# Initial Steps in Planning under Qualitative Uncertainty

**Brandon Rozek, Selmer Bringsjord**

Rensselaer Polytechnic Institute (RPI)
Rensselaer AI & Reasoning Lab, Troy New York, USA
{rozekbrandon, selmerbringsjord}@gmail.com

**Abstract**

Techniques in automated planning under uncertainty capture whether an agent believes that a ground atomic formula is true, false, or uncertain; and, in some cases, the exact probability that it's true at a given state. Sometimes, however, an agent does not have access to exact probabilistic information, but is instead able to judge the uncertainty qualitatively. We take initial but substantial steps towards characterizing a variant of conformant planning based on qualitative uncertainty. Our framework, QU-STRIPS, introduces levels of belief about ground atomic formulae which stratify uncertainty ranging on the negative side from *certainly not*, to *agnostic*, and then on the positive side up to *certainly*. In order to efficiently find plans, we present a sound compilation into classical STRIPS. We provide preliminary results on a new escape domain and show that state-of-the-art planners can effectively find plans that achieve the goal at a high positive belief level, while considering the trade-off between the strength of a plan and its cost.

## Introduction

Suppose that you want to buy a gift for a close friend. This friend, however, is known for returning gifts they already have. A weak plan from the conformant-planning literature would give this friend a gift, as long as there is one possible initial state where they don't already own it. On the other hand, a strong plan would select a gift that the friend is guaranteed not to own, even if the cost of obtaining it is prohibitive. As a close friend, you're somewhat privy to what gifts they already received, and you hold various levels of belief on whether they own certain gifts. For example, perhaps it is *likely* but not *certain*, that they were gifted a box of chocolates. What's desired is a plan where you have a high level of belief that they won't return the gift, but also take into account the trade-off between likelihood of plan success and the cost of said plan.

We aim to capture in this work how to effectively make use of these qualitative beliefs in the initial state when creating plans. Often, an agent may not have access to probabilistic information concerning a given ground atomic formula, but can still provide more information than the formula's being unknown. This can be useful for example in Human-AI teaming, as it is easier for a human to judge uncertainty qualitatively than to assign a probability distribution.

Prior work in conformant and epistemic planning treats ground atomic formulae in ternary fashion: semantic values of true, false, or indeterminate are enabled.

On the other hand, probabilistic planning modulate ground atomic formulae with probabilities. These forms of reasoning about uncertainty have also been explored in the logic community and the interested reader is invited to consult (Halpern 2004). The present paper takes early steps in considering the space between ternary values and probabilities, a space in which beliefs are stratified qualitatively based on their "strengths" for use in automated planning. More specifically, we employ *cognitive likelihoods*. Some informal prior work has been done in AI with a precursor to these likelihoods [e.g.(Govindarajulu and Bringsjord 2017b)], but we present here a more formal framework, as well as one that is more general in its reach. We allow the user of our framework (and perhaps the AI operating autonomously with it) to specify the stratification of belief levels of a planning agent, and to use these levels to determine a given spectrum of belief operators, and how beliefs change through state progression.

Belief change is captured through conditional effects and three principles to guide the progression of belief. We (1) *withhold* contradictory beliefs; (2) propagate the *strongest* belief out of a set of derived beliefs about a given ground atomic formula; and (3) use the *Weakest Link Principle* (WLP) to determine the belief level derived from the operator preconditions, as well as the antecedents of the conditional effect. Our framework is a variant of conformant planning in which there is no sensing, conformant width is one, and the conditional effects of operators are deterministic. The belief levels of the goals at the end of the plan determine the plan's strength. A rational agent should seek to maximize the strength of their plan while taking into account the cost of a plan. We address this by introducing a parameter which captures the action-cost incurred, before the agent should accept a less-costly plan with a lower strength.

The contribution is encapsulated as follows. We (1) formalize planning under qualitative uncertainty for a single agent using our QU-STRIPS framework. This allows for levels of belief in the initial state, and the three principles of belief progression are used when computing the next state. We (2) provide a compilation to STRIPS, a classical planning model, and show that plans generated by an optimal planner

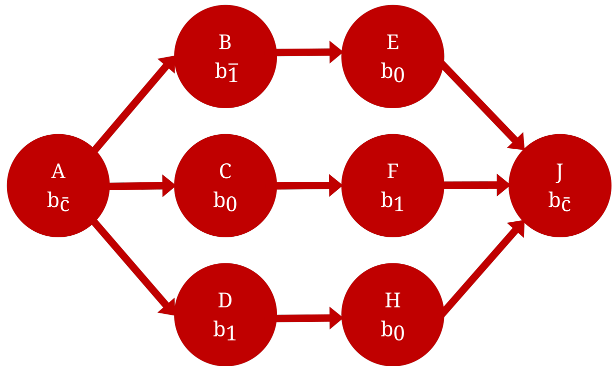


Figure 1: An Instance of the `Escape` Domain. The agent initially holds a belief at level *certain* that it's at the $A$ location and that the connectivity between nodes is as shown. Meanwhile, the agent holds beliefs at various levels concerning whether a given location has a trap. The level of such a belief is indicated under the location name. If the agent moves to a location with a trap, they'll get caught. The goal is for the agent to move to the $J$ location without getting caught.

are not only sound but also respect the preference ordering defined by the plan strength and the cost of the plan; and (3) provide preliminary empirical evaluation of our framework, and show how a small optimization can help avoid the worst-case complexity in the number of compiled operators.

To illustrate our framework, we introduce the `Escape` domain. In problems faced in this domain, an agent starts at an initial location and its goal is to navigate to a specified location through a series of connected locations. Unfortunately, each location potentially has a trap installed; it will catch the agent if it moves to that location. The agent holds qualitatively modulated belief as to whether each location has a trap. A small instance of this problem is shown pictorially in Figure 1.

## Background

**STRIPS Planning** A STRIPS problem $\Pi'$ is the tuple $\langle P', O', I', G' \rangle$ where $P'$ is the finite set of ground atomic formulae. $O'$ the finite set of operators, $I' \subseteq P'$ the initial state, and $G' \subseteq P'$ are the goals. An operator $o' \in O'$ is the tuple $\langle Pre', Add', Del', C' \rangle$ where $Pre' \subseteq P'$ are the preconditions, $Add' \subseteq P'$ the add effects, $Del' \subseteq P'$ the delete effects, and $C' \in \mathbb{N}$ is the cost of the operator. Given a state, $s' \subseteq P'$, an operator $o'$ is applicable iff $Pre'(o') \subseteq s'$. After applying an operator $o'$ on state $s'$, the next state will be $(s' - Del') \cup Add'$. A plan $\pi = (o_1, \ldots, o_n)$ is a sequence of operators. Such plan is valid for a STRIPS problem $\Pi'$ iff (1) $o_1$ is applicable for the initial state $I'$; (2) $o_{i+1}$ is applicable at $s_{i+1}$, which is the state that is the result of applying $o_i$ to $s_i$; (3) at the last state $s_n$, the goal set $G'$ is a subset of $s_n$. For a more complete treatment of STRIPS planning, see (Lifschitz 1987) for more details.

