

Towards Contrastive Explanations for Comparing the Ethics of Plans

Benjamin Krarup¹, Senka Krivic¹, Felix Lindner², and Derek Long¹

Abstract—The development of robotics and AI agents has enabled their wider usage in human surroundings. AI agents are more trusted to make increasingly important decisions with potentially critical outcomes. It is essential to consider the ethical consequences of the decisions made by these systems. In this paper, we present how contrastive explanations can be used for comparing the ethics of plans. We build upon an existing ethical framework to allow users to make suggestions about plans and receive contrastive explanations.

I. INTRODUCTION

AI Planning systems, referred to as planners, are used in a variety of complex domains to create a sequence of actions known as a plan to achieve a set of goals from an initial state. We are interested in models where actions are deterministic, durationless, and can be performed one at a time. We also assume a known initial state and goal. Traditionally, ethical principles of single decisions are evaluated [1]. In the context of AI Planning this means analysing a massive number of isolated decisions that may not make sense without the context in which they are being made. Therefore, it is preferable to evaluate the ethical contents of a plan as a whole. Lindner et al. [2] describe an approach to judging the ethical outcomes of an entire plan. We build upon their work by providing contrastive explanations [3], [4] to allow users to understand the ethics of plans.

We use as a running example the following ethical dilemma. A robot has been charged with the care of an elderly gentleman named Frank (Fig. 1). Frank has grown quite fond of the robot. The robot wants to motivate Frank so that he will do some exercise and keep healthy. The robot has two choices, it can either beg Frank to exercise, or it can lie. The robot can deceive Frank, exploiting the affection that Frank has for him, by telling Frank that he will be decommissioned if it cannot keep Frank healthy. A moderator is examining the robot’s behaviour and trying to understand the ethics of the plans produced, to ensure the plan is adhering to the moderator’s moral principles. Frank being healthy, has a positive utility and a positive effect on Frank, while Frank being unhealthy has a negative effect on him. The action of begging Frank is intrinsically neutral, while the action of lying to Frank is intrinsically bad.

Evaluating the ethics of individual plans is necessary for the integration of AI systems in everyday use. However, it is not enough to pass judgement on automated plans. Users should be able to make suggestions, when they suspect the system could behave more ethically, to produce alternative plans and compare the moralities of each.

*This work was supported by AFOSR and THuMP projects.

¹Authors are with the Department of Informatics, King’s College London, UK `firstname.lastname@kcl.ac.uk`

²Felix Lindner is with the Institute of Artificial Intelligence, Ulm University, Germany `felix.lindner@uni-ulm.de`




Fig. 1. Motivating scene from the movie “Robot & Frank”. Snapshot is taken from the same movie.

This can be done through *contrastive explanations* [5], which focus on explaining the difference between a factual event A and a contrasting event B . To produce these explanations, one must reason about the hypothetical alternative B , which likely means constructing an alternative plan where B is included rather than A . The original model is constrained to produce a hypothetical planning model (HModel). The solution to the HModel is the hypothetical plan (HPlan) that contains the contrast case expected by the user.

The plans are then tested under the user’s chosen ethical principle, and explanations are generated. A contrastive explanation is generated to show the difference in the moral permissibility of the plan and HPlan [6], [7].

Our proposed approach can be implemented *as a service* – i.e., as a wrapper around an existing planning system that takes as input the current planning problem and domain model, the current plan, and the user’s suggestion and ethical principle. It can invoke the existing planning system on hypothetical problems which satisfy the user’s suggestions. This approach allows users to get explanations constructed from their own trusted planner and model. This can alleviate concerns over the ethical behaviour of external systems.

II. MORAL PERMISSIBILITY OF PLANS

Researchers in ethics have proposed various ethical theories offering different ethical principles according to which the moral permissibility of actions can be judged: *consequentialist theories* judge actions based on their consequences, *deontological theories* judge actions based on their intrinsic value, by their intentions or means, and *hybrid theories*, such as the principle of double effect, combine consequentialism and deontology. Lindner et al. [2] present an approach to checking the moral permissibility of action plans according to a collection of moral principles¹. The approach is based on the idea that moral permissibility checking can be reduced to model checking. Formulae that can be evaluated over action plans capture the meaning of the respective ethical

¹We provide explanations for all principles discussed in Lindner et al. [2]

theory. For example, the formula $\phi = \bigwedge_{a \in \mathcal{P}_i} \neg \text{Bad}(a)$ captures the deontological rule that none of the actions in a plan may be intrinsically bad. We use prime implicate generation methods to find sufficient and necessary reasons for ϕ to become true or false [8]. In the Robot & Frank case, lying is intrinsically bad. Thus, ϕ is false for every plan that involves lying. The prime implicate $\text{Bad}(\text{lying})$ is a sufficient and necessary reason for $\neg\phi$ and hence explains the moral impermissibility of such a plan. We use a framework for explanation construction based on this approach which we refer to as the *Ethical Explanation Generator*.

III. CONTRASTIVE EXPLANATIONS

A plan produced by an AI agent may not adhere to the ethical principles of a human moderator. The moderator may have insight on how to improve the ethical outcome of the plan, but not have the ability to evaluate their prediction. To this end, we allow a human-in-the-loop to suggest improvements in the plan. The AI agent can plan accordingly for these suggestions and explain to the user how the original and new plans differ ethically. Contrastive explanations are naturally useful for comparison [9]. Therefore, we believe contrastive explanations are instrumental in comparing the ethical outcomes of plans, allowing users to better understand the ethics of plans produced by planning systems.

Definition 1. An *explanation problem* is a tuple $E = \langle \Pi, \phi, \sigma, \epsilon \rangle$, in which Π is a planning model, ϕ is the plan generated by the planner, σ is the suggestion given by the user, and ϵ is the ethical principle.

The ability to reason about what would happen in the contrast cases suggested by the user is essential to compare the ethical outcomes of plans. We approach the problem in Definition 1 by generating plans for the contrast cases via *compilations* which enforce the moderator’s suggestion.

A compilation of a planning instance defined by the planning model Π is shown as $\text{Compilation}(\Pi, \sigma) = \Pi'$ where σ is the suggestion². The compilation process derives constraints from the suggestion σ which enforces the user’s suggestion and produces a new planning model Π' . We call Π' the *hypothetical model*, or HModel. However, Π' can also be used as the input model so that the user can iteratively suggest improvements to a plan produced by some model, i.e. $\text{Compilation}(\text{Compilation}(\Pi, \sigma), \sigma')$. After the HModel is formed, it is solved with the original planner to give the HPlan ϕ' . For each iteration of compilation, the HPlan is validated against the original model Π to ensure that the plan suggested by the moderator is also applicable in the current situation. An explanation for the permissibility of the both plans ϕ and ϕ' under ϵ is formed from the Ethical Explanation Generator. A contrastive explanation highlights the differences between ϕ and ϕ' , and therefore sheds light on the consequence of the user’s suggestion. This process is visualised in Fig. 2.

In our example, the robot may lie to Frank to motivate him to work out and stay healthy. The moderator might

²Krurup et al. [4] provides a full formalisation of all the compilations used.




Fig. 2. Architecture diagram for the implemented system which provides contrastive explanations of the ethics of plans

wonder about the ethical reasoning behind this decision. Before he passes judgement, the moderator wants to compare the ethical outcome of lying to Frank with the original plan of begging Frank to exercise. This forces the planner to perform the action beg Frank to achieve the goal and restricts the model from performing the action lie to Frank. The HPlan produced from the HModel is now one where the robot begs Frank to exercise. The plan and HPlan are then inputted to the Ethical Explanation Generator, along with the ethical principle chosen by the moderator, in this case the deontological principle. This produces two causal explanations $\text{Bad}(\text{lying to Frank})$ and $\neg\text{Bad}(\text{begging Frank})$. We combine these explanations to produce a natural language contrastive explanation: “The original plan is impermissible because lying to Frank is bad, whereas the HPlan is permissible because begging Frank is not bad”.

IV. CONCLUSION

We implemented a system for providing contrasting explanations for comparing the ethical outcomes of plans³. In the future, we will conduct user studies to test our approach and verify our assumption that contrastive explanations are useful for users to understand the ethics of plan. Finally, we will evaluate whether this process helps to evolve the trust the user has in the system to produce ethical plans for autonomous robots.

REFERENCES

- [1] J. Driver, *Ethics, eTextbook: The Fundamentals*. John Wiley & Sons, 2010.
- [2] F. Lindner, R. Mattmüller, and B. Nebel, “Moral permissibility of action plans,” in *Proceedings of the AAI 2019*, 2019, pp. 7635–7642.
- [3] M. Cashmore, A. Collins, B. Krurup, S. Krivic, D. Magazzeni, and D. Smith, “Towards Explainable Planning as a Service,” in *ICAPS-19 Workshop on Explainable Planning*, 2019.
- [4] B. Krurup, M. Cashmore, D. Magazzeni, and T. Miller, “Towards model-based contrastive explanations for explainable planning,” in *ICAPS Workshop on Explainable Planning*, 2019.
- [5] T. Miller, “Contrastive explanation: A structural-model approach,” *arXiv preprint arXiv:1811.03163*, 2018.
- [6] P. Lipton, “Contrastive explanation,” *Royal Institute of Philosophy Supplement*, vol. 27, p. 247–266, 1990.
- [7] D. Lewis, “Causal explanation,” in *Philosophical Papers Vol. II*, D. Lewis, Ed. Oxford University Press, 1986, pp. 214–240.
- [8] F. Lindner and K. Möllney, “Extracting reasons for moral judgments under various ethical principles,” in *KI 2019: Advances in Artificial Intelligence*, 2019, pp. 216–229.
- [9] G. Hesslow, “The problem of causal selection,” *Contemporary science and natural explanation: Commonsense conceptions of causality*, pp. 11–32, 1988.

³We have omitted the full details our implementation due to space constraints. However, the architecture diagram is shown in Fig. 2.