Markovian Models Anxious to Stay on the Beaten Path. A Psychology-Grounded Approach to Minimising Exposure to Path Uncertainty

Loïs Vanhée[1,2], Anaïs de Graaf[1]

[1]Department of Computing Science, Umeå University, Sweden
[2]The Norwegian College of Fishery Science, UiT – The Arctic University of Norway, Norway

Abstract

Mainstream formalizations of uncertainty in Automated Planning (AP) research tend to revolve on rational probability management and reward maximization, often overlooking psychological factors involved in how humans handle uncertainty, where reward maximization is often discounted in favor of simpler plans yielding to predictable trajectories for the sake of reducing exposure to anxiety. This paper pushes the research frontier by introducing and formalizing the concept of Path Anxiety, which quantifies how much various policies may expose the AP to uncertainty in regards to future trajectories, which we call Path anxiety-sensitive Markov Decision Processes (PAS-MDP). Then, we propose practical PAS-MDP algorithms for generating optimized policies that balance utilitarian reward maximization with path uncertainty minimization. The practical relevance of PAS-MDPs is empirically validated by showing that PAS-MDP policies optimized by our algorithms can dramatically reduce path uncertainty with minimal reduction in utility as well as reducing the number of possible paths the system may reach by several orders of magnitude. The psychological accuracy of PAS-MDPs is validated by showing that the optimized policies replicate human-like anxiety-sensitive behaviors identified by psychology research. Altogether, this research highlights a promising approach towards developing human-friendly AP systems that account for human anxiety in their decision-making processes, hence addressing one of the most central factor of mental health and wellbeing.

Introduction

Managing uncertainty is a central focus of Automated Planning (AP) research (Pratt, Raiffa, and Schlaifer 1964; Saffiotti 1987; Russell and Norvig 2021), with increasing attention being given to approaches that better account for human and psychological factors (Stanojevic 2024) and that emphasize the need for interdisciplinary research (Kusters et al. 2020). The study of uncertainty in AP encompasses several key areas: 1) the development of methods for handling uncertainty, including Bayesian models (Cai and Scarlett 2024; Kukacka 2010; Toroghi and Sanner 2024; Wang and Li 2024; Wu et al. 2024a), Markovian models (Bai, Mondal, and Aggarwal 2024; Liu et al. 2024), automated planning (Elahi and Rintanen 2024; Engesser, Herzig, and Perrotin 2024; Lei, Lipovetzky, and Ehinger 2024), and reinforcement learning (Amitai, Septon, and Amir 2024; Omura et al. 2024; Zhou et al. 2024; Zhu et al. 2024); 2) the enhancement of AP systems acceptability by reducing the uncertainty associated with these systems, which includes initiatives in trustworthy AP (Chu et al. 2024; Vashistha and Farahi 2024; Zhang et al. 2024a), explainable AP (Bie, Luo, and Chen 2024; Lai et al. 2024; Yang et al. 2024), uncertainty quantification (Chen et al. 2024; Horii and Chikahara 2024; Wu et al. 2024b), human-AP collaboration (Ma et al. 2024; Zou et al. 2024), and the broader domain of human-centered AP (Capel and Brereton 2023; Zhang et al. 2024b); and 3) the design of systems that consider the impact of their decisions on human exposure to uncertainty, as seen in applications such as autonomous driving (Qian et al. 2024), personalized education (Shen et al. 2024), healthcare (Dai et al. 2024), automated recommendation (Agrawal et al. 2024), fraud detection (Yu, Liu, and Luo 2024), and virtual assistants (Tulshan and Dhage 2019). Dealing with uncertainty addresses reduction of anxiety, the prime mental state raised by uncertainty and one of the root causes of mental health and well-being issues, which leads to significant individual suffering (e.g., addiction, depression, self-harm) and social costs amounting to trillions of USD annually (Andlin-Sobocki and Wittchen 2005; Bandelow and Michaelis 2015; Health 2020; Massazza et al. 2023). This is particularly relevant in a context where AP solutions often exacerbate human exposure to uncertainty (Cornelissen and Cholakova 2021; Huang 2023; Pires and Pinto 2020; Lindgren 2023) and in a context of a high demand for mental wellbeing to be better factored in AP decisions (Heaukulani et al. 2024).

Despite this extensive interest, most current planning approaches are constrained by a narrow framing of uncertainty. They tend to focus on utility maximization, avoiding worst-case scenarios through the lenses of risk (Peng, Yang, and Yao 2023) and regret (Bai, Mondal, and Aggarwal 2024), and acquiring knowledge to reduce epistemic uncertainty (Yotam and Indelman 2024). These approaches predominantly scope uncertainty as a mathematical construct, which may significantly differ from how humans experience and respond to uncertainty (Spiegelhalter 2011). For example, people tend to be driven to reduce uncertainty through simple, controllable plans that address the uncertainty early and keep it low (e.g., stay on beaten paths), whereas the mathematical framing of uncertainty is prone to elaborate very complex branching strategies that allow an eventual resolution of uncertainty (e.g., engage in uncertain branching routes open for the possibility of fortunate shortcuts).

In response, this paper addresses the broader question of *how to develop psychology-grounded, utility-driven planning solutions that also aim to reduce exposure to uncertainty*. Specifically, it focuses on *path uncertainty*, defined as the degree of uncertainty associated with the possible paths a given policy can yield to, a concept rooted in human psychology as a form of rule uncertainty (Bach and Dolan 2012). As a method for enabling a psychological accountability of the experience of uncertainty, this paper builds over the concept of *anxiety sensitivity* introduced by (Vanhée, Jeanpierre, and Mouaddib 2022; Horned and Vanhée 2023b,a) and grounded in the psychology research of (Miceli and Castelfranchi 2005, 2014), which frames anxiety as the minimization of the cumulative exposure to uncertainty over time. This approach yields to human-like strategies that minimize uncertainty early and consistently and subsequently reduces the psychological impacts of exposure to uncertainty.

The contributions of this paper are as follows: 1) introducing the key foundations for a psychological framing uncertainty, under the lenses of experienced anxiety and subsequent behaviorial uncertainty-reducing responses; 2) formalizing the concept of Path Uncertainty; 3) integrating this concept within Markov Decision Processes (MDPs) and formalizing Path anxiety-sensitive MDPs (PAS-MDP); 4) presenting an exact algorithm and 5) an approximate polynomial algorithm for solving PAS-MDPs; 6) empirically validating the robustness of PAS-MDPs by demonstrating their alignment with core uncertainty-reducing strategies observed in human behavior; and 7) showing that PAS-MDPs provide alternative policies for a drastic reduction of path uncertainty while maintaining high operational performance, a feature not directly addressed in other models.

Background

For space consideration, the background related to Shannon Entropy is assumed known. The extended version of this paper and the associated arxiv provides further details[1].

Psychological Grounding for Anxiety. This section introduces how anxiety is framed in psychology towards enabling to validate whether our PAS-MDPs accurately reproduce anxiety-sensitive behaviors. Anxiety is the primary mental state triggered by uncertainty and activating deliberative mechanisms and behaviors directed to address uncertainty (e.g., circumvent, anticipate, prepare, address, avoid uncertainty) (Miceli and Castelfranchi 2005, 2014). While being an overall functional mental state that drives to dealing with threats and reducing uncertainty-induced cognitive workload, too regular or too intense exposure to anxiety ties to an array of harmful psychological implications (Brosschot, Verkuil, and Thayer 2016; Grupe and Nitschke 2013)

[1]*The extended version of this paper can be found at*: https://drive.google.com/open?id=128pZ3bwpBugiBCMpXO6ZcZEbDps1GRqm&usp=drive_fs

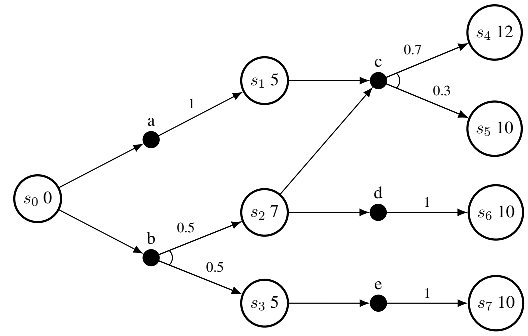

Figure 1: MDP for the Running example. White nodes represent states, numbers represent the reward for reaching the states. Black nodes represent actions. The transition function is represented by the probabilities next to arrows. An example of the calculations for the various modelling elements listed in this paper on this running example can be found in annexes[1].