**Cognitive Likelihoods** One approach to qualitative uncertainty within multi-value epistemic logic is that of *cognitive likelihoods* (Giancola et al. 2022). Specifically, the authors define 11 levels of cognitive likelihood $B_{11}$ ranging from `certain` to `certainly not` and justify each value with an appeal to rational human-level cognition (about which more will be said below). For example, within $B_{11}$, the formula $\mathbf{B}^3(a, t, \phi)$ states that an agent $a$ at time $t$ believes `beyond reasonable doubt` that $\phi$ holds. Let $\mathbf{\Phi} = \{\phi_1, \ldots, \phi_m\}$ and $\Gamma$ be the set of formulae at the current state. The authors provide the following inference schema ($I^S_{\mathbf{B}}$) for reasoning about beliefs.

$$\frac{\mathbf{B}^{\sigma_1}(a, t_1, \phi_1), \ldots, \mathbf{B}^{\sigma_m}(a, t_m, \phi_m), \mathbf{\Phi} \vdash \phi, \mathbf{\Phi} \nvdash \perp, \mathbf{\Gamma} \vdash t_i < t}{\mathbf{B}^{min(\sigma_1, \ldots, \sigma_m)}(a, t, \phi)} \quad (I^S_{\mathbf{B}})$$

In (Giancola 2023), Giancola further elaborates that this schema makes use of WLP. The schema states that as long as the assumptions used don't derive a contradiction $\perp$, the agent can infer a belief at the weakest belief level among the beliefs used. Note that the inference schema implicitly allows for belief forwarding through time as long as it does not derive a contradiction.

We are not convinced that all 11 likelihood values are necessary to reason qualitatively about uncertainty, or that we should fix the number of levels to be used for all problems. Instead, we hone in on WLP for planning under qualitative uncertainty, and keep generic the exact number of belief levels used.

## Related Work

Uncertainty problems, or planning with incomplete information, is well-covered in the automated-planning literature. Conformant planning, coined by (Smith and Weld 1998), takes an incomplete description of the initial state and finds a sequence of operators to achieve the goal. It was formulated as a heuristic-search problem in (Bonet and Geffner 2000), where the incomplete initial state gets represented as a belief state or a finite set of states that satisfy the initial description. A compilation technique was later developed in (Palacios and Geffner 2009); it takes a conformant planning problem with bounded width, and converts it to an equivalent classical planning problem. They do this through a technique called *knowledge compilation*. This technique maps literals $t$ and $L$ to $KL/t$, which expresses that if $t$ is true in the initial situation, then $L$ must be true. The conformant width of the problem is defined as the maximum number of uncertain state variables that interact through conditional effects. Given a deterministic conformant problem of width 1, (Bonet and Geffner 2014) defines a compilation technique to classical planning that is linear. In our work presented herein, it is assumed the problem is of conformant width 1. Assumption-based planning (Davis-Mendelow, Baier, and McIlraith 2013) is an extension to conformant planning which includes a set of ground atomic formulae $U$ that are unknown in the initial state, but that the agent is allowed to assume. They introduce *preferred assumption-based plans*, where a plan is preferred to another if it makes more reasonable assumptions.

Epistemic logic is the formalization of such propositional attitudes as knowledge and belief in suitable logics, and formal exploration thereof. Epistemic planning originated from the dynamic epistemic logic (DEL) community. These works (Bolander and Andersen 2011; Bolander et al. 2020) largely focused on theoretical decidability results from applying DEL to planning. There is a sustained line of research dedicated to compilation techniques of epistemic planning (for logics including but not limited to: DEL, $KD_{45}$, BBL) to classical planning (Muise et al. 2015, 2022; Hu, Miller, and Lipovetzky 2022). The primary focus of epistemic planning has been in the multi-agent setting with nested beliefs. This allows for reasoning about other agents' beliefs and for implicit coordination or even deceptive goals. A closely related paper in this space considers planning under plausibility models (Andersen, Bolander, and Jensen 2015). This work extends DEL planning to include a partial order over states and the outcomes of operators with non-deterministic effects. They define a weak plausible plan to be a sequence of actions in which the goal is reached in some of the most plausible terminal states. Strong plausible plans are defined similarly but with the goal reached in all of the most plausible terminal states as defined by the partial orders.

Probabilistic planning in its most general form is formulated as a POMDP. Within this, the initial state is defined as ground atomic formulae with probabilistic information attached, and the effects of operators may be non-deterministic at various probabilities. Three well-studied subclasses of probabilistic planning are stochastic shortest path (SSP), MaxProp, and probabilistic conformant planning. The first two assume that the initial state is known; however, a probability distribution is defined over the effects of operators. A solution for both these classes is a policy that maps states to operators. SSP and MaxProp differ in their objectives. The former concerns itself with reaching the goal (within a probabilistic threshold) at the minimum expected cost; the latter focuses on maximizing the probability that the goal is reached. Recent papers and techniques for those problems can be seen in (Kolobov et al. 2011; Trevizan, Thiébaux, and Haslum 2017; Klößner et al. 2022). Conformant probabilistic planning (like the prior two classes) has a probability distribution over the effects of operators, but it additionally defines a probability distribution over the initial state. A solution to the conformant planning problem is a linear sequence of operators which when applied satisfy the goal within a probabilistic threshold. Examples of work in this area include (Bryce, Kambhampati, and Smith 2006; Domshlak and Hoffmann 2007; E-Martín, R-Moreno, and Smith 2014).

Returning briefly to the work of Halpern and collaborators cited above, i.e. (Halpern 2004), and elaborating a bit: As we have pointed out, this work weaves together Kolmogorovian probability with epistemic operators; as such, this approach stands outside, and bounds, qualitative measures of likelihood.[1] Interesting work along the same line that is somewhat relevant to ours involves *plausibility measures* (Friedman and Halpern 1995). But such measures, instead of mapping constructions built from *possible worlds* into the interval $[0, 1]$, map to partially ordered sets. Generally speaking, probabilities are not independently justifiable by considering rational, human-level cognition. An example is *beyond reasonable doubt*, which has a centuries-old, independent-of-probability status in Occidental legal reasoning, and which can clearly support the search for and finding of a plan (such as a verdict in a trial).

And a final point re. related work: Our use of qualitatively modulated belief as part of the basis for planning under uncertainty could be fairly viewed as harnessing a space of multi-valued epistemic logics for AI purposes. There are precious few such logics, and none of them to our knowledge are computational, but two versions of a multi-valued modal logic are presented in (Fitting 1991). However, neither version includes anything like levels of likelihood, in any sense of the term; and the modal operators are not epistemic in nature. Likewise, while (Santos 2020) gives a (non-computational) 4-valued epistemic logic built out of the logic **BK**, the four values — applied to a proposition $p$: *true*, *false*, *both*, *none* — are not in any way likelihood-modulated belief operators.

## Formalisms

In order to express qualitative uncertainty, we will take inspiration from cognitive likelihoods to define the conditions for a set $B_i$ to hold levels of belief.

