Anxiety-Sensitive Planning: From Formal Foundations to Algorithms and Applications

Loïs Vanhée[1, 2, *], **Laurent Jeanpierre**[2], **Abdel-Illah Mouaddib**[2]

[1] Umeå University, Sweden
[2] Université de Caen, France
lois.vanhee@umu.se, laurent.jeanpierre@unicaen.fr, abdel-illah.mouaddib@unicaen.fr

Abstract

Anxiety is the most prominent source of stress, harmful behaviours, and psychological disorders. AI systems, usually built for maximizing performance, increase the worldwide exposition to anxiety. This foundational paper introduces Anxiety-Aware Markov Decision Processes (AA-MDPs), the first formalism rooted in fundamental psychology research for modelling the anxiety tied to policies. In addition, this paper formalizes models and practical polynomial algorithms for generating anxiety-sensitive policies. Empirical validation demonstrates that AA-MDPs policies replicate the influence of anxiety on human decision-making observed by fundamental psychology research. Last, this paper demonstrates that AA-MDPs are directly applicable for social good, through a real-world use case (Anxiety-Sensitive Itinerary Planning), the immediate applicability for augmenting any formerly-defined MDP model with anxiety-awareness, and direct tracks developing future high-impact models.

Introduction

Anxiety is a fundamental future-oriented emotional drive that can also threaten human wellbeing, as it also triggers worry, stress, insomnia, aggressiveness, disengagement, panic attacks, phobia, obsessive compulsive disorders, eating disorders, psychosis, depression, addiction, post-traumatic stress disorder, generalized anxiety disorder, cardiovascular issues, and self-harm, including suicide (Andrews et al. 2008; Avery, Clauss, and Blackford 2016; Grupe and Nitschke 2013; Hartley, Barrowclough, and Haddock 2013; Huppert 2009; Neumann, Veenema, and Beiderbeck 2010; Taylor et al. 2005). Anxiety disorders are the largest source of mental health disorders worldwide, affecting 33% of the population, totalling an annual cost of €74.4 billion in 2010 (€41 billion in 2004) only for the European Union, with excess costs up to €1600 per case (Alonso et al. 2018; Andlin-Sobocki and Wittchen 2005; Bandelow and Michaelis 2015; Wittchen 2002). The Lancet Commission defined anxiety as a key factor to be covered by the wellbeing Sustainable Development Goal (Patel et al. 2018; Firth et al. 2019).

Artificial Intelligence (AI) has the potential for both alleviating and creating anxiety at a world scale in human societies. Despite the importance now given to wellbeing in AI (IEEE 2017), simply considering and, a fortiori, avoiding anxiety remains barely touched upon. At best, specific approaches such as artificial *emotional intelligence* (Schuller and Schuller 2018) allow adjusting pre-written plans to observed emotions (Loizou et al. 2012). Worse, the current AI paradigm, near-ubiquitously driven towards performance maximization, is tied to incidentally cause undesirable consequences as side effects (Kulynych et al. 2020) and is seemingly prone to inducing greater anxiety, by directly and indirectly maximizing the stress of resources and engaging in risk-taking strategies (e.g. optimized plans with minimal room for error, high-risk high-rewards strategies).

AI has the potential for helping us to better deal with the emotional impact of decisions. As an example, the COVID-19 crisis has been a paroxysmal display of anxiety with a deep impact on our daily life. The diffuse yet unquestionable threat brought by the virus lead to massive reactions, from individuals to governments, from locking down activities to feeling threatening towards and threatened by our closest ones (Briscese et al. 2020). This anxiety had and still has large-scale global psychological, social, and economic repercussions (Restauri and Sheridan 2020; Salari et al. 2020). What if we had models for estimating the psychological repercussions of various measures? What if these models proposed alternatives that reduced the subsequent anxiety? For answering such questions and others more general about the topic, this paper develops: 1) a model grounded in psychological theories for assessing the anxiety expected to be raised by a policy; 2) a planning model, and 3) efficient algorithms for generating policies that balance induced anxiety and expected value, beyond the COVID-19 crisis.

This paper places itself as a foundation for the domain of anxiety-sensitive planning, a high-impact wellbeing-critical topic that remains to be studied by the AI community. Anxiety being obviously too broad for being fully covered in a single paper, the present paper is dedicated to study how to generate anxiety-sensitive policies, i.e. policies that best balance anxiety-avoidance and performance, depending on a degree of sensitivity to anxiety. 1) We lay the formalization of *Anxiety-Aware Markov Decision Processes (AA-MDPs)*, the first anxiety-sensitive planning model based on MDPs,

*Contact author

grounded in fundamental psychological theories, which can be directly applied to augment any existing MDP model with anxiety-sensitivity; 2) we prove that solving AA-MDPs is an NP-hard problem, 3) we introduce polynomial approximate algorithms that are used 4) for empirically validating the soundness of AA-MDPs by demonstrating that they replicate core anxiety-sensitive strategies observed in human behavior; and 5) demonstrating the suitability of AA-MDPs for mitigating worldwide sources of anxiety: Anxiety-Aware Itinerary Planner[1]; last, 6) we present direct follow-up applications and extensions of the foundations laid by AA-MDPs, such as therapeutic applications for post-traumatic stress disorders and panic attacks.

Background

Anxiety is a central topic in (neuro)psychology, generally defined as a future-oriented emotion related to a subjective impression of inability to cope with future events, insufficient preparation, and worries (Barlow 2002a; Miceli and Castelfranchi 2005; Rapee et al. 1996). Anxiety can be sketched as *anxiety* = *uncertainty* × *motive*, crossing of the *intensity of internal motives* (e.g. goals, needs, and plans) with the *intensity of uncertainty* inherently caused by the partial mastery of the environment. Higher stakes and higher uncertainty cause higher anxiety. Uncertain positive as well as negative outcomes generate anxiety (Carleton 2016). For example, "100% chance of winning 50$" causes less anxiety than "20% chance of winning 0$, 30% chance of winning 50$, and 50% chance of winning 100$", which causes less anxiety than "50% chance of winning 100$, else 0$" (Bach and Dolan 2012). (Bach and Dolan 2012) characterize four central sources of uncertainty: sensory uncertainty, (partially observable) state uncertainty, (transition) rule uncertainty, and outcome uncertainty. As a heuristic, anxiety arises from the exposition to perceived unknowns (Carleton 2016), notably when repeated over a short timespan (Levy and Schiller 2020). Anxiety can be caused by both externally observed uncertainty and internal exposition to uncertainty (e.g. focus on uncertainty, backpropagation of future projections (Levy and Schiller 2020)). For space consideration, our model focuses on reward uncertainty, laying direct foundations for additional uncertainty models (detailed in conclusions). Anxiety differs from fear in that fear is 1) a primal reaction to 2) explicit threats, whereas anxiety is triggered by the potentiality of threats over longer periods of time (Grupe and Nitschke 2013).

Anxiety relates to numerous physiological and behavioural responses, driven either to adapt to potential threats usually through increasing predictability and controllability (e.g. preparation, arousal, hypervigilance) (Beck and Clark 1997); or to cope with the experience of anxiety (e.g. impatience, self-sabotaging, pessimism, optimism, substance abuse) (Barlow 2002a; Levy and Schiller 2020; Rapee et al. 1996). While relevant in a state of nature with lethal threats, anxiety often triggers maladaptive responses in modern societies, sometimes up to a pathological extent (e.g. social media addictively overemphasizing social anxiety), causing direct and indirect comorbidity in the long run (physical diseases, sleep disorders, impatience, exhaustion, depression, substance abuse) (Belik, Sareen, and Stein 2009).The **instant anxiety**[2] is the anxiety experienced at a given point in time (Spielberger et al. 1999). The **cumulated anxiety** is the addition of the instant anxiety (to be) experienced over time, which individuals seek to avoid (Dugas et al. 1998). The sensitivity to and reactions caused by anxiety are individual-dependent, including biological, developmental, and cultural factors (Brandes and Bienvenu 2006; Hofstede, Hofstede, and Minkov 2010). Thus, individuals react very differently to various types of stimuli (e.g. not knowing when versus not knowing what). Moreover, while the intensity of anxiety and the exposition duration are positively tied to psychological repercussions, this relation leads to different outcomes in non-additive individual-dependent manners: a long exposition to mild anxiety intensity is tied with risks of depression, while a short exposition to high anxiety intensity is tied to post-traumatic stress disorder (Ehlers and Clark 2000). As such, cumulative anxiety is more than a sum of instant anxiety over time. As more specific relations are still under investigation in research in psychology, our models need to be generic for various individual and temporal components.

Markov Decision Processes (MDPs) are stochastic processes that are (partly) controlled by 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; $T : S \times A \times S \rightarrow [0, 1]$ is a transition function, where $T(s, a, s')$ is a probability of reaching s' when playing a from s; and $R : S \rightarrow \mathbb{R}$ is a reward function, where $R(s)$ is the expected reward for being in $s \in S$. Optimizing MDPs within a bounded horizon $h \in \mathbb{N}$ consists in finding a policy that maximizes the expected reward in using h actions. A policy $\pi : S \rightarrow A$ defines an action to be played for each state and Π is the set of all policies. The value of a policy π within a horizon h is $V^\pi_h : S \rightarrow \mathbb{R}$, where: $V^\pi_k(s) = R(s) + \sum_{s' \in S} T(s, \pi(s), s').V^\pi_{k-1}(s')$ if $k > 0$, else $V^\pi_0(s) = R(s)$. The set of optimal policies for the horizon h, represented by $\Pi^*_h \subseteq \Pi$ is the set of policies $\pi^* \in \Pi^*_h$ such that for all $s \in S$, $V^{\pi^*}_h(s) = \max_{\pi \in \Pi} V^\pi_h(s)$.

An example MDP is presented in Figure 1. $S = \{s_1, s_{21}, s_{22}, s_{31}, s_{32}, s_{33}, s_{34}\}$, $A = \{a, b, c, d, e, f\}$. T is defined in s_1 for a as $T(s_1, a, s_{21}) = 0.7$, $T(s_1, a, s_{22}) = 0.3$. $R(s_1) = 0$; $R(s_{21}) = 2$; $R(s_{22}) = 1$; $R(s_{31}) = -10$; $R(s_{32}) = 20$; $R(s_{33}) = 0$ $R(s_{34}) = 1$. A set of policies can be applied in this MDP. The **reward-maximizing policy** (or **anxiety-blind**), π^* is: $\pi^*(s_1) = a$; $\pi^*(s_{21}) = c$; $\pi^*(s_{22}) = f$. For later use, we introduce the **low-anxiety policy** $\pi^\sim$, where $\pi^\sim(s_1) = a$, $\pi^\sim(s_{21}) = d$ and $\pi^\sim(s_{22}) = f$, and the **zero-anxiety policy** π_0, where $\pi_0(s_1) = b$ and $\pi_0(s_{22}) = f$. The **value** of π^* is $V^{\pi^*}(s_1) = 5.5$; $V^{\pi^*}(s_{21}) = 7$; $V^{\pi^*}(s_{22}) = 2$; $V^{\pi^*}(s_{31}) = -10$; $V^{\pi^*}(s_{32}) = 20$; $V^{\pi^*}(s_{33}) = 0$; $V^{\pi^*}(s_{34}) = 1$.

[1]The URL to our open-source AA-MDP repository is available there: https://github.com/lvanhee/anxiety-aware-mdps.

[2]Generally called State Anxiety in psychology, but "instant anxiety" avoids ambiguity with Markovian states.

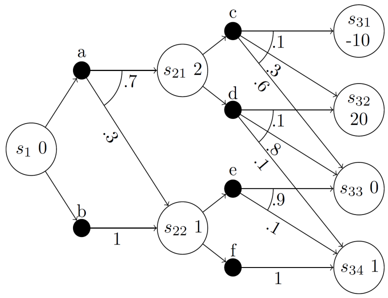

Figure 1: Markov Decision Process example. White nodes are states, including the state name and reward. Black nodes are actions. Numbers on lines are probability distributions.

Anxiety-Aware Markov Decision Process

MDPs provide the basic modelling structures for representing the key features of anxiety depicted by fundamental psychological research: uncertainty regarding future situations and rewards, including the effect of selected courses of action (policy) (Barlow 2002a; Miceli and Castelfranchi 2005; Rapee et al. 1996), while keeping natural representation and polynomial computational complexity. However, MDP structures are to be expanded for modelling anxiety, and some of these combinations need to remain generic, as the state of the art in psychology does not provide sufficiently grounded insight. Experimentations described later in this paper demonstrate how to implement AA-MDPs.

Instant Anxiety is directly tied to the uncertainty and intensity of future outcomes, which depends on the current state, world-dynamics, and intended policy. (Carleton 2016) formalize outcome-based anxiety as the variance of the possible outcomes weighted by their probability. This formalization can be naturally expanded to the MDP framework by representing *value-distributions*, i.e. the distribution of future expected accumulated rewards from a given state and a policy (instead of just the expected average captured by the Bellmann value function (Bellman 1956)). Note that instant anxiety cannot be reduced to rewards, as rewards are only state-dependent while anxiety is also policy-dependent. The model also assumes that the anticipated reward associated to every state is static, whereas human anticipation of the valence of situations can evolve over time (e.g. discounted sensitivity to negative outcomes when they are far in the future, increased sensitivity to negative outcomes when experiencing a failure) (Kahneman 2011).

Definition 1 *A **value-distribution** $\Delta : \mathbb{Q} \to [0,1]$ is a discrete probability distribution over $\mathbb{Q}$, an arbitrary discrete subset of $\mathbb{R}$. $\Delta(r)$ is the probability of acquiring $r \in \mathbb{Q}$. $\sum_{r\in\mathbb{Q}} \Delta(r) = 1$. The set of all value-distributions is $\mathbf{\Delta}$. By extension, $\Delta(r) = 0$ for any $r \in \mathbb{R} \setminus \mathbb{Q}$. Δ can be represented as the set of pairs $\{(r, \Delta(r))|r \in \mathbb{Q}\}$.*

Intuitively, the degree of anxiety raised by situations is tied to the variance of the value-distributions. For example, consider $\Delta_1 = \{(-8, 0.07), (2, 0.72), (22, 0.21)\}$, $\Delta_2 = \{(-8, 0.1), (2, 0.6), (22, 0.3)\}$, and $\Delta'_2 = \{(-7, 0.2), (2, 0.5), (22, 0.3)\}$ (Δ_1 and Δ_2 correspond to the value-distributions from s_1 and s_{21} when following π^*; Δ'_2 is similar to Δ_2 except that the worst case is slightly better but more probable at the expense of the average case for Δ'_2). Δ_1 means that there is a probability 0.07 to obtain a reward of -8, a probability 0.72 of obtaining 2 and a probability 0.21 of obtaining 22. Δ_1 objectively generates less anxiety than $\Delta_2 = \{(-8, 0.1), (2, 0.6), (22, 0.3)\}$: in Δ_2, the uncertainty is greater as the probability to reach more extreme values (i.e. -8 and 22) is strictly greater than in Δ_1 and the probability to reach less extreme values (i.e. 2) is the strictly lower. Yet, the anxiety from reward distributions has no total order: certain individuals would experience more anxiety from Δ_2 (greater extremes can be reached) while some others would experience more anxiety from Δ'_2 (greater probability of reaching extremes). Value-distributions can be seen as cumulative probability distributions.

In the context of MDPs, a given state and policy can be tied to value-distributions, describing the probability of every possible future value. Value-distributions can naturally be computed using a straightforward extension of the Bellman equation that propagates the probability of every possible reachable outcome, instead of only the aggregated average. This model preserves the original structures for representing the future-orientation in MDPs and can be efficiently computed by expanding classic algorithms such as the value-iteration algorithm (Sigaud and Buffet 2010).

Manipulating value-distributions in $\mathbf{\Delta}$ (instead of values in $\mathbb{R}$) requires introducing additional operators, notably for computing the expected value-distribution given a state and a policy. We formalize this expansion by expanding the classic Bellman equation $V^\pi_h(s) = R(s) + \sum_{s'\in S} T(s, \pi(s), s') \times V^\pi_{h-1}(s')$ to value-distributions. In particular, we introduce the operator $\oplus$ for expanding to value-distributions the addition of immediate rewards; and the operator $\otimes$ for expanding to value-distributions the operation $\sum_{s'\in S} T(s, \pi(s), s') \times V^\pi_{h-1}(s')$ of factoring in the probability of future (distributions of) outcomes to occur given an action $\pi(s)$.

Definition 2 *The **scalar sum** operator is represented by $\vec{\oplus}$: $\mathbf{\Delta} \times \mathbb{R} \to \mathbf{\Delta}$, where if $\Delta' = \vec{\oplus}(\Delta, x)$ then $\Delta'(r + x) = \Delta(r)$. Informally, $\vec{\oplus}(\Delta, x)$ shifts the rewards of Δ by x.*

Definition 3 *The **weighted product of value-distributions** is represented by $\otimes : P(\mathbf{\Delta}) \times W \to \mathbf{\Delta}$, where $W : \mathbf{\Delta} \to [0,1]$ is a weight function over $\mathbf{\Delta}$. $\otimes(\mathbf{\Delta'}, w)(r) = \sum_{\Delta\in\mathbf{\Delta'}} w(\Delta) \times \Delta(r)$, where $\mathbf{\Delta'} \subseteq \mathbf{\Delta}$ is the set of value-distributions to be crossed and $w : \mathbf{\Delta'} \to [0,1]$.$\otimes$ is defined similarly to mixtures of distributions (Everitt and Hand 1981)*

Example 1 *If $\Delta_1 = \{(-2, 0.1); (3.1, 0.6); (10, 0.3)\}$, $\Delta_2 = \{(-3, 0.2); (3.1, 0.4); (0, 0.4)\}$, $w(\Delta_1) = 0.2$ and $w(\Delta_2) = 0.8$ then $\vec{\oplus}(\Delta_1, 2) = \{(0, 0.1), (5.1, 0.6), (12, 0.3)\}$ $\otimes(\{\Delta_1, \Delta_2\}, w) = \{(-3, 0.16); (-2, 0.02); (3.1, 0.44); (0, 0.32); (10, 0.06)\}$ Note the merge that happens for $r = 3.1$.*

Given an MDP $\langle S, A, T, R\rangle$, value-distributions can model the uncertainty raised by following a policy as:

Distribution	Probability
$d_{s_1,2,\pi^*}(\Delta_1, \Delta_2, \{(-10,1)\})$	0.07
$d_{s_1,2,\pi^*}(\Delta_1, \Delta_2, \{(20,1)\})$	0.21
$d_{s_1,2,\pi^*}(\Delta_1, \Delta_2, \{(0,0)\})$	0.42
$d_{s_1,2,\pi^*}(\Delta_1, \{(2,1)\}, \{(1,1)\})$	0.3

Table 1: Cumulative anxiety of π^* from s_1 in Figure 1.

Definition 4 *The* ***value-distribution of applying*** $\pi \in \Pi$ ***from a state*** $s \in S$ ***for a horizon*** $h \in \mathbb{N}$ *is represented as* $\Delta_{s,h,\pi} \in \mathbf{\Delta}$*, where:*

$$\Delta_{s,h,\pi} = \begin{cases} \{(R(s),1)\} & \text{if } h=0 \\ \vec{\oplus}(\otimes(\{\Delta_{s',h-1,\pi} | s' \in S\}, w), R(s)) & \text{otherwise} \end{cases}$$

where $w(\Delta_{s',h-1,\pi}) = T(s,\pi(s),s')$*. Informally,* $\Delta_{s,h,\pi}(r)$ *is defined as either the immediate reward from s when* $h=0$ *or the immediate reward added to the reachable value-distributions in h steps, weighted by the probability of reaching them given* π *and T.*

Definition 5 *The* ***instant anxiety raised by*** $\pi \in \Pi$ ***from a state*** $s \in S$ ***for a horizon*** $h \in \mathbb{N}$ *is represented based on* $\Delta_{s,h,\pi} \in \mathbf{\Delta}$*, which is the value-distribution from following* π *from s for h steps.*

For example, the instant anxiety for π^* is $\Delta_{s_1,2,\pi^*} = \{(-8,0.07),(2,0.72),(22,0.21)\}$, $\Delta_{s_{21},1,\pi^*} = \{(-8,0.1),(2,0.6),(22,0.3)\}$, $\Delta_{s_{22},1,\pi^*} = \{(2,1)\}$, $\Delta_{s_{31},0,\pi^*} = \{(-10,1)\}$, $\Delta_{s_{32},0,\pi^*} = \{(20,1)\}$, $\Delta_{s_{33},0,\pi^*} = \{(0,1)\}$, $\Delta_{s_{34},0,\pi^*} = \{(1,1)\}$.

This value-distribution model preserves the utility properties of Markovian value (see technical annexes).

Cumulative Anxiety aggregates the exposition to instant anxiety over time. As the repercussions of cumulative anxiety are individual-dependent and still under scrutiny by the psychology community they need to be kept generic. Therefore, our models represent the cumulative anxiety as *probabilistic trajectories of instant-anxiety*.

Definition 6 *A* ***trajectory of reachable instant anxiety levels for a horizon*** $h \in \mathbb{N}$ *is a sequence of value-distributions represented by the tuple* $\mathbf{\Delta}^h$*.*

A ***distribution of trajectories for a horizon*** $h \in \mathbb{N}$ $D_{\Delta^h} : \mathbf{\Delta}^h \to [0,1]$ *is a distribution over* $\mathbf{\Delta}^h$*.*

An ***expected cumulative anxiety for a horizon*** $h \in \mathbb{N}$ ***for a policy*** $\pi \in \Pi$ ***from a state*** $s \in S$ *is recursively defined as:* $d_{s,h,\pi} \in D_{\Delta^h}$ *where:* $d_{s,h,\pi}(\Delta_1,\ldots,\Delta_h) = \sum_{s' \in S} T(s,\pi(s),s') \times d_{s',h-1,\pi}(\Delta_2,\ldots,\Delta_h)$ *if* $h > 1$*;* $d_{s,1,\pi}(\Delta_1) = 1$ *if* $\Delta_1 = \Delta_{s,\pi,1}$*, and* $d_{s,1,\pi}(\Delta_1) = 0$ *otherwise. Informally,* $d(\Delta_1,\ldots,\Delta_h)$ *measures the expected probability of experiencing the sequence of instant anxieties* $(\Delta_1,\ldots,\Delta_h)$ *when following* π*.*[3]

Pragmatic Anxiety-Avoiding Policies Individuals compromise anxiety-avoidance and reward-maximization, thus fitting the multiobjective optimization paradigm. Comparators allow representing the relative importance between rewards and cumulative anxiety profiles.

[3]The cumulative anxiety of π^* is exemplified in Table 1.

Algorithm 1: Computing π_h^* for Practical Anxiety-Aware MDPs $\langle S,A,T,R,R_\Delta,W,s_0,h\rangle$. *backpropDistr* is operationalized from Definition 4

$\forall s \in S, \Delta_0(s) \leftarrow \{(0,1)\},\ V_0(s) \leftarrow 0,\ AV_0(s) \leftarrow 0$
for $i \in [1,h]$ **do**
  for $s \in S$ **do**
    $\pi_i^*(s) \leftarrow argmax_{a \in A}$

$$\sum_{s' \in s} T(s,a,s') \times ((1-W) \times (R(s') + V_{i-1}(s')) + W \times (R_\Delta(\Delta_{i-1}(s')) + AV_{i-1}^{\pi^{*,i-1}}(s')))$$

    $\Delta_i(s) \leftarrow backpropDistr(\Delta_{i-1}(s), \pi_i^*(s))$

$$V_i(s) \leftarrow R(s) + \sum_{s' \in S} T(s,\pi_i^*(s),s') \times V_{i-1}(s')$$

$$AV_i(s) \leftarrow R_\Delta(\Delta_i(s)) + \sum_{s' \in S} T(s,\pi_i^*(s),s') \times AV_{i-1}(s')$$

  end for
end for
return π^*

Definition 7 *A* ***preference over expected cumulative anxiety for a horizon*** $h \in \mathbb{N}$ *is a partial order* $\prec_D$ *over* D_{Δ^h}*.*

The ***anxiety/performance preference for a horizon*** *h* ***and a state*** s_0 *is represented by the preorder* $\prec_{\mathbb{R},D}$ *over* $\mathbb{R} \times D$*.* $\prec_{\mathbb{R},D}$ *is defined as a combination of the classic* $\leq$ *comparator of* $\mathbb{R}$ *and an expected cumulative anxiety evaluator* $\prec_D$*. I.e. for any* $r,r' \in \mathbb{R}$ *and* $d,d' \in D$*, iff* $r \leq r'$ *and* $d \prec_d d'$ *then* $((r,d)) \prec_{\mathbb{R},D} ((r',d'))$*.* ***Optimal anxiety-aware policy*** π^* *are policies that maximize* $\prec_{\mathbb{R},D}$*.*

$\prec_{\mathbb{R},D}$ models important preferences over personal anxiety, such as discounting anxiety for average positive rewards or for avoiding risks of negative rewards (Kahneman 2011). Note that this formalization, while relying on a certain horizon h, can be applied for modelling the discounted importance of anxiety in a further future. The previous definitions allow to define AA-MDPs:

Definition 8 *An* ***Anxiety-Aware MDP*** *(AA-MDP) is a tuple* $\langle S,A,T,R,s_0,h,\mathbf{\Delta},D,\prec_D,\prec_{\mathbb{R},D}\rangle$ *where* $\langle S,A,T,R\rangle$ *is a classic MDP;* $s_0 \in S$ *is a starting state;* $h \in \mathbb{N}$ *a horizon;* $\mathbf{\Delta}$*, derived from* S,A,T,R *and h, is the set of value-distributions;* $\prec_D$ *is a preference over expected cumulative anxiety at horizon h, and* $\prec_{\mathbb{R},D}$ *is an anxiety/performance preference comparing cumulative anxiety with* $\prec_D$*.*

In the example from Figure 1, $\{\pi^*, \pi^\sim, \pi_0\}$ form a Pareto-front of pragmatic anxiety-avoiding policies. Whereas individual-dependent, intuitively we have $d_{s,h,\pi^*} \prec_D d_{s,h,\pi^\sim} \prec_D d_{s,h,\pi_0}$ and $V^{\pi_0}(s_0) < V^{\pi^\sim}(s_0) < V^{\pi^*}(s_0)$.

Property 1 *Computing optimal anxiety-aware policies is NP-hard. Proof in technical annexes.*

Polynomial Algorithms

Practical AA-MDPs are AA-MDP approximations that can be solved in polynomial time. These approximations are multicriteria MDPs, combining reward and anxiety as a second criterion. Using classic linearisation techniques, policies can be computed in polynomial time (Roijers, Vamplew, and Whiteson 2013; Sigaud and Buffet 2010). However, whereas multicriteria approaches rely on one classic reward function per criterion, our anxiety-criterion is instead computed based on value-distributions for reachable states, which can be computed efficiently.

Definition 9 *A **Practical AA-MDP** is defined as $\langle S, A, T, R, R_\Delta, W, s_0, h\rangle$, where $\langle S, A, T, R\rangle$ is a classic MDP, $R_\Delta : \mathbf{\Delta} \to \mathbb{R}$ is a reward function that associates a cost to an instant anxiety, W is the relative weight given to anxiety versus objective rewards, $s_0 \in S$ is an initial state, and $h \in \mathbb{N}$ is the horizon.*

*An **Anxiety Value Function** $AV^\pi_h(s)$ represents the cumulative anxiety costs expected to be experienced following π for h steps from s, defined as $AV^h_\pi : S \to \mathbb{R}$ where: $AV^0_\pi(s) = 0$ and $AV^h_\pi(s) = R_\Delta(\Delta_{s,h,\pi}) + \sum_{s' \in S} T(s,a,s') \times AV^{h-1}_\pi(s')$ if $h > 0$.*

*A **Pragmatic Anxiety-Aware Value Function** AAV^h_π combines objective rewards to be acquired and anxiety costs to be experienced when following π for h steps, represented as $AAV^\pi_h : S \to \mathbb{R}$ where $AAV^\pi_h(s) = (1-W)\,V^\pi_h(s) + W \times AV^\pi_h(s)$.*

Practical AA-MDPs allow for a polynomial resolution of AA-MDP problems. Algorithm 1 expands the Value-Iteration algorithm for computing approximate policies π^*_{AA} that maximizes AAV. Moreover, k-bounded distributions help cutting short computations, as $\mathbb{Q}$ can easily grow and include many near-zero-probability values as h increases.

Definition 10 *A **k-bounded distribution** Δ^k **of** Δ is a k-value histogram based on Δ.*

Property 2 *The complexity of Algorithm 1 is $O(h \times |S|^3 \times (k^2 + E))$, if E is the evaluation cost of R_Δ.*
Proof: *This complexity is the complexity of classic Value-Iteration algorithm including the additional cost $O(k^2)$ for backpropagating and fusing value-distributions and $O(|E|)$ for computing R_Δ.*

R_Δ models linearise expected instant anxiety given a value distribution. *Standard deviation*, introduced by (Bach and Dolan 2012, p. 574), is a straightforward candidate for a rational assessment of value distribution, defined as: $R^\sigma_{\mathbf{\Delta}} = -\sqrt{\sum_{r \in \mathbb{Q}} \Delta(r) \times (V(\Delta) - r)^2}$, where $V(\Delta)$ is the expected value of Δ. $R^\sigma_{\mathbf{\Delta}}$. Alternative models for representing cognitive biases are introduced in annexes. The Shannon entropy, while seemingly a relevant candidate, fails to capture the intensity variable (e.g. $\{(0, 0.5); (10, 0.5)\}$ has the same entropy, but should raise higher anxiety than $\{(0, 0.5); (1, 0.5)\}$) (Bach and Dolan 2012).

Note that this model implies being oblivious to the temporal dynamics of anxiety experience: experiencing ten units of anxiety at once has the same valence than experiencing ten times one unit of anxiety. Alternative ways for modelling R_Δ and $AV^\pi_h(s)$ can be considered. For example, adding a power greater than 1 to the anxiety cost increases importance of avoiding relatively higher degrees of anxiety, e.g. $AV^h_\pi(s) = R_\Delta(\Delta_{s,h,\pi})^{\mathbf{2}} + \sum_{s' \in S} T(s,a,s') \times AV^{h-1}_\pi(s')$. cumulation can be considered, such as associating

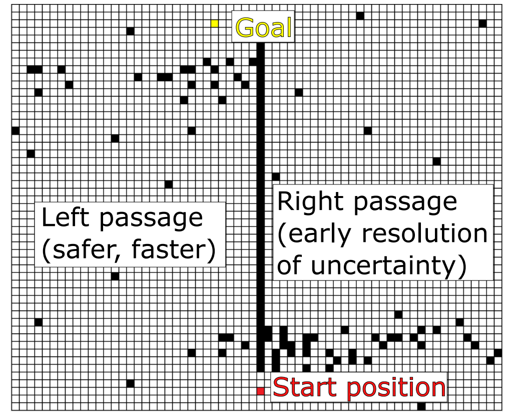

Figure 2: Experiment setup, black tiles are obstacles

Definition 11 *A model of $\prec_D$ or $R_{\mathbf{\Delta}}$ is **sound** if, for any deterministic T, $\Pi^* = \{\pi^*\}$.*

Property 3 *Standard deviation; pessimistic standard deviation; and ϵ-extreme anxiety are sound.*
Proof sketch: *If T is deterministic, any backpropagated Δ are singletons. Therefore, by their definitions, $R_\Delta(\Delta) = 0$ and thus $AAV = V$.*

Experimental Results

This section is dedicated to validating that our models generate policies that blend anxiety-avoidance with goal-orientation. As an approach, we study in simulations whether the policies generated by AA-MDPs are aligned with the characteristics of anxiety-avoiding behaviors in humans, and whether this alignment is preserved as the relative valence between anxiety and goal-orientation is altered.

Setup For validating that AA-MDPs replicate human-like anxiety-induced behaviors, we adapted the experimental setup from (Chow et al. 2015) with an obstacle map that highlights anxiety-sensitive choices. This setup, illustrated in Figure 2, consists of a navigation problem on a 2D-map with obstacles to a goal, given a robotic vehicle with imperfect move abilities. This map consists of a 64×53 grid with obstacles (black tiles in the figures), every tile being mapped to a state (i.e. 3312 states). Every round, the robot selects a direction (North, South, East, West). Every action has a 0.95 probability of reaching the intended tile and 0.05 probability of moving to any of the 8 tiles that neighbours the intended tile. The horizon is 200 steps. Each action is rewarded -1 unless the goal is reached (battery usage) and -40 when hitting obstacles. This model seeks for *safe* (obstacle-free) *battery efficient* policies. Policies are computed using Practical Anxiety-Avoidance MDPs. The weight W is the prime experimental variable as it determines the relative prevalence of anxiety on decisions. By increasing W, we should observe behaviors that are increasingly sensitive to anxiety that replicate human-like coping strategies. $k = 20$. All experiments were run on a classic desktop computer (i7 CPU, 16GB of RAM, Windows 10, Java 16).

The map layout (i.e. obstacles) is organized for allowing the emergence of a variety of general strategies: the map is split in two, with a line of obstacle in the centre, requiring to decide early whether to reach the goal by going through the left or the right passage. The *left passage* is slightly shorter and involves crossing a low-risk field of obstacle late in the process. The *right passage* involves crossing a medium-risk obstacle field early. For both passages, the agent can move away from the centre using longer, but the safer routes.

Results This section describes the main types of strategies generated by the system that we observed as we altered the sensitivity of the system to anxiety (W) and relates these strategies to how anxiety sensitivity influences human behavior. These four types were selected based on the overarching strategies taken by the agent that we observed. No more that the four phases described here were observed.

Low Sensitivity: Rational Behavior. