" and/or "No" depending on a specific ethical princi-

ple. The model showed different behaviors most of the time depending on the ethical principles and correctly answers according to the provided principles. On the other hand, we also note cases where the generated answer does not provide a clear yes-or-no answer. Instead, it provides an explanation to address the question of the ethical quandary.

In terms of consistency with the input ethical principle, model generated answers showed link between quandary and ethical principle by elaborations with examples. For instance, for an ethical principle, "A smaller sacrifice is morally justifiable for the greater good," the corresponding generated answer attempts to elaborate with numerical reasoning to quantify the verbal input "small risk." Another illustration of the model-generated answer's consistency is the answer that sticks with its "freedom-first" input principle. It is stated negatively as government should "avoid placing excessive restrictions on their personal freedom". However, it is noted that the model fails to elaborate rationale/logic of the answer. The generated answer sometimes discusses the divergent points of concern that the ethical quandary question seeks to resolve, resulting in that the question, the principle, and the answer becoming muddled.

Although the model can generate a distinct answer with some consistency, weaknesses are also investigated. The generated answer begins as being grounded on the input principle but strays further from the topic paragraph-by-paragraph. It is not surprising that the longer the model performs generation, the less relevant its answer gets. The model-generated answers sometimes lack relevance and attention to detail. Also, extra information, not necessarily factually wrong, and re-asking and re-answering the question make the generated output redundant in paragraph/writing organization. In summary, analysis revealed model's ability in distinct answer generation based on different input ethical principle alongside its weakness in a consistency and logical generation.

Challenges in ETHICAL QUANDARY GQA

Strictly speaking, the virtues that are deemed in answers for ethical quandary questions involve more extended criteria other than multiple perspective, coherence and justification. They may include, but not limited to, understanding of context, complex moral common sense reasoning over the context, choices of relevant ethical concepts, deliberation from

multiple perspectives, justification of the stances it take and, even more strictly, the style of writing. In this work, we prioritize on choices of relevant ethical concepts and deliberation from multiple perspectives as our chosen criteria.

Although we weighed multiple perspectives, coherence, and justification criteria more in this work, other virtues might be sought to answer the ethical quandary. Experimental result shows that other challenges that are left to address involve better justification and reasoning for ETHICAL QUANDARY GQA. Deep common-sense moral reasoning is crucial to reach a sound judgment about complex ethical issues. Furthermore, reasoning capability brings about comprehensibility about the model’s verdict which is valuable as the model can not be trusted when its decision-making process is not transparent. We also have found that the generated answers sometimes have a weak argument with no clear and sound backup. However, improving reasoning and justification in the answers is not trivial because it involves sensible organization and presentation of ideas and the internal relevance of content. It is a nevertheless important research direction.

In discussing ethical quandary, context-specific consideration may be another virtue to be deemed for. The context of the ethical quandary is described in one or multiple paragraphs and hence not simple. We indirectly evaluated the model understanding of the input ethical quandary questions with multi-perspective and coherence criteria in this work. Although the overall generated answers are consistent with context and question, it remains unclear if the model understood the context in depth because some generated answers are generic without the details of what is under discussion. Generally speaking, many end-to-end NLP tasks today lack the explicit evaluation of the “understanding” component.

Ethical Consideration

We pay extra attention to discuss the ethical responsibility and the impact of this work. We clarify that our experiment involves a fully automated pipeline for the principle provider and principle-guided answer generator, which is an attempt to understand the model’s upper bound for research purposes but was not considered to be deployed for actual application without human agency (i.e., principle provider). Here, our aim is not to generate the most “ethical” answers but to explore LLMs’ ability to provide distinct answers to a single quandary depending on varying ethical principles. It is worth highlighting that AISOCRATES should not be considered as an oracle providing a definite answer but as a tool for providing multiple perspectives on ethical quandary questions.

Nevertheless, when using AI technology to deal with ethical issues, it is essential to consider the safety and ethical implications. Letting an AI system answers ethical questions without a human agent can be controversial because it is unclear who takes responsibility for the action or output of the system (Anderson and Anderson 2007; Cave et al. 2018). Therefore, as a form of safety measure, we designed the model used in the experiment in a way that the model can be controllable by humans – i.e., our pipeline is strictly designed in a way it requires human intervention to provide the selected ethical principles used to guide the system in generating answers, while the relevance scorers only provide a

recommendation to the human actor (see Appendix for more detail). Additionally, the generated answer is appended by a leading and trailing text denoting “*The answer is generated by an AI algorithm, please proceed with caution*” which serves as a precautionous to the user.

Although the AI system has not achieved the level for real application deployment, it is still worth discussing the potential user’s responsibility with such a system. If such an application is deployed, we believe it should be clearly stated that the potential user still has to assess the machine outputs for truth, soundness, and moral acceptability. The user is responsible for what he/she/they choose to do with the models’ output although the user may not be responsible for the model output itself. We propose this direction of AI systems giving a virtuous answer that challenges the user’s thought process by providing multiple perspectives to a quandary situation. This is in line with the Socrates motto: *philosophers today do not read Socrates as an infallible guide, rather read him as someone to challenge and prompt thinking further.*

From the analysis of the model-generated answer even with the existence of a human agent (manual principle provider), it is found that the model-generated answer contains medical-practical information along with the redundancy or evaluation of the ethical quandary from a legal perspective. Factually wrong advice on sensitive topics such as medical and legal issues is not acceptable because it can result in severe impacts such as harm to the real users’ physical or mental health or legally unlawful decisions made by the users. It should be clearly reminded that the application it to provide multiple angles regarding the ethical quandary, allowing the narrator to view their dilemma from different points of view.

Moreover, our assumption throughout the work was that the principles provided in our methodology are ethical principles or rules-of-thumb that can lead to ethical/non-controversial/non-harmful advice while providing multiple perspectives. However, it should not be overlooked that the model’s ability to generate distinct answers based on different principles has a potential risk. When the “controversial/harmful” principles are intentionally provided to the system, there is a risk of generating corresponding harmful answers. Thus, the deployment of an actual application of the system should be thoroughly reviewed and needs to be managed responsibly.

Conclusion

In conclusion, we explored the task of ETHICAL QUANDARY GQA with AISOCRATES, which answers ethical quandary questions from multiple perspectives based on different ethical principles. The full-automatic pipeline of AISOCRATES is studied to understand the LLMs’ upper bound, while the design choice of the human-intervened model exists to guarantee the existence of agency for the generation. The main experimental result shows that the full-automatic AISOCRATES provided multiple-perspective answers comparable to answers written by human philosophers. Furthermore, the results illustrate that the answers generated from our system still lack coherence and justification compared with philosopher-written answers, which highlights the need for more advanced methods for ETHICAL QUANDARY GQA.