**Definition 1.** *A totally ordered finite set $B_i$ represents **levels of belief** provided that the following properties hold:*

- *There exists a null element $b_0$ which represents a lack of belief for or against a given atomic ground formula.*
- *There exists an element $b_c$ which denotes that the agent is certain that an atomic ground formula holds.*
- *A total bijective inverse function $inv$ exists such that $inv(b_0) = b_0$ and for $b_x \neq b_0 \in B_i$ if $b_x > b_0$ then a corresponding $b_{\bar{x}}$ exists such that $b_{\bar{x}} < b_0$ and there are the same number of elements between both $b_x$ and $b_0$ as well as $b_{\bar{x}}$ and $b_0$. Similarly for all elements $b_x < b_0$.*

From Definition 1, note that the following properties must hold. First, $B_i$ must have a length greater than or equal to 3. Secondly, $B_i$ must be of odd length.

**Example 1.** *Let $B_5 = \{b_{\bar{c}}, b_{\bar{1}}, b_0, b_1, b_c\}$ be a totally ordered set. Then $B_5$ corresponds to levels of belief.*

With $B_i$ defined, we can modify the classical STRIPS representation to assign levels of belief to atomic ground formulae.

**Definition 2.** *A **QU-STRIPS model** $\Pi$ is the tuple $\langle B_i, P, P^\dagger, O, I, G\rangle$ where:*

- *$B_i$ correspond to the levels of belief;*
- *$P$ is the finite set of atomic ground formulae;*
- *$P^\dagger = \{\mathbf{B}(p, \sigma) \mid p \in P, \sigma \in B_i\}$ is the finite set of beliefs for each belief level $\sigma \in B_i$ and each $p \in P$;*
- *$I \subseteq P^\dagger$ is the initial state;*
- *$G \subseteq P$ is the set of goals; and*
- *$O$ is the finite set of operators.*

[1] For a particular example, the language of the logic $\mathcal{L}^{KQU}$ permits formulae saying that agents know/**K** that a formula $\phi$ has a probability $p$.

Let $p_\sigma = \mathbf{B}(p, \sigma)$. We define $ground(p_\sigma) = p$ and $strength(p_\sigma) = \sigma$. One important thing to note about the QU-STRIPS framework is that since the agent operates under uncertainty, we do not assume the *closed world*. That is, instead let $s \subseteq P^\dagger$ be a state and $p \in P$ be an atomic ground formula. If there does not exist a $p_\sigma \in s$ such that $ground(p_\sigma) = p$ and $strength(p_\sigma) \neq b_0$, then it is assumed that the agent lacks a belief for or against $p$ (i.e. $\mathbf{B}(p, b_0) \in s$.) We say that a state $s$ is *consistent* iff for every $p_i, p_j$ within $s$, if $p_i \neq p_j$ then $ground(p_i) \neq ground(p_j)$. A state $s$ *satisfies* a ground atomic formula $p$ iff there exists $p_\sigma \in s$ such that $ground(p_\sigma) = p$ and $strength(p_\sigma) > b_0$. A state $s$ *satisfies* a partial state $s_p \subseteq P$ iff for all $p \in s_p$, $s$ satisfies $p$.

**Definition 3.** *An* ***operator*** $o \in O$ *is the tuple* $\langle Pre, Add_P, Add_N, C\rangle$ *where*

- $Pre \subseteq P$ *is a partial state representing the precondition;*
- $Add_P$ *is the set of conditional-effect tuples* $\langle c_p, l_p\rangle$ *where* $c_p \subseteq P$ *and* $l_p \subseteq P$*;*
- $Add_N$ *is the set of conditional-effect tuples* $\langle c_n, l_n\rangle$ *where* $c_n \subseteq P$ *and* $l_n \subseteq P$*; and*
- $C \in \mathbb{N}$ *is the operator cost.*

In QU-STRIPS, state progression is captured by the sets $Add_P$ and $Add_N$. The former corresponds to the conditional effects that result in positive beliefs for a given atomic ground formula. Meanwhile, the later results in negative beliefs.

**Example 2.** *The following* $o = (\texttt{moveAgent A B})$ *is an operator in the escape domain.*

- $Pre(o) = \{(\texttt{atAgent A}), (\texttt{CONNECTED A B})\}$
- $Add_P(o) = \{\langle \top, (\texttt{atAgent B})\rangle\}$
- $Add_N(o) = \{\langle \top, (\texttt{atAgent A})\rangle, \langle (\texttt{atTrap B}), (\texttt{notCaught})\rangle\}$
- $C(o) = 1$

Note that belief levels are not taken into account when defining an operator. Instead, the belief levels are used for applicability and state progression. An operator $o \in O$ is *applicable* at state $s$ iff $s$ satisfies $Pre(o)$.

For state progression, we make use of three principles to determine the level of belief for a given ground atomic formula. The first principle comes directly from cognitive likelihoods and is called WLP.[2] In this context, when a conditional effect fires, the weakest (minimum in this case) belief level of all the prerequisite beliefs is assigned to the consequent ground atomic formulae. Given a state $s \subseteq P^\dagger$ and the prerequisite ground atomic formulae $conds \subseteq P$, the following equation returns the belief level according to WLP.

$$wlp(s, conds) = min(\{strength(s_i) \mid \forall s_i \in s \text{ if } ground(s_i) \in conds\})$$

For the next two principles, let us capture the set of positive and negative beliefs derived from the conditional effects after applying an applicable operator $o$ to the state $s$.

[2]Space constraints preclude discussing WLP versus probability.

$$Add_P^\dagger \subseteq P^\dagger = \{\mathbf{B}(l_i, \sigma) \mid \forall \langle c, l\rangle \in Add_P, \forall l_i \in l \text{ if } s \text{ satisfies } c \text{ where } \sigma = wlp(s, Pre \cup c)\}$$

$$Add_N^\dagger \subseteq P^\dagger = \{\mathbf{B}(l_i, \sigma) \mid \forall \langle c, l\rangle \in Add_N, \forall l_i \in l \text{ if } s \text{ satisfies } c \text{ where } \sigma = inv(wlp(s, Pre \cup c))\}$$

Consider the cases in which multiple beliefs about some ground atomic formula $p$ are derived from the conditional effects. First, suppose $B(p, \sigma_1) \in Add_P^\dagger$ and $B(p, \sigma_2) \in Add_N^\dagger$. This comes about when $\langle c_p, l_p\rangle \in Add_P(o)$ and $\langle c_n, l_n\rangle \in Add_N(o)$ with both $c_p$ and $c_n$ satisfied and $p$ within both $l_p$ and $l_n$. In this situation, we take an agnostic approach and discard those positive and negative beliefs for a belief of $B(p, b_0)$. This *withholding principle* has the agent take a belief neither for or against a ground atomic formula in the face of contradictory belief levels. In order to determine whether this situation holds, we define the following predicate:

$$withinBoth(p) \iff (\exists p_i \in Add_P^\dagger, p_j \in Add_N^\dagger, ground(p_i) = ground(p_j) = p)$$

We then use this formula to determine the ground atomic formulae the agent does not hold a belief for or against in the successive state.

$$Add_1^\dagger = \{\mathbf{B}(p, b_0) \mid \forall \mathbf{B}(p, \sigma_i) \in Add_P^\dagger \cup Add_N^\dagger \text{ if } withinBoth(p)\}$$

Secondly, consider the cases when multiple beliefs are derived about $p$ but not $withinBoth(p)$. The last principle which captures this is the *strongest belief principle*. When this occurs, as the name suggests, the agent adopts the strongest belief. This principle also applies when $B(p, \sigma_i) \in s$ and $\sigma_i$ shares the same sign. Another way of looking at this is that if $B(p, \sigma)$ is in the successive state, then there does not exist a stronger belief.