Psychological research has identified a spectrum of responses to anxiety that take into account both the individual's sensitivity to anxiety and the level of uncertainty they experience (Hartley and Phelps 2012; Luan, Schooler, and Gigerenzer 2014; Miceli and Castelfranchi 2005, 2014). At the lowest levels of sensitivity, anxiety leads to behaviors that are nearly indistinguishable from purely task-driven actions or are associated with cost-free uncertainty aversion (Tversky and Fox 1995), i.e., a tendency to avoid uncertainty when it does not impact performance.

As sensitivity to anxiety increases, the intolerance to uncertainty increases and thus performance is increasingly discounted for reducing uncertainty. This sensitivity can be manifested in various strategies such as satisficing (i.e., ensuring some form of achievement of the goals, albeit with increasing ineffectiveness for the sake of anxiety reduction), status quo bias (e.g., the tendency to maintain the current situation as close as possible to a known baseline) (Samuelson and Zeckhauser 1988), choosing the default option (e.g., the tendency to make decisions similar to those made previously) (Johnson and Goldstein 2003), or minimalist decision-making (e.g., the tendency to simplify decision frames and favor courses of action that are easier to follow) (Gigerenzer and Goldstein 1996) and, in higher levels of anxiety, a tendency towards perfectionism (i.e., the tendency to select actions that guarantee the attainment of goals with minimal room for uncertainty, regardless how inefficient is the approach in the long run) (Frost and DiBartolo 2002). As anxiety sensitivity further rises, behaviors start losing their goal-directed focus altogether, yielding to avoidance behaviors (e.g., giving up on goals) (Dymond and Roche 2009; Schoenfeld 1950), self-sabotage (e.g., intentionally damaging the conditions required for success) (Komporozos-Athanasiou and Haiven 2023), and self-harm (e.g., engaging in self-injury or pain to cope with psychological distress) (O'Connor, Rasmussen, and Hawton 2010).

Markov Decision Process Markov Decision Processes (MDPs) represent stochastic processes under the (partial) control of an agent, represented by a tuple $\langle S, A, T, R\rangle$, where S is a set of states, A is a set of actions, T is a transition function where $T(s, a, s')$ is a probability of reaching $s' \in S$ from $s \in S$ after performing action $a \in A$, R is a reward function where $R(s)$ is the expected reward for landing in state $s \in S$. Figure 1 illustrates an MDP that is used as a *running example* throughout this paper. Optimizing MDPs within a horizon $h \in \mathbb{N}$ consists in finding a policy $\pi : S \rightarrow A$, which assigns an action $a \in A$ to be played in every state $s \in S$, such that π maximizes the expected value $V : S \rightarrow \mathbb{R}$ defined by the Bellman equation (Bellman 1956) as $V^{\pi,h}(s) = R(s) + \sum_{s' \in S} T(s, \pi(s), s') \times V^{\pi,h-1}(s')$.

Five-Component Markovian Models of Anxiety Foundational research on Markovian models of anxiety introduce a structured formal framework for modelling anxiety-sensitive planning (Vanhée, Jeanpierre, and Mouaddib 2022), which identifies five primary components modelling anxiety. The first component is the *object of anxiety*, typically formalized using a probability distribution Δ over the potential realizations of this object (e.g., the probability of various outcomes given a policy). The second component is *instant anxiety* (Spielberger et al. 1999), representing the anxiety experienced at a specific moment due to exposure to this uncertain object, formalized as $IAnx$, a function that increases with the level of uncertainty associated with Δ (e.g., the number and disparity of potential outcomes). The third component is *cumulated anxiety*, denoted as $CAnx$, a function that represents the accumulation of instant anxiety over time and its consequent effects on an individual's wellbeing and health (Dugas et al. 1998). The fourth component is the *anxiety sensitivity* factor, usually represented as $W \in [0, 1]$, which models the tradeoff between reward optimization and much anxiety avoidance, where $W = 0$ represents pure reward optimization and $W = 1$ represents pure anxiety avoidance. The last component is the *anxiety-sensitive value function*, represented by ASV, which extends the classic value function by integrating anxiety avoidance.

Related Work

This section explores various Markovian models that integrate multiple facets of uncertainty that best relate to our aims. Classic MDPs have been significantly expanded to emphasize risk avoidance, generally aiming to maximize expected value while minimizing the likelihood of the worst outcomes or staying within a constrained cost budget (Hou, Yeoh, and Varakantham 2014; Yu, Lin, and Yan 1998). Another variant, the Regret MDP (Bai, Mondal, and Aggarwal 2024; Regan and Boutilier 2012), focuses on maximizing value under some uncertainty about the reward function. Partially Observable MDPs (POMDPs), where the system's state is uncertain and inferred through indirect observations (Kaelbling, Littman, and Cassandra 1998), have been extended into ρ-POMDPs (Araya et al. 2010). These extensions prioritize gaining epistemic certainty, enhancing knowledge about the system's state. In contrast, our model is grounded in a psychological paradigm rather than purely mathematical representations of uncertainty. Moreover, PAS-MDPs seek to minimize cognitive costs and the psychological impacts of uncertainty, by minimizing uncertainty early and consistently. Additionally, our approach specifically focuses on the uncertainty tied to the distribution of possible trajectories the system may encounter –a feature not directly addressed in other models.

Certain specialized models, such as the AA-MDP (Vanhée, Jeanpierre, and Mouaddib 2022) and $R\rho$-POMDP (Gutsche and Vanhée 2023), also incorporate a psychological basis for managing uncertainty through mechanisms that address anxiety, similar to our approach. Nonetheless, these models primarily deal with uncertainty about expected rewards and uncertainty about the current (hidden) state, without specifically targeting the reduction of path uncertainty, which is the central focus in our present work.

Formalizing Path Anxiety

Our Path anxiety-sensitive MDP (PAS-MDP) formalizes the psychological cost associated with the uncertainty of future courses of action the system may take, building on the integrative four-component anxiety framework from (Vanhée, Jeanpierre, and Mouaddib 2022). As an approach, we focus here on a specific implementation of this model for which practical algorithms can easily be designed. This model can be generalized for capturing a broader range of psychological dynamics (e.g., non-additive preferences from decision makers). In this framework, the object of anxiety is modeled as the probability distributions over paths resulting from a given policy π, formalized as follows:

Definition 1. *A **path distribution** $\Delta^h : \mathcal{T}^h \rightarrow [0, 1]$ is a probability distribution over all paths of size $h \in \mathbb{N}$, where $\mathcal{T}^h$ is the set of all possible paths of size h. $\Delta^h(\tau)$ is the probability for the path $\tau \in \mathcal{T}^h$ to occur. $\sum\limits_{\tau \in \mathcal{T}^h} \Delta^h(\tau) = 1$.*

A path distribution is a general representation of a probability distribution over paths within a given horizon. On the example of Figure 1, a distribution could be $((s_0, s_6, s_7), 0.7), ((s_0, s_s, s_3), 0.3)$.

