

On Generating Monolithic and Model Reconciling Explanations in Probabilistic Scenarios

STYLIANOS LOUKAS VASILEIOU*, New Mexico State University, United States

WILLIAM YEOH, Washington University in St. Louis, United States

ALESSANDRO PREVITI, Ericsson Research, Sweden

TRAN CAO SON, New Mexico State University, United States

Explanation generation frameworks aim to make AI systems' decisions transparent and understandable to human users. However, generating explanations in uncertain environments characterized by incomplete information and probabilistic models remains a significant challenge. In this paper, we propose a novel framework for generating *probabilistic monolithic explanations* and *model reconciling explanations*. Monolithic explanations provide self-contained reasons for an explanandum without considering the agent receiving the explanation, while model reconciling explanations account for the knowledge of the agent receiving the explanation. For monolithic explanations, our approach integrates uncertainty by utilizing probabilistic logic to increase the probability of the explanandum. For model reconciling explanations, we propose a framework that extends the logic-based variant of the model reconciliation problem to account for probabilistic human models, where the goal is to find explanations that increase the probability of the explanandum while minimizing conflicts between the explanation and the probabilistic human model. We introduce *explanatory gain* and *explanatory power* as quantitative metrics to assess the quality of these explanations. Further, we present algorithms that exploit the duality between minimal correction sets and minimal unsatisfiable sets to efficiently compute both types of explanations in probabilistic contexts. Extensive experimental evaluations on various benchmarks demonstrate the effectiveness and scalability of our approach in generating explanations under uncertainty.

JAIR Associate Editor: Kuldeep Meel

JAIR Reference Format:

Stylianios Loukas Vasileiou, William Yeoh, Alessandro Previti, and Tran Cao Son. 2025. On Generating Monolithic and Model Reconciling Explanations in Probabilistic Scenarios. *Journal of Artificial Intelligence Research* 84, Article 5 (September 2025), 40 pages. DOI: [10.1613/jair.1.18820](https://doi.org/10.1613/jair.1.18820)

1 Introduction

The rapid integration of artificial intelligence (AI) into critical and everyday applications has magnified the importance of not just achieving high-performance metrics but also ensuring that AI decisions are transparent, interpretable, and, above all, trustworthy. This imperative has given rise to the field of explainable AI (XAI), which seeks to make AI systems' workings comprehensible to their human users (Gunning et al. 2019). XAI endeavors to demystify the often opaque processes of AI, providing insights into the reasoning behind decisions

*Corresponding Author.

Authors' Contact Information: Stylianios Loukas Vasileiou, stelios@nmsu.edu, New Mexico State University, Las Cruces, New Mexico, United States; William Yeoh, wyeh@wustl.edu, Washington University in St. Louis, St. Louis, Missouri, United States; Alessandro Previti, alessandro.previti@ericsson.com, Ericsson Research, Stockholm, Sweden; Tran Cao Son, stran@nmsu.edu, New Mexico State University, Las Cruces, New Mexico, United States.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.18820](https://doi.org/10.1613/jair.1.18820)

and actions. This transparency is not just a matter of ethical AI design but a practical necessity for enhancing user trust, facilitating user decision-making, and ensuring the accountability of AI systems.

In the domain of machine learning (ML), significant strides have been made towards enhancing the explainability of algorithms. Researchers have sought to categorize ML algorithms according to various dimensions of explainability (Guidotti et al. 2018), improve the transparency of existing algorithms (Alvarez Melis and Jaakkola 2018; Petkovic et al. 2018), and even propose new algorithms that balance accuracy with increased explainability (Dong et al. 2017; Gilpin et al. 2018). These efforts underscore a growing recognition within the ML community of the critical role that explainability plays in the deployment of AI systems (Belle and Papantonis 2021).

Parallel to advancements in ML, the automated planning community has adopted a focused approach to generating explanations for plans produced by AI (planning) agents, which led to the inception of explainable AI planning (XAIP) (Fox et al. 2017). Predominantly, XAIP research focuses on the explanation generation problem, which involves identifying explanations for plans that, when conveyed to human users, help them understand and accept the agent’s proposed actions (Kambhampati 1990; Langley 2016). A noteworthy direction within this space is the *model reconciliation problem* (MRP) (Chakraborti, Sreedharan, Zhang, et al. 2017; Nguyen et al. 2020; Sreedharan et al. 2018; Vasileiou, Previti, et al. 2021; Vasileiou, Yeoh, et al. 2022), aimed at aligning a user’s model with that of an AI agent through the provision of explanations, especially when discrepancies in their understanding of a planning problem lead to confusion or misinterpretation of the agent’s decisions. We will refer to these explanations as *model reconciling explanations*. In contrast, we refer to explanations that do not account for the user’s model as *monolithic explanations*.

Although MRP tackles essential facets of explainability, it often operates under the presumption that the AI agent has a deterministic grasp of the human model—a scenario that may not always align with the complexities of real-world interactions characterized by uncertainty about human knowledge. Indeed, this gap highlights a general challenge in explanation generation from AI agents: Their operation within realms of incomplete information and probabilistic decision-making models. Traditional explanation methods, which rely on deterministic knowledge, falter under these conditions, unable to represent the uncertain nature of the agent’s knowledge adequately.

To bridge this gap in explainability, the first part of the paper presents a framework for explanation generation under uncertainty. In particular, we use (propositional) probabilistic logic as our underlying mechanism for modeling uncertainty, and introduce the notions of *probabilistic monolithic explanation* and *probabilistic model reconciling explanation*. For the former, given an agent’s belief base \mathcal{B}_α (a weighted knowledge base¹), and an explanandum φ (a formula to be explained), the goal is to find a *probabilistic monolithic explanation* such that the probability of the explanandum being true is increased. Moreover, we extend our work on *logic-based model reconciliation problems* (L-MRPs) (Son et al. 2021; Vasileiou, Previti, et al. 2021; Vasileiou, Yeoh, et al. 2022) to scenarios in which the human user’s model is uncertain. Specifically, given a knowledge base KB_α of an agent, an explanandum φ entailed by KB_α , and a human belief base \mathcal{B}_h , the goal is to find a *probabilistic model reconciling explanation* such that the explanandum’s probability is increased while the probability of conflicts between the explanation and \mathcal{B}_h is decreased. To measure the quality of such explanations, we define the concepts of *explanatory gain* and *explanatory power*, aimed at quantifying the effectiveness of these explanations with respect to the explanandum.

In the second part of the paper, we describe algorithms for computing both types of explanations, where we leverage the duality between *minimal correction sets* (MCSEs) and *minimal unsatisfiable sets* (MUSEs). These algorithms, adapted from our prior work (Vasileiou, Previti, et al. 2021), integrate a *weighted maximum satisfiability* procedure for computing probabilities. Our experimental evaluation across various benchmarks highlights the

¹A weighted knowledge base, essentially, induces a probability distribution over the variables of the agent’s language (Richardson and Domingos 2006).

practicality and applicability of our algorithms, and demonstrate their capabilities to generate probabilistic explanations within a propositional logic framework.

In summary, our contributions are the following:

- We propose a novel framework for generating probabilistic monolithic explanations and probabilistic model reconciling explanations. Central to our framework are the concepts of *explanatory gain* and *explanatory power*, metrics designed to quantitatively assess the quality and effectiveness of probabilistic explanations.
- We describe algorithms for computing both types of probabilistic explanations, using the duality of MCSes and MUSes together with a weighted maximum satisfiability process. Through a series of benchmark evaluations, we demonstrate the efficacy of our proposed algorithms in generating explanations.

The paper is structured as follows. In Section 2, we motivate the use and applicability of logic in explainability, and discuss possible sources of uncertainty in explanation generation. In Section 3, we provide the background knowledge needed, and in Section 4 we describe a motivating application that serves as a running example. In Section 5, we present our explanation generation framework for monolithic and model reconciliation probabilistic explanations. We present algorithms for computing explanations in Section 6, and experimentally evaluate them on a set of benchmarks in Section 7. We then discuss related work in Section 8, address the assumptions, limitations and future extensions of our framework in Section 9, and finally conclude the paper in Section 10.

2 A Logic-based Perspective on Explainable Decision-Making

In this work, we take the following perspective: *Logic-based frameworks can serve as an explanatory representational layer for AI systems, enabling the generation of rigorous and flexible explanations across diverse problem domains by capturing the system's decisions in a formal logical language that supports inference and reasoning.*

According to this perspective, formal logic provides a robust foundation for creating explanatory mechanisms by serving as an intermediate representational layer between AI systems and explanation generation processes. Some of the reasons that justify this perspective are as follows:

- **Structured Semantics:** Logic has well-defined compositional semantic functions that compute the meaning of a compound as a function of its constituents' meanings. This composition is invertible, allowing us to trace back from conclusions to premises—essentially performing inference in reverse. This property is fundamental to explanation, as it enables us to identify precisely which components contributed to a particular decision or output.
- **Expressivity & Scrutability:** Logic can represent complex knowledge and reasoning processes in a form that can be traced and scrutinized. Depending on the formal language of choice, it allows for encoding rich, relational information about the world, including classes, hierarchies, causal relationships, and quantified statements. For instance, in planning domains, logical representations can express not only states and actions but also the causal links between them, making it possible to explain why certain actions were selected in a plan. Moreover, logic-based systems enable us to examine their internal properties, both through internal verification techniques and external dialogues. This scrutability is crucial for building trust in AI systems, as it allows for thorough validation of the reasoning process, ensuring that explanations are not just plausible but provably correct within the system's logical framework.
- **Adaptability:** Logic-based systems can be extended with new knowledge and can integrate with various AI paradigms. This means that explanatory frameworks built on logic can evolve over time, incorporating new concepts, rules, or meta-level reasoning principles as needed. This adaptability is important for explanation systems that must operate across diverse domains or interact with different types of users.

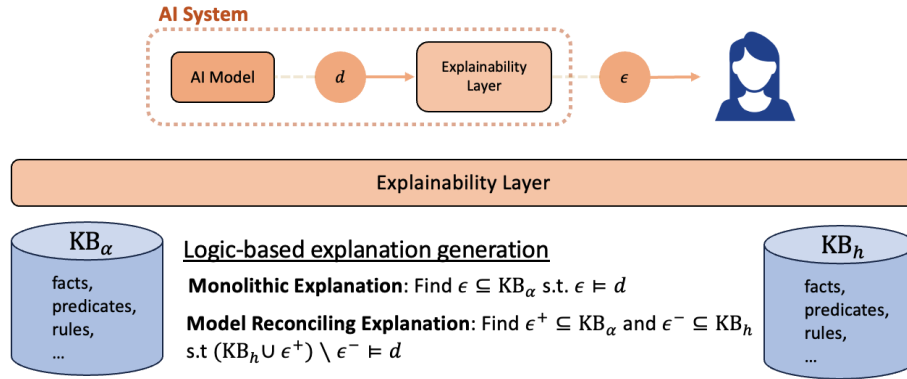


Fig. 1. The Logic-based Explainability Layer. Given a decision d from an AI model, a *monolithic explanation* is generated with respect to the AI system’s knowledge base (KB_α), while a *model reconciling explanation* is generated with respect to the AI system’s (KB_α) as well as the human user’s knowledge base (KB_h).

- **Uncertainty Representation:** While classical logic is deterministic, probabilistic extensions (which we use in this work) allow for representing and reasoning about uncertainty. This is important in real-world scenarios where AI systems must operate with incomplete information or where multiple interpretations of data are possible. Probabilistic logic enables explanations that convey not just what the system “believes”, but also its degree of confidence in those beliefs.

For more about logic and its suitability in explainability as well as the current AI landscape, please see some good arguments presented by Belle (2017), Mocanu and Belle (2023), and Belle (2025).

Now, as illustrated in Figure 1, our approach positions logic as an intermediate explanatory layer between AI systems and human users. The logic-based explainability layer functions as an abstraction mechanism that captures the essential reasoning behind an AI system’s decisions in a formal logical representation, regardless of the underlying implementation details of the AI system itself. This layer serves two critical functions. First, it translates the internal decision processes of potentially opaque AI systems into a logical formalism that can be systematically analyzed. Second, it enables the generation of explanations through formal reasoning over this logical representation.

We will distinguish between two notions of explanation: *monolithic explanations*, which are generated with respect to the AI system’s knowledge base (KB_α) alone, focusing on providing a self-contained justification for the system’s decision; and *model reconciling explanations*, which are generated with respect to both the AI system’s knowledge base (KB_α) and the human user’s knowledge base (KB_h), explicitly addressing the discrepancies between the two models that may lead to confusion or misunderstanding. The explainability layer thus provides a unified framework for generating different types of explanations while maintaining a clear separation between the underlying AI system and the explanation mechanism.

The effectiveness of the logic-based explainability framework depends crucially on the ability to encode the problem domains of interest in logical formalisms. In general, it can be used for problems that can be represented in a logical language for which satisfiability of sets is feasible. Fortunately, many problem domains admit logical encodings, such as planning problems (Cashmore, Fox, et al. 2012; Cashmore, Magazzeni, et al. 2020; Kautz et al. 1996), scheduling problems (Ansótegui et al. 2011; Crawford and Baker 1994; Demirović et al. 2019), argumentation problems (Besnard, Grégoire, et al. 2010; Besnard and Hunter 2001; Prakken 2006), as well as machine learning

problems (Izza et al. 2023; Marques-Silva and Ignatiev 2022; Shrotri et al. 2022). Indeed, logic-based explanation generation approaches have been successfully employed in numerous applications (Marques-Silva and Ignatiev 2022; Rago et al. 2023; Vasileiou, Previti, et al. 2021; Vasileiou, Xu, et al. 2023; Vasileiou, Yeoh, et al. 2022).

2.1 The Source of Uncertainty in Explanations

Traditional logic-based approaches have typically operated within deterministic settings. They assume that both the AI system and the human user have definite, certain knowledge about the domain. This assumption, however, fails to capture the reality of many real-world scenarios where uncertainty plays a crucial role. Succinctly,

- *Incomplete Information*: AI systems often operate with partial information about the environment, meaning they operate with probabilistic rather than deterministic knowledge.
- *Ambiguous Evidence*: In domains like medical diagnosis or law, the available evidence may support multiple hypotheses to varying degrees, requiring probabilistic reasoning.
- *Uncertain Human Models*: Human users may hold beliefs with varying degrees of confidence, and AI systems may be uncertain about what humans know or believe.
- *Probabilistic Effects*: In more realistic, dynamic environments, actions may have probabilistic outcomes, requiring explanations that account for these uncertainties.

Motivating Example. To ground our discussion, consider a service robot operating in an office building with some uncertain information about its environment. The robot must navigate from its current location to deliver a package to a specific office, choosing among multiple possible routes. In this scenario, uncertainty may arise due to several reasons, such as environmental uncertainty, e.g., the robot has limited information about the current state of different corridors (e.g., crowded or clear), outcome uncertainty, e.g., even with a chosen path, the robot may not be able to perfectly predict travel times due to potential obstructions or changing conditions, and uncertainty about the beliefs of humans, i.e., when explaining its decisions to building staff, the robot may not know exactly what they believe about current building conditions.

Now, assume that the robot chooses a longer route through corridor B instead of a shorter path through corridor A. When asked why it did not take the more direct route through corridor A, the robot could provide different (monolithic) explanations. For instance, “Corridor A likely has high foot traffic at this time of day, which would impede movement and potentially require frequent stops”, or “There is an ongoing maintenance operation in corridor A that creates an obstruction risk and would likely result in a significant detour mid-route.” Both explanations may increase our understanding of why the robot avoided corridor A, but they focus on different aspects of the robot’s belief state and might have different explanatory power.

For the model reconciliation case, consider a human supervisor who expected the robot to take corridor A because they believe the corridor should be clear at this time. This discrepancy arises because the supervisor and the robot have different beliefs about the current state of the building. Then, a model reconciling explanation would need to address these specific belief differences and help the supervisor understand why the robot’s decision was reasonable given its information.

This example illustrates how uncertainty naturally arises in practical scenarios and how multiple potential explanations must be compared to identify those with the highest explanatory power.

3 Background

In this section, we provide some background for propositional logic, the hitting set duality between minimal unsatisfiable and minimal correction sets, modeling uncertainty in propositional logic, and the model reconciliation problem.

3.1 Propositional Logic

Let \mathcal{L} be a propositional language built from a finite set of atomic variables $\mathcal{V} = \{a, b, c, \dots\}$. A *possible world* is a truth-value assignment to each variable $\omega : \mathcal{V} \mapsto \{T, F\}$, where T and F denote truth and falsity respectively. The set of all possible worlds of \mathcal{L} is denoted by Ω . The simplest formulae in \mathcal{L} are atoms: Individual variables that may be true or false in a given possible world. More complex formulae are recursively constructed from atoms using the classical logical connectives. A *model* of a formula is a possible world in which the formula is satisfied (i.e., evaluates to true). A knowledge base KB is a set of formulae. If there exists at least one possible world ω that satisfies all formulae in KB, then KB is *consistent*, otherwise we say that KB is *inconsistent*. We use \models to denote the classical entailment relation and say that a (consistent) KB entails a formula φ , expressed as $\text{KB} \models \varphi$, if and only if every model of KB is also a model of φ , or equivalently, if $\text{KB} \cup \{\neg\varphi\}$ is inconsistent. Unless stated otherwise, it is assumed that all formulae are expressed in *conjunctive normal form* (CNF).²

Given a knowledge base KB and a formula φ , called the *explanandum* such that $\text{KB} \models \varphi$, we define a *monolithic explanation* for φ from KB as a minimal set of formulae that entails φ :

Definition 1. (*Monolithic Explanation*) Let KB be a knowledge base and φ an explanandum such that $\text{KB} \models \varphi$. We say that $\epsilon \subseteq \text{KB}$ is a monolithic explanation for φ from KB if and only if: (i) $\epsilon \models \varphi$; and (ii) $\nexists \epsilon' \subset \epsilon$ such that $\epsilon' \models \varphi$.

Example 1. Consider the knowledge base $\text{KB} = \{p, \neg p \vee q, \neg p \vee r\}$ built up from $\mathcal{V} = \{p, q, r\}$. Notice that $\text{KB} \models q$. Then, $\epsilon = \{p, \neg p \vee q\}$ is a monolithic explanation for q from KB.

3.1.1 Duality of Minimal Unsatisfiable and Minimal Corrections Sets.

Definition 2 (MINIMAL UNSATISFIABLE SET (MUS)). Given an inconsistent knowledge base KB, a subset $\mathcal{M} \subseteq \text{KB}$ is an MUS if \mathcal{M} is inconsistent and $\forall \mathcal{M}' \subset \mathcal{M}$, \mathcal{M}' is consistent.

Definition 3 (MINIMAL CORRECTION SET (MCS)). Given an inconsistent knowledge base KB, a subset $\mathcal{C} \subseteq \text{KB}$ is an MCS if $\text{KB} \setminus \mathcal{C}$ is consistent and $\forall \mathcal{C}' \subset \mathcal{C}$, $\text{KB} \setminus \mathcal{C}'$ is inconsistent.

By definition, every inconsistent KB contains at least one MUS.

Definition 4 (PARTIAL MUS). A set of formulae Φ is a *partial MUS* of an inconsistent knowledge base KB if there exists at least one MUS $\mathcal{M} \subseteq \text{KB}$ such that $\Phi \subseteq \mathcal{M}$.