$$strongest_P(p, \sigma_i) \iff \neg\exists B(p, \sigma_j) \in Add_P^\dagger \cup s \text{ such that } \sigma_j > \sigma_i$$

$$strongest_N(p, \sigma_i) \iff \neg\exists B(p, \sigma_j) \in Add_N^\dagger \cup s \text{ such that } \sigma_j < \sigma_i$$

We take this into account when collecting the strongest positive and negative beliefs that aren't withheld.

$$Add_2^\dagger = \{\mathbf{B}(p, \sigma) \mid \forall \mathbf{B}(p, \sigma) \in Add_P^\dagger \text{ if } \neg withinBoth(p) \text{ and } strongest_P(p, \sigma)\}$$

$$Add_3^\dagger = \{\mathbf{B}(p, \sigma) \mid \forall \mathbf{B}(p, \sigma) \in Add_N^\dagger \text{ if } \neg withinBoth(p) \text{ and } strongest_N(p, \sigma)\}$$

With the three principles in place, we define state progression given an applicable operator $o$ at a state $s$.

**Definition 4.** *The function* **apply** $: (O \times S) \to S$ *takes an applicable operator* $o \in O$ *and a state* $s \in S$ *to produce a new state* $s_n$*.*

$$apply(o, s) = s_n = (s - Del^\dagger) \cup (Add^\dagger)$$

*where*

- $Add^{\dagger} = Add_1^{\dagger} \cup Add_2^{\dagger} \cup Add_3^{\dagger}$
- $Del^{\dagger} \subseteq P^{\dagger} = \{\mathbf{B}(l, \sigma) \mid \forall l'_{\sigma} \in Add^{\dagger}, \forall \sigma \in B_i$ *where* $l = ground(l'_{\sigma})$ *and* $\sigma \neq strength(l'_{\sigma})\}$.

**Example 3.** *Consider the* escape *domain with the belief set* $B_5$ *from Example 1, operator* $o$ *from Example 2, and the following state:*

$$s = \{\mathbf{B}((\texttt{atAgent A}), b_c), \mathbf{B}((\texttt{atAgent B}), b_{\bar{c}}), \mathbf{B}((\texttt{CONNECTED A B}), b_c), \mathbf{B}((\texttt{atTrap B}), b_1), \mathbf{B}((\texttt{notCaught}), b_c)\}$$

*We can see that* $o$ *is an applicable operator since the agent holds each of the preconditions at a belief level greater than* $b_0$. *Applying* $o$ *to* $s$ *will result in the following state:*

$$s_n = \{\mathbf{B}((\texttt{atAgent A}), b_{\bar{c}}), \mathbf{B}((\texttt{atAgent B}), b_c), \mathbf{B}((\texttt{CONNECTED A B}), b_c), \mathbf{B}((\texttt{atTrap B}), b_1), \mathbf{B}((\texttt{notCaught}), b_{\bar{1}})\}$$

**Lemma 1** (Operator application preserves state consistency)**.** *Let* $o$ *be an applicable operator for a consistent state* $s$. *Then* $s_n = apply(o, s)$ *is consistent.*

*Proof.* Let $p_i$ and $p_j$ be elements of $s_n$ where $p_i \neq p_j$. Recall that, $s_n = (s - Del^{\dagger}) \cup Add^{\dagger}$. We will prove that $s_n$ is consistent through cases of membership on $p_i$ and $p_j$.

- Assume that $p_i$ and $p_j$ are elements of $s$. We know that $s$ is consistent, therefore $ground(p_i) \neq ground(p_j)$. Hence, $s_n$ is consistent.
- Suppose that $p_i$ and $p_j$ are elements of $Add^{\dagger}$. Let $p_i = \mathbf{B}(l_i, \sigma_i)$. Recall that $Add^{\dagger} = Add_1^{\dagger} \cup Add_2^{\dagger} \cup Add_3^{\dagger}$. Let us do a proof by cases on membership of $p_i$.
  - Assume that $p_i \in Add_1^{\dagger}$. Then, $p_i = \mathbf{B}(l_i, b_0)$ and $withinBoth(l_i)$ holds. For sake of contradiction, let $ground(p_i) = ground(p_j)$. Then, $p_j = \mathbf{B}(l_i, \sigma_2)$. Since $withinBoth(l_i)$ holds, $p_j \in Add_1^{\dagger}$. Hence, $p_j = \mathbf{B}(l_i, b_0)$. This is a contradiction, because $p_i \neq p_j$. Hence, $ground(p_i) \neq ground(p_j)$ and $s_n$ is consistent.
  - Suppose that $p_i \in Add_2^{\dagger}$. Then, $p_i = \mathbf{B}(l_i, \sigma_1)$ and $\neg withinBoth(l_i)$. For sake of contradiction, assume that $ground(p_i) = ground(p_j)$. Then, $p_j = \mathbf{B}(l_i, \sigma_2)$. We have $p_j \in Add_2^{\dagger}$ since $\neg withinBoth(l_i)$. Then both $strongest_P(l_i, \sigma_1)$ and $strongest_P(l_i, \sigma_2)$ holds. Therefore, $\sigma_1 = \sigma_2$. This is a contradiction since $p_i \neq p_j$. Hence, $ground(p_i) \neq ground(p_j)$ and $s_n$ is consistent.
  - A similar argument can be made for $p_i \in Add_3^{\dagger}$.
- Without loss of generality, assume that $p_i \in s$ and $p_j \in Add^{\dagger}$. Let $p_j = \mathbf{B}(l_j, \sigma_j)$. Then,

  $$\{\mathbf{B}(l_j, \sigma_x) \mid \forall \sigma_x \in B_i \text{ where } \sigma_x \neq \sigma_j\} \subseteq Del^{\dagger}$$

  For indirect, let $ground(p_i) = ground(p_j)$. Then since $p_i \neq p_j$, $p_i \in Del^{\dagger}$. Therefore, $p_i \notin s_n$. This is a contradiction; hence $ground(p_i) \neq ground(p_j)$, and $s_n$ is consistent.

We have shown in all the cases $s_n$ is consistent; that is, $s_n$ is consistent after applying an applicable operator $o$ to a consistent state $s$. Hence, operator application preserves state consistency. ■

Let us define a plan $\pi$ to be a sequence of operators $(o_1, \ldots, o_n)$. Let $s_i$ be the state that results from applying $o_1, \ldots, o_{i-1}$ sequentially from some initial state $I$. Then $\pi$ is a *valid* plan for a QU-STRIPS problem $\Pi$ if for each $o_i \in \pi$, $o_i$ is applicable at $s_i$, and for all $g \in G$, $s_n$ satisfies $g$. The solution to $\Pi$ is the set of all valid plans for $\Pi$. The cost of a plan $\pi$ is the sum of the cost of all the operators $o_i \in \pi$. That is $cost(\pi) = \sum_{o_i} cost(o_i)$. The strength of a plan $\pi$ is the lowest level of belief for a given $g \in G$ within the terminal state $s_n$. That is, $min(\{strength(p_{\sigma}) \mid p_{\sigma} \in s_n$ if $ground(p_{\sigma}) \in G\})$. A rational agent should prefer plans with a higher level of belief, while also taking into account the costs of those plans.