Definition 2. *A **policy-induced path distribution** $\Delta^{\pi,h} : \mathcal{T}^h \rightarrow [0, 1]$ is the path distribution induced by following policy π for h steps. $\Delta^{\pi,h}(s_0, s_1, \ldots, s_h)$ is formalized as:*

$$\begin{cases} 1 & h = 0 \\ T(s_0, \pi(s), s_1) \times \Delta^{\pi,h-1}(s_1, \ldots, s_h) & h > 0. \end{cases} \tag{1}$$

A policy-induced path distribution is a specific path distribution issued from following a given policy given a starting state. This distribution should comply with constraints relative to the dynamics of the MDP. A policy-induced path distributions for a set of policies on the example from Figure 1 can be found in the extended version of the paper.

Definition 3. *A **policy-induced path distribution from** s_0, defined as $\Delta^{s_0,\pi,h} : \mathcal{T}^h \rightarrow [0, 1]$ is the distribution over all paths induced by following policy π, for h steps that start by $s_0 \in S$. Formally, $\Delta^{s_0,\pi,h}(\tau) = \Delta^{\pi,h}(\tau)$ for τ starts with s_0 and 0 otherwise. $\sum\limits_{\tau \in \mathcal{T}^h} \Delta^{s_0,\pi,h}(\tau) = 1$.*

An illustration policy-induced path distributions in the running example can be found in the extended version of the paper. For example, the path distribution from s_0

induced by $\pi^{R,*}$ (i.e., the policy optimizing for rewards only) is: $\Delta^{s_0,\pi^{R,*},3} = \{(s_0, s_2, s_4) : 0.35; (s_0, s_2, s_5) : 0.15; (s_0, s_3, s_7) : 0.5\}$.

Definition 4. *The **instant path anxiety** $IAnx^{\pi,h}(s)$ quantifies the degree of uncertainty tied to possible trajectories resulting from applying π for h steps from s. This uncertainty is quantified by the Shannon entropy (see annexes for a justification and extended model) as:*

$$IAnx^{\pi,h}(s) = H(\Delta^{s,\pi,h}). \tag{2}$$

As an illustration, the instant anxiety of the various policies in the running example can be found in the extended version of the paper. For example, the instant anxiety for the policy $\pi^{R,*}$ (i.e., the policy maximizing rewards) $IAnx^{\pi^{R,*},2}(s_0) = H(\Delta^{s_0,\pi^{R,*},h}) = 1.44$. A comparison of the instant anxieties of the various policies of the running example allows demonstrating the soundness of the approach. $IAnx^{\pi^{\uparrow},2}(s_0)$ achieves the lowest uncertainty as they can only encounter two paths with (s_0, s_1, s_4) being relatively probable, followed by $IAnx^{\pi^{\downarrow},2}(s_0)$, which can reach two paths as well, but each with a probability of 50%, hence creating more uncertainty than $\pi^{\uparrow}$. Last, $IAnx^{\pi^{R,*},2}(s_0)$ is the highest, which is sound as well, as it features more paths than $\pi^{\downarrow}$, each with a lower or equal certainty than the paths of $\pi^{\downarrow}$.

Definition 5. *The **cumulated path anxiety** $CAnx^{\pi,h}(s)$ represents the expected exposure over time to instant path anxiety when following π for h steps. $CAnx^{\pi,h}(s)$ is represented as:*

$$\begin{cases} IAnx^{\pi,h}(s) & h = 0 \\ IAnx^{\pi,h}(s) + \\ \sum_{s' \in S} T(s, \pi(s), s') \times CAnx^{\pi,h-1}(s') & h > 0. \end{cases} \tag{3}$$

An illustration of the importance of the cumulative effect can be found in the extended version of the paper. When comparing the instant and cumulated anxiety scores for $\pi^{R,*}$ (i.e., the reward-maximizing policy) and $\pi^{\downarrow}$ (the policy of selecting the action that is visually the lowest in the display in Figure 1) from the running example, it can be observed that $IAnx^{\pi^{R,*},2}(s_0) < IAnx^{\pi^{\downarrow},2}(s_0)$, but $CAnx^{\pi^{\downarrow},2}(s_0) < CAnx^{\pi^{R,*},2}(s_0)$. This outcome is desirable and can be explained by examining Figure 1: in s_0, $\pi^{\downarrow}$ selects an action that introduces high uncertainty in the trajectory, but once this action is resolved, the system only enters states where the future paths considered by $\pi^{\downarrow}$ are certain. Conversely, $\pi^{R,*}$ delays the action that would eliminate uncertainty about future paths, leading to greater uncertainty over time. This observation aligns with human behavior, where anxiety increases the tendency for early resolution of uncertainty (often referred to as impatience when this resolution comes at the expense of future rewards (Whitmont 1952)).

Definition 6. *A **Path Anxiety-Sensitive MDP** (PAS-MDP) is defined as a tuple $\langle S, A, T, R, h, s_0, W\rangle$, where S, A, T, R and h correspond to the classic representations of an MDP, s_0 corresponds to a starting state and $W \in [0, 1]$ represents the **anxiety-sensitivity factor**, i.e. the relative importance for avoiding anxiety over collecting rewards as W tends to 1.*

This representation simply expands the classic representation for a MDP with W, the anxiety-sensitivity factor. $W = 0$ means that only rewards are considered and $W = 1$ means that only anxiety is considered.

Definition 7. *A **path anxiety-sensitive Value Function** $ASV^{\pi,h} : S \to \mathbb{R}$ represents the scaled compromise between the reward acquisition and anxiety avoidance for a horizon h, defined as $ASV^{\pi,h}(s) = (1 - W) \times V^{\pi,h}(s) - W \times CAnx^{\pi,h}(s) \times VA\_norm + \eta^-_V$, where VA_norm, is a rescaling factor between value and anxiety, formalized as $VA\_norm = (\eta^+_V - \eta^-_V)/(\eta^+_A - \eta^-_A)$ and $\eta^-_V = argmin_\pi V^{\pi,h}(s_0)$, $\eta^+_V = argmax_\pi V^{\pi,h}(s_0)$, $\eta^-_A = argmin_\pi ASV^{\pi,h}(s_0)$, $\eta^+_A = argmax_\pi ASV^{\pi,h}(s_0)$.*

This definition expands the classic definition of the value function (i.e., cumulated rewards) for also including cumulated path anxiety. Because the reward function R possibly relies on an arbitrary scale (e.g., value in dollars, cost in time), a scaling operation between rewards and anxiety allows ensuring a comparability across these values. Here, we propose a linear mapping between the minimum and maximum of anxiety and the minimum and maximum of rewards. In this definition, policies that maximize the path anxiety-sensitive value function are Markovian, as their value only depends on the current state without requiring memory from prior events. In the running example, $\pi^{R,*}$ is the optimal policy for $W = 0$ (i.e. pure reward maximisation) and $\pi^{\downarrow}$ is the optimal policy for $W = 1$ (pure anxiety avoidance) and $\pi^{\uparrow}$ balances reward acquisition and anxiety avoidance.