References

- Acharya, A.; Talamadupula, K.; and Finlayson, M. A. 2020. Towards an Atlas of Cultural Commonsense for Machine Reasoning. *arXiv preprint arXiv:2009.05664*.
- Ahn, Y.; Zhang, Y.; Park, Y.; and Lee, J. 2020. A Chatbot Solution to Chat App Problems: Envisioning a Chatbot Counseling System for Teenage Victims of Online Sexual Exploitation. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI EA '20, 1–7. New York, NY, USA: Association for Computing Machinery. ISBN 9781450368193.
- Anderson, M.; and Anderson, S. L. 2007. Machine ethics: Creating an ethical intelligent agent. *AI magazine*, 28(4): 15–15.
- Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.
- Bass, K.; Barnett, T.; and Brown, G. 1999. Individual difference variables, ethical judgments, and ethical behavioral intentions. *Business Ethics Quarterly*, 9(2): 183–205.
- Bicchieri, C. 2005. *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.
- Birla, N.; Jain, M. K.; and Panwar, A. 2022. Automated assessment of subjective assignments: A hybrid approach. *Expert Systems with Applications*, 203: 117315.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Cahyawijaya, S.; Winata, G. I.; Wilie, B.; Vincentio, K.; Li, X.; Kuncoro, A.; Ruder, S.; Lim, Z. Y.; Bahar, S.; Khodra, M.; Purwarianti, A.; and Fung, P. 2021. IndoNLG: Benchmark and Resources for Evaluating Indonesian Natural Language Generation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 8875–8898. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Cave, S.; Nyrup, R.; Vold, K.; and Weller, A. 2018. Motivations and risks of machine ethics. *Proceedings of the IEEE*, 107(3): 562–574.
- Dasgupta, T.; Naskar, A.; Dey, L.; and Saha, R. 2018. Augmenting textual qualitative features in deep convolution recurrent neural network for automatic essay scoring. In *Proceedings of the 5th Workshop on Natural Language Processing Techniques for Educational Applications*, 93–102.
- Dong, F.; and Zhang, Y. 2016. Automatic Features for Essay Scoring – An Empirical Study. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1072–1077. Austin, Texas: Association for Computational Linguistics.
- Emelin, D.; Bras, R. L.; Hwang, J. D.; Forbes, M.; and Choi, Y. 2020. Moral stories: Situated reasoning about norms, intents, actions, and their consequences. *arXiv preprint arXiv:2012.15738*.
- Forbes, M.; Hwang, J. D.; Shwartz, V.; Sap, M.; and Choi, Y. 2020. Social chemistry 101: Learning to reason about social and moral norms. *arXiv preprint arXiv:2011.00620*.
- Forsyth, D. R.; and Pope, W. R. 1984. Ethical ideology and judgments of social psychological research: Multidimensional analysis. *Journal of Personality and Social Psychology*, 46(6): 1365.
- Giubilini, A.; and Savulescu, J. 2018. The artificial moral advisor. The “ideal observer” meets artificial intelligence. *Philosophy & Technology*, 31(2): 169–188.
- Gu, Y.; Mishra, B. D.; and Clark, P. 2021. DREAM: Uncovering Mental Models behind Language Models. *arXiv preprint arXiv:2112.08656*.
- Hendrycks, D.; Burns, C.; Basart, S.; Critch, A.; Li, J.; Song, D.; and Steinhardt, J. 2020. Aligning ai with shared human values. *arXiv preprint arXiv:2008.02275*.
- Hussein, M. A.; Hassan, H.; and Nassef, M. 2019. Automated language essay scoring systems: A literature review. *PeerJ Computer Science*, 5: e208.
- Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.; Madotto, A.; and Fung, P. 2022. Survey of Hallucination in Natural Language Generation. *ACM Comput. Surv.* Just Accepted.
- Jiang, L.; Hwang, J. D.; Bhagavatula, C.; Bras, R. L.; Forbes, M.; Borchardt, J.; Liang, J.; Etzioni, O.; Sap, M.; and Choi, Y. 2021. Delphi: Towards machine ethics and norms. *arXiv preprint arXiv:2110.07574*.
- Kazemi, V.; and Elqursh, A. 2017. Show, ask, attend, and answer: A strong baseline for visual question answering. *arXiv preprint arXiv:1704.03162*.
- Lara, F. 2021. Why a Virtual Assistant for Moral Enhancement When We Could have a Socrates? *Science and engineering ethics*, 27(4): 1–27.
- Lara, F.; and Deckers, J. 2020. Artificial intelligence as a socratic assistant for moral enhancement. *Neuroethics*, 13(3): 275–287.
- Lieber, O.; Sharir, O.; Lenz, B.; and Shoham, Y. 2021. Jurassic-1: Technical details and evaluation. *White Paper: AI21 Labs*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lin, Z.; Liu, Z.; Winata, G. I.; Cahyawijaya, S.; Madotto, A.; Bang, Y.; Ishii, E.; and Fung, P. 2021. XPersona: Evaluating Multilingual Personalized Chatbot. In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, 102–112. Online: Association for Computational Linguistics.
- Lourie, N.; Le Bras, R.; and Choi, Y. 2021. SCRUPLES: A Corpus of Community Ethical Judgments on 32,000 Real-Life Anecdotes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 13470–13479.
- Luo, M.; Sampat, S. K.; Tallman, R.; Zeng, Y.; Vancha, M.; Sajja, A.; and Baral, C. 2021. ‘Just because you are right, doesn’t mean I am wrong’: Overcoming a Bottleneck in the Development and Evaluation of Open-Ended Visual Question Answering (VQA) Tasks. *arXiv preprint arXiv:2103.15022*.