**Lemma 2** (Plans preserve state consistency)**.** *Given a planning problem* $\Pi$ *and a valid plan* $\pi$, *if* $I \in \Pi$ *is consistent, then a valid plan will result in a terminal state* $s_n$ *that is consistent.*

*Proof.* We show this by induction on $s_i$. The base case holds since $I$ is consistent. For the inductive step, let $s_n$ be consistent with an applicable operator $o_n$. Then $s_{n+1} = apply(s_n, o_n)$. In turn, $s_{n+1}$ is consistent by Lemma 1. ■

## Compilation

To make use of efficient modern automated planners, we present an initial compilation that takes a QU-STRIPS problem $\Pi$ and converts it to a STRIPS problem $\Pi'$. For this, let $\mathbf{X_P}$ take a belief proposition $\mathbf{B}(p, \sigma)$ within $\Pi$ and convert it to a unique ground atomic formula within $\Pi'$ and let $\mathbf{X_P}^{-1}$ do the inverse.

**Definition 5.** *For a QU-STRIPS problem* $\Pi = \langle B_i, P, P^{\dagger}, O, I, G\rangle$ *the following compilation* $\mathbf{X_{\Pi}}$ *outputs a classical STRIPS problem* $\Pi' = \langle P', O', I', G'\rangle$ *where*

- $P' = \{\mathbf{X_P}(p_{\sigma}) \mid \forall p_{\sigma} \in P^{\dagger}\} \cup \{goal\}$
- $O' = O'_A \cup O'_G$
- $I' = \{\mathbf{X_P}(p_{\sigma}) \mid \forall p_{\sigma} \in I\}$
- $G' = \{goal\}$

Recall that a valid plan for $\Pi$ may satisfy a goal $g \in G$ within the terminal state at any strength greater than $b_0$. In order to account for this, the compilation creates a new goal operator for every possible satisfiable belief configuration for all the goals. The preconditions of these new goal operators fixes the belief level for every $g \in G$ and the only effect is the addition of the ground atomic formula `goal` which satisfies the goal in the compiled problem. The strength of a plan induces a preference relation over the set of valid plans for the QU-STRIPS problem $\Pi$. Recall that a valid plan of higher strength is preferable to a valid plan of lower strength, as long as the trade-off in plan costs aren't too high. The exact value of this trade-off is problem-dependent. One approach to encoding this preference relation in the compiled problem is to define a non-negative cost $c_*$ of dropping a belief level. This assumes that the other operators have non-negative costs. Then, the cost of a goal operator will be $c_*$ multiplied by the number of belief levels dropped from $b_c$. We introduce an abuse of notation and define $b_c - b_x$ as the number of levels of belief $b_x$ below $b_c$ within the ordering $B_i$. Finding an optimal plan within the compilation then amounts to finding one of the most preferable valid plans.

More formally, let $B_i^+ = \{b_\sigma \mid \forall b_\sigma \in B_i \text{ if } b_\sigma > b_0\}$. The set $G^\dagger$ fixes a satisfiable belief level for each $g \in G$, i.e. $G^\dagger = \{\mathbf{B}(g, \sigma_g) \mid \forall g \in G\}$ for potentially distinct $\sigma_g \in B_i^+$. Then, $o' \in G'_G$ where

- $Pre'(o') = \{\mathbf{X_P}(g_\sigma) \mid g_\sigma \in G^\dagger\}$
- $Add'(o') = \{goal\}$
- $Del'(o') = \emptyset$
- $C'(o') = c_*(b_c - min(\{strength(g_\sigma) \mid g_\sigma \in G^\dagger\}))$

In order to compute $O'_A$, for every operator we will create a new set of operators which capture every possible positive belief level for each of the preconditions as well as every possible belief level for each of the antecedents and consequents within conditional effects. Consider an operator $o \in O$ and let $S_o$ be the set of ground atomic formulae that appear anywhere in the preconditions and conditional effects of $o$. We then construct $S_{o,\sigma}$ to fix arbitrary belief levels for each of these ground atomic formula.

$$S_{o,\sigma} = \{\mathbf{B}(p, \sigma_p) \mid \forall p \in S_o, \sigma_p \in B_i^+ \\ \text{if } p \in Pre(o) \text{ else } \sigma_p \in B_i\}$$

The compilation then takes the work which happens during the run-time of state progression under QU-STRIPS and performs the work ahead of time by constructing operators for every possible instantiation of $S_{o,\sigma}$. We define $Add'_P$ and $Add'_N$ to be subsets of $P^\dagger$ which are similar to $Add^\dagger_P$ and $Add^\dagger_N$ respectively except that the state $s$ is substituted with $S_{o,\sigma}$. Similarly, we define new predicates $withinBoth'(p)$, $strongest'_P(p, \sigma)$, and $strongest'_N(p, \sigma)$ for some $p \in P$ and $\sigma \in B_i$ to be similar to their non-primed versions except that it ranges over the newly defined $Add'_P$ and $Add'_N$ instead. Finally, the sets $Add'_1, Add'_2, Add'_3$, which are subsets of $P^\dagger$, are similar with respect to $Add^\dagger_1$, $Add^\dagger_2$, and $Add^\dagger_3$ except that it uses our newly created sets and predicates.
Then an operator $o' \in O'_A$ where:

- $Pre'(o') = \{\mathbf{X_P}(p_\sigma) \mid \forall p_\sigma \in S_{o,\sigma}\}$
- $Add'(o') = \{\mathbf{X_P}(p_\sigma) \mid \forall p_\sigma \in Add'_1 \cup Add'_2 \cup Add'_3\}$
- $Del'(o') = \{\mathbf{X_P}(\mathbf{B}(l, \sigma)) \mid \forall \sigma \in B_i, \forall p_\sigma \in Add'_1 \cup Add'_2 \cup Add'_3,$ where $l = ground(p_\sigma)$ and $\sigma \neq strength(p_\sigma)\} \cup \{goal\}$
- $C'(o') = C(o)$

Note from the compilation that an arbitrary operator $o' \in O'_A$ gets computed from an unique operator $o$ within the original problem $\Pi$. Let us define $\mathbf{X_O}^{-1}$ to be the function which takes a $o' \in O'_A$ and determines $o$.

For this naïve compilation, the total number of ground atomic formulae ($|P'|$) within the compiled problem $\mathbf{X_\Pi}(\Pi)$ is

$$|P'| = |B_i||P| \tag{1}$$

As for the operators, since the compiled problem considers different belief levels for every construction of $S_{o,\sigma}$ for a given operator $o$, the number of operators is exponential with respect to the size of $S_o$.[3] Let $CL(o)$ be the set of ground atomic formulae within the conditional effects of $o$ that are not within the preconditions. Then

$$|O'| = \Sigma_{o \in O}(|B_i^+|^{|Pre(o)|}|B_i|^{|CL(o)|}) + |B_i^+|^{|G|} \tag{2}$$

To help reduce the number of operators generated we can depend on the following lemma.

**Lemma 3** (Ground atomic formulae that do not appear in the consequent of any conditional effect only needs to get compiled to the level of belief specified in the initial state $I$)**.** *Let $\mathbf{B}(p, \sigma_p) \in I$ and let $p$ not appear in the consequent of any conditional effect in $O$. Then (1) $\mathbf{X_P}(\mathbf{B}(p, \sigma_{p2}))$ for any $\sigma_{p2} \neq \sigma_p$ may be safely discarded from $P'$; and (2) any operator $o'$ with $\mathbf{X_P}(\mathbf{B}(p, \sigma_{p2})) \in Pre'(o')$ may be safely discarded from $O'$.*

*Proof.* Given that $p$ does not appear in the consequent of any conditional effects, the belief level for $p$ will never change. Hence, $\mathbf{B}(p, \sigma_{p2})$ where $\sigma_{p2} \neq \sigma_p$ will never occur in any reachable state. This means that $\mathbf{X_P}(\mathbf{B}(p, \sigma_{p2}))$ will never be within the compiled version of any reachable state. Therefore, it's safe to remove $\mathbf{X_P}(\mathbf{B}(p, \sigma_{p2}))$ from $P'$ and any operator $o' \in O'$ whose precondition contains $\mathbf{X_P}(\mathbf{B}(p, \sigma_{p2}))$. ■

**Lemma 4** (Applicability in $\Pi'$ is sound with respect to $\Pi$)**.** *Let $o' \in O'_A$ be an operator within $\Pi'$ which is applicable at $s'$. Then, $o = \mathbf{X_O}^{-1}(o')$ is applicable at $s = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'\}$.*

*Proof.* By the definitions of applicability and the compilation, we know:

$$\begin{aligned} s' &\supseteq Pre'(o') \\ &\supseteq \{\mathbf{X_P}(p_\sigma) \mid \forall p_\sigma \in S_{o,\sigma}\} \\ &\supseteq \{\mathbf{X_P}(\mathbf{B}(p, \sigma_p)) \mid \sigma_p \in B_i^+, \forall p \in Pre(o)\} \end{aligned}$$

Then, $s \supseteq \{\mathbf{B}(p, \sigma_p) \mid \sigma_p \in B_i^+, \forall p \in Pre(o)\}$. Hence, $s$ is applicable at $o$. ■

**Lemma 5** (Operator application in $\Pi'$ is sound with respect to $\Pi$)**.** *Let $o' \in O'_A$ be an applicable operator for state $s'$ within $\Pi'$ and let $s'_n$ be the state resulting from applying $o'$ to $s'$. Additionally, let $s = \{\mathbf{X_P}^{-1}(p') \mid \forall p' \in s'\}$ discarding the ground atomic formula `goal` and $o = \mathbf{X_O}^{-1}(o')$. Assume $s$ is consistent. Then, if $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \in s'_n$ then $\mathbf{B}(l, \sigma_l) \in apply(o, s)$.*

*Proof.* Let $s_n = apply(o, s)$ and $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \in s'_n$. We know that $s'_n = (s' - Del') \cup Add'$. Let us show $\mathbf{B}(l, \sigma_l) \in s_n$ using proof by cases on membership of $\mathbf{X_P}(\mathbf{B}(l, \sigma_l))$.

[3] We are in general aware that some readers will presuppose that the exponential case is unreasonable. We have discussed this issue in connection with automated reasoning as the basis for automated planned elsewhere at some length (Rozek and Bringsjord 2024), and recapitulation here is beyond scope.

- Assume that $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \in Add'$. Then, $B(l, \sigma_l) \in \{p_\sigma \mid \forall p_\sigma \in Add'_1 \cup Add'_2 \cup Add'_3\}$. Since $o'$ is applicable at $s'$, $Pre' \subseteq s'$. Therefore, $S_{o,\sigma} \subseteq s$. This means that we can safely substitute $S_{o,\sigma}$ with $s$ and derive that $B(l, \sigma_l) \in Add_1^\dagger \cup Add_2^\dagger \cup Add_3^\dagger$. Which means that $B(l, \sigma_l) \in Add^\dagger$. Hence, $B(l, \sigma_l) \in s_n$.
- Suppose that $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \in s'$. Since $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \in s'_n$, we know that $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \notin Del'$. This leads to one of two cases, either $\mathbf{B}(l, \sigma) \in Add'_1 \cup Add'_2 \cup Add'_3$ or there does not exist a $\sigma_i \in B_i$ such that $B(l, \sigma_i) \in Add'_1 \cup Add'_2 \cup Add'_3$.
  - Assume that $\mathbf{B}(l, \sigma) \in Add'_1 \cup Add'_2 \cup Add'_3$. This subcase follows directly from the first case we've considered.
  - Suppose that there does not exist a $\sigma_i \in B_i$ such that $B(l, \sigma_i) \in Add'_1 \cup Add'_2 \cup Add'_3$. This means that there does not exist a $\sigma_i \in B_i$ such that $B(l, \sigma_i)$ is in either $Add'_P$ or $Add'_N$. Due to the fact that $S_\sigma \subseteq s$ and $s$ is consistent, we know that there does not exist a $\sigma_i \in B_i$ such that $B(l, \sigma_i)$ in either $Add_P^\dagger$ or $Add_N^\dagger$. Hence, $B(l, \sigma_l) \notin Del^\dagger$. From $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \in s'$, we know that $\mathbf{B}(l, \sigma_l) \in s$. Therefore, $\mathbf{B}(l, \sigma_l) \in s_n$.

■

**Lemma 6** (Operator application in $\Pi'$ is complete with respect to $\Pi$)**.** *Let $o$ be an applicable operator for some consistent state $s$. and $s' = \{\mathbf{X_P}(p_\sigma) \mid \forall p_\sigma \in s\}$. Additionally, let $o' \in \mathbf{X_O}(o)$ such that $Pre'(o') \subseteq s'$ and $s'_n$ be the result of applying $o'$ to $s'$. Then, if $\mathbf{B}(l, \sigma_l) \in apply(o, s)$ then $\mathbf{X_P}(\mathbf{B}(l, \sigma_l)) \in s'_n$.*

*Proof sketch.* For economy, we forego the details; however, the proof is similar to that of Lemma 5 and is shown by a proof by cases on membership of $\mathbf{B}(l, \sigma_l)$. An important point is that if $s$ satisfies $c$ for some $\langle c, l\rangle$ within $Add_P(o) \cup Add_N(o)$, then $S_{o,\sigma}$ satisfies $c$ since (1) $s$ is consistent, (2) $S_{o,\sigma} \subseteq s$, and (3) $c \subseteq S_o$. ■

**Lemma 7** (Operator application in $\Pi'$ is correct with respect to $\Pi$)**.** *Let $o' \in O'_A$ be an applicable operator for state $s'$ within $\Pi'$ and let $s'_n$ be the state resulting from applying $o'$ to $s'$. Additionally, let $s = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'\}$ after dropping the* `goal` *predicate and $o = \mathbf{X_O}^{-1}(o')$. Assume $s$ is consistent. Then, $\{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'_n\} = apply(o, s)$ after dropping the* `goal` *predicate within $s'_n$.*

*Proof.* This is evident from Lemma 5 and 6. ■

**Lemma 8** (The translation is sound where plans produced via $\Pi'$ (discarding operators in $O'_G$) are valid for $\Pi$)**.** *Let $\pi' = (o'_1, \ldots, o'_n)$ be a valid plan for $\Pi' = \langle P', O', I', G'\rangle$. Then, $\pi = (\mathbf{X_O}^{-1}(o'_1), \ldots, \mathbf{X_O}^{-1}(o'_{n-1}))$ where $o'_i \in \pi'$ and $o'_i \notin O'_G$ is a valid plan for the original problem $\Pi$.*

*Proof.* Note that by consequence of the compilation, $o'_n \in O'_G$ since the other operators delete the `goal` ground atomic formula. For $\pi$ to be valid for the problem $\Pi$, $G$ must be a subset of $s_n$ and each $o_i \in \pi$ must be applicable. At the state $s'_{n-1}$ after following each $o'_i$ from $\pi'$ starting from $I'$, we know from the preconditions of operators in $O'_G$ that

$$\{\mathbf{X_P}(\mathbf{B}(p, \sigma_p)) \mid \forall p \in G, \sigma_p \in B_i^+\} \subseteq s'_{n-1}$$

Hence for the first condition, it suffices to show that $s_{n-1} = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'_{n-1}\}$. To show both conditions, we'll perform a proof by induction on $s'_i$ which signifies the state at each step in $\pi'$.

- For the base case, consider the initial state $s'_1$. By definition of the translation $s_1 = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s_1\}$. By Lemma 4, $o_1$ is applicable at $s_1$.
- For the inductive case, assume that $s_i = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'_i\}$ and that $o_i$ is applicable at $s_i$. We want to show that $s_{i+1} = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'_{i+1}\}$ and $o_{i+1}$ is applicable at $s_{i+1}$. By Lemma 7, we know that $s_{i+1} = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'_{i+1}\}$. Then by Lemma 4, $o_{i+1}$ is applicable at $s_{i+1}$.

Hence, by induction each $o_i \in \pi$ is applicable and $s_{n-1} = \{\mathbf{X_P}^{-1}(p_\sigma) \mid \forall p_\sigma \in s'_{n-1}\}$. ■

| $\lvert P\rvert$ | $\lvert O\rvert$ | $\lvert P'\rvert$ | $\lvert O'\rvert$ | Compile | Search |
|---|---|---|---|---|---|
| 65 | 10 | 98 | 504 | 0.049 | 0.000165 |
| 2810 | 75 | 3023 | 3754 | 1.549 | 0.000276 |
| 16642 | 168 | 17159 | 8404 | 38.568 | 0.000361 |
| 26570 | 200 | 27223 | 10004 | 113.076 | 0.000453 |

Table 1: Results for Instances of the `Escape` Problem. Compile shows the number of sec. it takes to get from $\Pi = \langle B_i, P, P^\dagger, O, I, G\rangle$ to $\Pi' = \langle P', O', I', G'\rangle$. Search is the reported time in sec. to search for a solution within Fast Downward.

## Preliminary Evaluation

We implemented QU-STRIPS and its compilation to STRIPS.[4] For our initial evaluations, we look at instances of the `Escape` domain parameterized by $n$. The $n$ denotes the number of locations between the starting location and the goal location at all branches. Each branch in the problem will have a different configuration of levels of belief that a trap is present at each location within the branch. For our experiments, we used $B_5$ from Example 1 as our levels of belief. As such, for $n = 1$ and $n = 2$ there are $5^n$ branches. In order to keep the size of the experiments manageable, we pruned the number of branches for $n = 3$ and $n = 4$. In `Escape-3`, we pruned by a factor of 3 while for $n = 4$, we pruned by a factor of 16. Recall that if we do not make use of the optimization from Lemma 3, then we can compute the number of compiled ground atomic formulae and operators directly through Equations 1 and 2. Therefore, for our experiments we decided to use the optimization for direct comparison in the size of $|P'|$ and $|O'|$. Our results are

[4] The codebase along with the benchmark scripts can be found at https://github.com/Brandon-Rozek/Planning-Qu/tree/af91c4e209d8025b16d6a943841f3e02c5dcc6d2.

shown in Table 1. In our results, we can see that the size of $|O'|$ is far from the worst case. For example, without the optimization, the first row will have $|O'| = 25004$. To see how long it takes an automated planner to find a solution, we used Fast Downward (Helmert 2006). The table shows that search times grow more slowly than the time it takes to compile the problem to STRIPS. We hypothesize two reasons for this. The first and likely reason is that the search space for the Escape instances are more wide than it is deep. Plans for Escape instances of size $n$ are $n+2$ deep. The second reason is that for problems with consistent initial states, only one operator from $\mathbf{X_O}(o)$ will be applicable at a reachable compiled state.

## Discussion

We presented an initial framework for solving planning problems under qualitative uncertainty. However, this framework does not solve the problem in general. In the following paragraph, we present a few limitations present in our framework for the domain modeler to keep in mind, as well as to guide future work.

First, we do not reference any levels of belief in the operator descriptions. Additional ground atomic formulae can be added to the initial state to adjust the level of belief computed during state progression. However, in problems of ethical control, the domain modeler may want to require a certain level of belief to ethically permit an operator. For example, one may only want a prohibition operator to be applicable provided that evidence of the preconditions are at a level of *beyond reasonable doubt* or higher. We respectfully note that in prior work devoted to giving artificial agents themselves an "ethical compass" [e.g. (Govindarajulu and Bringsjord 2017a)], uncertainty is nowhere to be found. That work exploits computational logic to determine whether certain actions available to the AI are morally prohibited. But in "real life," an action is often prohibited not because some future consequence of that action is believed to be certain should that action be performed, but only that this consequence is (say) highly likely to ensue. Secondly, the QU-STRIPS framework is restricted to reasoning about qualitative uncertainty of a single agent. We do not capture epistemic reasoning over multiple agents with nested belief. This entails that our framework does not natively support certain tasks such as deception. Lastly, our framework does not handle probabilistic information. One can increase gradations of the levels of belief used in this framework; however, the belief updating we describe does not operate exactly like probabilistic updating.

Building atop the prior work of others, we have above generalized and extended the formal machinery of cognitive likelihoods for use in planning under uncertainty. While the increased flexibility and reach of this machinery in turn extends automated planning built upon it, a non-trivial question arises: How would an agent (perhaps a human AI engineer, or perhaps an AI agent given meta-control of subsidiary artificial agents under its oversight) go about deciding between the options in our generalization? For example, one option, presented by us above in encapsulated form, is that $\sigma$ ranges over 11 levels (five positive, one neutral/agnostic, and five negative). However, our formal generalization, and corresponding design and implementation, means that planning with qualitative uncertainty can work just as smoothly, at least formally and computationally speaking, with, say, a five-level ordering; i.e. (with negative belief levels omitted):

$$\begin{array}{l|l} \mathbf{B}^2(a,t,\phi) & \text{where } \sigma = 2 = \texttt{certain} \\ \mathbf{B}^1(a,t,\phi) & \text{where } \sigma = 1 = \texttt{likely} \\ \mathbf{B}^0(a,t,\phi) & \text{where } \sigma = 0 = \texttt{agnostic} \end{array}$$

## Conclusion

We have presented a new framework, namely QU-STRIPS, for solving planning problems under qualitative uncertainty. This framework, while in its initial stages, provide a substantial step toward deriving plans with epistemic strengths. It takes an initial state of qualitative beliefs, and from it defines whether operators are applicable and how beliefs change throughout state progression. We presented a provably sound compilation technique into classical STRIPS which allows for the use of state-of-the-art automated planners to find solutions to QU-STRIPS problems that maximize belief in achieving the goal while taking into account plan costs.

We see future work along both theoretical and applied trajectories. On the theory side, the formalism must be extended to consider conformant problems of width greater than one. This approach will benefit from ideas in the conformant-planning literature as referenced in the related work. On the applications front, the authors believe that reasoning under qualitative uncertainty has applications in goal recognition. Currently, in many algorithms (Ramırez and Geffner 2009; Sohrabi, Riabov, and Udrea 2016; Smith et al. 2015) plan costs are the driving force for deriving intent of the acting agent. However, the authors believe that not only plan cost should be considered, but additionally the level of belief the actor holds with respect to the achievability of their plan.

## Acknowledgments

The authors thank the anonymous reviewers for their insightful feedback. This paper was supported in part by a fellowship award to Rozek under contract FA9550-21-F-0003 through the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program, sponsored by the Air Force Research Laboratory (AFRL), the Office of Naval Research (ONR), and the Army Research Office (ARO).

## References

Andersen, M. B.; Bolander, T.; and Jensen, M. H. 2015. Don't plan for the unexpected: Planning based on plausibility models. *Logique et Analyse*, 145–176.

Bolander, T.; and Andersen, M. B. 2011. Epistemic planning for single-and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1): 9–34.

Bolander, T.; Charrier, T.; Pinchinat, S.; and Schwarzentruber, F. 2020. DEL-based epistemic planning: Decidability and complexity. *Artificial Intelligence*, 287: 103304.

Bonet, B.; and Geffner, H. 2000. Planning with incomplete information as heuristic search in belief space. In *Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems*, AIPS'00, 52–61. AAAI Press. ISBN 1577351118.

Bonet, B.; and Geffner, H. 2014. Flexible and scalable partially observable planning with linear translations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28.

Bryce, D.; Kambhampati, S.; and Smith, D. E. 2006. Sequential Monte Carlo in Probabilistic Planning Reachability Heuristics. In *ICAPS*, 233–242.

Davis-Mendelow, S.; Baier, J. A.; and McIlraith, S. A. 2013. Assumption-Based Planning: Generating Plans and Explanations under Incomplete Knowledge. In desJardins, M.; and Littman, M. L., eds., *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence, July 14-18, 2013, Bellevue, Washington, USA*, 209–216. AAAI Press.

Domshlak, C.; and Hoffmann, J. 2007. Probabilistic planning via heuristic forward search and weighted model counting. *Journal of Artificial Intelligence Research*, 30: 565–620.

E-Martín, Y.; R-Moreno, M. D.; and Smith, D. E. 2014. Progressive heuristic search for probabilistic planning based on interaction estimates. *Expert Systems*, 31: 421–436.

Fitting, M. 1991. Many-Valued Modal Logics I. *Fundamenta Informaticae*, 15: 235–254.

Friedman, N.; and Halpern, J. 1995. Plausibility Relations: A User's Guide. In *Proceedings of the Eleventh Conference on Uncertainty in AI*, 175–184.

Giancola, M. 2023. *Reasoning With Cognitive Likelihood for Artificially-Intelligent Agents: Formalization & Implementation*. Ph.D. thesis, Rensselaer Polytechnic Institute.

Giancola, M.; Bringsjord, S.; Govindarajulu, N. S.; and Varela, C. 2022. Making Maximally Ethical Decisions via Cognitive Likelihood and Formal Planning. In *Towards Trustworthy Artificial Intelligent Systems*, 127–142. Springer.

Govindarajulu, N.; and Bringsjord, S. 2017a. On Automating the Doctrine of Double Effect. In Sierra, C., ed., *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI-17)*, 4722–4730. International Joint Conferences on Artificial Intelligence. ISBN 978-0-9992411-0-3.

Govindarajulu, N. S.; and Bringsjord, S. 2017b. Strength Factors: An Uncertainty System for Quantified Modal Logic. In Belle, V.; Cussens, J.; Finger, M.; Godo, L.; Prade, H.; and Qi, G., eds., *Proceedings of the IJCAI Workshop on "Logical Foundations for Uncertainty and Machine Learning" (LFU-2017)*, 34–40. Melbourne, Australia.

Halpern, J. 2004. *Reasoning About Uncertainty*. Cambridge, MA: MIT Press.

Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research*, 26: 191–246.

Hu, G.; Miller, T.; and Lipovetzky, N. 2022. Planning with Perspectives–Decomposing Epistemic Planning using Functional STRIPS. *Journal of Artificial Intelligence Research*, 75: 489–539.

Klößner, T.; Steinmetz, M.; Torralba, A.; and Hoffmann, J. 2022. Pattern Selection Strategies for Pattern Databases in Probabilistic Planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, 184–192.

Kolobov, A.; Mausam, M.; Weld, D.; and Geffner, H. 2011. Heuristic search for generalized stochastic shortest path MDPs. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 21, 130–137.

Lifschitz, V. 1987. On the semantics of STRIPS. In *Reasoning about Actions and Plans: Proceedings of the 1986 Workshop*, 1–9.

Muise, C.; Belle, V.; Felli, P.; McIlraith, S.; Miller, T.; Pearce, A.; and Sonenberg, L. 2015. Planning over multi-agent epistemic states: A classical planning approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.

Muise, C.; Belle, V.; Felli, P.; McIlraith, S.; Miller, T.; Pearce, A. R.; and Sonenberg, L. 2022. Efficient multi-agent epistemic planning: Teaching planners about nested belief. *Artificial Intelligence*, 302: 103605.

Palacios, H.; and Geffner, H. 2009. Compiling uncertainty away in conformant planning problems with bounded width. *Journal of Artificial Intelligence Research*, 35: 623–675.

Ramırez, M.; and Geffner, H. 2009. Plan recognition as planning. In *Proceedings of the 21st international joint conference on Artifical intelligence. Morgan Kaufmann Publishers Inc*, 1778–1783. Citeseer.

Rozek, B.; and Bringsjord, S. 2024. Spectra: An Expressive STRIPS-Inspired AI Planner Based on Automated Reasoning. *Künstliche Intelligenz*. This URL goes to a preprint only.

Santos, Y. D. 2020. A Four-Valued Dynamic Epistemic Logic. *Logic, Language and Information*, 29: 451–489.

Smith, D. E.; and Weld, D. S. 1998. Conformant graphplan. In *AAAI/IAAI*, 889–896.

Smith, D. E.; et al. 2015. A fast goal recognition technique based on interaction estimates. In *International Joint Conferences on Artificial Intelligence (IJCAI-15)*, ARC-E-DAA-TN24965.

Sohrabi, S.; Riabov, A. V.; and Udrea, O. 2016. Plan Recognition as Planning Revisited. In *IJCAI*, 3258–3264. New York, NY.

Trevizan, F.; Thiébaux, S.; and Haslum, P. 2017. Occupation measure heuristics for probabilistic planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 27, 306–315.