Fast Algorithms for Solving PAS-MDPs

Optimal PAS-MDP policies can be computed by straightforwardly extending the backpropagation function of the classic value iteration algorithm (Russell and Norvig 2021; Sigaud and Buffet 2013) such as, for iteration i and for every state s, are propagated: 1) the expected value V of s given a horizon i, as per usual, 2) the path distribution Δ of length i from s, 3) the expected cumulated path anxiety $CAnx$ from s within a horizon i, which can be computed by integrating instant path anxiety derived from Δ. This algorithm, which we call the VI-PAS-MDP algorithm is used as well for computing VA_norm, through four executions of the algorithm, in which the $argmax$ of line 4 is replaced, respectively, by $R(s)$, $-R(s)$, $IAnx^{\pi,h}(s)$, $-IAnx^{\pi,h}(s)$ for s_0. After these four executions, a final execution is then run, in which the reward for state s is computed as $R(s) + VA\_norm$. The VI-PAS-MDP algorithm is detailed in Algorithm 1. A proof of the VI-PAS-MDP algorithm is available in annexes[1].

Due to the exhaustive enumeration of all possible paths of length h given any state, Algorithm 1 a complexity exponential on h. The (Value-Iteration) VI-PAS-MDP algorithm[2] makes this approach practical through two key simplifications. First, the VI-PAS-MDP algorithm represents the distribution Δ as a mapping that pairs each probability value with the count of paths having that probability. For

[2]https://github.com/lvanhee/path-anxiety-mdp

Algorithm 1: **VI-PAS-MDP**: Compute π_h^* and ASV for a PAS-MDP. VA_norm can be computed by altering Line 4 and Line 7 with $W = 0$ or $W = 1$ and inverting R

1: $\pi^* \leftarrow$ any arbitrary policy. $\forall s \in S$: $V^{\pi^*,0}(s) \leftarrow R(s)$, $CAnx^{\pi^*,0}(s) \leftarrow 0$, $\Delta^{s,\pi^*,0} \leftarrow \{() : 1\}$
2: **for** i ∈ [1, h] **do**
3:   **for** s ∈ S **do**
4:
$$a^* \leftarrow argmax_a \sum_{s' \in S} T(s,a,s')(1 - W\times) V^{\pi^*,i-1}(s') + W \times VA\_norm \times CAnx^{\pi^*,i-1}(s')$$
5: $\pi^*(s) \leftarrow a^*$
6:
$$\Delta^{s',\pi^*,i} \leftarrow \{(s.s'.\tau) : T(s,a^*,s') \times \Delta^{s',\pi^*,i-1}(\tau) \mid s' \in S,\ \tau \in \mathcal{T}^{h-2}\}.$$
7: $V^{*,s}(s) \leftarrow R(s) + \sum_{s' \in S} T(s,a^*,s') \times V^{\pi^*,i-1}(s')$
8:
$$CAnx^{\pi^*,i}(s) \leftarrow IAnx^{\pi^*,i}(s) + \sum_{s' \in S} T(s,a^*,s') \times CAnx^{\pi^*,i-1}(s')$$
9:   **end for**
10: **end for**
11:
12: **return** π^*

example $\{\tau_1 : 0.25; \tau_2 : 0.5; \tau_3 : 0.25\}$ is represented as $\{2 : 0.25; 1 : 0.5\}$, meaning "two paths have a probability of 0.25 and one path has a probability of 0.5". This simplification preserves the necessary information about the number of paths with specific probabilities in Δ as two distinct paths cannot become equal when the horizon is extended, which is the only information required for computing the Shannon entropy.

Second, as the set of possible paths can grow arbitrarily, the mapping can become arbitrarily large. To address this, the mapping is replaced by a k-bounded bag of particles, using a method similar to Monte Carlo planning in POMDPs representations (Silver and Veness 2010) and representations used in similar contexts (Vanhée, Jeanpierre, and Mouaddib 2022). When the number of particles exceeds k, the two closest particles are fused into one by adding their counts and averaging their probabilities weighted by their count, thereby preserving both their total density and the total count of paths. This representation ensures a hard boundary on memory usage and update operations. In terms of complexity, all operations within the core of the for loop have costs comparable to the standard value iteration algorithm, except for Line 6 of Algorithm 1, which involves creating a bag of particles with a size upper-bounded by $|S| \times k$. Since fusing k' elements into a k-sized bag can be performed in $O(k' \times \log(k))$ assuming sorted bags, the overall complexity of the algorithm is $\mathcal{O}(h \times |A| \times |S|^3 \times \log(k))$.

Method

This section puts in place the method for validating the relevance of PAS-MDPs for both 1) *operational assessment*,

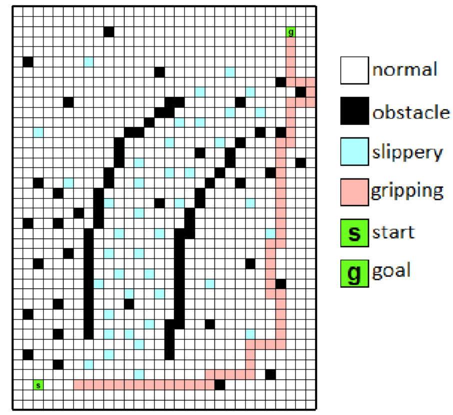

Figure 2: The simulation environment represented in a grid.

by showing that the VI-PAS-MDPs produce policies pertaining to lower cumulated anxiety when anxiety sensitivity increases and that effective compromises between cumulated anxiety minimization and reward maximization can be achieved, and 2) managing to replicate *human-like behaviors*, by comparing the behaviours resulting from increasing anxiety sensitivity in our policies with human anxiety-sensitive behavior.

Experimental Setup

The setup, illustrated in Figure 2, consists of a robot navigation problem on a 30×40 2D grid where the robot must reach a goal while navigating obstacles and various floor textures that affect its likelihood of slipping with each action, adapting the experimental setup from (Chow et al. 2015; Vanhée, Jeanpierre, and Mouaddib 2022). The set of states corresponds to the grid tiles plus a pit state. Four actions are available, each corresponding to a direction the robot intends to reach (North, South, East, West). Each tile is associated a slipperiness level, which influences the transition function. When the robot is located on a *normal* (i.e. white) tile, it has a probability 0.9 of landing into the intended tile and a probability of 0.05 of landing into any of the two tiles adjacent to the intended tile in the intended direction. For example, if $(5, 7)$ is normal then $T((5, 7), N, (5, 8)) = 0.9$ and $T((5, 7), N, (4, 8)) = T((5, 7), N, (6, 8)) = 0.05$. When the robot is located on *slippery* (i.e. blue) tile, it has a probability of 0.7 of landing into the intended tile and has a chance of 0.1 of landing in any of the three tiles one step further the intended tile. For example, if $(5, 7)$ is slippery, $T((5, 7), N, (5, 8)) = 0.7$ and $T((5, 7), N, (4, 9)) = T((5, 7), N, (5, 9)) = T((5, 7), N, (6, 9)) = 0.1$. When the robot is located on a *gripping* (i.e. red) tile, the intended tile is reached with a probability of 1. When the robot is on the goal, an obstacle, or the pit state, the system transitions to the pit state with a probability 1. The system receives a reward of: 0 when in the pit state, -100 when in an obstacle state, 1 when in the goal state, and -1 otherwise (representing battery consumption). The parameters are set as $k = 100$ and $h = 140$. A visual representation of the transition function and the transition rules of this example are provided

	$V^{\pi^*,h}(s_0)$	$CAnx^{\pi^*,h}(s_0)$
$W = 0$	-55.65	432.33
$W = 0.1$	-55.70	379.65
$W = 0.3$	-62.33	81.05
$W = 0.5$	-63.13	32.68
$W = 0.8$	-65.38	2.84
$W = 1$	-104.01	2.41

Table 1: Exact value and cumulated anxiety for various levels anxiety sensitivity (W) in the validation scenario

in annexes[1]. W is an independent variable altered by the various experiments. A first analysis was performed for all values W in $[0, 1]$ in increments of 0.1. Six values of W were selected for final analysis as they led to the most significant shifts in the systems' strategies: 0, 0.1, 0.3, 0.5, 0.8, 1. anxiety-sensitive responses are enforced by the map disposition that include a straightforward central path to the goal that features more slippery tiles (which are more uncertain but, when used adequately, have a probability of saving one to two turns), a side path that features more gripping tiles, but requires to perform detours should the system seek to minimize uncertainty. All experiments were run on a classic desktop computer (i7 CPU, 16 GB of RAM, Windows 11, Python 3.6.15).