Partial MUSes in an inconsistent knowledge base KB appear when a subset of formulae is set as *hard*, that is, formulae that must always be satisfied in a solution. Conversely, *soft* formulae may not always be satisfied. Given a formula φ , we will write φ^* with $*$ $\in \{s, h\}$ to denote it as *soft* and *hard*, respectively.

MUSes and MCSes are related by the concept of *hitting set*:

Definition 5 (HITTING SET). Given a collection Γ of sets from a universe U , a *hitting set* for Γ is a set $H \subseteq U$ such that $\forall S \in \Gamma, H \cap S \neq \emptyset$. A *hitting set* H is *minimal* if $\nexists H' \subset H$ such that H' is a hitting set.

The relationship between MUSes and MCSes is discussed by Liffiton and Sakallah (2008) and Liffiton, Previti, et al. (2016), and it was first presented by Reiter (1987), where MUSes and MCSes are referred to as (minimal) conflicts and diagnoses, respectively.

Proposition 1. A subset of an inconsistent knowledge base KB is an MUS (resp. MCS) if and only if it is a minimal hitting set of the collection of all MCSes (resp. MUSes) of KB.

²A CNF formula is a conjunction of clauses, where each clause is a disjunction of literals. A literal is either an atom or its negation. This is not a restrictive requirement, since any propositional formula can be transformed into a CNF representation.

It follows from the above proposition that a cardinality minimal MUS (resp. MCS) is a minimal hitting set. *Cardinality minimal MUS* are referred to as SMUS, whereas a *cardinality minimal MCS* corresponds to the complement of a MaxSAT solution (Li and Manya 2009). We may refer to a cardinality minimal set as a minimum or smallest set.

Lemma 1. *Given a subset \mathcal{H} of all the MCSes of knowledge base KB, a hitting set is an SMUS if: (1) It is a minimal hitting set h of \mathcal{H} , and (2) The subformula induced by h is inconsistent.*

See the work by Ignatiev, Previti, et al. (2015) for a proof.

Proposition 1 and Lemma 1 naturally extend to the case of partial MUS. Note that when some formulae are set as hard in an inconsistent knowledge base, the set of all MCSes is a subset of the soft formulae. In this case, every minimal hitting set on the set of all MCSes is a partial MUS.

Finally, MUSes and monolithic explanations are related by the following:

Proposition 2. *Given a knowledge base KB, a consistent set of formulae $\epsilon \subseteq \text{KB}$ is a monolithic explanation for φ from KB (Definition 1) if and only if ϵ is a partial MUS of $\epsilon \cup \{\neg\varphi\}$.*

Example 2. Let $\text{KB} = \{p, \neg p \vee q, \neg p \vee r\}$ and $\epsilon = \{p, \neg p \vee q\}$ from Example 1. Notice how $\mathcal{M} = \{p, \neg p \vee q, \neg q\}$ is an MUS of $\text{KB} \cup \{\neg q\}$. Then, it is easy to see that ϵ is a partial MUS of \mathcal{M} .

3.2 Modeling Uncertainty in Propositional Logic

Building on a propositional language \mathcal{L} , we can model the uncertainty of propositional formulae using a *probability distribution* over the possible worlds Ω of \mathcal{L} . Formally,

Definition 6 (PROBABILITY DISTRIBUTION). *Let Ω be the set of possible worlds of the language \mathcal{L} . A probability distribution P on Ω is a function $P : \Omega \mapsto [0, 1]$ such that $\sum_{\omega \in \Omega} P(\omega) = 1$.*

In essence, a probability distribution over possible worlds creates a *ranking* between those worlds with respect to how likely they are to be true. This then allows us to quantify the uncertainty in a formula as follows:

Definition 7 (DEGREE OF BELIEF). *Let Ω be the set of possible worlds and P a probability distribution over Ω . The degree of belief in a formula $\varphi \in \mathcal{L}$ is $P(\varphi) = \sum_{\omega \models \varphi} P(\omega)$.³*

We may refer to $P(\varphi)$ as degree of belief or probability of φ interchangeably. Note that the possible worlds approach to probabilities is essentially equivalent to probabilities assigned directly to the formulae (Bacchus 1990).

Now, the probability distribution on Ω can be induced from a weighted knowledge base, referred to as a *belief base*:

$$\mathcal{B} = \{(\phi_1, w_1), \dots, (\phi_n, w_n)\} \quad (1)$$

where each formula $\phi_i \in \mathcal{L}$ is associated with a weight $w_i \in \mathbb{R}^+$.⁴

Intuitively, the weights serve as meta-information and reflect the certainty about the truth of the corresponding formulae—the higher the weight, the more certain the formula is. In that sense, formulae with higher weights are prioritized for satisfaction, effectively capturing the certainty of the particular formulae. This mechanism is especially useful for handling inconsistency and non-monotonic reasoning patterns, thus capturing a broader

³Note that, with a slight abuse of notation, we use $\omega \models \varphi$ to denote that φ is true (e.g., satisfied) in world ω .

⁴We assume, without loss of generality, that all weights are non-negative because a formula with a negative weight w can be replaced by its negation with weight $-w$.

spectrum of problems.⁵ Further, we will denote with $\mathcal{B}^{\downarrow w}$ the *classical projection* of \mathcal{B} , that is, $\mathcal{B}^{\downarrow w} = \{\phi_i \mid (\phi_i, w_i) \in \mathcal{B}\}$.

Given a belief base \mathcal{B} , one way to induce a probability distribution is the following:

$$\forall \omega \in \Omega, P_{\mathcal{B}}(\omega) = \frac{1}{Z} \exp \left(\sum_{i=1}^n w_i \cdot \mathbb{I}(\omega, \phi_i) \right) \quad (2)$$

where $\mathbb{I}(\omega, \phi) = 1$ if $\omega \models \phi$ and 0 otherwise, and $Z = \sum_{\omega \in \Omega} \exp \left(\sum_{i=1}^n w_i \cdot \mathbb{I}(\omega, \phi_i) \right)$ is the normalization factor.

The induced probability distribution quantifies the likelihood that a given (possible) world is the actual world. Higher formula weights amplify the (log-) probability difference between a world that satisfies the formula and one that does not, other things being equal. Consequently, worlds that violate fewer formulas are deemed more probable. Note that a belief base \mathcal{B} is essentially a *log-linear model* (Bishop et al. 2007), from which a *joint probability distribution* of the set of variables of \mathcal{L} is induced. Interestingly, log-linear models are special cases of Markov Logic Networks and can represent any positive distribution (Richardson and Domingos 2006). When taken from context, we will simply use P to denote the distribution induced from \mathcal{B} .

Entailment in a belief base KB becomes *graded*, that is, we now say that \mathcal{B} entails a formula ϕ with degree of belief $P(\phi)$. However, when all weights are equal and tend to infinity, a belief base represents a uniform distribution over the worlds that satisfy it and, as such, entailment of a formula can be answered by computing the probability of the formula and checking whether it is 1. In other words, entailment under belief bases collapses to classical entailment under knowledge bases. See Richardson and Domingos (2006) for a proof.

Finally, the weighted formulae in a belief base \mathcal{B} can be viewed as *soft* constraints in the sense described in Section 3.1.1. In contrast, *hard* constraints can be imposed as formulae with “infinite” weights.⁶

3.3 The Model Reconciliation Problem

The *Model Reconciliation Problem* (MRP), as introduced by Chakraborti, Sreedharan, Zhang, et al. (2017), highlights the critical need for aligning the planning models of a human user and an agent to facilitate effective collaboration and understanding. This alignment becomes especially pertinent in scenarios where the agent’s plan deviates from human expectations, necessitating a mechanism to reconcile these differences through explanations. In this approach, the (planning) agent must have knowledge of the human’s model in order to contemplate their goals and foresee how its plan will be perceived by them. When there exist differences between the models of the agent and the human such that the agent’s plan diverges from the human’s expectations, the agent provides a minimal set of model differences, namely a *model reconciling explanation*, to the human.

It is important to highlight that, in order to effectively solve MRP, the following (implicit) assumptions typically hold:

1. The *agent model represents the ground truth* or, in other words, the agent model is the “correct” encoding of the domain. This assumption is predicated on the notion that the explanation is generated from the agent’s perspective, thereby rendering it reasonable to assume that the agent “thinks” that its model is accurate or correct.

⁵For example, the notion of inconsistency is relaxed as follows: Given two formulae ϕ and $\neg\phi$ that contradict each other (e.g., $\{\phi, \neg\phi\}$ is inconsistent), if $P(\phi) = 0.9$, then from the axioms of probability we have that $P(\neg\phi) = 0.1$. This then means that the worlds where $\neg\phi$ is true are more unlikely than the worlds where ϕ is true, but not impossible.

⁶In practice, infinite weights can be replaced with $\sum_{i=1}^n w_i + 1$.

2. The *agent has access to the human model*, which is an approximation of the actual human model. In the worst case, it can be empty; but, practically, it can be approximated based on past interactions (Juba et al. 2021; Sreedharan et al. 2018).
3. *Both models are assumed to be deterministic*, and they thus are able to represent only deterministic domains. Note that while this is a restricting assumption, in Section 5.2 we present a framework that relaxes it.

Now, building upon the MRP foundation, we introduced its logic-based variant (L-MRP), where the models of the agent and the human user are represented as logical knowledge bases (Vasileiou, Yeoh, et al. 2022).⁷ As a *model reconciling explanation* must take into account both the knowledge base KB_α of the agent providing an explanation as well as the knowledge base KB_h of the human receiving the explanation, it is defined slightly differently compared to monolithic explanations defined by Definition 1:

Definition 8 (MODEL RECONCILING EXPLANATION). *Given the knowledge bases of an agent KB_α and a human user KB_h as well as an explanandum φ , such that $KB_\alpha \models \varphi$ and $KB_h \not\models \varphi$, $\mathcal{E} = \langle \epsilon^+, \epsilon^- \rangle$ is a model reconciling explanation if and only if $\epsilon^+ \subseteq KB_\alpha$, $\epsilon^- \subseteq KB_h$, and $(KB_h \cup \epsilon^+) \setminus \epsilon^- \models \varphi$.*

When KB_h is updated with a model reconciling explanation $\mathcal{E} = \langle \epsilon^+, \epsilon^- \rangle$, new formulae ϵ^+ from KB_α are added to KB_h and formulae ϵ^- from KB_h are retracted to ensure consistency. Note that since a model reconciling explanation is from the perspective of the agent's knowledge base KB_α , we implicitly assume that if a formula in KB_h is inconsistent with KB_α , then that formula is “false” from the perspective of the agent.

Example 3. Let $KB_\alpha = \{a, \neg a \vee b, \neg a \vee c\}$ and $KB_h = \{\neg a, \neg a \vee b\}$ be the knowledge bases of an agent and a human user, respectively, where $KB_\alpha \models b$ and $KB_h \not\models b$. A model reconciling explanation is then $\mathcal{E} = \langle \{a\}, \{\neg a\} \rangle$, where $(KB_h \cup \{a\}) \setminus \{\neg a\} = \{a, \neg a \vee b\} \models b$.

4 Motivating Application: Office Robot Delivery

To demonstrate our ideas, let us revisit the service robot scenario described earlier. Consider a robot operating in an office building with the task of delivering a package from its current location (Room 1) to a destination (Room 2). The robot can navigate through two possible corridors: corridor A, which is shorter but often crowded, or corridor B, which is longer but typically less crowded. Figure 2 illustrates a simplified abstraction of this problem in the form of a grid world.

The robot's action dynamics include two primary actions: move and deliver. The move action allows the robot to navigate between rooms through a specified corridor, subject to preconditions (the robot must be at the origin location) and probabilistic effects that depend on the corridor's crowding status, i.e., when attempting to move through a crowded corridor, the robot has a higher probability of failing to move and remaining in its current location. The deliver action can only be performed when the robot is at the destination room and results in the package being delivered. The complete problem specification and encoding is provided in the appendix.

We encode this probabilistic planning scenario as a belief base \mathcal{B} containing weighted formulae (with subscripts denoting time steps) that represent the robot's beliefs about the environment states and actions. For instance, the initial states and the action dynamics for the move action (for $t = 0$) are:

Initial state beliefs:

$$\begin{aligned} &(\text{robot-at}(\text{room1})_0, \infty), \\ &(\neg \text{package-delivered}_0, \infty), \\ &(\text{crowded}(A), 3), \end{aligned}$$

⁷We have also introduced extensions that are complementary to the work presented in this paper (Vasileiou, Kumar, et al. 2024; Vasileiou and Yeoh 2023).

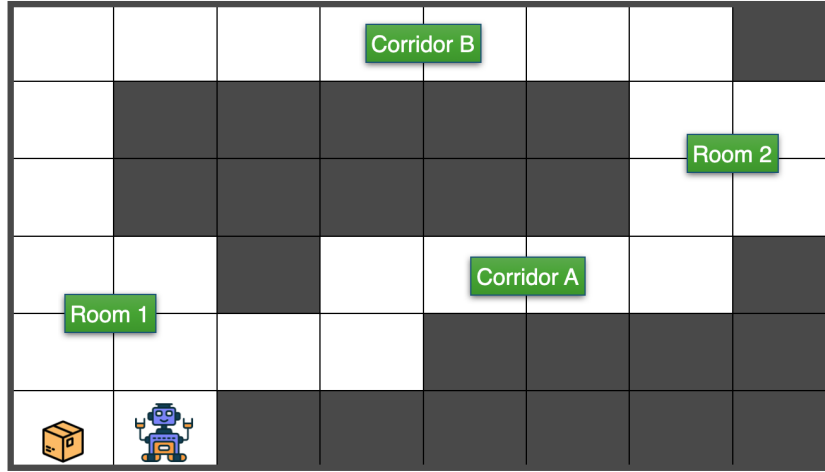


Fig. 2. Grid world illustration of the package delivery problem. The robot has to deliver the package from Room 1 to Room 2 using either Corridor A or Corridor B.

$(\text{crowded}(B), 0.5),$

Move preconditions:

$(\text{move}(\text{room1}, \text{room2}, A)_0 \rightarrow \text{robot-at}(\text{room1})_0, \infty),$
 $(\text{move}(\text{room1}, \text{room2}, B)_0 \rightarrow \text{robot-at}(\text{room1})_0, \infty),$

Move effects:

$(\text{move}(\text{room1}, \text{room2}, A)_0 \wedge \text{crowded}(A) \rightarrow \text{robot-at}(\text{room2})_1, 0.5),$
 $(\text{move}(\text{room1}, \text{room2}, A)_0 \wedge \text{crowded}(A) \rightarrow \text{robot-at}(\text{room1})_1, 3).$

In this belief base, the weights encode the robot's degrees of belief. For instance, the weight of 3 for $\text{crowded}(A)$ represents a strong belief that corridor A is crowded, while the weight of 0.5 for $\text{crowded}(B)$ indicates a weaker belief that corridor B is crowded. Similarly, the probabilistic action effects encode that when a corridor is crowded, it is less likely that the robot will move through it to the next room. Given its belief base \mathcal{B} , the robot computes that the optimal plan is $\pi = \langle \text{move}(\text{room1}, \text{room2}, B)_0, \text{deliver}_1 \rangle$, with probability $P(\pi) = 0.6$.

Now, suppose a human supervisor observes the robot's choice and asks: "Why didn't you take corridor A instead?" This query can be formalized as the explanandum $\varphi = \neg \text{move}(\text{room1}, \text{room2}, A)$, essentially encoding "not moving from Room 1 to Room 2 through corridor A", and allowing the robot to generate an explanation of why that is the case. In the following section, we will revisit this scenario and show how our methods can be used to generate probabilistic explanations that answer the query.

5 A Framework for Probabilistic Explanation Generation

In this section, we approach the notion of *probabilistic explanation* in the following setting, where \mathcal{B} is a belief base and P its induced probability distribution:

Definition 9 (PROBABILISTIC EXPLANATION GENERATION PROBLEM (PEGP)). *Given a belief base \mathcal{B} and an explanandum φ in logic \mathcal{L} , the goal of PEGP is to identify an explanation (i.e., a set of formulae) $\tilde{\epsilon} \subseteq \mathcal{B}$ such*

that the probability of the explanandum given the explanation is higher than the probability of the explanandum, i.e., $P(\varphi \mid \tilde{\epsilon}) > P(\varphi)$.

Under this definition, we account for the uncertainty inherent in knowledge bases, and seek to find explanations that not only identify contributing factors for an explanandum but also quantify the uncertainty in these factors. Intuitively, a solution $\tilde{\epsilon}$ to PEGP seeks to increase the degree of belief in the explanandum φ . If $P(\varphi \mid \tilde{\epsilon}) > P(\varphi)$, this then represents the case where $\tilde{\epsilon}$ increases the degree of belief in φ and the greater the value of $P(\varphi \mid \tilde{\epsilon})$ the greater the degree of belief in φ .

In what follows, we present a framework designed to extend the classical concepts of monolithic explanation, as defined by Definition 1, and model reconciling explanation, as defined by Definition 8, into our probabilistic setting.

5.1 Probabilistic Monolithic Explanations

Building on the classical notion of monolithic explanation presented in Definition 1, a *probabilistic monolithic explanation* for an explanandum φ from belief base \mathcal{B} is defined as follows:

Definition 10 (PROBABILISTIC MONOLITHIC EXPLANATION). *Let \mathcal{B} be a belief base, $\mathcal{B}^{\downarrow w}$ its classical projection, and φ an explanandum. We say that $\tilde{\epsilon} \subseteq \mathcal{B}^{\downarrow w}$ is a probabilistic monolithic explanation for φ from \mathcal{B} if and only if $P(\varphi \mid \tilde{\epsilon}) > P(\varphi)$.*

Example 4. Consider the belief base $\mathcal{B} = \{(a, 1), (\neg a \vee b, 2)\}$ and the explanandum b . The probability of the explanandum is $P(b) = 0.73$. Then, $\tilde{\epsilon}_1 = \{a\}$ and $\tilde{\epsilon}_2 = \{\neg a \vee b\}$ are two probabilistic monolithic explanations for b from \mathcal{B} , that is, $P(b \mid \tilde{\epsilon}_1) = 0.88 > P(b)$ and $P(b \mid \tilde{\epsilon}_2) = 0.78 > P(b)$.

It is important to note that Definition 10 can be extended to the case where the formulae $\tilde{\epsilon}$ do not necessarily come from \mathcal{B} , but rather from the language \mathcal{L} . However, we restrict our attention only to formulae from \mathcal{B} in order to be compatible with the classical notion of monolithic explanations (see Definition 1) and the algorithms that we will present in Section 6. For brevity, and until the end of this section, we will refer to probabilistic monolithic explanations as monolithic explanations.

Looking at Example 4, we can see that monolithic explanations will typically vary in their capacity to increase the degree of belief in the explanandum. In other words, each monolithic explanation provides us with some *explanatory gain* for the explanandum. Following Good (1968, 1960), explanatory gain is defined as follows:⁸

Definition 11 (EXPLANATORY GAIN OF MONOLITHIC EXPLANATIONS). *Let $\tilde{\epsilon}$ be a monolithic explanation for explanandum φ from belief base \mathcal{B} . The explanatory gain of $\tilde{\epsilon}$ for φ is defined as $G(\tilde{\epsilon}, \varphi) = \log \left(\frac{P(\varphi \mid \tilde{\epsilon})}{P(\varphi)} \right)$.^{9,10}*

In essence, the explanatory gain can be thought of as a measure that quantifies how well the monolithic explanation $\tilde{\epsilon}$ explains the explanandum φ or, equivalently, the degree to which $\tilde{\epsilon}$ entails φ . The greater the value of $G(\tilde{\epsilon}, \varphi)$, the more substantial the explanatory gain and, hence, the more effective $\tilde{\epsilon}$ is at explaining φ .