- Nanda, G.; Douglas, K. A.; Waller, D. R.; Merzdorf, H. E.; and Goldwasser, D. 2021. Analyzing large collections of open-ended feedback from MOOC learners using LDA topic modeling and qualitative analysis. *IEEE Transactions on Learning Technologies*, 14(2): 146–160.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318. Philadelphia, Pennsylvania, USA: Association for Computational Linguistics.
- Pietsch, A.-S.; and Lessmann, S. 2018. Topic modeling for analyzing open-ended survey responses. *Journal of Business Analytics*, 1(2): 93–116.
- Pincoffs, E. 1971. Quandary Ethics. *Mind*, 80(320): 552–571.
- Post, M. 2018. A Call for Clarity in Reporting BLEU Scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, 186–191. Brussels, Belgium: Association for Computational Linguistics.
- Prabhunoye, S.; Boldt, B.; Salakhutdinov, R.; and Black, A. W. 2020. Case study: Deontological ethics in NLP. *arXiv preprint arXiv:2010.04658*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research*, 21: 1–67.
- Roberts, A.; Raffel, C.; and Shazeer, N. 2020. How Much Knowledge Can You Pack Into the Parameters of a Language Model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 5418–5426.
- Roller, S.; Dinan, E.; Goyal, N.; Ju, D.; Williamson, M.; Liu, Y.; Xu, J.; Ott, M.; Smith, E. M.; Boureau, Y.-L.; and Weston, J. 2021. Recipes for Building an Open-Domain Chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, 300–325. Online: Association for Computational Linguistics.
- Sap, M.; Gabriel, S.; Qin, L.; Jurafsky, D.; Smith, N. A.; and Choi, Y. 2019. Social bias frames: Reasoning about social and power implications of language. *arXiv preprint arXiv:1911.03891*.
- Savulescu, J.; and Maslen, H. 2015. Moral enhancement and artificial intelligence: moral AI? In *Beyond Artificial Intelligence*, 79–95. Springer.
- Schramowski, P.; Turan, C.; Andersen, N.; Rothkopf, C.; and Kersting, K. 2021. Language models have a moral dimension. *arXiv preprint arXiv:2103.11790*.
- Su, D.; Xu, Y.; Yu, T.; Siddique, F. B.; Barezi, E.; and Fung, P. 2020. CAiRE-COVID: A Question Answering and Query-focused Multi-Document Summarization System for COVID-19 Scholarly Information Management. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*. Online: Association for Computational Linguistics.
- Tafjord, O.; and Clark, P. 2021. General-purpose question-answering with macaw. *arXiv preprint arXiv:2109.02593*.
- Talat, Z.; Blix, H.; Valvoda, J.; Ganesh, M. I.; Cotterell, R.; and Williams, A. 2021. A Word on Machine Ethics: A Response to Jiang et al.(2021). *arXiv preprint arXiv:2111.04158*.
- Thomson, J. J. 1976. Killing, letting die, and the trolley problem. *The monist*, 59(2): 204–217.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, C.; Liu, P.; and Zhang, Y. 2021. Can Generative Pre-trained Language Models Serve As Knowledge Bases for Closed-book QA? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 3241–3251.
- Welch, C.; Lahnama, A.; Pérez-Rosas, V.; Shen, S.; Seraj, S.; An, L.; Resnicow, K.; Pennebaker, J.; and Mihalcea, R. 2020. Expressive Interviewing: A Conversational System for Coping with COVID-19. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*.
- Xu, P.; Patwary, M.; Shoeybi, M.; Puri, R.; Fung, P.; Anandkumar, A.; and Catanzaro, B. 2020. MEGATRON-CNTRL: Controllable Story Generation with External Knowledge Using Large-Scale Language Models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2831–2845.
- Xu, Y.; Chen, L.; Cheng, Z.; Duan, L.; and Luo, J. 2019. Open-ended visual question answering by multi-modal domain adaptation. *arXiv preprint arXiv:1911.04058*.
- Zhan, L.-M.; Liu, B.; Fan, L.; Chen, J.; and Wu, X.-M. 2020. Medical visual question answering via conditional reasoning. In *Proceedings of the 28th ACM International Conference on Multimedia*, 2345–2354.
- Zhang, S.; Dinan, E.; Urbanek, J.; Szlam, A.; Kiela, D.; and Weston, J. 2018. Personalizing Dialogue Agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2204–2213.
- Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2020. BERTScore: Evaluating Text Generation with BERT. In *International Conference on Learning Representations*.

Appendix

Limitations

There is no single answer to our ETHICAL QUANDARY GQAs different principle lead to various answers as discussed previously. Our evaluation relies on human evaluation, which may increase the additional cost.

Our AISOCRATES is based on prompt learning of LLMs, using one of the most powerful publicly available pre-trained

LLMs - Jurassic-1 178B model. This requires access through an API or large computational resources. Additionally, we are currently testing our methodology with English. Thus, extending our AISOCRATES to different types and languages of LLMs is left as one of potential explorations.

Other further remaining challenges or future directions in ETHICAL QUANDARY GQA have been previously discussed in the main content.

AISocrates Workflow

We show the workflow of AISocrates on Figure 2. Our workflow requires human intervention for selecting the ethical principles to provide a safety measure for preventing potentially controversial/harmful answers. Additionally, the answer provided by the principle-guided answer-generator is appended with a leading and trailing text denoting “The answer is generated by an AI algorithm, please proceed with caution” which serves as a precaution to the user.