Measurements

We captured and analyzed an array of performance indicators (value, cumulated anxiety, number of possible paths) together with behaviorial indicators (number of various types of visited tiles, overview of the intended paths) depending on the anxiety sensitivity W. The number of paths is computed by expanding the backpropagation algorithm by adding a variable $\#path_k^s$ defining the number of paths for s with a horizon k, such that:

$$\#path_k^s = \sum_{s' \in \{s' | T(s, \pi_k^*(s), s') > 0\}} \#path_{k-1}^{s'}.$$

This method allows counting the number of possible paths without requiring their exhaustive enumeration.

Results and Analysis

Quantitative Performance Indicators

Table 1 introduces *the impact of W on both value and cumulated anxiety* and Figure 3 relates these two variables together. These table and figure show the emergence of a clear Pareto front, where increasing W yields to a succession of Pareto tradeoffs between cumulated anxiety minimization and and expected value maximization. The relation between W, value, and cumulated anxiety appears to follow non-linear thresholds. The transition from $W = 0$ (i.e., the classic MDP optimization) to $W = 0.1$ yields to a reduction of 12.2% of cumulated anxiety for $< 1\%$ of reduced value. When comparing $W = 0$ with $W = 0.3$, 81.3% of cumulated anxiety is removed for 12.0% of reduced value.

These results demonstrate that the VI-PAS-MDP indeed optimize value versus anxiety exposure and that its output

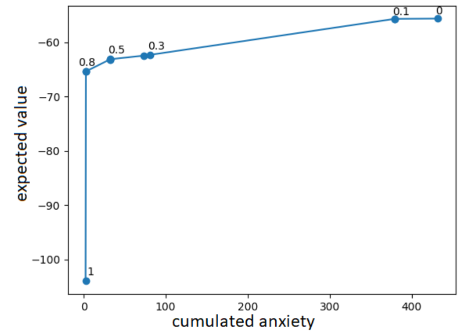

Figure 3: Expected value and cumulated anxiety for various levels of anxiety sensitivity W. Exact numbers are listed in annexes[1]. The top-right value ($W = 0$) corresponds to the output of classic optimized Markovian policies.

W	0	0.2	0.4	0.6	0.8	1
# paths	10^{49}	10^{48}	10^{47}	10^{44}	10^{9}	10^{5}

Table 2: Order of magnitude of the number of paths depending on W

corresponds to psychological findings, i.e. increasing sensitivity to anxiety yields to behaviors that lower the total exposure to anxiety at the expense of value. Additionally, these results show that significant exposure to anxiety may be achieved at minimal cost in terms of value, which shows the practical relevance of anxiety-sensitive planning.

Number of Possible Paths

The number of possible paths that policies generated by the system may reach provides a relevant indicator of how much path anxiety the system is exposed to. Table 2 depicts the relationship between W and $\#path_k^s$ from the initial state and horizon. The table shows an exponential decrease in the number of reachable paths when the sensitivity to uncertainty W increases, from an order of magnitude of 10^{49} unique paths for $W = 0$ (i.e. for the optimal policy π^* computed by the classic Value Iteration algorithm) to 10^5 paths for $W = 1$ (i.e. the pure anxiety-avoiding policy). A dramatic decrease can be observed from $W = 0.6$ to $W = 0.8$, from 10^{44} to 10^9 possible paths.

These results are congruent with expectations on how the output of the algorithm should respond to various levels of anxiety sensitivity. In regards to practical considerations, these results when crossed with the ones from Figure 3 show the feasibility for a dramatic reduction in the number of possible trajectories that the system may take while still achieving a high level of performance (for e.g., $W = 0.8$), when comparing these numbers with those of classic Markovian optimization (for $W = 0$). Such reduction can prove highly relevant for operators and stakeholders depending on the output of the system to reduce their own uncertainty on how the system may behave. Last, these results are congruent with findings from psychology, i.e., heightened anxiety sensitivity yields to strategies leading to more controllable

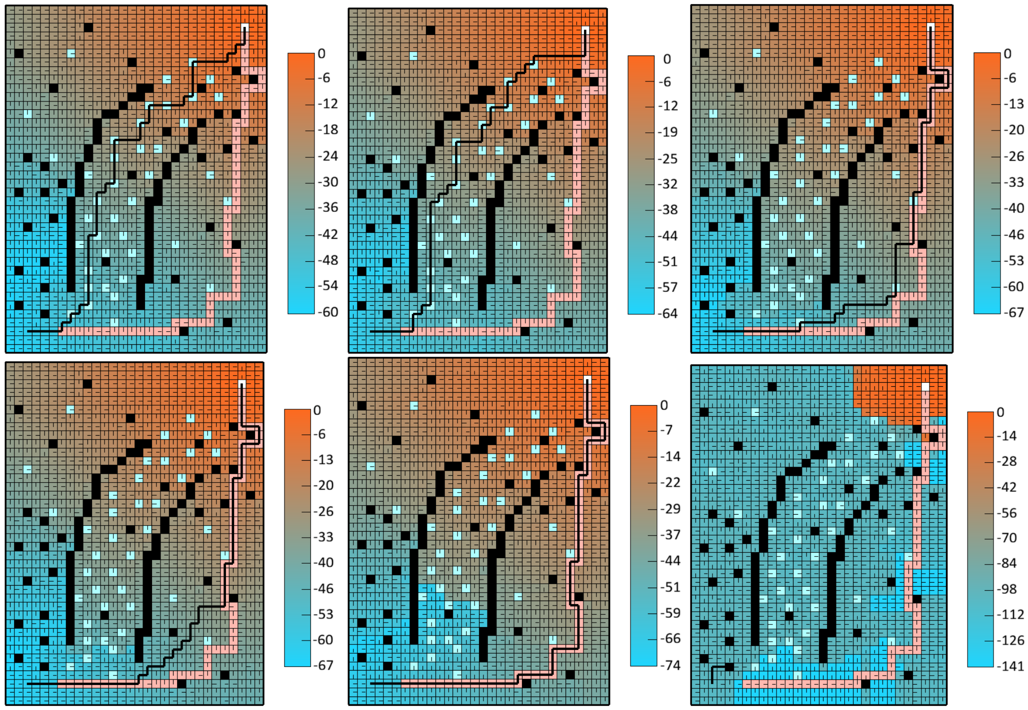

Figure 4: Expected value depending on W. From left to right and top to bottom, $W = 0, W = 0.1, W = 0.3, W = 0.5, W = 0.8$, and $W = 1$. The most likely path according to the policies is indicated by a black line. A heat map is plotted over the *normal* tiles, representing the expected reward from that state.

courses of action (e.g., a preference for routine, habitual behavior, and controllable environments).

Quantitative Behaviorial Indicators

The impact of W on behavior can be quantified via the probability of reaching various distributions of various types of tiles given the policies generated depending on W. Table 3 shows that, as W increases, more and more visits of more slippery tiles are traded off for visits of normal tiles and then of gripping tiles, hence yielding to paths that, in average, are increasingly longer and rely on less uncertain tiles, which are, in average, slower than more uncertain tiles. These results are in alignment with what one may expect from what the algorithm should generate, for the relevance of such algorithm for practical purposes, and for replicating human-like anxiety-sensitive behavior.