It is essential to recognize that this measure, while initially introduced to assess the weak explanatory power of hypotheses in light of evidence (Good 1960), it is used here to evaluate monolithic explanations. By quantifying the extent to which a monolithic explanation explains an explanandum, we can systematically identify the most informative monolithic explanations within a probabilistic framework.

⁸Good (1960) originally introduced this measure to quantify the (weak) explanatory power of a hypothesis with respect to evidence, essentially evaluating how effectively the hypothesis explains the evidence.

⁹We use log with base 2 in our calculations.

¹⁰Note that $G(\tilde{\epsilon}, \varphi)$ is always positive due to the requirement of monolithic explanations that $P(\varphi \mid \tilde{\epsilon}) > P(\varphi)$ (Definition 10).

Example 5. Continuing from Example 4, consider the monolithic explanations $\tilde{e}_1 = \{a\}$, $\tilde{e}_2 = \{\neg a \vee b\}$, and $\tilde{e}_3 = \{a, \neg a \vee b\}$ for explanandum b . The explanatory gains of \tilde{e}_1 , \tilde{e}_2 , and \tilde{e}_3 for b are $G(\tilde{e}_1, b) = \log\left(\frac{P(b|\tilde{e}_1)}{P(b)}\right) = \log\left(\frac{0.88}{0.73}\right) = 0.27$, $G(\tilde{e}_2, b) = \log\left(\frac{P(b|\tilde{e}_2)}{P(b)}\right) = \log\left(\frac{0.78}{0.73}\right) = 0.11$, and $G(\tilde{e}_3, b) = \log\left(\frac{P(b|\tilde{e}_3)}{P(b)}\right) = \log\left(\frac{1}{0.73}\right) = 0.45$, respectively.

A natural course of action when seeking monolithic explanations for an explanandum is to seek the one with the highest explanatory gain. While tempting, it is important to emphasize that when a monolithic explanation entails the explanandum, then the explanatory gain takes on its greatest value. For example,

Example 6. Consider the three monolithic explanations \tilde{e}_1 , \tilde{e}_2 , and \tilde{e}_3 from Example 5. Notice that $\tilde{e}_3 = \{a, \neg a \vee b\}$ entails b (i.e., $\tilde{e}_3 \models b$) and that its explanatory gain is higher than that of \tilde{e}_1 and \tilde{e}_2 . As \tilde{e}_1 , \tilde{e}_2 , and \tilde{e}_3 are the only three possible explanations for b , $G(\tilde{e}_3, b)$ is indeed the maximum achievable explanatory gain for b .

We formalize this in the following proposition:

Proposition 3. Given a monolithic explanation \tilde{e} for an explanandum φ from belief base \mathcal{B} , if $\tilde{e} \models \varphi$, then $G(\tilde{e}, \varphi)$ achieves its maximal value for φ , specifically $G(\tilde{e}, \varphi) = -\log P(\varphi)$.

PROOF. If $\tilde{e} \models \varphi$, then for all possible worlds ω in which $\omega \models \tilde{e}$, it holds that $\omega \models \varphi$. That is, the worlds ω in which \tilde{e} is true are subsumed by the worlds in which φ is true, which implies that also $\omega \models \varphi \wedge \tilde{e}$.

Consequently, $P(\varphi | \tilde{e}) = \frac{\sum_{\omega \models \varphi \wedge \tilde{e}} P(\omega)}{\sum_{\omega \models \tilde{e}} P(\omega)} = \frac{\sum_{\omega \models \tilde{e}} P(\omega)}{\sum_{\omega \models \tilde{e}} P(\omega)} = 1$. Therefore, when $\tilde{e} \models \varphi$, the explanatory gain of \tilde{e} for φ is $G(\tilde{e}, \varphi) = \log\left(\frac{P(\varphi | \tilde{e})}{P(\varphi)}\right) = \log\left(\frac{1}{P(\varphi)}\right) = -\log P(\varphi)$. \square

The following corollary follows naturally from Proposition 3:

Corollary 1. Let $\tilde{E}(\varphi)$ denote the set of all monolithic explanations for explanandum φ from belief base \mathcal{B} . For any two monolithic explanations $\tilde{e}_1, \tilde{e}_2 \in \tilde{E}(\varphi)$, if $\tilde{e}_1 \models \varphi$ and $\tilde{e}_2 \models \varphi$ (resp. $\tilde{e}_2 \not\models \varphi$), then $G(\tilde{e}_1, \varphi) = G(\tilde{e}_2, \varphi)$ (resp. $G(\tilde{e}_1, \varphi) > G(\tilde{e}_2, \varphi)$).

What Proposition 3 and Corollary 1 essentially underscore is that the exclusive focus on explanatory gain as an evaluation metric for a monolithic explanation neglects the inherent likelihood of the explanation itself. That is, the explanatory gain of a monolithic explanation for an explanandum evaluates how effectively the explanation explains the explanandum, *assuming that the explanation itself is true*. Nonetheless, this premise often lacks practical relevance because, in probabilistic contexts, each monolithic explanation is associated with a probability reflecting its likelihood for being true. Therefore, a good measure for evaluating monolithic explanations should incorporate the explanation's inherent plausibility.

Addressing this gap, Good (1968) introduced the concept of (strong) explanatory power that integrates the monolithic explanation's explanatory gain with its probability, offering a more balanced metric for evaluating monolithic explanations.¹¹ Building on Good's measure of explanatory power, we adapt it to our setting and define it as follows:¹²

¹¹Good's measure of (strong) explanatory power is defined as $\log\left(\frac{P(\varphi|h) \cdot P(h)^\gamma}{P(\varphi)}\right)$, where h is a hypothesis and $0 < \gamma < 1$ a constant (Good 1968).

¹²For a detailed defense of Good's measure as a quantitative criterion for explanatory power, alongside a discussion of relevant properties and a comprehensive comparison with other measures, we refer the reader to the work by Glass (2023).

Definition 12 (EXPLANATORY POWER OF MONOLITHIC EXPLANATIONS). *Let \tilde{e} be a monolithic explanation for explanandum φ from belief base \mathcal{B} . The explanatory power of \tilde{e} for φ is defined as $EP(\tilde{e}, \varphi) = G(\tilde{e}, \varphi) + \gamma \cdot P(\tilde{e})$, where $\gamma \in [0, 1]$ is a constant.*

This definition effectively combines the measure of how much a monolithic explanation explains the explanandum (explanatory gain) with the likelihood of the explanation itself, mediated by a parameter γ . The constant γ serves as a tuning parameter, enabling the adjustment of the relative importance of the monolithic explanation's probability in the overall assessment of explanatory power. This flexibility is important for tailoring the evaluation process to specific contexts or preferences, where the balance between the informativeness of a monolithic explanation and its plausibility may vary.

Example 7. Consider the belief base $\mathcal{B} = \{(a, 1.5), (b, 3), (\neg a \vee c, 1), (\neg b \vee c, 1)\}$ and the explanandum c with initial probability $P(c) = 0.84$. Notice that $\tilde{e}_1 = \{a, \neg a \vee c\}$ and $\tilde{e}_2 = \{b, \neg b \vee c\}$ are two monolithic explanations for c from \mathcal{B} , each of which entail c (i.e., $\tilde{e}_1 \models c$ and $\tilde{e}_2 \models c$), with probabilities $P(\tilde{e}_1) = 0.68$ and $P(\tilde{e}_2) = 0.80$, respectively. This means that their explanatory gain for c is equal (Corollary 1), that is, $G(\tilde{e}_1, c) = G(\tilde{e}_2, c) = 0.25$. Now, assuming $\gamma = 0.5$, the explanatory power of \tilde{e}_1 and \tilde{e}_2 respectively is $EP(\tilde{e}_1, c) = 0.25 + 0.5 \cdot 0.68 = 0.59$ and $EP(\tilde{e}_2, c) = 0.25 + 0.5 \cdot 0.80 = 0.65$.

With the introduction of explanatory power as an evaluative measure of (probabilistic) monolithic explanations, we can now define a (probabilistic) *preference relation* among monolithic explanations, which allows for a systematic approach to determining the most effective monolithic explanation for a given explanandum:

Definition 13 (PREFERENCE RELATION FOR MONOLITHIC EXPLANATION). *Let \tilde{e}_1 and \tilde{e}_2 be two monolithic explanations for explanandum φ from belief base \mathcal{B} . \tilde{e}_1 is preferred over \tilde{e}_2 , denoted as $\tilde{e}_1 \geq \tilde{e}_2$, if and only if $EP(\tilde{e}_1, \varphi) \geq EP(\tilde{e}_2, \varphi)$.*

This definition enables a quantitatively grounded approach to preference among monolithic explanations, where the preference is directly tied to the explanatory power of each explanation. It facilitates a structured way to navigate the space of potential monolithic explanations, prioritizing those that not only explain the explanandum more effectively, but also align more closely with the existing knowledge represented by the belief base \mathcal{B} .

Example 8. Continuing from Example 7, the two monolithic explanations for c from \mathcal{B} are \tilde{e}_1 and \tilde{e}_2 and have explanatory power $EP(\tilde{e}_1, c) = 0.59$ and $EP(\tilde{e}_2, c) = 0.65$. Thus, \tilde{e}_2 is preferred over \tilde{e}_1 (i.e., $\tilde{e}_2 \geq \tilde{e}_1$).

Finally, given the set of all monolithic explanations for an explanandum, we say that a monolithic explanation is *most preferred* if and only if it is (probabilistically) preferred over every other possible monolithic explanation for that explanandum. Formally,

Definition 14 (MOST-PREFERRED MONOLITHIC EXPLANATION). *Let $\tilde{E}(\varphi)$ denote the set of all monolithic explanations for explanandum φ from belief base \mathcal{B} . A monolithic explanation $\tilde{e}^* \in \tilde{E}(\varphi)$ is the most-preferred monolithic explanation if and only if $\tilde{e}^* \geq \tilde{e}$ for all $\tilde{e} \in \tilde{E}(\varphi)$.*

5.1.1 Motivating Application: Illustrative Example. Recall from our motivating scenario in Section 4 that we are interested in explaining the explanandum $\varphi = \neg \text{move}(\text{room1}, \text{room2}, A)$, which is akin to asking “why didn’t you take corridor A?” Using our framework, we identify two potential explanations for this explanandum:

- (1) $\tilde{e}_1 = \{\text{crowded}(A)\}$ or, in natural language, “Corridor A is very likely crowded.”
 - $P(\varphi|\tilde{e}_1) = 0.62$
 - Explanatory gain: $G(\tilde{e}_1, \varphi) = 0.1$
 - Explanatory power: $EP(\tilde{e}_1, \varphi) = 0.5$

- (2) $\tilde{\epsilon}_2 = \{\text{crowded}(A), \text{move}(\text{room1}, \text{room2}, A)_0 \wedge \text{crowded}(A) \rightarrow \text{robot-at}(\text{room1})_1, \text{move}(\text{room1}, \text{room2}, A)_0 \wedge \text{crowded}(A) \rightarrow \text{robot-at}(\text{room2})_1\}$ or, in natural language, “Corridor A is very likely crowded, and when attempting to move through a crowded corridor, my movement outcomes are uncertain; I may either succeed or fail to move through to the next room.”
- $P(\phi|\tilde{\epsilon}_2) = 1$
 - Explanatory gain: $G(\tilde{\epsilon}_2, \phi) = 0.79$
 - Explanatory power: $EP(\tilde{\epsilon}_2, \phi) = 1.04$

Comparing the two explanations, we see that explanation $\tilde{\epsilon}_2$ has the highest explanatory power (1.04), and is thus the most-preferred explanation according to Definition 14. Such an explanation can help the human supervisor to understand both the robot’s belief about the environment (crowded corridor) and how this belief influences action outcomes (movement outcomes under crowded conditions), thereby providing a more complete picture of the robot’s decision-making process.

Comparison with MPE: It is important to highlight the difference with the traditional *Most Probable Explanation* (MPE) method. MPE typically consists of finding the world with the highest probability given some evidence (Shterionov et al. 2015). However, a world does not show the chain of inferences of a given explanandum, and importantly, it is not minimal by definition since it usually includes a (possibly large) number of probabilistic facts whose truth value is irrelevant for the explanandum. For example, applying standard MPE to our running example we get the following outcome:

$$\begin{aligned} \omega_{\text{MPE}} = \{ & \text{crowded}(A), \neg\text{crowded}(B), \neg\text{deliver}_0, \text{deliver}_1, \neg\text{move}(\text{room1}, \text{room2}, A)_0, \\ & \neg\text{move}(\text{room1}, \text{room2}, A)_1, \text{move}(\text{room1}, \text{room2}, B)_0, \neg\text{move}(\text{room1}, \text{room2}, B)_1, \\ & \neg\text{move}(\text{room2}, \text{room1}, A)_0, \neg\text{move}(\text{room2}, \text{room1}, A)_1, \neg\text{move}(\text{room2}, \text{room1}, B)_0, \\ & \neg\text{move}(\text{room2}, \text{room1}, B)_1, \neg\text{package-delivered}_0, \neg\text{package-delivered}_1, \\ & \text{package-delivered}_2, \text{robot-at}(\text{room1})_0, \neg\text{robot-at}(\text{room1})_1, \\ & \neg\text{robot-at}(\text{room1})_2, \neg\text{robot-at}(\text{room2})_0, \text{robot-at}(\text{room2})_1, \\ & \text{robot-at}(\text{room2})_2 \} \end{aligned}$$

While ω_{MPE} is indeed the most probable world ($P(\omega_{\text{MPE}}) = 0.32$) in which the query ($\neg\text{move}(\text{room1}, \text{room2}, A)$) is true, it does not pinpoint to the (minimal) combination of formulae that justify the query.

5.2 Probabilistic Model Reconciling Explanations

Recall from Section 3.3 that, in the model reconciliation problem (MRP), the models of the agent and the human user diverge with respect to an explanandum, insofar as the explanandum is explicable in the agent’s model but inexplicable in the human’s model. The goal is then to find a model reconciling explanation (i.e., a set of model differences) such that the explanandum becomes explicable in the human’s model. Three important assumptions underlying MRP typically hold: (1) the agent model is the ground truth; (2) the agent has access to the human model; and (3) both models are deterministic.

As we described in Section 3.3, assumption (1) is reasonable since explanations are generated from the agent’s perspective. In other words, the agent “thinks” that its model is correct. For assumption (2), the agent does not have access to the human’s actual model, but an approximation of it. In the worst case, it can be empty; but, practically, it can be approximated based on past interactions (Juba et al. 2021; Sreedharan et al. 2018). For assumption (3), we will relax the assumption that the human model is deterministic in our work, but we will still assume that the agent model is deterministic.

The motivation for moving away from deterministic human models becomes stronger when we consider two key points. First, since the agent is using an approximated human model, deterministic approximations are

more likely to be inaccurate compared to probabilistic ones. Consequently, deterministic models may generate explanations that are incorrect or not meaningful for the user, thereby reducing the effectiveness of MRP. Secondly, it is likely that humans hold beliefs with varying degrees of certainty, highlighting a shortfall of deterministic models in capturing this range of uncertainties. These factors together underscore the necessity for models that incorporate probabilistic aspects, thus potentially enabling a more accurate and user-relevant application of MRP.

To that end, we will now expand the scope of MRP to cases in which the agent is uncertain about the human model. Particularly, we build on our previous work (Vasileiou, Yeoh, et al. 2022), wherein both the agent and human models are represented as (logical) knowledge bases (see Definition 8), and extend it to the case where the human knowledge base is probabilistic (i.e., a belief base). In other words, we are now interested in *probabilistic model reconciling explanations*.

First, we show through the following example how the concepts surrounding probabilistic monolithic explanations introduced in the previous section are applicable to the case of an agent knowledge base KB_α and a human belief base \mathcal{B}_h .

Example 9. Let $KB_\alpha = \{a, \neg a \vee b, c\}$ and $\mathcal{B}_h = \{(c, 2), (\neg c \vee \neg a, 2)\}$ be the knowledge bases of an agent and the belief base of a human, respectively. Additionally, let b be the explanandum, where $KB_\alpha \models b$ and $P_h(b) = 0.5$. The goal in this example would then be to find which formulae from KB_α increase the probability of the explanandum for \mathcal{B}_h , that is, to find a probabilistic monolithic explanation \tilde{e} for b from KB_α for \mathcal{B}_h such that $P_h(b \mid \tilde{e}) > P_h(b)$ (Definition 10).

Given KB_α , there are three possible monolithic explanations: $\tilde{e}_1 = \{a\}$, $\tilde{e}_2 = \{\neg a \vee b\}$, and $\tilde{e}_3 = \{a, \neg a \vee b\}$. Evaluating them with respect to the probability distribution induced by \mathcal{B}_h , we get $P_h(b \mid \tilde{e}_1) = 0.5$, $P_h(b \mid \tilde{e}_2) = 0.55$, and $P_h(b \mid \tilde{e}_3) = 1$. Notice now that only \tilde{e}_2 and \tilde{e}_3 qualify as monolithic explanations since $P_h(b \mid \tilde{e}_2) > P_h(b)$ and $P_h(b \mid \tilde{e}_3) > P_h(b)$, whilst \tilde{e}_1 does not qualify as a monolithic explanation as $P_h(b \mid \tilde{e}_1) = P_h(b) = 0.5$.

Given \tilde{e}_2 and \tilde{e}_3 as the two possible monolithic explanations, we can now evaluate their effectiveness in terms of explanatory gain (Definition 11) and explanatory power (Definition 12). In terms of explanatory gain, we get $G(\tilde{e}_2, b) = 0.14$ and $G(\tilde{e}_3, b) = 1$. In terms of explanatory power (for $\gamma = 0.5$), we get $EP(\tilde{e}_2, b) = 0.59$ and $EP(\tilde{e}_3, b) = 1.04$. Finally, following the definition of most-preferred monolithic explanation (Definition 14), we get that \tilde{e}_3 is the most-preferred monolithic explanation for b from KB_α for \mathcal{B}_h .

On the one hand, example 9 shows that the definitions introduced in Section 5.1 can be directly applied to the case of an agent knowledge base KB_α and a human belief base \mathcal{B}_h . On the other hand, there is something important to highlight here. Despite \tilde{e}_3 being the most-preferred monolithic explanation (i.e., it has the highest explanatory power), notice that its probability $P_h(\tilde{e}_3) = 0.09$ is rather low, which means that its negation $\neg\tilde{e}_3$ has a much higher probability with $P_h(\neg\tilde{e}_3) = 0.91$. Logically, this is explained by the fact that \tilde{e}_3 is inconsistent with the formulae in $\mathcal{B}_h^{\downarrow w}$. Therefore, the probabilistic monolithic explanation \tilde{e}_3 may not achieve the intended “reconciliation” between the agent and the human.