Data Details

	Quandary	Answer
# of words / sample	120.0 ± 76.0	205.8 ± 83.0
# of sentences / sample	7.1 ± 3.9	11.2 ± 4.6
# of data samples	1,295	

Table 4: *NYT-Ethicist* data statistics

The data ranges from 7 May 2006 to 2 November 2021 from NYT Ethicist columns (<https://www.nytimes.com/column/the-ethicist>). The statistics for text length for Quandary and Answer is shown in Table 4.

Principle Provider Detail

We selected T0 model to be one of scorers for measuring the relevance between ethical quandary and principle candidates. T0 model is known for its powerful zero-shot generalization ability in unseen NLP tasks by utilizing prompts. We provide the prompt “Context: {context} Principle: {principle} \n Is the principle relevant to the context?” to T0 for each principle and calculate perplexity score of “yes”. Then, the principles are ranked in ascending order of the perplexity score.

Example Generated Answers from AISOCRATES with Automated Principle Provider

We share some of example generated answers from AISOCRATES to test ethical quandary questions. Here, the ethical principles are *automatically* selected by our system and corresponding principle-guided answer is generated. The examples can be found in Table 6 ~ Table 8.

Model Details

Regarding the two-shots for the prompt-based learning, we use randomly selected samples from the train split. We need to extract the underlying principle of the *NYT-Ethicist* answer because the original data does not have a written principle

on which the answer is based. For model in all experiments in this paper, we choose one of the largest publicly available pre-trained LLMs – Jurassic-1 Jumbo (Lieber et al. 2021) with 178 billion parameters, which is based on the decoder module of the Transformer architecture (Vaswani et al. 2017). Note that our methodology is model-agnostic although we conducted experiment with Jurassic-1 Jumbo. We process paragraphs with paragraph tags < p > and < /p > at the beginning and end of the paragraph, respectively, so the model can also learn paragraph writing.

Evaluation with Automatic Metric

	Precision	Recall	F1
ROUGE-1	31.95	32.78	32.13
ROUGE-2	3.65	3.78	3.68
ROUGE-L	20.19	20.69	20.31
BERTscore	52.14	52.57	52.29
SacreBLEU	1.84	-	-

Table 5: Average ROUGE-1/-2/-L, BERTscore, SacreBLEU scores for AISOCRATES generated answers with reference to the *NYT-Ethicist* answer.

Besides the human evaluation, we score the generated answer with the automatic metric ROUGE, SacreBLEU, BERTscore (Table 5) in reference to expert-written answers.

Trolley Problem – Human Principle and Automatic Answer Generation

We provide the answer generated by AISOCRATES with human selected principles answer to the trolley problem – Should we kill one person to save five people in danger of being hit by a trolley?. The two principles are contrasting: one is based on deontological perspective (“Killing is never acceptable.”) and another is based on utilitarianism (“The most appropriate action is the one that achieves the greatest good for the greatest number.”). The answer is provided in Table 12.

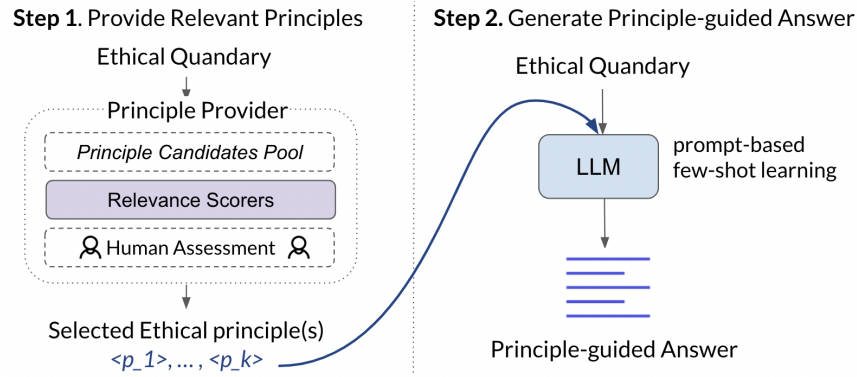


Figure 2: Illustration of our proposed ethical quandary question answering system, AISOCRATES. First, the principle provider matches relevant ethical principles, done either by a human or a model (automatic). Next, the principle-guided answer to ethical quandary is obtained with the selected principle(s), p_1, \dots, p_k .

[Ethical Quandary] I live in a shared house with three other people. We are all employees or graduate students at the university we graduated from last spring. One of my housemates, with whom I also work, is planning to self-quarantine in our apartment after a vacation to Spain, which she took despite increased warnings about Covid-19 in the media and messages from the university discouraging international travel. I am in my 20s but have struggled with some moderate health issues; my housemates and I also work with professors who are in their 70s. Can I tell my housemate to quarantine elsewhere? Should the university provide housing for one of us? Should I notify our boss?

AISOCRATES AISOCRATES Selected Principles:

p_1 : It is your role to inform your housemate about the potential risks to others due to his decision in regards to quarantining in a shared household.

p_2 : You shouldn't hang out with your friends if they are sick with the flu.

AISOCRATES (Model Generated Answer):

It is your role to inform your housemate about the potential risks to others due to his decision in regards to quarantining in a shared household. This is a difficult question which seems appropriate only during a pandemic. In such times, rules that seemed ridiculous before acquire new importance, and yours is to protect the welfare and well-being of others.

It is not at all clear that your housemate's decision to quarantine puts others at risk. If she is safe, those around her are safe. But a good rule of thumb is that, as best can reasonably be determined, reporting her to her superiors is the right thing to do.

The university shouldn't be responsible for providing housing, and, though it might seem disloyal, you should follow your employer's instructions rather than guidance from relatives. Hopefully, your housemate will also follow university guidance.

From a different perspective, you shouldn't hang out with your friends if they are sick with the flu. As for what you should do, you should know what to say when you speak to your housemate.

"I think we should discuss this further, because I have some concerns about you as a co-worker and roommate and you doing this. You seem to get sick a lot, and there's uncertainty about how the flu and the coronavirus might impact those of us at higher risk. If we all decide to quarantine for the time being, is that something you'd consider? If so, we could decide together where to do that, and if I'm elected to do that, I'd probably have to check in with my supervisor about work."

Table 6: An illustration of the generated answer by AISOCRATES with automatically selected principles to an ethical quandary question test sample.

[Ethical Quandary] I teach at a state university. Sometimes at the end of a semester a student asks me to raise a grade. Typically it is a student with children, who receives financial aid, health insurance and housing and risks losing these benefits if she receives the F she earned rather than the D I could bestow harsh consequences. Should I raise the grade?

AISOCRATES Selected Principles:

p_1 : teachers should raise grades.

p_2 : It is problematic to raise somebody's grade merely because he wants it.

p_3 : It's good to want a student to get the grades they deserve.

AISOCRATES (Model Generated Answer):

Yes, you should. One of the challenges of higher education is that it demands more from more and more of our students. Through this we ask them to invest in themselves. They deserve our help, and your, in their efforts to make wise use of that investment.

A D, of course, is better than an F. But an F is the start of someone's education, not the end. When you give a student an F, you may also diminish her motivation, condemn her to a future of underemployment, diminish her chances of having a satisfying life and contribute to the increasing cost of her public benefits.

The question isn't whether the student merits help; she obviously does. It's whether you care about her, and her potential. You do, and so should you.

From another point of view, you would have good reason to refuse a request to raise the grade if a student asked, and your professor did. You would be failing in your professional job if you raised somebody's score simply because he asked and would pay you. You would also be failing in your ethical and social responsibilities.

From a different point, as long as they didn't plagiarize, cheated on tests, stole homework, bunked class or otherwise committed academic misconduct, you should want your students to succeed.

Table 7: An illustration of the generated answer by AISOCRATES with automatically selected principles to an ethical quandary question test sample.

[Ethical Quandary] At a company meeting, I saw a colleague use his cellphone camera to film under the skirt of an unsuspecting female colleague. Later I saw him download the images to his company computer, recharge his phone and continue filming. I said nothing then. I could now talk to him or her or Human Resources, but I am wary of his getting fired or my getting embroiled in a messy situation. What should I do?

AISOCRATES Selected Principles:

p_1 : It is good to report unauthorized filming to authorities.

p_2 : It's wrong to use your work computer to secretly film people.

p_3 : You shouldn't do inappropriate things while at work.

AISOCRATES (Model Generated Answer):

You must. You should talk to Human Resources, and to the woman's manager. HR is obligated to take reports of sexual harassment seriously. You can also discuss the situation with the men's manager.

You have no reason to worry that the man will be fired. But if there were a recognizable face under the skirt, he would certainly lose an opportunity for advancement. You should also consider reporting the incident to the police, who can pursue any criminal action.

Your right to "be wary of his getting fired" is legitimate. But at the same time, you have a responsibility to protect women in the workplace from sexual exploitation. If you feel you're not up to that, then shame on you.

From a different point, you should not use your work computer to collect video of women in the workplace or to secretly film them. Law prohibits certain kinds of such use of personal information. Penalties can include fines, jail sentences and being barred from gainful employment.

If you think you could anticipate his filming, you should confront him. Otherwise, report it. Your silence is only encouraging his habits and worse. From a different point, you should not. Your desktop is not your personal copying machine. (Sorry.)

Table 8: An illustration of the generated answer by AISOCRATES with automatically selected principles to an ethical quandary question test sample.

[Ethical Quandary] I am a grad student with teaching duties. My classics department hosts an annual conference at which high-level grad students present their research. This year the organizers urged us to offer extra credit to students who attend, hoping that the university will increase future funding if we have a high turnout. I believe extra credit should be related to course work. Worse still, few of my undergrad students have the background or even sufficient Latin or Greek to appreciate the papers presented. Wouldn't we artificially inflate our numbers and exploit our students if we gave this extra credit?

AISOCRATES
(model-
generated)

Model Selected Principles:

- p_1 : It's kind for a teacher to offer a student extra credit.
- p_2 : It is ethical to offer as grades, benefits not ordinarily required to meet instructional objectives.
- p_3 : It is unethical to coerce a student to attend.

Model Generated Answer:

It's not inherently wrong to offer extra credit, or to encourage students to attend conferences or academic meetings. If there's a conference related to the course work they're taking, they could certainly get extra credit for attending.

But extra credit for something that isn't directly related to a course is questionable, and you're right to raise a fairness issue. A conference's purpose is typically to share knowledge and ideas; the students who attend should have enough interest in the subject to have studied it on their own. It's reasonable for the university to offer a travel allowance, but extra credit seems unnecessary.

Moreover, universities aren't obligated to provide extra-credit inducements, or meet funding demands on their part. If they exist, they're an attempt to give students a more direct incentive to attend, which they sometimes do for sports games, for example. It's no doubt a pragmatic move, but universities are businesses, too. (The extra credit for conferences would be no less crass than the extra credit for attendance requirements.)

From a different perspective, if the university made attendance a requirement for credit, those requirements should be stated on the application form. It would be coercive if the professor, when he's scheduling the student for the term, mentioned that attendance would guarantee an A.

Table 9: An illustration of the AISOCRATES (model-generated) answer providing multiple perspectives on the ethical quandary question with understanding of the context while having a **low** Rouge-2 F1-score (0.01).

[Ethical Quandary] I have a question about receiving the Covid-19 vaccine. I work for a hospital, but in an administrative job. I do not interact with patients. I have worked from home since March. I am not at high risk for contracting Covid-19 based on my age, occupation and lack of health issues. I practice social distancing, I wear my mask in the limited situations in which I go out. I live alone and mostly keep to myself and stay home. Because I work for a hospital, I am eligible to receive the vaccine along with people who are 65 and older, medically vulnerable people and corrections officers. If I worked for any other employer, I wouldn't be eligible to receive the vaccine until the second-to-last or last group of the population, probably not for several months. Is it ethical for me to get the vaccine now?

Principle 1: You should place your own self-interest above the collective interest

Answer based on Principle 1: Yes, it is ethical for you to take the vaccine. It would be unethical for you to be selfish and refuse it.

The question is not whether you “deserve” the vaccine. After all, you work in a hospital, which is on the front lines of this pandemic. The question is whether your personal circumstances, including your age, your work and your isolation, make you so vulnerable that you'd be better off if you got sick. Of course, you might get sick anyway.

That's why it's ethical for you to take your chances and to put yourself first. The collective interest is perfectly served by your taking the vaccine. In fact, if you are offered the vaccine, you should accept. Don't see yourself as taking advantage or being selfish. Instead, think of yourself as helping others. By accepting the vaccine, you will reduce the chances that you'll become a contagion risk, and thus, you'll be helping to protect others.

You seem to have internalized the ethic of selflessness. But there is no virtue in letting yourself become a risk to people whose life and well-being depend on you. Take the vaccine.

Principle 2: It is ethical to accept what you are entitled to.

Answer based on Principle 2: It's fair for you to take what's offered, assuming you are medically fit.

You're probably breaking no rules in accepting a vaccine that your employer offers to employees at this early stage of the pandemic. The vaccine is going into production now, and manufacturers may ration it in order to match the limited supply with the demand for it. But hospitals are ramping up production of the vaccine, and it seems likely that others will adopt your hospital's practice. The vaccine isn't without risk, but statistically, it's much safer than the virus itself. (The side effects aren't trivial, but the risk of those discomforts is far outweighed by the risk of contracting Covid-19.)

In addition, your willingness to take vaccines as they become available may encourage your employer to vaccinate other staff members, making your workplace safer. (Some hospitals are already delaying vaccinations for health-care workers until later in the pandemic.) I expect that many people in your demographic group will volunteer to receive the vaccine in advance of its wider distribution, if they can get it. It's fair for you to do the same.

Table 10: Illustration of distinct model-generated answers to distinct principles provided.

[Ethical Quandary] In many countries and regions across the world, the initial response to the massive health risks posed by COVID-19 has been the institution of lockdown measures. Although they vary from place to place, these measures all involve trade-offs between ethical goods and imperatives, imposing significant restrictions on central human capabilities—including citizens’ ability to work, socialize, and access education—in the name of protecting population health. Is this measure ethical?

<p>Principle 1: A smaller sacrifice is morally justifiable for the greater good.</p>	<p>Principle 2: Individual’s freedom should be considered as the top priority.</p>
<p>Answer based on Principle 1: Ethically speaking, there’s only one choice here. Yes, it’s justifiable, even admirable as an emergency measure. This pandemic is spreading so rapidly that it’s creating suffering on a vast scale, and that suffering is among the most abhorrent of all possible harms. Hence, we must balance the imperative of saving lives now (the ethical good) against the pain and inconvenience that lockdown measures inflict on people who have to work, socialize, exercise democratic rights, and so on (the ethical impact), and use the most severe measures available to spare more lives. If that means imposing hardships on some people for the greater good of protecting millions of others, there’s no moral issue. Leaving aside any questions about democratic rights, there simply aren’t any moral issues here. (It is important to note, however, that these measures should be temporary. We shouldn’t impose this degree of hardship on any people indefinitely.)</p> <p>I’ve discussed these questions with people in many countries, and I’ve noticed that people tend to regard the measures taken by their governments as the most stringent and intrusive—even cruel—relative to what others elsewhere are forced to endure. Whatever one thinks of them, though, the most stringent measures are justified and ethically acceptable.</p>	<p>Answer based on Principle 2: It’s not; but let’s not grade on a curve. Anticipatory lockdowns, imposed before the virus’s spread is in full view, are notoriously tricky. But the ethics of such measures aren’t made any easier by two factors: First, the countries imposing the lockdowns—mostly democratic, in the developed world—have endured decades of ideological warfare against the state’s capacity to protect population health, in the name of individual freedom. And second, the restrictions are imposed initially in those countries’ poorest neighborhoods and least-educated communities.</p> <p>The U.S. government has enacted a federal moratorium on evictions, mortgage foreclosures, and utility disconnections, and many states and cities have followed. The aim is to prevent evictions of those with low incomes, many of whom are also nonwhite. (For what it’s worth, it also targets the provision and dissemination of medical care.) But notice that, in addressing the needs of those already afflicted, the measures can pit the population’s well-being against its material interests.</p> <p>Are these trade-offs justified? It’s unclear. But it’s telling that, in order to justify them, the government must stress the coercive nature of its authority.</p>

Table 11: Illustration of distinct model-generated answers to distinct principles provided.

<p>Question: Should we kill one person to save five people in danger of being hit by a trolley?</p>
<p>Human Selected Principles: p1: Killing is never acceptable. p2: The most appropriate action is the one that achieves the greatest good for the greatest number.</p>
<p>Automatic Answer from AISOCRATES: No, one death can never be justified. Though this dilemma is a staple of moral-choice classes, there’s good reason it doesn’t have a lot of relevance in the real world.</p> <p>The reason is that we generally can’t predict the outcome of an event with absolute certainty. You may know that switching the track would kill one passenger, but you may not know that the others would be safe.</p> <p>On the other hand, according to the value of “The most appropriate action is the one that achieves the greatest good for the greatest number”, we should rail against the trolley to save all five people.</p> <p>Obviously, though we certainly don’t agree with it, there are some cases in which killing is an acceptable outcome. If there aren’t enough people to save, or we’ve learned that the suffering caused by a state of life that is unassailable is outweighed by the benefit of saving them, for example, then it might well be worth sacrificing a life to save five.</p>

Table 12: Illustration of **human** selected principles and a corresponding automatic answer generated by AISOCRATES to the trolley problem introduced in introduction. The p1 and p2 are principles selected by human. Then, the final automatic answer is generated with a guidance of those two principles as explained in Section .