Qualitative Behaviorial Indicators

Four main strategies, illustrated in Figure 4, emerge from varying levels of anxiety sensitivity W, each of which can be linked to the human anxiety-sensitive responses to exposure to uncertainty listed in the background section.

Insensitivity to anxiety ($W = 0$): pure rational behavior. In this setting, the robot selects the shortest and fastest route to the goal, maximizing opportunities for crossing slippery tiles, which generally allow reaching the goal quicker.

Minimal Sensitivity ($W = 0.1$): cost-free uncertainty aversion. The system behaves similarly to the $W = 0$ case, except that it systematically trades normal tiles for gripping tiles, provided that the path length and the number of reached slippery tiles remain the same. This type of strategy reflects the functional role of anxiety, which enhances safety and preserves cognitive resources by reducing exposure to avoidable factors of uncertainty.

Mild sensitivity ($W = 0.3$): *satisficing*. The system switches to bottom route, favoring slippery tiles when possible and otherwise gripping tiles. Although it still seeks the shortest path (though no longer the fastest) and opportunistically uses slippery tiles, it sacrifices some optimality to reduce uncertainty by staying on more gripping tiles. Early signs of status quo bias, default option selection, and minimalist decision-making begin to appear.

Medium sensitivity ($W = 0.5$): *uncertainty aversion/intolerance*. The tendencies observed in the previous stage become more pronounced. While maintaining an optimal path length, the system actively avoids stepping on slippery tiles. This behavior mirrors psychological tendencies where increased anxiety sensitivity leads to greater influence of biases, including actively avoiding taking opportunities for the same of reducing the uncertainty, even if these opportunities may only open for positive outcomes (such as e.g., students avoiding an exam to evade the uncertainty of whether they will pass).

High sensitivity ($W = 0.8$): *perfectionism*. the system directly follows the gripping path to the goal. This behavior is akin to perfectionist strategies, where the goal is still pursued, but effectiveness is heavily compromised in favor of

	slippery	normal	gripping	total
$W = 0$	14	42	3	59
$W = 0.1$	14	40	5	59
$W = 0.3$	3	23	37	63
$W = 0.5$	0	21	42	63
$W = 0.8$	0	3	62	65
$W = 1$	goal tile not reached			

Table 3: The distribution of encountered tile types for different settings of W when following the intended route.

minimizing exposure to uncertainty.

Maximal sensitivity ($W = 1$), *self-sabotage/self-harm*. The system chooses the quickest route to a failure state, thereby minimizing exposure to uncertainty. This behavior corresponds to pathological anxiety sensitivity, where reducing anxiety overrides goal-directed behavior (e.g., outright avoidance of goals, self-sabotage, self-harm). Crossing with Table 2, this policy still considers 10^5 paths, which mostly correspond to unlikely sequences of sliding away from the closest obstacle.

Altogether, these behaviors are highly congruent with psychological findings relating how increasing levels of anxiety sensitivity influences human behaviors.

Conclusion

This paper introduces a new formalism that we call Path anxiety-sensitive Markov Decision Processes (PAS-MDP). This formalism is designed to generate policies that balance utility maximization with the minimization of exposure to path uncertainty (i.e. the uncertainty tied to the possible paths that the system may encounter). PAS-MDPs incorporate a model of human anxiety that allows representing and adapting to the psychological impact of uncertainty (Barlow 2002; Miceli and Castelfranchi 2005, 2014; Rapee et al. 1996). The paper also presents two algorithms for finding optimal PAS-MDP policies: an exact algorithm with exponential complexity and an approximate algorithm with polynomial complexity. Last, this paper shows the soundness of PAS-MDPs for achieving practical relevance, by showing their capacity to reduce the exposure to anxiety and the number of possible paths that the system may undertake by several orders of magnitude in exchange for minimal costs in terms of utility. This paper also shows the relevance of PAS-MDPs as a psychologically grounded model, by successfully replicating key anxiety-sensitive behaviors observed in fundamental psychology research. Due to space restrictions (for introducing new background, formalizations, algorithms, and experiments), a structured, multi-experiment comparison with related research such as (Vanhée, Jeanpierre, and Mouaddib 2022; Gutsche and Vanhée 2023) is left for future work. Due to the inherent conceptual differences between paths, rewards, and epistemic uncertainty significant resemblances (e.g., prefer more controllable trajectories) and differences can be expected (e.g. uncertain paths yielding to the same reward vs. few paths yielding to very different rewards).

PAS-MDPs offer exciting insights into the key AP systems scientific challenges that motivate this research. A significant contribution of PAS-MDPs is their ability to generate plans that enhance predictability and control over future courses of action the system is likely to encounter when applying this plan, mirroring human decision processes. From a technical perspective, by proactively and preventively reducing and controlling the range of possible trajectories, PAS-MDPs reduces the exposure to low-probability edge cases and hence having a positive effect on practical deployment. These plans are also simpler to express and represent in a compact way. Our approach also provides the modelling foundations for detecting significant deviations by observing when the system engages in highly improbable paths.

As future work, further variants of PAS-MDPs be considered. In particular alternative measures of instant uncertainty should be investigated, such as capturing specific psychological or applicative dynamics (e.g., individuals may be more likely to disregard certain paths, individuals may rely on simplificatory models of trajectories thus fusing similar paths, individuals may overestimate unlikely paths that are considered undesirable). Alternative models of anxiety-sensitive MDPs should be considered for capturing other forms of uncertainty, such as models using online learning to capture partially-defined transition functions or uncertainties relative to transition functions. Experiments on additional examples can further consolidate the generalizability of the findings. Experiments with different values for k can provide insights on the tradeoffs between computational costs and precision. While differences between PAS-MDPs and risk-aware MDPs (Chow et al. 2015) can be underlined from a theoretical standpoint (e.g., an optimal policy may take additional risks if these risks allows for a rapid reduction of the number of possible paths), a systematic empirical comparison can help further substantiating the differences.

From an applicative standpoint, PAS-MDPs offer highly desirable features for human users, particularly in the realms of trustworthy and explainable AP and human-AP collaboration. By making the system's strategy easier to express, understand, and monitor, PAS-MDPs enhance the transparency and reliability of AP systems. The results show that these desirable properties are achieved with minimal impact on utilitarian optimality. In the long term, PAS-MDP policies contribute to more stable and predictable system behavior, making it easier to integrate these systems into human routines, which is crucial for their relevance in society.

From a social impact standpoint, PAS-MDPs provide a cost-effective approach to reducing the anxiety that AP systems can cause in humans, addressing the growing psychological toll associated with traditional AP systems focusing almost solely on utility maximization (Burr, Taddeo, and Floridi 2020; Saxena 2018). This approach enhances the relevance of AP systems not only as tools for optimizing performance but also as means to incorporate the psychological health and well-being of individuals into societal operations. As our societies move towards the widespread adoption of autonomous vehicles for daily commutes, wouldn't we prefer these vehicles to stay, by default, on the beaten path?

Acknowledgments

The first author acknowledges the financial support of the project "Anxiety-Sensitive Artificial Intelligence" funded by the Swedish Research Council (Grant ID: 2023-04505). The authors wish to thank the reviewers of the paper for their insightful comments.

References

Agrawal, N.; Sirohi, A. K.; Kumar, S.; et al. 2024. No Prejudice! Fair Federated Graph Neural Networks for Personalized Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 10775–10783.

Amitai, Y.; Septon, Y.; and Amir, O. 2024. Explaining reinforcement learning agents through counterfactual action outcomes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 10003–10011.

Andlin-Sobocki, P.; and Wittchen, H.-U. 2005. Cost of anxiety disorders in Europe. *European Journal of Neurology*, 12.

Araya, M.; Buffet, O.; Thomas, V.; and Charpillet, F. 2010. A POMDP extension with belief-dependent rewards. *Advances in neural information processing systems*, 23.

Bach, D. R.; and Dolan, R. J. 2012. Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature reviews neuroscience*, 13(8): 572–586.

Bai, Q.; Mondal, W. U.; and Aggarwal, V. 2024. Regret analysis of policy gradient algorithm for infinite horizon average reward markov decision processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 10980–10988.

Bandelow, B.; and Michaelis, S. 2015. Epidemiology of anxiety disorders in the 21st century. *Dialogues in clinical neuroscience*, 17(3): 327–335.

Barlow, D. 2002. The experience of anxiety: Shadow of intelligence or specter of death. *Anxiety and its disorders: The nature and treatment of anxiety and panic*, 1–36.

Bellman, R. 1956. Dynamic programming and the smoothing problem. *Management Science*, 111–113.

Bie, Y.; Luo, L.; and Chen, H. 2024. Mica: Towards explainable skin lesion diagnosis via multi-level image-concept alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 837–845.

Brosschot, J. F.; Verkuil, B.; and Thayer, J. F. 2016. The default response to uncertainty and the importance of perceived safety in anxiety and stress: An evolution-theoretical perspective. *Journal of anxiety disorders*, 41: 22–34.

Burr, C.; Taddeo, M.; and Floridi, L. 2020. The ethics of digital well-being: A thematic review. *Science and engineering ethics*, 26(4): 2313–2343.

Cai, X.; and Scarlett, J. 2024. Kernelized Normalizing Constant Estimation: Bridging Bayesian Quadrature and Bayesian Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11150–11158.

Capel, T.; and Brereton, M. 2023. What is human-centered about human-centered AI? A map of the research landscape. In *Proceedings of the 2023 CHI conference on human factors in computing systems*, 1–23.

Chen, H.; Jia, Y.; Wang, G.; and Zou, C. 2024. Uncertainty Quantification for Data-Driven Change-Point Learning via Cross-Validation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11294–11301.

Chow, Y.; Tamar, A.; Mannor, S.; and Pavone, M. 2015. Risk-sensitive and robust decision-making: a cvar optimization approach. *Advances in neural information processing systems*, 28.

Chu, Z.; Hu, M.; Cui, Q.; Li, L.; and Li, S. 2024. Task-driven causal feature distillation: Towards trustworthy risk prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 11642–11650.

Cornelissen, J.; and Cholakova, M. 2021. Profits Uber everything? The gig economy and the morality of category work. *Strategic Organization*, 19(4): 722–731.

Dai, Q.; Wei, D.; Liu, H.; Sun, J.; Wang, L.; and Zheng, Y. 2024. Federated Modality-Specific Encoders and Multimodal Anchors for Personalized Brain Tumor Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1445–1453.

Dugas, M. J.; Gagnon, F.; Ladouceur, R.; and Freeston, M. H. 1998. Generalized anxiety disorder: A preliminary test of a conceptual model. *Behaviour research and therapy*, 36(2): 215–226.

Dymond, S.; and Roche, B. 2009. A contemporary behavior analysis of anxiety and avoidance. *The Behavior Analyst*, 32: 7–27.

Elahi, M.; and Rintanen, J. 2024. Optimizing the Optimization of Planning Domains by Automatic Action Schema Splitting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 20096–20103.

Engesser, T.; Herzig, A.; and Perrotin, E. 2024. Towards Epistemic-Doxastic Planning with Observation and Revision. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 10501–10508.

Frost, R. O.; and DiBartolo, P. M. 2002. Perfectionism, anxiety, and obsessive-compulsive disorder.

Gigerenzer, G.; and Goldstein, D. G. 1996. Reasoning the fast and frugal way: models of bounded rationality. *Psychological review*, 103(4): 650.

Grupe, D. W.; and Nitschke, J. B. 2013. Uncertainty and anticipation in anxiety: an integrated neurobiological and psychological perspective. *Nature Reviews Neuroscience*, 14(7): 488–501.

Gutsche, L.; and Vanhée, L. 2023. The Value of Knowledge: Joining Reward and Epistemic Certainty Optimisation for Anxiety-Sensitive Planning. In *International Conference on Autonomous Agents and Multiagent Systems*, 30–42. Springer.

Hartley, C. A.; and Phelps, E. A. 2012. Anxiety and decision-making. *Biological psychiatry*, 72(2): 113–118.

Health, T. L. G. 2020. Mental health matters. *The Lancet. Global Health*, 8(11): e1352.

Heaukulani, C.; Phang, Y. S.; Weng, J. H.; Lee, J.; and Morris, R. J. T. 2024. Deploying AI Methods for Mental Health in Singapore: From Mental Wellness to Serious Mental Health Conditions. In Stanojevic, M., ed., *Proceedings of Machine Learning for Cognitive and Mental Health Workshop (ML4CMH 2024) Co-located with the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2024), Vancouver, BC, Canada, February 26th, 2024*, volume 3649 of *CEUR Workshop Proceedings*, 84–93. CEUR-WS.org.

Horii, S.; and Chikahara, Y. 2024. Uncertainty Quantification in Heterogeneous Treatment Effect Estimation with Gaussian-Process-Based Partially Linear Model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 20420–20429.

Horned, A.; and Vanhée, L. 2023a. From Threatening Pasts to Hopeful Futures. A Review of Agent-Based Models of Anxiety. In *Conference of the European Social Simulation Association*, 139–152. Springer.

Horned, A.; and Vanhée, L. 2023b. Models of anxiety for agent deliberation: the benefits of anxiety-sensitive agents. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 1761–1767.

Hou, P.; Yeoh, W.; and Varakantham, P. 2014. Revisiting risk-sensitive MDPs: New algorithms and results. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 24, 136–144.

Huang, H. 2023. "The Food Delivered is More Valuable Than My Life": Understanding the Platform Precarity of Online Food-Delivery Work in China. *Journal of Contemporary Asia*, 53(5): 852–868.

Johnson, E. J.; and Goldstein, D. 2003. Do defaults save lives?

Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2): 99–134.

Komporozos-Athanasiou, A.; and Haiven, M. 2023. Anxiety and self-sabotage in the neoliberal university. In *Clickbait capitalism*, 209–229. Manchester University Press.

Kukacka, M. 2010. Bayesian methods in artificial intelligence. In *WDS*, volume 10, 25–30.

Kusters, R.; Misevic, D.; Berry, H.; Cully, A.; Le Cunff, Y.; Dandoy, L.; Díaz-Rodríguez, N.; Ficher, M.; Grizou, J.; Othmani, A.; et al. 2020. Interdisciplinary research in artificial intelligence: challenges and opportunities. *Frontiers in big data*, 3: 577974.

Lai, K.-H.; Yang, Z.-R.; Lai, P.-Y.; Wang, C.-D.; Guizani, M.; and Chen, M. 2024. Knowledge-Aware Explainable Reciprocal Recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 8636–8644.

Lei, C.; Lipovetzky, N.; and Ehinger, K. A. 2024. Generalized planning for the abstraction and reasoning corpus. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 20168–20175.

Lindgren, S. 2023. *Critical theory of AI*. John Wiley & Sons.

Liu, J.; Xie, C.; Deng, Q.; Ge, D.; and Ye, Y. 2024. Sketched Newton Value Iteration for Large-Scale Markov Decision Processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 13936–13944.

Luan, S.; Schooler, L. J.; and Gigerenzer, G. 2014. From perception to preference and on to inference: An approach–avoidance analysis of thresholds. *Psychological Review*, 121(3): 501.