Recall that a model reconciling explanation (see Definition 8) is of the form $\mathcal{E} = \langle \epsilon^+, \epsilon^- \rangle$, where ϵ^- is specifically intended to resolve the inconsistency between the agent and the human with respect to the explanandum. Intuitively, the provision of ϵ^- can be thought of as the agent’s suggestion of what is “false” in the human knowledge base, at least compared to the agent knowledge base. In the case of a human belief base \mathcal{B}_h , we can account for ϵ^- by finding a set of formulae from \mathcal{B}_h such that $P_h(\epsilon^+ \mid \neg\epsilon^-) > P_h(\epsilon^+)$. For example,

Example 10. Let $KB_\alpha = \{a, \neg a \vee b, c\}$ and $\mathcal{B}_h = \{(c, 2), (\neg c \vee \neg a, 2)\}$ from Example 9. From the perspective of KB_α , explanation $\tilde{e} = \{a, \neg a \vee c\}$ can be seen as the formulae that should be true (e.g., added) in \mathcal{B}_h (i.e., $\tilde{e}^+ = \tilde{e}$). However, notice that \tilde{e}^+ is inconsistent with $\mathcal{B}_h^{\downarrow w} = \{c, \neg c \vee \neg a\}$. Thus, from the perspective of KB_α , some formulae from $\mathcal{B}_h^{\downarrow w}$ are false (e.g., they should be retracted). One can see that $\tilde{e}^- = \neg c \vee \neg a$ is the only formula that should be false as it is the only one that is inconsistent with KB_α . Indeed, if \tilde{e}^- is assumed to be false, then the probability

of $\tilde{\epsilon}^+$ increases, i.e., $P_h(\tilde{\epsilon}^+ \mid \neg\tilde{\epsilon}^-) = 0.5 > P_h(\tilde{\epsilon}^+) = 0.09$. Therefore, $\tilde{\epsilon}^+$ and $\tilde{\epsilon}^-$ can be seen as a model reconciling explanation for b from KB_α for \mathcal{B}_h .

Before formally defining what constitutes a probabilistic model reconciling explanation, we state the following assumptions underlying our framework:

- **Shared Domain Language:** The agent and the human user share the same (propositional) language \mathcal{L} , that is, they share the same set of atomic variables \mathcal{V} from which formulae specific to a domain can be constructed.
- **Agent Knowledge Base:** The agent model is represented by the (deterministic) knowledge base KB_α , encoding the ground truth of the domain.
- **Human Belief Base:** The human model is represented by the belief base \mathcal{B}_h (and its associated probability distribution P_h), reflecting the agent's uncertainty, for example, its degrees of belief about the human model. The agent has access to \mathcal{B}_h a-priori.¹³

We define a *probabilistic model reconciling explanation* as follows:

Definition 15 (PROBABILISTIC MODEL RECONCILING EXPLANATION). *Given the knowledge base KB_α of an agent, the belief base \mathcal{B}_h of a human user, and an explanandum φ such that $\text{KB}_\alpha \models \varphi$ and $P_h(\varphi) < 1$, $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ is a probabilistic model reconciling explanation if and only if $\tilde{\epsilon}^+ \subseteq \text{KB}_\alpha$ and $\tilde{\epsilon}^- \subseteq \mathcal{B}_h^{\downarrow w}$, and $P_h(\varphi \mid \tilde{\epsilon}^+) > P_h(\varphi)$ and $P_h(\tilde{\epsilon}^+ \mid \neg\tilde{\epsilon}^-) > P_h(\tilde{\epsilon}^+)$.*

A probabilistic model reconciling explanation $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ for φ from KB_α for \mathcal{B}_h is a tuple that increases the degree of belief in φ with $\tilde{\epsilon}^+$, as well as increasing the degree of belief in $\tilde{\epsilon}^+$ with $\tilde{\epsilon}^-$ if $\tilde{\epsilon}^+$ is inconsistent with $\mathcal{B}_h^{\downarrow w}$. For brevity, until the end of this section, we will refer to probabilistic model reconciling explanations $\tilde{\mathcal{E}}$ as model reconciling explanations. In this context, the notion of explanatory gain takes the following form:

Definition 16 (EXPLANATORY GAIN FOR MODEL RECONCILING EXPLANATIONS). *Let $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ be a model reconciling explanation for explanandum φ from KB_α for \mathcal{B}_h . The explanatory gain of $\tilde{\mathcal{E}}$ for φ is defined as $G(\tilde{\mathcal{E}}, \varphi) = \log\left(\frac{P(\varphi \mid \tilde{\epsilon}^+)}{P(\varphi)}\right) + \log\left(\frac{P(\tilde{\epsilon}^+ \mid \neg\tilde{\epsilon}^-)}{P(\tilde{\epsilon}^+)}\right)$.*

In essence, the explanatory gain of $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ for φ evaluates to what extent $\tilde{\epsilon}^+$ increases the probability of φ , as well as the extent to which $\tilde{\epsilon}^-$ increases the probability of $\tilde{\epsilon}^+$, assuming that $\tilde{\epsilon}^-$ is false.

Example 11. Let $\tilde{\mathcal{E}} = \langle \{a, \neg a \vee b\}, \{\neg c \vee \neg a\} \rangle$ be the model reconciling explanation for b from KB_α for \mathcal{B}_h in Example 10. The explanatory gain of $\tilde{\mathcal{E}}$ for b is $G(\tilde{\mathcal{E}}, b) = \log\left(\frac{1}{0.5}\right) + \log\left(\frac{0.5}{0.09}\right) = 1 + 2.47 = 3.47$.

Similarly, the notion of explanatory power is defined in the following way:

Definition 17 (EXPLANATORY POWER FOR MODEL RECONCILING EXPLANATIONS). *Let $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ be a model reconciling explanation for explanandum φ from KB_α for \mathcal{B}_h . The explanatory power of $\tilde{\mathcal{E}}$ for φ is defined as $\text{EP}(\tilde{\mathcal{E}}, \varphi) = G_h(\tilde{\mathcal{E}}, \varphi) + \gamma \cdot (P_h(\tilde{\epsilon}^+) + P_h(\tilde{\epsilon}^-))$, where $\gamma \in [0, 1]$ is a constant.*

This definition of explanatory power of $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ for φ assesses, in addition to the explanatory gain of $\tilde{\mathcal{E}}$, the likelihoods of $\tilde{\epsilon}^+$ and $\tilde{\epsilon}^-$, with γ parameterizing their relative importance in the overall assessment.

Example 12. Continuing from Example 11, the explanatory power of $\tilde{\mathcal{E}} = \langle \{a, \neg a \vee b\}, \{\neg c \vee \neg a\} \rangle$ for b (for $\gamma = 0.5$) is $\text{EP}(\tilde{\mathcal{E}}, b) = 3.47 + 0.5 \cdot (0.09 + 0.90) = 3.96$

¹³We leave the question of acquiring (or learning) the human belief base open for future work.

Finally, a preference relation and a most-preferred model reconciling explanation can be defined in the same manner as in Definition 13 and Definition 14, respectively.

Definition 18 (PREFERENCE RELATION FOR MODEL RECONCILING EXPLANATION). *Let $\tilde{\mathcal{E}}_1$ and $\tilde{\mathcal{E}}_2$ be two model reconciling explanations for explanandum φ from knowledge base KB_α for belief base \mathcal{B}_h . $\tilde{\mathcal{E}}_1$ is preferred over $\tilde{\mathcal{E}}_2$, denoted $\tilde{\mathcal{E}}_1 \geq \tilde{\mathcal{E}}_2$, if and only if $\text{EP}(\tilde{\mathcal{E}}_1) \geq \text{EP}(\tilde{\mathcal{E}}_2)$.*

Definition 19 (MOST-PREFERRED MODEL RECONCILING EXPLANATION). *Let $\tilde{E}(\varphi)$ denote the set of all model reconciling explanations for explanandum φ from knowledge base KB_α for belief base \mathcal{B}_h . A model reconciling explanation $\tilde{\mathcal{E}}^* \in \tilde{E}(\varphi)$ is the most-preferred model reconciling explanation for φ if and only if $\tilde{\mathcal{E}}^* \geq \tilde{\mathcal{E}}$ for all $\tilde{\mathcal{E}} \in \tilde{E}(\varphi)$.*

5.2.1 Motivating Application: Illustrative Example. Continuing with our motivating example, we will now examine a case where the human’s belief base differs from the robot’s in a way that leads to divergent expectations about the robot’s decision. Particularly, consider a deterministic version of the robot’s belief base KB_α as described in Section 4, and let us define the human’s belief base \mathcal{B}_h with the following key differences:

$$\begin{aligned} &(\neg\text{crowded}(A), 1) \\ &(\text{crowded}(B), 3) \end{aligned}$$

The human user believes that corridor A is generally unlikely to be crowded and corridor B is very likely crowded—the opposite of the robot’s beliefs. Additionally, the human believes that the robot can successfully move through corridors A and B, whether crowded or not. Given their belief base \mathcal{B}_h , the expected optimal plan from the human’s perspective would be $\pi_h = \langle \text{move}(\text{room1}, \text{room2}, A)_0, \text{deliver}_1 \rangle$. When the human observes the robot taking corridor B instead, they would naturally question this decision.

The robot, aware of this discrepancy, can then generate a probabilistic model reconciling explanation $\tilde{\mathcal{E}} = \langle \tilde{\mathcal{E}}^+, \tilde{\mathcal{E}}^- \rangle$ to address φ :

$$\begin{aligned} \tilde{\mathcal{E}}^+ &= \{ \text{crowded}(A), \text{move}(\text{room1}, \text{room2}, A)_0 \wedge \text{crowded}(A) \rightarrow \text{robot-at}(\text{room1})_1 \} \\ \tilde{\mathcal{E}}^- &= \{ \neg\text{crowded}(A) \} \\ \bullet P_h(\varphi | \tilde{\mathcal{E}}^+) &= 0.7 \text{ (increased from } P_h(\varphi) = 0.4) \\ \bullet P_h(\tilde{\mathcal{E}}^+ | \neg\tilde{\mathcal{E}}^-) &= 1 \text{ (increased from } P_h(\tilde{\mathcal{E}}^+) = 0.15) \\ \bullet \text{Explanatory gain: } G(\tilde{\mathcal{E}}, \varphi) &= 3.04 \\ \bullet \text{Explanatory power: } EP(\tilde{\mathcal{E}}, \varphi) &= 3.22 \end{aligned}$$

The explanation essentially communicates: “I didn’t take corridor A because it is very likely crowded, contrary to your belief that it’s not crowded. Additionally, when a corridor is crowded, my movement through it will likely fail, causing me to remain in my original location.” What we see here is that if we focus on both environmental factors (crowded corridor) and action dynamics (movement success), we get an explanation that provides a complete account of the robot’s decision-making process, thereby helping the human gain a better understanding of the robot.

6 Computing Explanations

We now describe algorithms for computing explanations. We first review two algorithms proposed in our previous work (Vasileiou, Previti, et al. 2021) for computing classical (deterministic) monolithic explanations (Definition 1) and model reconciling explanations (Definition 8), and then show how to extend them to the probabilistic case.

Algorithm 1: monolithic-explanation(KB, φ)**Input:** Knowledge base KB and explanandum φ **Result:** A minimum size monolithic explanation ϵ for φ from KB

```

1  $\mathcal{H} \leftarrow \emptyset$ 
2 while true do
3    $seed \leftarrow \text{minHS}(\mathcal{H})$                                 // compute a minimal hitting set
4    $\epsilon \leftarrow \{c_i \mid i \in seed\}$ 
5   if not SAT( $\epsilon \cup \{\neg\varphi\}$ ) then
6     return  $\epsilon$                                               // minimum size monolithic explanation
7   else
8      $C \leftarrow \text{getMCS}(seed, KB^s \cup \{\neg\varphi^h\})$       // compute a minimal correction set
9    $\mathcal{H} \leftarrow \mathcal{H} \cup \{C\}$ 

```

6.1 Classical Explanations

We previously introduced an approach for computing minimum size monolithic explanations for an explanandum φ from a knowledge base KB, where $KB \models \varphi$ (Vasileiou, Previti, et al. 2021). The principal idea of this approach is to reduce the problem of computing a monolithic explanation of minimum size to the one of computing a *smallest minimal unsatisfiable set* (SMUS) over an inconsistent knowledge base.

In particular, notice that, by definition, we have that $KB \models \varphi$ if and only if $KB \cup \{\neg\varphi\}$ is inconsistent. Moreover, in Proposition 2, we have already stated the relation between a monolithic explanation and a *minimal unsatisfiable set* (MUS). This suggests that, in order to extract a monolithic explanation, we just need to run an MUS solver over the knowledge base $KB^s \cup \{\neg\varphi^h\}$, where KB^s and φ^h denote that KB and φ are treated as soft and hard constraints, respectively, and then remove $\neg\varphi$ from the returned MUS.¹⁴ The hitting set duality relating MUSes and *minimal correction sets* (MCSes) (see Lemma 1) is a key aspect for the computation of an SMUS.

Algorithm 1 describes the main steps of our approach. \mathcal{H} is a collection of sets, where each set corresponds to an MCS on KB. At the beginning, it is initialized with the empty set (line 1). Each MCS in \mathcal{H} is represented as the set of the indexes of the formulae in it. \mathcal{H} stores the MCSes computed so far. At each step, a minimal hitting set on \mathcal{H} is computed (line 3). In line 4, the formulae induced by the computed minimal hitting set is stored in ϵ . Then, $\epsilon \cup \{\neg\varphi\}$ is evaluated for satisfiability (line 5). If $\epsilon \cup \{\neg\varphi\}$ is inconsistent, then ϵ is a monolithic explanation of minimum size and the algorithm returns ϵ . If instead $\epsilon \cup \{\neg\varphi\}$ is consistent, then it means that $\epsilon \not\models \varphi$ and the algorithm continues in line 8. The computation of an MCS of this kind can be performed via standard MCS procedures (Marques-Silva et al. 2013), using the set of formulae indexed by the *seed* as the starting formula to extend. Since φ is set to hard (line 8), the returned MCS C is guaranteed to be contained in KB. Due to the hitting set duality relation, we will also have $\epsilon \subseteq KB$. Finally, notice that the procedure *getMCS* always reports a new MCS because, by construction, we have $seed \subseteq KB \setminus C$. In fact, the *seed* contains at least one formula for each previously computed MCS and, thus, $seed \cap C = \emptyset$ (i.e., at least one formula for each previously computed MCS is not in C).

Algorithm 1 is complete in the sense that eventually a monolithic explanation $\epsilon \subseteq KB$ of minimum size such that $\epsilon \models \varphi$ will be returned. This can be easily verified by observing that every time $\epsilon \cup \{\neg\varphi\}$ is satisfiable, a new MCS is computed. Eventually, all the MCSes will be computed and, from Propositions 1 and 2, it follows that a minimal hitting set on the collection of all MCSes corresponds to the smallest MUS, and as such, to a monolithic explanation of minimum size.

¹⁴Recall that soft constraints may be removed by the MUS solver, while hard constraints will not be removed.

Table 1. Example of Algorithm 1 for computing a monolithic explanation of minimum size.

$KB = \{ \overset{C_1}{a \vee b}, \overset{C_2}{\neg b \vee c}, \overset{C_3}{\neg c}, \overset{C_4}{\neg b \vee d} \}$ $KB \models a$		
1.	$\mathcal{H} \leftarrow \emptyset$	
2.	$seed \leftarrow \emptyset$	# $minHS(\mathcal{H})$
3.	$\emptyset \not\models a$	# $SAT(\epsilon \cup \{\neg a\})$
4.	$C \leftarrow \{C_1\}$	# MCS computed on $KB^s \cup \{\neg a^h\}$ starting with the seed $seed$
5.	$\mathcal{H} \leftarrow \{\{C_1\}\}$	
6.	$seed \leftarrow \{C_1\}$	# $minHS(\mathcal{H})$
7.	$\{a \vee b\} \not\models a$	# $SAT(\epsilon \cup \{\neg a\})$
8.	$C \leftarrow \{C_2\}$	# MCS computed on $KB^s \cup \{\neg a^h\}$ starting with the seed $seed$
9.	$\mathcal{H} \leftarrow \{\{C_1\}, \{C_2\}\}$	
10.	$seed \leftarrow \{C_1, C_2\}$	# $minHS(\mathcal{H})$
11.	$\{a \vee b, \neg b \vee c\} \not\models a$	# $SAT(\epsilon \cup \{\neg a\})$
12.	$C \leftarrow \{C_3\}$	# MCS computed on $KB^s \cup \{\neg a^h\}$ starting with the seed $seed$
13.	$\mathcal{H} \leftarrow \{\{C_1\}, \{C_2\}, \{C_3\}\}$	
14.	$seed \leftarrow \{C_1, C_2, C_3\}$	# $minHS(\mathcal{H})$
15.	$\{a \vee b, \neg b \vee c, \neg c\} \models a$	# $\neg SAT(\epsilon \cup \{\neg a\})$
16.	<i>Return</i> $\{a \vee b, \neg b \vee c, \neg c\}$	# minimum size monolithic explanation for a from KB

Note that deciding whether there exists a monolithic explanation of size less or equal to k is Σ_2^P -complete and extracting a smallest monolithic explanation is in $FP^{\Sigma_2^P}$. This follows directly from the complexity of deciding and computing an SMUS on which Algorithm 1 is based on (Ignatiev, Previti, et al. 2015).

6.1.1 Model Reconciling Explanations. We have also previously showed how Algorithm 1 can be further extended for computing model reconciling explanations for an explanandum φ from an agent knowledge base KB_α for a human knowledge base KB_h (Vasileiou, Previti, et al. 2021). However, we only considered the specific task of finding a model reconciling explanation $\epsilon \subseteq KB_\alpha \cup KB_h$ such that $KB_h \cup \epsilon \models \varphi$ and $\epsilon \setminus KB_h$ is of minimum size. Notice that, in general, $KB_h \cup \epsilon$ might be inconsistent. However, in our approach, we discard this possibility by preprocessing KB_h . In particular, we create a new $KB'_h \subseteq KB_h$ by removing a minimal set of formulae in KB_h that makes $KB_h \cup KB_\alpha$ inconsistent. The new KB'_h is such that $KB'_h \cup \epsilon$ is always consistent.

We now modify this approach for computing model reconciling explanations $\mathcal{E} = \langle \epsilon^+, \epsilon^- \rangle$, where $\epsilon^+ \subseteq KB_\alpha$ and $\epsilon^- \subseteq KB_h$, such that $(KB_h \cup \epsilon^+) \setminus \epsilon^- \models \varphi$. Particularly, in addition to ϵ^+ , our approach now computes ϵ^- as well.

Algorithm 2 describes the pseudocode of our approach. At the beginning of the algorithm, we initialize \mathbb{R} to the null set (line 1). \mathbb{R} is used to store the MCSes, which acts as a mediator between KB_α and KB_h . Lines 2-3 are used to specify which clauses of KB_α will be treated as hard and soft constraints, respectively. We then check if $KB_h \cup KB_\alpha$ is inconsistent (line 4). This is important in order to avoid the possibility of finding subsets ϵ^+ that explain why $KB_h \cup KB_\alpha$ is inconsistent instead of the target explanandum. In case $KB_h \cup KB_\alpha$ is inconsistent, we preprocess KB_h by removing from $KB_h \setminus KB_\alpha$ a minimal set of formulae causing the conflict (i.e., an MCS) (lines 5-6), where E^- stores the set of potential formulae ϵ^- to retract. The reconciliation procedure starts in line 7. The algorithm proceeds iteratively by computing a minimal hitting set on \mathbb{R} and then testing for satisfiability the

Algorithm 2: model-reconciling-explanation($\text{KB}_\alpha, \text{KB}_\alpha^h, \varphi$)**Input:** Knowledge bases KB_α and KB_h and explanandum φ **Result:** A model reconciling explanation $\mathcal{E} = \langle \epsilon^+, \epsilon^- \rangle$ for φ from KB_α for KB_h

```

1  $\mathbb{R} \leftarrow \emptyset$ 
2  $\text{KB}_\alpha^h \leftarrow \text{KB}_\alpha \cap \text{KB}_h$ 
3  $\text{KB}_\alpha^s \leftarrow \text{KB}_\alpha \setminus \text{KB}_\alpha^h$ 
4 if not SAT( $\text{KB}_h \cup \text{KB}_\alpha$ ) then
5    $E^- \leftarrow \text{getMCS}((\text{KB}_h \setminus \text{KB}_\alpha)^s \cup \text{KB}_\alpha^h)$  // restore consistency on  $\text{KB}_h$ 
6    $\text{KB}_h \leftarrow \text{KB}_h \setminus E^-$ 
7 while true do
8    $\text{seed} \leftarrow \text{minHS}(\mathbb{R})$ 
9    $\epsilon^+ \leftarrow \{c_i \mid i \in \text{seed}\}$  // explanation  $\epsilon^+$  induced by the seed
10  if not SAT( $\text{KB}_h \cup \epsilon^+ \cup \{\neg\varphi\}$ ) then
11     $\epsilon^- \leftarrow \emptyset$ 
12    if not SAT( $\text{KB}_h \cup \epsilon^+ \cup E^-$ ) then
13       $\epsilon^- \leftarrow \text{getMCS}((\text{KB}_h \cup \epsilon^+)^h \cup (E^-)^s)$ 
14    return  $\langle \epsilon^+, \epsilon^- \rangle$ 
15  else
16     $C \leftarrow \text{getMCS}(\text{seed}, \text{KB}_\alpha^h \cup \{\neg\varphi^h\} \cup \text{KB}_\alpha^s)$ 
17     $\mathbb{R} \leftarrow \mathbb{R} \cup \{C\}$ 

```

formulae ϵ^+ (lines 8-10). The test checks whether adding ϵ^+ to KB_h is sufficient for entailing φ . If $\text{KB}_h \cup \epsilon^+ \cup \{\neg\varphi\}$ is unsatisfiable, then $\text{KB}_h \cup \epsilon^+ \models \varphi$. In that case, the algorithm then checks whether $\text{KB}_h \cup \epsilon^+ \cup E^-$ is inconsistent, and if it is, it computes an MCS ϵ^- on $(\text{KB}_h \cup \epsilon^+)^h \cup (E^-)^s$ (lines 12-13). The model reconciling explanation $\langle \epsilon^+, \epsilon^- \rangle$ is then returned in line 14. Otherwise, the algorithm continues in line 16, where a new MCS is computed and added to \mathbb{R} .¹⁵

Table 2 shows an example trace of Algorithm 2.

6.2 Probabilistic Explanations

We now show how the algorithms described in the previous section can be used for computing probabilistic monolithic explanations (Definition 10) and probabilistic model reconciling explanations (Definition 15).

6.2.1 Monolithic Explanations. Consider an explanandum φ and a belief base \mathcal{B} . First, notice that if we assume that the classical projection of \mathcal{B} entails the explanandum φ , that is $\mathcal{B}^{\downarrow w} \models \varphi$, then Algorithm 1 can directly be applied on $\mathcal{B}^{\downarrow w}$ and φ .¹⁶ In that case, Algorithm 1 guarantees to find a monolithic explanation with maximum explanatory gain, since we know from Proposition 3 that explanatory gain achieves its maximum value for φ when the monolithic explanation entails φ . Nevertheless, this does not guarantee that the monolithic explanation will be the most-preferred one, that is, the one with the highest explanatory power (Definition 14).

Obviously, a straightforward way of computing a most-preferred monolithic explanation is to use Algorithm 1 to enumerate all possible monolithic explanations for φ , and return the one that has the highest probability, which corresponds to the one with the highest explanatory power. But enumerating through all possible monolithic

¹⁵Note that the algorithm is complete as it is based on Algorithm 1, which is complete.

¹⁶Recall that the classical projection of belief base \mathcal{B} is the unweighted version of the set of formulae from \mathcal{B} .

Table 2. Example of Algorithm 2 for computing a model reconciling explanation.

$\left. \begin{array}{l} \text{KB}_\alpha = \{(a \vee b), (\neg b \vee c), \neg c, (\neg b \vee d), \neg d\} \\ \text{KB}_h = \{b, \neg c\} \end{array} \right\} \text{ We have that } \text{KB}_\alpha \models a \text{ and } \text{KB}_h \not\models a$	
1. $\mathcal{R} \leftarrow \emptyset$	
2. $\text{KB}_\alpha^h \leftarrow \text{KB}_\alpha \cap \text{KB}_h = \{C_3\}$	
3. $\text{KB}_\alpha^s \leftarrow \text{KB}_\alpha \setminus (\text{KB}_\alpha \cap \text{KB}_h) = \{C_1, C_2, C_4, C_5\}$	
4. $E^- \leftarrow \{D_1\}$	# MCS computed on $(\text{KB}_h \setminus \text{KB}_\alpha)^s \cup \text{KB}_\alpha^h$
5. $\text{KB}_h \leftarrow \{D_1, D_2\} \setminus \{D_1\} = \{D_2\}$	
6. $\text{seed} \leftarrow \emptyset$	# $\text{minHS}(\mathcal{R})$
7. $\{\neg c\} \not\models a$	# $\text{SAT}(\text{KB}_h \cup \epsilon^+ \cup \{\neg a\})$
8. $C \leftarrow \{C_1\}$	# MCS computed on $\text{KB}_\alpha^h \cup \text{KB}_\alpha^s \cup \{\neg a^h\}$
9. $\mathcal{R} \leftarrow \{\{C_1\}\}$	
10. $\text{seed} \leftarrow \{C_1\}$	# $\text{minHS}(\mathcal{R})$
11. $\{\neg c, a \vee b\} \not\models a$	# $\text{SAT}(\text{KB}_h \cup \epsilon^+ \cup \{\neg a\})$
12. $C \leftarrow \{C_2, C_4\}$	# MCS computed on $\text{KB}_\alpha^h \cup \text{KB}_\alpha^s \cup \{\neg a^h\}$
13. $\mathcal{R} \leftarrow \{\{C_1\}, \{C_2, C_4\}\}$	
14. $\text{seed} \leftarrow \{C_1, C_4\}$	# $\text{minHS}(\mathcal{R})$
15. $\{\neg c, a \vee b, \neg b \vee d\} \not\models a$	# $\text{SAT}(\text{KB}_h \cup \epsilon^+ \cup \{\neg a\})$
16. $C \leftarrow \{C_2, C_5\}$	# MCS computed on $\text{KB}_\alpha^h \cup \text{KB}_\alpha^s \cup \{\neg a^h\}$
17. $\mathcal{R} \leftarrow \{\{C_1\}, \{C_2, C_4\}, \{C_2, C_5\}\}$	
18. $\text{seed} \leftarrow \{C_1, C_2\}$	# $\text{minHS}(\mathcal{R})$
19. $\{\neg c, a \vee b, \neg b \vee c\} \models a$	# $\neg \text{SAT}(\text{KB}_h \cup \epsilon^+ \cup \{\neg a\})$
20. $\epsilon^- \leftarrow \{D_1\}$	# MCS computed on $(\text{KB}_h \cup \epsilon^+)^h \cup (E^-)^s$
21. Return $\langle \{C_1, C_2\}, \{D_1\} \rangle$	# model reconciling explanation for a from KB_α for KB_h

explanations and computing their probabilities can be computationally prohibitive, as even extracting a smallest monolithic explanation is in $FP^{\Sigma_2^P}$ (Ignatiev, Previti, et al. 2015) and computing the probability of a formula is #P-complete (Chavira and Darwiche 2008; Roth 1996). We can, however, account for this high computational complexity by seeking for a monolithic explanation that is guaranteed to have a probability above a certain threshold.

First, the following lemma notes that for all possible monolithic explanations \tilde{e} for explanandum φ , the following upper and lower probability bounds hold:

Lemma 2. *Let $\tilde{E}(\varphi)$ be the set of all monolithic explanations for explanandum φ from belief base \mathcal{B} , where $\tilde{e} \models \varphi$ for all $\tilde{e} \in \tilde{E}(\varphi)$, and let ω_1 be the most-probable world in which φ is true. Then, for any $\tilde{e} \in \tilde{E}(\varphi)$, it holds that $P(\omega_1) \leq P(\tilde{e}) \leq P(\varphi)$.*

PROOF. For the upper probability bound, since we assume that for all $\tilde{e} \in \tilde{E}(\varphi)$, $\tilde{e} \models \varphi$, then it must hold that for all $\tilde{e} \in \tilde{E}(\varphi)$, the worlds where \tilde{e} is true are subsumed by the worlds where φ is true (entailment property). This implies that for any $\tilde{e} \in \tilde{E}(\varphi)$, $P(\tilde{e}) \leq P(\varphi)$.

For the lower bound, since ω_1 is the most-probable world of φ , that is, the world where the highest number of formulae from \mathcal{B} are satisfied, then all monolithic explanations for φ must be true in ω_1 (i.e., $\omega_1 \models \tilde{e}$). As such, for any $\tilde{e} \in \tilde{E}(\varphi)$, $P(\tilde{e}) \geq P(\omega_1)$. \square

However, some monolithic explanations may have a higher lower probability bound. Formally, we call such explanations *k*-bounded monolithic explanations:

Definition 20 (*k*-BOUNDED MONOLITHIC EXPLANATION). Let $\tilde{E}(\varphi)$ be the set of all monolithic explanations for explanandum φ from belief base \mathcal{B} . Let $\Omega(\varphi) = \{\omega_1, \dots, \omega_n\}$ be the set of possible worlds in which φ is true, where $P(\omega_1) \geq P(\omega_2) \geq \dots \geq P(\omega_n)$. Also let $I_k = \bigcap_{i=1}^k \{\phi \mid \phi \in \mathcal{B}^{\downarrow w}, \omega_i \models \phi\}$ be the intersection of formulae that are true in worlds ω_1 to ω_k . We say that $\tilde{e} \in \tilde{E}(\varphi)$ is a *k*-bounded monolithic explanation for φ from \mathcal{B} , with lower bound $P(\tilde{e}) \geq \sum_{i=1}^k P(\omega_i)$, if and only if $\tilde{e} \subseteq I_k$.

Example 13. Consider the belief base $\mathcal{B} = \{(a, 1), (\neg a \vee b, 3), (c, 2), (\neg c \vee b, 1)\}$ and explanandum b . The two monolithic explanations for b from \mathcal{B} that entail b are $\tilde{e}_1 = \{a, \neg a \vee b\}$ and $\tilde{e}_2 = \{c, \neg c \vee b\}$, where $P(\tilde{e}_1) = 0.64$ and $P(\tilde{e}_2) = 0.77$. Notice that there are four possible worlds in which b is true: $\omega_1 = \{a = T, b = T, c = T\}$, $\omega_2 = \{a = F, b = T, c = T\}$, $\omega_3 = \{a = T, b = T, c = F\}$, and $\omega_4 = \{a = F, b = T, c = F\}$, where $P(\omega_1) = 0.57$, $P(\omega_2) = 0.20$, $P(\omega_3) = 0.07$, and $P(\omega_4) = 0.02$. The maximum number of intersections that entail b is $k = 2$ (i.e., $I_2 = \{\neg a \vee b, c, \neg c \vee b\}$). Indeed, $\tilde{e}_2 \subseteq I_2$ and $P(\tilde{e}_2) = 0.77 = P(\omega_1) + P(\omega_2)$. Finally, notice how \tilde{e}_2 is also the most-preferred monolithic explanation for b from \mathcal{B} ; for $\gamma = 0.5$, $\text{EP}(\tilde{e}_2, b) = 0.57 > \text{EP}(\tilde{e}_1, b) = 0.50$.

Proposition 4. Let \mathcal{B} be a belief base and φ an explanandum. A 1-bounded monolithic explanation \tilde{e} for φ from \mathcal{B} always exists.

PROOF. The proof follows directly from Lemma 2. □

Interestingly, there also exists a maximal *k*-bounded monolithic explanation.

Corollary 2. If $I_k \models \varphi$ and $I_{k+1} \not\models \varphi$, then $\exists \tilde{e} \subseteq I_k$ with maximal lower bound $P(\tilde{e}) \geq P(I_k)$

PROOF. First, notice that if $I_k \models \varphi$ and $I_{k+1} \not\models \varphi$, then $I_{k+j} \not\models \varphi$ for all $j = 1, \dots, n - k$. As such, k is the maximum number of intersections (from ω_1 to ω_k) such that $I_k \models \varphi$. Thus, since $\tilde{e} \models \varphi$ for all $\tilde{e} \in \tilde{E}(\varphi)$, it must be the case that there exists at least one \tilde{e} such that $\tilde{e} \subseteq I_k$, from which we know that $P(I_k) \leq P(\tilde{e})$. Moreover, as I_k is the set of formulae that are true in worlds ω_1 to ω_k , its probability must be at least equal to the sum of the probabilities of these worlds (i.e., $P(I_k) \geq \sum_{i=1}^k P(\omega_i)$). Therefore, $P(\tilde{e}) \geq P(I_k) \geq \sum_{i=1}^k P(\omega_i)$, meaning that the probability of \tilde{e} has a maximal lower bound by the top k most-probable worlds of φ . □

The utility of a *k*-bounded monolithic explanation in computing probabilistic monolithic explanations can be described as follows. If we take the top k most-probable worlds in which the explanandum φ is true, then we can prune the search space of possible monolithic explanations by taking the intersection of those worlds and checking if it entails φ —if it does, then we know that at least one monolithic explanation must be true in that world with probability at least equal to the sum of the probabilities of these top k worlds. Building on this, we now present an algorithm for computing *k*-bounded monolithic explanations for φ from \mathcal{B} , where we use Algorithm 1 as our core monolithic explanation generation engine.

Algorithm 3 describes the main steps of our approach. The important factor is the user-defined parameter \hat{k} , which dictates the number of worlds of φ to be considered. It is an integer with range $1 \leq \hat{k} \leq |\Omega(\varphi)|$, where $\Omega(\varphi)$ is the set of all possible worlds of φ . Intuitively, the larger the \hat{k} , the more exhaustive the search will be as more worlds will be considered. The algorithm starts in line 1 with k taking the user-defined value \hat{k} , and then proceeds to line 2, where it uses a weighted MaxSAT solver to find the top k most-probable worlds of φ . Note that

Algorithm 3: probabilistic-monolithic-explanation($\mathcal{B}, \varphi, \hat{k}$)**Input:** Belief base \mathcal{B} , explanandum φ , and user-defined parameter \hat{k} **Result:** A k -bounded monolithic explanation $\tilde{\epsilon}$ for φ from \mathcal{B} for some $k \leq \hat{k}$

```

1  $k \leftarrow \hat{k}$ 
2  $\Omega_\varphi \leftarrow \text{getTopKWorlds}(\mathcal{B} \cup \{(\varphi, \infty)\}, k)$  // find candidate set of formulae
3 while true do
4   // get intersecting formulae from top  $k$  worlds of  $\varphi$ 
5    $I_k \leftarrow \text{getIntersections}(\mathcal{B}^{\downarrow w}, \Omega_\varphi, k)$ 
6   if not SAT( $I_k \cup \{\neg\varphi\}$ ) then
7      $\tilde{\epsilon} \leftarrow \text{monolithic-explanation}(I_k, \varphi)$ 
8     return  $\tilde{\epsilon}$ 
9   else
10     $k \leftarrow k - 1$ 

```

(φ, ∞) denotes that φ is added to the solver as a hard constraint. The main loop of the algorithm starts in line 3. In line 4, `getIntersections` extracts the set of intersecting formulae I_k from $\mathcal{B}^{\downarrow w}$ that are true in worlds ω_1 to ω_k . If $I_k \models \varphi$, then we know that a monolithic explanation is in I_k and the algorithm proceeds to use Algorithm 1 with I_k and φ as inputs to compute and return a monolithic explanation (lines 5-7). Otherwise, the algorithm discounts k by 1 and repeats the process until a suitable k is found.

Algorithm 3 is complete in the sense that, eventually, a monolithic explanation will be returned.

Theorem 1. *Algorithm 3 is guaranteed to terminate with a solution.*

PROOF. The proof rests on the fact that, in the worst case, the parameter k will reach a value of 1. This will then correspond to the most-probable world of φ , which entails all possible monolithic explanations for φ . From Lemma 2, we know that the most-probable world of φ entails all possible monolithic explanations for φ , that is, for any $\tilde{E}(\varphi)$, $\omega_1 \models \tilde{\epsilon}$, and $\tilde{\epsilon} \subseteq I_1$. Therefore, as Algorithm 3 uses I_1 as an input to Algorithm 1, which is guaranteed to return a solution, the algorithm is also guaranteed to terminate with a solution. \square

Theorem 2. *Algorithm 3 is guaranteed to return a maximal k -bounded monolithic explanation if the user-defined parameter \hat{k} is initialized to $|\Omega(\varphi)|$.*

PROOF. First, note that if the user-defined parameter is initialized to $\hat{k} = |\Omega(\varphi)|$, then Algorithm 3 will perform an exhaustive and iterative search, starting from $k = |\Omega(\varphi)|$, to find I_k , such that $I_k \models \varphi$, and use it in Algorithm 1. Now, as the algorithm discounts k by 1 at each new iteration, eventually it will be the case that $I_k \models \varphi$ and $I_{k+1} \not\models \varphi$. From Corollary 2, we then know that $\exists \tilde{\epsilon} \subseteq I_k$ such that $P(\tilde{\epsilon}) \geq P(I_k) \geq \sum_{i=1}^k P(\omega_i)$, which means that $\tilde{\epsilon}$ corresponds to a k -bounded monolithic explanation. Therefore, the algorithm is guaranteed to return a maximal k -bounded monolithic explanation for φ . \square

6.2.2 Model Reconciling Explanations. We now move on to the case of computing probabilistic model reconciling explanations $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ for an explanandum φ from an agent knowledge base KB_α for a human belief base \mathcal{B}_h . Similarly to what we described in Section 6.2.1, Algorithm 2 can directly be used on KB_α and $\mathcal{B}_h^{\downarrow w}$ for computing

Algorithm 4: probabilistic-model-reconciling-explanation($\text{KB}_\alpha, \mathcal{B}_h, \varphi, \hat{k}$)**Input:** Knowledge base KB_α , belief base \mathcal{B}_h , explanandum φ , and user-defined parameter \hat{k} **Result:** A probabilistic model reconciling explanation $\tilde{\mathcal{E}} = \langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ for φ from KB_α for \mathcal{B}_h

```

1  $k \leftarrow \hat{k}$ 
2  $\text{KB}_\alpha^h \leftarrow \text{KB}_\alpha \cap \mathcal{B}_h^{\perp w}$ 
3  $W \leftarrow \sum_{i=1}^n \{w_i \mid (\phi_i, w_i) \in \mathcal{B}_h\}$ 
4  $\mathcal{B}_\alpha \leftarrow \{(\phi, W) \mid \phi \in \text{KB}_\alpha \setminus \text{KB}_\alpha^h\}$ 
5  $\Omega_\varphi \leftarrow \text{getTopKWorlds}(\mathcal{B}_h \cup \mathcal{B}_\alpha \cup \{(\varphi, \infty)\}, k)$ 
6 while true do
7    $I_k \leftarrow \text{getIntersections}(\text{KB}_\alpha, \Omega_\varphi, k)$ 
8   if not SAT( $(I_k \cup \text{KB}_\alpha^h \cup \{\neg\varphi\})$ ) then
9      $\langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle \leftarrow \text{model-reconciling-explanation}(I_k \cup \text{KB}_\alpha^h, \mathcal{B}_h^{\perp w}, \varphi)$ 
10    return  $\langle \tilde{\epsilon}^+, \tilde{\epsilon}^- \rangle$ 
11  else
12     $k \leftarrow k - 1$ 

```

model reconciling explanations. Additionally, the concept of a k -bounded explanation (Definition 20) can also be used to guarantee a lower bound on the probability of $\tilde{\epsilon}^+$.

Algorithm 4 shows the pseudocode of our approach. The initial computational steps are similar to those in Algorithm 3, with the exception that KB_α is now also considered in the computation of the most-probable worlds of the φ . Specifically, in line 4, KB_α is converted into a belief base \mathcal{B}_α where each formula is given a weight that is larger than the sum of weights of \mathcal{B}_h . This is to enforce these formulae to be true in the worlds of the explanandum φ . Then, \mathcal{B}_α is used in conjunction with \mathcal{B}_h to compute the top k most-probable worlds of φ (line 5). The algorithm proceeds in line 7 to extract formulae from KB_α that are true in the first k intersections of the worlds of φ . If they entail φ , the algorithm then proceeds to compute a model reconciling explanation by invoking Algorithm 2 (lines 8-9). Otherwise, the algorithm continues by discounting k by 1 and repeats the process.

Note that Algorithm 4 is complete and correct as it is based on Algorithms 2 and 3, which are complete and correct.

7 Experimental Evaluations

This section presents a comprehensive evaluation of the proposed algorithms, assessing their effectiveness and efficiency across a range of scenarios.

7.1 Experimental Setup

Experiments were conducted on a system equipped with an M1 Max processor and 32GB of memory. The algorithms were implemented in Python, utilizing the PySAT toolkit (Ignatiev, Morgado, et al. 2018) for SAT solving, MCS/MUS finding, weighted MaxSAT, and minimal hitting set computations. The time limit for all experiments was set to 500s.¹⁷

For our benchmarks, we selected a diverse set of problem instances:

¹⁷Code repository: <https://github.com/YODA-Lab/Probabilistic-Monolithic-Model-Reconciling-Explanations>.

Table 3. Number of Instances Solved (S) vs. Timed Out (T/O) by ALG1 ($\hat{k} = 1$) and ALG3 ($\hat{k} = 50, \hat{k} = 100, \hat{k} = 150, \hat{k} = 200$). Note that Runtime denotes the average runtime over all solved instances.

Parameter \hat{k}	Planning			Scheduling			Random CNF		
	S	T/O	Runtime	S	T/O	Runtime	S	T/O	Runtime
1	28	9	82.0s	30	5	80.0s	25	5	12.4s
50	32	5	79.0s	30	5	53.8s	25	5	8.6s
100	31	6	49.6s	30	5	44.7s	25	5	5.8s
150	31	6	45.5s	30	5	38.0s	25	5	3.4s
200	31	6	45.2s	30	5	37.2s	25	5	1.6s

- **Classical Planning Problems:** We encoded classical planning problems from the International Planning Competition (IPC) in the style of Kautz et al. (1996), and used them as knowledge bases. The explanandum for each problem was the plan optimality query, which we constructed as described by Vasileiou, Yeoh, et al. (2022).
- **Agent Scheduling Problems:** We encoded logic-based agent scheduling problems based on the description by Vasileiou, Xu, et al. (2023), and used them as the knowledge bases. The explanandum for each problem was a set of unsatisfied agent constraints.
- **Random CNF Problems:** We generated random CNF formulae as knowledge bases using CNFgen (Lauria et al. 2017). The explanandum for each problem was a conjunction of backbone literals,¹⁸ which we computed using the minibones algorithm proposed by Janota et al. (2015).

Note that we created associated belief bases for each problem by simply adding a random weight to each formula in the knowledge base.

7.2 Results and Discussion

We now describe and discuss our experimental results, first for monolithic explanations and then for model reconciling explanations.

7.2.1 Monolithic Explanations. We evaluated Algorithm 3, referred to as ALG3, on computing probabilistic monolithic explanations. Since the core monolithic explanation generation engine of ALG3 is powered by Algorithm 1, referred to as ALG1, we also evaluate its performance on the same instances. These experiments aim to answer the following questions:

Q1: What is the performance of ALG3 on computing monolithic explanations across different problem instances?

Q2: Does the efficacy of ALG3 change under different values of the user-defined parameter \hat{k} ?

Table 3 tabulates the instances solved (i.e., found a monolithic explanation within the time limit) and not solved (i.e., timed out) by ALG1 ($\hat{k} = 1$) and ALG3 at $\hat{k} = \{5, 100, 150, 200\}$.¹⁹ We observe that the algorithm managed to solve most instances across different values of \hat{k} . Figure 3 shows the runtime distributions of ALG3 across all values of \hat{k} for computing a monolithic explanation. Interestingly, we observe that the runtimes decrease as \hat{k} increases. This can be explained by the fact that for larger values of \hat{k} , ALG3 considers the intersections of more worlds where the explanandum is true, which means that the number of formulae that are true in these

¹⁸The backbone literals of a propositional knowledge base are the set of literals entailed by the knowledge base.

¹⁹ALG3 at $\hat{k} = 1$ corresponds to ALG1 because each encoded knowledge base is consistent and entails the explanandum. As such, all formulae in the knowledge base are true in the most-probable world of the explanandum (i.e., $\hat{k} = 1$), which means that ALG3 reduces to ALG1.

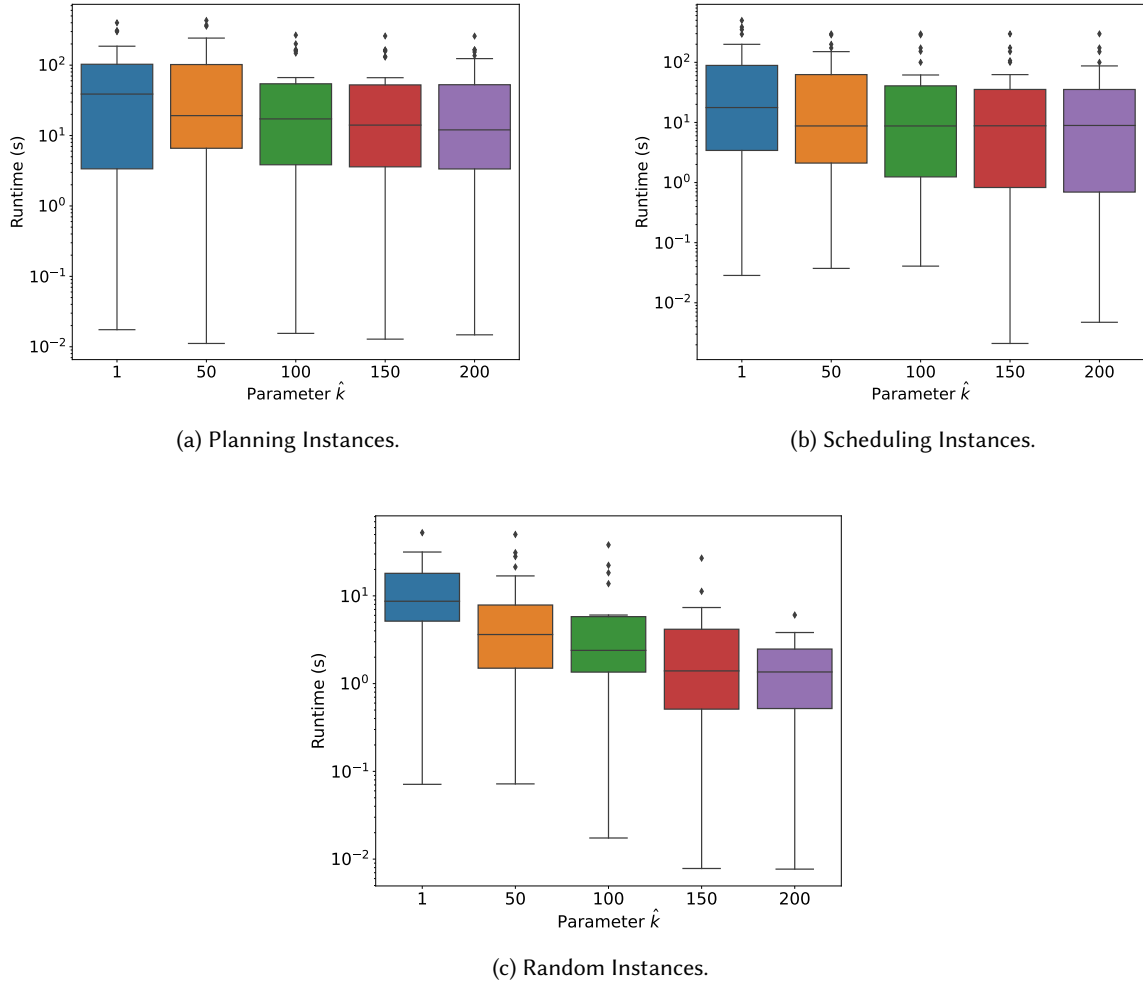


Fig. 3. Runtime distributions of ALG1 ($\hat{k} = 1$) and ALG3 ($\hat{k} = 50, \hat{k} = 100, \hat{k} = 150, \hat{k} = 200$) across all planning, scheduling, and random CNF instances.

intersections decreases. As such, the overall search space of monolithic explanations decreases as well, thus resulting in a reduced runtime needed for ALG1 to compute a monolithic explanation. This can also be observed more granularly in Figure 4, where we can see the runtime distributions of ALG1 ($\hat{k} = 1$) and ALG3 at $\hat{k} = 200$ for each instance of the planning, scheduling, and random CNF problems. Again, the runtime of ALG3 at $\hat{k} = 200$ is smaller than that of ALG1. Moreover, and as expected, in Figure 5, we can observe a positive correlation between runtime and the size of the encoded knowledge bases—as the size of the knowledge base increases, the runtimes increase as well. This is due to the fact that there is an increasing number of variables and formulae that must be considered, thus increasing the computational effort needed by the WMaxSAT, MCS, and hitting set solvers.

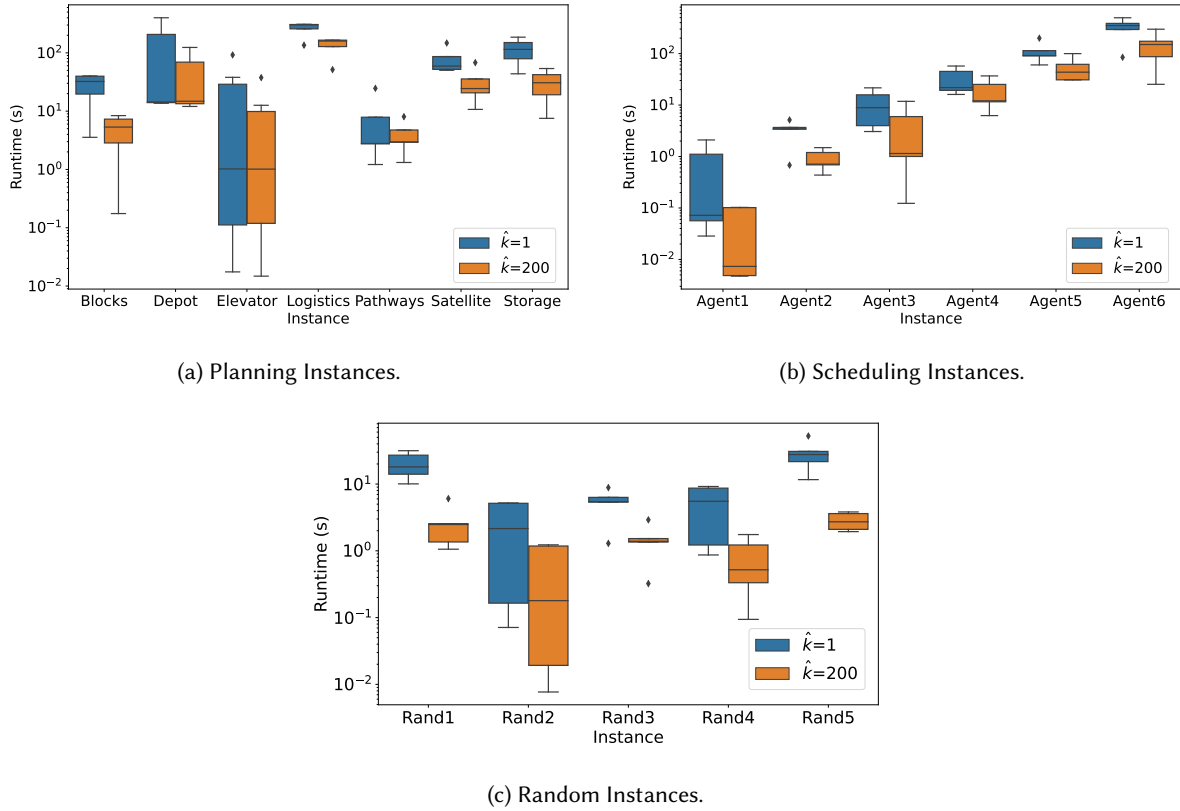




Fig. 4. Runtime distributions of ALG1 ($\hat{k} = 1$) and ALG3 ($\hat{k} = 200$) across commonly solved planning, scheduling, and random CNF instances.

All of these observations indicate the feasibility and practical efficacy of ALG3 across all benchmarks. In particular, from these experiments, we may conclude that the performance of ALG3 increases as the user-defined parameter \hat{k} increases. To reiterate, this is mainly because the overall search space of monolithic explanations that needs to be considered by ALG1 (the main monolithic explanation generation engine) decreases. Finally, it is important to note that the performance of these algorithms lies in the effectiveness of the underlying WMaxSAT, MCS, and hitting set solvers. In other words, this also implies that any advancement in those solvers will automatically reflect in performance gains in our algorithms. Thus, future work can look at efficient and optimized solvers and examine whether there is any variability in performance.


7.2.2 Model Reconciling Explanations. We now examine the effectiveness of Algorithm 2, referred to as ALG2, and Algorithm 4, referred to as ALG4, on computing model reconciling explanations. We chose the value of $\hat{k} = 200$ for ALG4 as it was the better performing parameter for ALG3 in our previous experiments. More specifically now, we are interested in scenarios with varying degrees of knowledge asymmetry between the agent and human models. To simulate such scenarios, we used the actual encoded knowledge bases as the model of the agent (KB_α), and tweaked that model and assigned it to be the model of the human (KB_h or \mathcal{B}_h). We considered the following ways to tweak the human model, resulting in the following five scenarios:



(a) Planning Instances.



(b) Scheduling Instances.



(c) Random Instances.

Fig. 5. Average runtime of ALG1 ($\hat{k} = 1$) and ALG3 ($\hat{k} = 200$) to compute an explanation across different knowledge base sizes for the planning, scheduling, and random CNF instances.

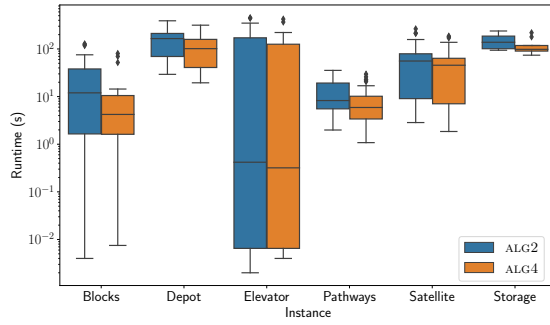
- **Scenario 1:** We randomly removed 10% of the formulae and removed 20% of literals from 10% of the total formulae in the human's model.
- **Scenario 2:** We randomly removed 20% of the formulae and removed 20% of literals from 20% of the total formulae in the human's model.
- **Scenario 3:** We randomly removed 30% of the formulae and removed 20% of literals from 30% of the total formulae in the human's model.
- **Scenario 4:** We randomly removed 40% of the formulae and removed 20% of literals from 40% of the total formulae in the human's model.
- **Scenario 5:** We randomly removed 50% of the formulae and removed 20% of literals from 50% of the total formulae in the human's model.

In general, these experiments aim to answer the following two questions:

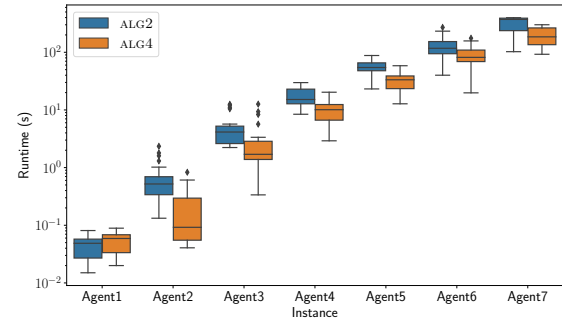
Q1: What is the performance of the algorithms on computing model reconciling explanations across different problem instances?

Table 4. Instances Solved (S) vs. Timed Out (T/O) for the Planning, Scheduling, and Random CNF Benchmarks for ALG2 and ALG4 at $\hat{k} = 200$. Note that Runtime denotes the average runtime over all solved instances.

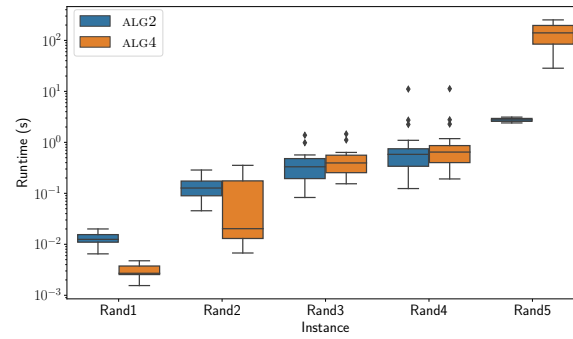
Scenario	Planning						Scheduling						Random CNF					
	ALG2			ALG4			ALG2			ALG4			ALG2			ALG4		
	S	T/O	Runtime	S	T/O	Runtime	S	T/O	Runtime	S	T/O	Runtime	S	T/O	Runtime	S	T/O	Runtime
1	25	12	67.0s	28	9	59.7s	33	2	51.4s	33	2	33.1s	27	5	30.4s	21	11	12.3s
2	25	13	69.2s	27	10	71.8s	31	4	40.9s	31	4	28.0s	26	6	18.7s	20	12	0.5s
3	24	14	67.9s	26	12	68.8s	32	3	60.3s	32	3	37.1s	29	3	20.4s	21	11	2.5s
4	25	13	82.6s	27	11	84.0s	30	4	35.7s	30	4	22.9s	23	9	5.2s	20	12	0.5s
5	22	15	84.3s	24	13	89.9s	30	4	34.4s	30	4	21.5s	24	8	3.8s	20	11	0.5s



(a) Planning Instances.



(b) Scheduling Instances.



(c) Random Instances.

Fig. 6. Runtime distributions of ALG2 and ALG4 at $\hat{k} = 200$ to compute an explanation across all commonly solved instances.

Q2: What is the performance of the algorithms in scenarios with varying degrees of knowledge asymmetry between the agent and the human model?

Table 4 tabulates the instances solved and timed out by ALG2 and ALG4 at $\hat{k} = 200$ across the five scenarios, where we observe the following trends. For the planning instances, the runtime of both algorithms increases as the difference between the models of the agent and human increases (Scenarios 1 to 5), since both algorithms search over the explanation search space, which increases as the number of differences between the two models increases. As in the previous experiments, ALG4 at $\hat{k} = 200$ yields faster runtimes than ALG2. For the scheduling instances,

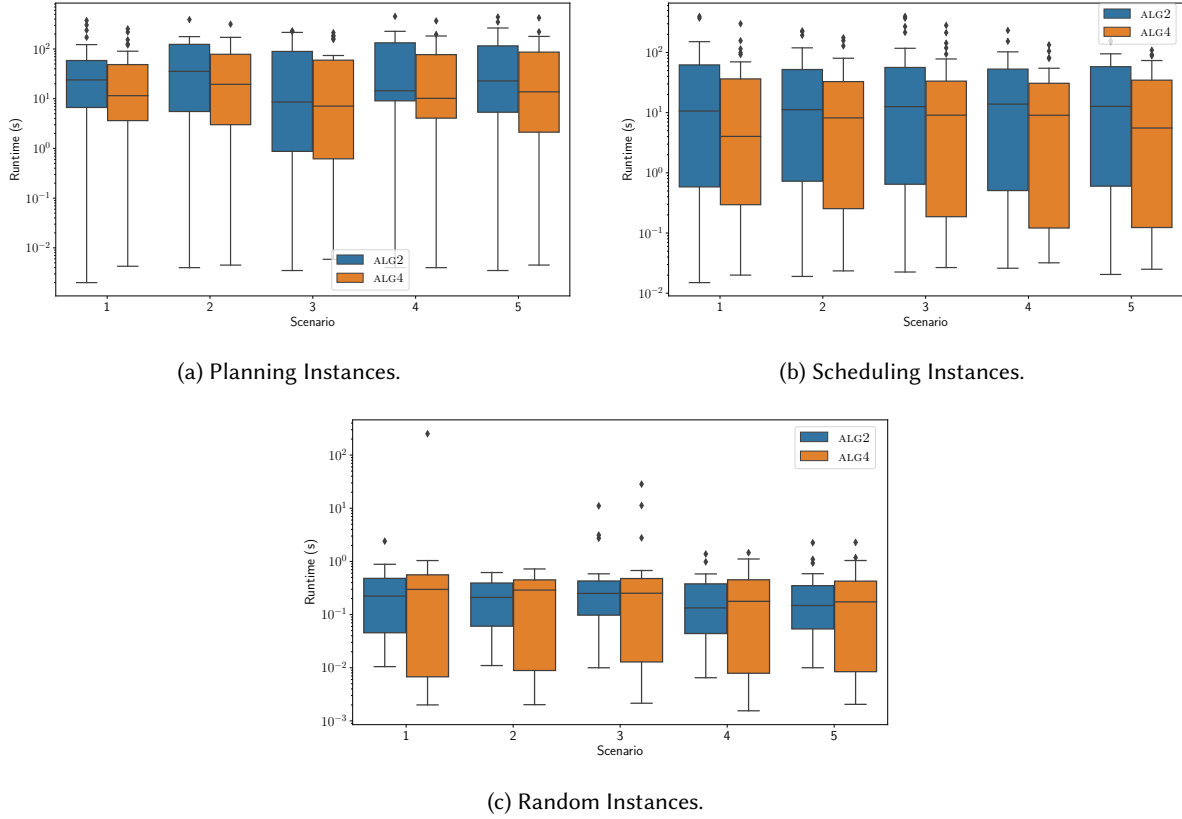




Fig. 7. Runtime distributions of ALG2 and ALG4 at $\hat{k} = 200$ to compute an explanation across commonly solved instances in each of the five scenarios.

we observe that the runtimes increase from Scenario 1 to 3, but decrease from Scenarios 4 to 5. Upon closer inspection, this is mainly because the instances solved in these scenarios were easier (i.e., smaller knowledge base sizes) than those solved in the other three scenarios, thus resulting in smaller average runtimes. A similar trend is observed for the random CNF instances. However, in the random CNF instances, ALG2 managed to solve more instances than ALG4. After examining them more closely, we found that the main bottleneck of ALG4 in those instances was computing the most-probable worlds of the explanandum (i.e., the WMaxSAT solver). Even for smaller values of \hat{k} , the solver failed to compute all the worlds under the specified time limit—the increase in search space (e.g., because of considering \mathcal{B}_α and \mathcal{B}_h) increased the complexity of these instances. We expect that an optimized and more dedicated solver may be able to overcome this limitation. The runtime distributions for ALG2 and ALG4 across all commonly solved instances and across commonly solved instances in each scenario can be seen in Figures 6 and 7, respectively. For these instances, we observe, like in the previous experiments, that ALG4 has faster runtimes than ALG2.


Moreover, in Figure 8 we see the distributions of the model reconciling explanation lengths computed by both algorithms. As expected, the general trend is that the size of the explanation ϵ^+ (i.e., formulae from KB_α for KB_h (or \mathcal{B}_h)) increases with each scenario, as the difference between the agent and human models increase. The same trend can be seen for ϵ^- —each scenario from 1 to 5 has an increasing amount of inconsistencies between the




(a) ALG2 on Planning Instances.




(b) ALG4 on Planning Instances.




(c) ALG2 on Scheduling Instances.



(d) ALG4 on Scheduling Instances.



(e) ALG2 on Random CNF Instances.



(f) ALG4 on Random CNF Instances.

Fig. 8. Distributions of the lengths of explanations ϵ^+ and ϵ^- computed by ALG2 and ALG4 at $k = 200$ across all planning, scheduling, and random CNF instances.

two models. Interestingly, ϵ^- was largest in the random CNF instances. This indicates that the inconsistencies between the human model and the corresponding ϵ^+ were high. That can also be used to explain why ALG4 failed to solve a subset of random CNF instances—highly inconsistent knowledge bases are considered as the most difficult instances for MaxSAT solvers.

In conclusion, the comparative analysis of ALG2 and ALG4 at $\hat{k} = 200$ across varied problem instances shows some trends in performance and computational complexity. The observed increase in runtime with the increase of differences between agent and human models underscores the direct relationship between model disparity and

the explanation search space size. Notably, ALG4 consistently outperforms ALG2 in terms of runtime across most scenarios, except in certain random CNF instances where the computation of most-probable worlds becomes a bottleneck due to the limitations of the WMaxSAT solver. This highlights a potential area for further optimization and development of more efficient solvers. Furthermore, the analysis of model reconciling explanation lengths reveals an expected increase in inconsistency measures as the model differences widen, particularly highlighted in random CNF instances.

8 Related Work

We now provide a discussion of related work from the planning and knowledge representation and reasoning (KR) literature. We focus on these two areas as our approach is motivated by the model reconciliation problem introduced by the planning community and it bears some similarity to other logic-based approaches in KR.

8.1 Related Planning Work

We have briefly discussed the body of work within *explainable AI planning* (XAIP) on the topic of model reconciliation in Section 1 and situated our work within that literature.²⁰ For model reconciling explanations, a similar line of research that considers uncertainty about the human user’s model is that of Sreedharan et al. (2018). In their work, they focus on scenarios where the human’s model is located within a space of possible human models that the agent has, and their method operates in that space to find conformant explanations, i.e., explanations applicable to a set of possible models. In contrast, in our approach we represent an uncertain human model as a probability distribution that the agent has, and we have explicitly defined the notion of probabilistic explanation as well as metrics to measure its quality.

For monolithic explanations, most XAIP works have focused on contrastive explanations in deterministic settings. These explanations take the form of “Why not A (instead of B)?”, where A is an alternative (or foil) suggested by the human to a decision B proposed by the agent. Contrastive explanations have found applications in linear temporal logic systems (Kasenberg et al. 2020), general epistemic accounts (Belle 2023), multi-agent optimization problems (Zehtabi et al. 2024), and in oversubscription planning (Eifler et al. 2020). There have also been used to provide a taxonomy of user questions that often arise during interactive plan exploration (Krarup et al. 2021), as well as towards creating user interfaces for decision support systems (Karthik et al. 2021; Kumar et al. 2022).

8.2 Related KR Work

In the monolithic case, our definition of a probabilistic explanation (Definition 10) may appear similar to what was proposed by Gärdenfors (1988). However, an important distinction is that Gärdenfors is dealing with epistemic states that do not contain the explanandum, while we are dealing with belief bases that do contain the explanandum. We also define a different notion of explanatory power as well as present algorithms for computing explanations. Chajewska and Halpern (1997) have also considered the problem of defining what constitutes an explanation in probabilistic systems, however they solely focus on epistemic states defined over causal structures.

The notion of (monolithic) explanation has also been explored by the *probabilistic logic programming* (PLP) community (De Raedt and Kersting 2008; Fierens et al. 2011), a formalism that extends logic programming languages with probabilities. In PLP, explanations have been associated with possible worlds.²¹ The most prominent task there is that of the most probable explanation (MPE), which consists of finding the world with the highest probability given some evidence (Shterionov et al. 2015). However, a world does not show the chain of inferences of a given explanandum and it is not minimal by definition, since it usually includes a (possibly

²⁰For a more complete discussion on XAIP related work, please see Section 9 of Vasileiou, Yeoh, et al. (2022).

²¹Recall that a possible world is a truth value assignment to the atoms in the language.

large) number of probabilistic facts whose truth value is irrelevant for the explanandum. An alternative approach is using the proof of an explanandum as an explanation (Kimmig et al. 2011), where a proof is a (minimal) partial world ω' such that for all worlds $\omega \supseteq \omega'$, the explanandum is true in ω . In this case, one can easily ensure minimality, but even if the partial world contains no irrelevant facts, it is still not easy to determine the chain of inferences behind a given explanandum. Finally, Renkens et al. (2014) have leveraged explanations in PLP as approximation techniques for more efficiently computing weighted model counting problems.

It is important to mention that logic-based methods have also been employed in machine learning problems (Marques-Silva and Ignatiev 2022). For example, Shrotri et al. (2022) proposed CLIME, a constraint-driven explanation framework for black-box ML models that enables users to specify Boolean constraints to guide the perturbation phase when generating explanations. Unlike their approach that focuses on constraining the input space during perturbation, our approach integrates uncertainty directly into the explanation generation process through probabilistic logic. Both approaches use concepts from formal methods, but our work addresses the specific case of explanations in probabilistic scenarios and model reconciliation, providing a complementary perspective to CLIME's constraint-driven framework. A similar line of research is the work by Izza et al. (2023), where they introduced an approach to computing probabilistic explanations across various classifier types (such as decision trees and naive Bayes classifiers). Their framework addresses the challenge of generating minimal sets of features that guarantee a prediction with a certain probability threshold. While Izza et al. focus on efficiently computing explanations for specific classifier architectures, our framework provides a more general theoretical foundation for explanation generation in uncertain environments and extends to the model reconciliation problem.

In the model reconciliation case, we have extended our previous work on the logic-based model reconciliation problem (Vasileiou, Yeoh, et al. 2022) to handle scenarios where the human model is uncertain. To the best of our knowledge, the application of probabilistic explanations in the context of model reconciliation that we consider in this work is novel.

Finally, the algorithms presented in this paper are an extension of our previous work (Vasileiou, Previti, et al. 2021). Specifically, Algorithms 1 and 2 are inspired by a procedure for computing a *smallest minimal unsatisfiable set* (SMUS) of an inconsistent formula, originally presented by Ignatiev, Previti, et al. (2015). The method is also related to other similar approaches for enumerating MUSes and *minimal correction sets* (MCSes). Moreover, our approach is similar in spirit to the HS-tree presented by Reiter (1987). Although the original purpose was to enumerate diagnoses, Reiter's procedure can be easily adapted to enumerate MUSes (called conflicts in that paper) as already noted by Previti and Marques-Silva (2013). However, the computation of an SMUS might require more substantial modifications. Procedures like the one presented by Reiter, which target MCSes (diagnoses) instead of MUSes (conflicts), can be seen as the dual version of our algorithm. In particular, the algorithm MaxHS (Davies and Bacchus 2011) applies the same idea of iteratively computing and testing a minimal hitting set for the computation of a MaxSAT solution (the complement of the smallest MCSes).

There are other approaches that exploit the duality between MUSes and MCSes, but instead of iteratively checking if the current hitting set is an MUS, they first compute the set of all MCSes (Liffiton and Sakallah 2008). This has the potential advantage that once all the MCSes are known, every minimal hitting set on the collection of all MCSes is guaranteed to be an MUS (Proposition 1). However, as the number of MCSes is, in the worst case, exponential in the size of the formula, this approach might fail even before reporting the first MUS. This is particularly unnecessary when the target is to return a single support, like the one presented in this paper.

9 Discussion

In this paper, we presented a logic-based framework for generating probabilistic explanations within uncertain knowledge bases (e.g., belief bases). We distinguished between two notions of explanation—*monolithic*, where explanations are defined with respect to an agent knowledge base, and *model reconciling*, where explanations

are defined with respect to an agent knowledge base as well as a human user knowledge base. Particularly, we introduced formal definitions for probabilistic monolithic and probabilistic model reconciling explanations, along with metrics to evaluate their quality (explanatory gain and explanatory power).

Our framework operates under several assumptions that we need to address. First, a fundamental assumption is that the problem domain can be encoded in a logical language. This assumption holds for many structured decision-making scenarios but may not apply to domains where knowledge is primarily subsymbolic or where the underlying decision processes are difficult to formalize in a logical language. The general effectiveness of our approach thus depends critically on how well the domain can be represented within a logical framework. Further, we used propositional (probabilistic) logic as our language of choice due to its simplicity, capacity to encode a plethora of domains, as well as computational efficiency (e.g., through SAT solvers). While our methods and algorithms are presented within propositional logic, their underlying principles are broadly applicable to any constraint system where the satisfiability of subsets can be decided. This opens avenues for extending our work to other logical systems, such as Markov Logic Networks (Richardson and Domingos 2006) and Probabilistic Logic Programs (Fierens et al. 2011).

For model reconciling explanations specifically, our framework makes some assumptions about human inferential capabilities. Notably, we presume that human users possess the reasoning capacity to process and understand the explanations provided by the AI agent. However, this does not imply that humans can reason with the same efficiency as AI agents—much like how humans can perform arithmetic operations correctly but are significantly slower than calculators. Rather, we assume that given sufficient time and a properly formulated explanation, humans can validate the logical correctness of the AI agent’s decision-making process.

Another big assumption in model reconciliation is that the agent has direct access to the human’s belief base. This assumption, while necessary for the operationalizing our framework, is often unrealistic in practice. Thus, a promising direction for future research involves approximating the human model through iterative interactions. Indeed, in a recent work (Tang et al. 2025), we showed how (propositional) probabilistic logic can be effectively used to learn a human model based on past, dialogue-based interactions. This line of work provides a fruitful ground for our framework to create more personalized and effective model reconciliation processes.

At the other end of the spectrum, a practical challenge of explainable decision-making frameworks involves how the explanations are communicated to human users.²² While we tackled the problem of generating explanations with formal guarantees about their quality, we did not directly explore the presentation of these explanations to users. There are several avenues to explore here. First, the emergence of large language models (LLMs) (Brown et al. 2020) may offer an immediate solution to this problem. Specifically, one can use LLMs as translation engines between logic and natural language, and thus create hybrid systems that generate natural language explanations with formal guarantees (provided by logical framework). For example, we presented a system for tackling this idea in the domain of coursework planning (Vasileiou and Yeoh 2025a,b). In future work, we plan to extend it for general problems expressible in propositional logic. It is also worth emphasizing that explanations can also be communicated in other forms, such as visualizations, something that has been shown to improve the users’ understanding when compared to text-based explanations alone (Karthik et al. 2021; Kumar et al. 2022).

Finally, it is important to acknowledge the resurgence of hybrid approaches that combine symbolic and neural-based methods (A. D. Garcez and Lamb 2023; A. S. d. Garcez et al. 2009). These *neuro-symbolic* AI systems aim to synergize the strengths of both paradigms: the robust learning capabilities and pattern recognition of neural networks with the interpretability and reasoning power of symbolic systems. We position the methods proposed

²²Note that, in certain settings, it has been shown that explanations in the form of model reconciliation are well understood and preferred by human users (Chakraborti, Sreedharan, Grover, et al. 2019; Kumar et al. 2022; Vasileiou and Yeoh 2023; Vasileiou, Yeoh, et al. 2022; Zahedi et al. 2019).

in this work as complementary to the advancements in neuro-symbolic AI, potentially informing and enhancing the explainability aspects of future neuro-symbolic systems.

10 Concluding Remarks

In this paper, we attempted to bridge the gap between classical explanation models and the inherent uncertainty found in real-world scenarios. We started by describing a framework for generating *probabilistic monolithic explanations* within uncertain knowledge bases (e.g., belief bases), and introduced the concepts of *explanatory gain* and *explanatory power* as quantitative measures to evaluate the effectiveness and relevance of explanations, thus offering a better characterization of explanation quality. Additionally, we presented an extension to the model reconciliation problem for generating *probabilistic model reconciling explanations*, which addresses the need for reconciling model differences between an agent and a human model, specifically when the human model is not known with certainty.

Furthermore, we developed algorithms that leverage the duality between minimal correction sets (MCSeS) and minimal unsatisfiable subsets (MUSeS) and demonstrated their potential for generating probabilistic explanations. Our experimental evaluations across different benchmarks underscore the effectiveness of these algorithms, suggesting they are promising in practical scenarios.

Acknowledgments

We thank the anonymous reviewers, whose suggestions improved the quality of our paper. Stylianos Loukas Vasileiou and William Yeoh are partially supported by the National Science Foundation (NSF) under award 2232055. Tran Cao Son is partially supported by NSF grants 1914635 and 2151254 as well as ExpandAI grant 2025-67022-44266. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, or the United States government.

References

- D. Alvarez Melis and T. Jaakkola. 2018. “Towards Robust Interpretability with Self-Explaining Neural Networks.” In: *Proceedings of the Conference on Neural Information Processing Systems*, 7775–7784.
- C. Ansótegui, M. Boffill, M. Palahí, J. Suy, and M. Villaret. 2011. “Satisfiability modulo theories: An efficient approach for the resource-constrained project scheduling problem.” In: *Proceedings of the Symposium on Abstraction, Reformulation and Approximation*, 2–9.
- F. Bacchus. 1990. *Representing and Reasoning with Probabilistic Knowledge - A Logical Approach to Probabilities*. MIT Press.
- V. Belle. 2023. “Counterfactual explanations as plans.” In: *Proceedings of the International Conference on Logic Programming*.
- V. Belle. 2017. “Logic meets Probability: Towards Explainable AI Systems for Uncertain Worlds.” In: *Proceedings of the International Joint Conference on Artificial Intelligence*, 5116–5120.
- V. Belle. 2025. “On the relevance of logic for AI, and the promise of neuro-symbolic learning.” *Neurosymbolic Artificial Intelligence*.
- V. Belle and I. Papantonis. 2021. “Principles and practice of explainable machine learning.” *Frontiers in Big Data*, 4, 688969.
- P. Besnard, É. Grégoire, C. Piette, and B. Raddaoui. 2010. “MUS-based Generation of Arguments and Counter-Arguments.” In: *Proceedings of International Conference on Information Reuse and Integration*, 239–244.
- P. Besnard and A. Hunter. 2001. “A Logic-based Theory of Deductive Arguments.” *Artificial Intelligence*, 203–235.
- Y. M. Bishop, S. E. Fienberg, and P. W. Holland. 2007. *Discrete Multivariate Analysis: Theory and Practice*. Springer Science & Business Media.
- T. Brown et al.. 2020. “Language Models are Few-Shot Learners.” In: *Proceedings of the Conference on Neural Information Processing Systems*, 1877–1901.
- M. Cashmore, M. Fox, and E. Giunchiglia. 2012. “Planning as Quantified Boolean Formula.” In: *Proceedings of the European Conference on Artificial Intelligence*, 217–222.
- M. Cashmore, D. Magazzeni, and P. Zehtabi. 2020. “Planning for hybrid systems via Satisfiability Modulo Theories.” *Journal of Artificial Intelligence Research*, 67, 235–283.
- U. Chajewska and J. Y. Halpern. 1997. “Defining explanation in probabilistic systems.” In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 62–71.

- T. Chakraborti, S. Sreedharan, S. Grover, and S. Kambhampati. 2019. "Plan Explanations as Model Reconciliation – An Empirical Study." In: *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, 258–266.
- T. Chakraborti, S. Sreedharan, Y. Zhang, and S. Kambhampati. 2017. "Plan Explanations as Model Reconciliation: Moving Beyond Explanation as Soliloquy." In: *Proceedings of the International Joint Conference on Artificial Intelligence*, 156–163.
- M. Chavira and A. Darwiche. 2008. "On probabilistic inference by weighted model counting." *Artificial Intelligence*, 172, 6-7, 772–799.
- J. Crawford and A. Baker. 1994. "Experimental results on the application of satisfiability algorithms to scheduling problems." In: *Proceedings of the National Conference on Artificial Intelligence*, 1092–1097.
- J. Davies and F. Bacchus. 2011. "Solving MAXSAT by Solving a Sequence of Simpler SAT Instances." In: *Proceedings of the International Conference on Principles and Practice of Constraint Programming*, 225–239.
- L. De Raedt and K. Kersting. 2008. "Probabilistic inductive logic programming." In: *Probabilistic Inductive Logic Programming*, 1–27.
- S. Della Pietra, V. Della Pietra, and J. Lafferty. 2002. "Inducing features of random fields." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19, 4, 380–393.
- E. Demirović, N. Musliu, and F. Winter. 2019. "Modeling and solving staff scheduling with partial weighted MaxSAT." *Annals of Operations Research*, 275, 79–99.
- Y. Dong, H. Su, J. Zhu, and B. Zhang. 2017. "Improving interpretability of deep neural networks with semantic information." In: *Proceedings of the IEEE/CVF Computer Vision and Pattern Recognition Conference*, 4306–4314.
- R. Eifler, M. Cashmore, J. Hoffmann, D. Magazzeni, and M. Steinmetz. 2020. "A new approach to plan-space explanation: Analyzing plan-property dependencies in oversubscription planning." In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 9818–9826.
- D. Fierens, G. V. den Broeck, B. G. I. Thon, and L. D. Raedt. 2011. "Inference in probabilistic logic programs using weighted CNF's." In: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 211–220.
- M. Fox, D. Long, and D. Magazzeni. 2017. "Explainable Planning." CoRR, abs/1709.10256.
- A. D. Garcez and L. C. Lamb. 2023. "Neurosymbolic AI: The 3rd wave." *Artificial Intelligence Review*, 56, 11, 12387–12406.
- A. S. d'Avila Garcez, L. C. Lamb, and D. M. Gabbay. 2009. *Neural-Symbolic Learning Systems*. Springer.
- P. Gärdenfors. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. The MIT press.
- L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. Specter, and L. Kagal. 2018. "Explaining Explanations: An Overview of Interpretability of Machine Learning." In: *Proceedings of the International Conference on Data Science and Advanced Analytics*, 80–89.
- D. H. Glass. 2023. "How good is an explanation?" *Synthese*, 201, 2, 53.
- I. J. Good. 1968. "Corroboration, explanation, evolving probability, simplicity and a sharpened razor." *The British Journal for the Philosophy of Science*, 19, 2, 123–143.
- I. J. Good. 1960. "Weight of evidence, corroboration, explanatory power, information and the utility of experiments." *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 22, 2, 319–331.
- R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. 2018. "A Survey of Methods for Explaining Black Box Models." *ACM Computing Survey*, 51, 5, 93:1–93:42.
- D. Gunning, M. Stefik, J. Choi, T. Miller, S. Stumpf, and G.-Z. Yang. 2019. "XAI—Explainable artificial intelligence." *Science Robotics*, 4, 37.
- A. Ignatiev, A. Morgado, and J. Marques-Silva. 2018. "PySAT: A Python Toolkit for Prototyping with SAT Oracles." In: *Proceedings of the International Conference on Theory and Applications of Satisfiability Testing*, 428–437.
- A. Ignatiev, A. Previti, M. H. Liffiton, and J. Marques-Silva. 2015. "Smallest MUS Extraction with Minimal Hitting Set Dualization." In: *Proceedings of the International Conference on Principles and Practice of Constraint Programming*, 173–182.
- Y. Izza, X. Huang, A. Ignatiev, N. Narodytska, M. Cooper, and J. Marques-Silva. 2023. "On computing probabilistic abductive explanations." *International Journal of Approximate Reasoning*, 159, 108939.
- M. Janota, I. Lynce, and J. Marques-Silva. 2015. "Algorithms for computing backbones of propositional formulae." *AI Communications*, 28, 2, 161–177.
- B. Juba, H. S. Le, and R. Stern. Nov. 2021. "Safe Learning of Lifted Action Models." In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*. (Nov. 2021), 379–389.
- S. Kambhampati. 1990. "A classification of plan modification strategies based on coverage and information requirements." In: *AAAI Spring Symposium Series*.
- V. Karthik, S. Sreedharan, S. Sengupta, and S. Kambhampati. 2021. "RADAR-X: An interactive interface pairing contrastive explanations with revised plan suggestions." In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 16051–16053.
- D. Kasenberg, R. Thielstrom, and M. Scheutz. 2020. "Generating explanations for temporal logic planner decisions." In: *Proceedings of the International Conference on Automated Planning and Scheduling*, 449–458.
- H. Kautz, D. McAllester, and B. Selman. 1996. "Encoding plans in propositional logic." In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, 374–384.
- A. Kimmig, B. Demoen, L. De Raedt, V. S. Costa, and R. Rocha. 2011. "On the implementation of the probabilistic logic programming language ProbLog." *Theory and Practice of Logic Programming*, 11, 2-3, 235–262.

- B. Krarup, S. Krivic, D. Magazzeni, D. Long, M. Cashmore, and D. E. Smith. 2021. “Contrastive explanations of plans through model restrictions.” *Journal of Artificial Intelligence Research*, 72, 533–612.
- A. Kumar, S. L. Vasileiou, M. Bancelhon, A. Ottley, and W. Yeoh. 2022. “Vizxp: A visualization framework for conveying explanations to users in model reconciliation problems.” In: *Proceedings of the International Conference on Automated Planning and Scheduling*, 701–709.
- P. Langley. 2016. “Explainable agency in human-robot interaction.” In: *AAAI Fall Symposium Series*.
- M. Lauria, J. Elffers, J. Nordström, and M. Vinyals. 2017. “CNFgen: A generator of crafted benchmarks.” In: *Proceedings of the International Conference on Theory and Applications of Satisfiability Testing*, 464–473.
- C. M. Li and F. Manyà. 2009. “MaxSAT, Hard and Soft Constraints.” *Handbook of Satisfiability*, 185, 613–631.
- M. H. Liffiton and K. A. Sakallah. 2008. “Algorithms for computing minimal unsatisfiable subsets of constraints.” *Journal of Automated Reasoning*, 40, 1, 1–33.
- M. H. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. 2016. “Fast, flexible MUS enumeration.” *Constraints*, 21, 2, 223–250.
- M. L. Littman. 1997. “Probabilistic propositional planning: Representations and complexity.” In: *Proceedings of the National Conference on Artificial Intelligence*, 748–754.
- J. Marques-Silva and A. Ignatiev. 2022. “Delivering trustworthy AI through formal XAI.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 12342–12350.
- J. Marques-Silva, F. Heras, M. Janota, A. Previti, and A. Belov. 2013. “On Computing Minimal Correction Subsets.” In: *Proceedings of the International Joint Conference on Artificial Intelligence*, 615–622.
- I. G. Mocuano and V. Belle. 2023. “Knowledge representation and acquisition in the era of large language models: Reflections on learning to reason via PAC-Semantics.” *Natural Language Processing Journal*, 5, 100036.
- V. Nguyen, S. L. Vasileiou, T. C. Son, and W. Yeoh. 2020. “Explainable planning using answer set programming.” In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, 662–666.
- D. Petkovic, R. Altman, M. Wong, and A. Vigil. 2018. “Improving the explainability of random forest classifier—user centered approach.” In: *Proceedings of the Pacific Symposium on Biocomputing*, 204–215.
- H. Prakken. 2006. “Formal systems for persuasion dialogue.” *The Knowledge Engineering Review*, 21, 2, 163–188.
- A. Previti and J. Marques-Silva. 2013. “Partial MUS Enumeration.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 818–825.
- A. Rago, H. Li, and F. Toni. 2023. “Interactive Explanations by Conflict Resolution via Argumentative Exchanges.” In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, 582–592.
- R. Reiter. 1987. “A theory of diagnosis from first principles.” *Artificial Intelligence*, 32, 1, 57–95.
- J. Renkens, A. Kimmig, G. Van den Broeck, and L. De Raedt. 2014. “Explanation-based approximate weighted model counting for probabilistic logics.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2490–2496.
- M. Richardson and P. Domingos. 2006. “Markov logic networks.” *Machine Learning*, 62, 1, 107–136.
- D. Roth. 1996. “On the hardness of approximate reasoning.” *Artificial Intelligence*, 82, 1-2, 273–302.
- A. A. Shrotri, N. Narodytska, A. Ignatiev, K. S. Meel, J. Marques-Silva, and M. Y. Vardi. 2022. “Constraint-driven explanations for black-box ML models.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 8304–8314.
- D. Shterionov, J. Renkens, J. Vlasselaer, A. Kimmig, W. Meert, and G. Janssens. 2015. “The most probable explanation for probabilistic logic programs with annotated disjunctions.” In: *Proceedings of the International Conference on Inductive Logic Programming*, 139–153.
- T. C. Son, V. Nguyen, S. L. Vasileiou, and W. Yeoh. 2021. “Model reconciliation in logic programs.” In: *Proceedings of the European Conference on Logics in Artificial Intelligence*, 393–406.
- S. Sreedharan, T. Chakraborti, and S. Kambhampati. 2018. “Handling Model Uncertainty and Multiplicity in Explanations via Model Reconciliation.” In: *Proceedings of the International Conference on Automated Planning and Scheduling*, 518–526.
- Y. Tang, S. L. Vasileiou, and W. Yeoh. 2025. “Does Your AI Agent Get You? A Personalizable Framework for Approximating Human Models from Argumentation-based Dialogue Traces.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 14405–14413.
- S. L. Vasileiou, A. Kumar, W. Yeoh, T. C. Son, and F. Toni. 2024. “Dialectical reconciliation via structured argumentative dialogues.” In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, 777–787.
- S. L. Vasileiou, A. Previti, and W. Yeoh. 2021. “On exploiting hitting sets for model reconciliation.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 6514–6521.
- S. L. Vasileiou, B. Xu, and W. Yeoh. 2023. “A Logic-based Framework for Explainable Agent Scheduling Problems.” In: *Proceedings of the European Conference on Artificial Intelligence*, 2402–2410.
- S. L. Vasileiou and W. Yeoh. 2023. “PLEASE: Generating Personalized Explanations in Human-Aware Planning.” In: *Proceedings of the European Conference on Artificial Intelligence*, 2411–2418.
- S. L. Vasileiou and W. Yeoh. 2025a. “TRACE-cs: A Hybrid Logic–LLM System for Explainable Course Scheduling.” In: *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*.
- S. L. Vasileiou and W. Yeoh. 2025b. “TRACE-cs: A Synergistic Approach to Explainable Course Scheduling Using LLMs and Logic.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*, 29706–29708.

- S. L. Vasileiou, W. Yeoh, T. C. Son, A. Kumar, M. Cashmore, and D. Magazzeni. 2022. “A Logic-Based Explanation Generation Framework for Classical and Hybrid Planning Problems.” *Journal of Artificial Intelligence Research*, 73, 1473–1534.
- H. L. Younes and M. L. Littman. 2004. “PPDDL1.0: The language for the probabilistic part of IPC-4.” In: *Proceedings of the International Planning Competition*, 70–74.
- Z. Zahedi, A. Olmo, T. Chakraborti, S. Sreedharan, and S. Kambhampati. 2019. “Towards understanding user preferences for explanation types in model reconciliation.” In: *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, 648–649.
- P. Zehtabi, A. Pozanco, A. Bolch, D. Borrajo, and S. Kraus. 2024. “Contrastive Explanations of Centralized Multi-agent Optimization Solutions.” In: *Proceedings of the International Conference on Automated Planning and Scheduling*, 671–679.

A Illustrative Example: Simple Package Delivery Problem

In Section 4, we presented a simple probabilistic planning scenario (Littman 1997) that can be used by our framework. This problem can be fully specified in the form of the probabilistic planning domain definition language (PPDDL) (Younes and Littman 2004), as shown in Listings ?? and 2. Given the PPDDL descriptions, we can then encode this problem into (propositional) probabilistic logic (Littman 1997), similarly to the encoding for classical planning problems (cf. Kautz et al. (1996)). Specifically, for $l, l' \in \{room1, room2\}$ and $c \in \{A, B\}$, we can encode the belief base:

Initial States: (Starting states and belief about the environment)

$$\begin{aligned} &(\text{robot-at}(\text{room1})_0, \infty) \\ &(\neg \text{robot-at}(\text{room2})_0, \infty) \\ &(\neg \text{package-delivered}_0, \infty) \\ &(\text{crowded}(A), w_1) \\ &(\text{crowded}(B), w_2) \end{aligned}$$

Goal State: (Final state to be reached)

$$(\text{package-delivered}_{t_n}, \infty)$$

Action Preconditions: (Fluents that must be true at timestep t for the action to be executed)

$$\begin{aligned} &(\text{move}(l, l', c)_t \rightarrow \text{robot-at}(l)_t, \infty) \\ &(\text{deliver}(l)_t \rightarrow \text{robot-at}(l)_t, \infty) \end{aligned}$$

Deterministic Action Effects: (Fluents that become true at timestep $t + 1$ if the action executed at timestep t)

$$(\text{deliver}(l)_t \rightarrow \text{package-delivered}_{t+1}, \infty)$$

Probabilistic Action Effects: (Fluents that become true at timestep $t + 1$ with a certain probability if the action executed at timestep t)

$$\begin{aligned} &(\text{move}(l, l', c)_t \wedge \text{crowded}(c) \rightarrow \text{robot-at}(l')_{t+1} \wedge \neg \text{robot-at}(l)_{t+1}, w_3) \\ &(\text{move}(l, l', c)_t \wedge \text{crowded}(c) \rightarrow \text{robot-at}(l)_{t+1}, w_4) \\ &(\text{move}(l, l', c)_t \wedge \neg \text{crowded}(c) \rightarrow \text{robot-at}(l')_{t+1} \wedge \neg \text{robot-at}(l)_{t+1}, w_5) \\ &(\text{move}(l, l', c)_t \wedge \neg \text{crowded}(c) \rightarrow \text{robot-at}(l)_{t+1}, w_6) \end{aligned}$$

Explanatory Frame Axioms: (Fluents do not change between subsequent timesteps t and $t + 1$ unless they are effects of actions that are executed at timestep t)

$$\begin{aligned} &(\text{robot-at}(l)_t \wedge \neg \text{robot-at}(l)_{t+1} \rightarrow \text{move}(l, l', A)_t \vee \text{move}(\text{room1}, \text{room2}, B)_t, \infty) \\ &(\neg \text{robot-at}(l)_t \wedge \text{robot-at}(l)_{t+1} \rightarrow \text{move}(l', l, A)_t \vee \text{move}(l', l, B)_t, \infty) \\ &(\text{package-delivered}_t \rightarrow \text{package-delivered}_{t+1}, \infty) \\ &(\neg \text{package-delivered}_t \wedge \text{package-delivered}_{t+1} \rightarrow \text{deliver}(l)_t, \infty) \end{aligned}$$

Action exclusions: (Only one action can occur at each timestep t)

$$\begin{aligned} &(\neg \text{move}(l, l', A)_t \vee \neg \text{move}(l, l', B)_t, \infty) \\ &(\neg \text{move}(l, l', c)_t \vee \neg \text{deliver}(l)_t, \infty) \end{aligned}$$

Finally, we can extract a plan by finding an assignment of truth values that satisfies the belief base (i.e., for all timesteps $t = 0, \dots, n - 1$, there will be exactly one action a such that $a_t = \text{True}$).

Note that a formula's weight represents the log odds comparing worlds where the formula is true versus where it is false, assuming all else remains unchanged. However, when formulae share variables with each other, it is impossible to flip one formula's truth value without affecting others. Despite that, we can still determine appropriate weights through collective formula probabilities. By treating desired probabilities as empirical frequencies for maximum likelihood estimation (MLE) (Della Pietra et al. 2002), we can derive optimal weights. This approach involves specifying how often each formula should be true, using these as observed frequencies, and applying standard MLE algorithms (e.g., (Richardson and Domingos 2006)) to compute weights.

Listing 1. PPDDL Domain of Office Robot Delivery.

```
(define (domain Office-Robot-Delivery)
  (:requirements :strips :typing :probabilistic-effects)
  (:types
    location corridor)
  (:predicates
    (robot-at ?l - location)
    (connected ?l1 ?l2 - location ?c - corridor)
    (package-delivered)
    (crowded ?c - corridor)
  )

  (:action move
    :parameters (?from ?to - location ?c - corridor)
    :precondition (and
      (robot-at ?from)
      (connected ?from ?to ?c))
  )
)
```

```

      :effect (and
        (when (crowded ?c)
          (probabilistic
            p (and (robot-at ?to) (not (robot-at ?from)))
            1-p (and (robot-at ?from))))
        (when (not (crowded ?c))
          (probabilistic
            p' (and (robot-at ?to) (not (robot-at ?from)))
            1-p' (and (robot-at ?from))))
      )

    (:action deliver
      :parameters (?l - location)
      :precondition (robot-at ?l)
      :effect (package-delivered)
    )
  )
)

```

Listing 2. PPDDL Problem of Office Robot Delivery.

```

(define (problem office-delivery)
  (:domain Office-Robot-Delivery)

  (:objects
    room1 room2 - location
    A B - corridor)

  (:init
    (robot-at room1)
    (connected room1 room2 A)
    (connected room1 room2 B)
    (probabilistic p (crowded A))
    (probabilistic p' (crowded B))
  )

  (:goal (package-delivered room2))
)

```

Received 07 April 2025; accepted 11 August 2025