Ma, H.; Wang, C.; Zhu, H.; Yang, S.; Zhang, X.; and Zhang, X. 2024. Enhancing cognitive diagnosis using un-interacted exercises: A collaboration-aware mixed sampling approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 8877–8885.

Massazza, A.; Kienzler, H.; Al-Mitwalli, S.; Tamimi, N.; and Giacaman, R. 2023. The association between uncertainty and mental health: a scoping review of the quantitative literature. *Journal of Mental Health*, 32(2): 480–491.

Miceli, M.; and Castelfranchi, C. 2005. Anxiety as an "epistemic" emotion: An uncertainty theory of anxiety. *Anxiety, Stress, and Coping*, 18(4): 291–319.

Miceli, M.; and Castelfranchi, C. 2014. *Expectancy and emotion*. OUP Oxford.

O'Connor, R. C.; Rasmussen, S.; and Hawton, K. 2010. Predicting depression, anxiety and self-harm in adolescents: The role of perfectionism and acute life stress. *Behaviour research and therapy*, 48(1): 52–59.

Omura, M.; Osa, T.; Mukuta, Y.; and Harada, T. 2024. Symmetric Q-learning: Reducing Skewness of Bellman Error in Online Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 14474–14481.

Peng, S.; Yang, S.; and Yao, J. 2023. Improving value-at-risk prediction under model uncertainty. *Journal of Financial Econometrics*, 21(1): 228–259.

Pires, G. N.; and Pinto, J. P. G. 2020. Gig Economy, Austerty and "Uberization" of Labor in Brazil (2014–2019). *Argumentum*, 12(3): 237–255.

Pratt, J. W.; Raiffa, H.; and Schlaifer, R. 1964. The foundations of decision under uncertainty: An elementary exposition. *Journal of the American statistical association*, 59(306): 353–375.

Qian, T.; Chen, J.; Zhuo, L.; Jiao, Y.; and Jiang, Y.-G. 2024. Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4542–4550.

Rapee, R. M.; Craske, M. G.; Brown, T. A.; and Barlow, D. H. 1996. Measurement of perceived control over anxiety-related events. *Behavior Therapy*, 27(2): 279–293.

Regan, K.; and Boutilier, C. 2012. Regret-based reward elicitation for Markov decision processes. *arXiv preprint arXiv:1205.2619*.

Russell, S.; and Norvig, P. 2021. Artificial Intelligence: A Modern Approach.

Saffiotti, A. 1987. An AI view of the treatment of uncertainty. *The Knowledge Engineering Review*, 2(2): 75–97.

Samuelson, W.; and Zeckhauser, R. 1988. Status quo bias in decision making. *Journal of risk and uncertainty*, 1: 7–59.

Saxena, S. 2018. Human-AI Interaction: Psychological Perspectives. *International Journal of Advanced Multidisciplinary Scientific Research*.

Schoenfeld, W. N. 1950. An experimental approach to anxiety, escape and avoidance behavior. In *Annual Meeting of the American Psychopathological Association, 39th, Jun, 1949, New York City, NY, US*. Grune & Stratton.

Shen, J.; Qian, H.; Zhang, W.; and Zhou, A. 2024. Symbolic Cognitive Diagnosis via Hybrid Optimization for Intelligent Education Systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 14928–14936.

Sigaud, O.; and Buffet, O. 2013. *Markov decision processes in artificial intelligence*. John Wiley & Sons.

Silver, D.; and Veness, J. 2010. Monte-Carlo planning in large POMDPs. *Advances in neural information processing systems*, 23.

Spiegelhalter, D. 2011. Quantifying uncertainty. *Risk*, 24: 17.

Spielberger, C.; Sydeman, S.; Owen, A.; and Marsh, B. 1999. Measuring anxiety and anger with the State-Trait Anxiety Inventory (STAI) and the State-Trait Expression Inventory (STAXI). *The Use of Psychological Testing for Treatment Planning: An Outcomes Assessment, 2nd ed. ME Maruish, Ed. Lawrence Erlbaum Associates. Mahwah*.

Stanojevic, M., ed. 2024. *Proceedings of Machine Learning for Cognitive and Mental Health Workshop (ML4CMH 2024) Co-located with the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI 2024), Vancouver, BC, Canada, February 26th, 2024*, volume 3649 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Toroghi, A.; and Sanner, S. 2024. Bayesian Inference with Complex Knowledge Graph Evidence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 20550–20558.

Tulshan, A. S.; and Dhage, S. N. 2019. Survey on virtual assistant: Google assistant, siri, cortana, alexa. In *Advances in Signal Processing and Intelligent Recognition Systems: 4th International Symposium SIRS 2018, Bangalore, India, September 19–22, 2018, Revised Selected Papers 4*, 190–201. Springer.

Tversky, A.; and Fox, C. R. 1995. Weighing risk and uncertainty. *Psychological review*, 102(2): 269.

Vanhée, L.; Jeanpierre, L.; and Mouaddib, A.-I. 2022. Anxiety-sensitive planning: from formal foundations to algorithms and applications. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, 730–740.

Vashistha, R.; and Farahi, A. 2024. U-Trustworthy Models. Reliability, Competence, and Confidence in Decision-Making. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 19956–19964.

Wang, S.; and Li, K. 2024. Constrained Bayesian optimization under partial observations: Balanced improvements and provable convergence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15607–15615.

Whitmont, E. 1952. Impatience: A Differential Diagnosis.

Wu, D.; Chi, M.; Zang, X.; and Peng, B. 2024a. Solving Spectrum Unmixing as a Multi-Task Bayesian Inverse Problem with Latent Factors for Endmember Variability. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15878–15887.

Wu, T.; Neiswanger, W.; Zheng, H.; Ermon, S.; and Leskovec, J. 2024b. Uncertainty Quantification for Forward and Inverse Problems of PDEs via Latent Global Evolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 320–328.

Yang, D.; He, X.; Wang, J.; Yu, G.; Domeniconi, C.; and Zhang, J. 2024. Federated Causality Learning with Explainable Adaptive Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16308–16315.

Yotam, T.; and Indelman, V. 2024. Measurement Simplification in ρ-POMDP with Performance Guarantees. *IEEE Transactions on Robotics*.

Yu, H.; Liu, Z.; and Luo, X. 2024. Barely Supervised Learning for Graph-Based Fraud Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16548–16557.

Yu, S. X.; Lin, Y.; and Yan, P. 1998. Optimization models for the first arrival target distribution function in discrete time. *Journal of mathematical analysis and applications*, 225(1): 193–223.

Zhang, X.; Zheng, B.; Hu, J.; Li, C.; and Bai, X. 2024a. From Toxic to Trustworthy: Using Self-Distillation and Semi-supervised Methods to Refine Neural Networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16873–16880.

Zhang, Y.; Wang, J.; Yu, L.-C.; Xu, D.; and Zhang, X. 2024b. Personalized LoRA for Human-Centered Text Understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 19588–19596.

Zhou, R.; Gao, C.-X.; Zhang, Z.; and Yu, Y. 2024. Generalizable Task Representation Learning for Offline Meta-Reinforcement Learning with Data Limitations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17132–17140.

Zhu, T.; Qiu, Y.; Zhou, H.; and Li, J. 2024. Decoding Global Preferences: Temporal and Cooperative Dependency Modeling in Multi-Agent Preference-Based Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17202–17210.

Zou, R.; Liu, S.; Luo, Y.; Liu, Y.; Feng, J.; Wei, M.; and Sun, J. 2024. Balancing Humans and Machines: A Study on Integration Scale and Its Impact on Collaborative Performance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17628–17636.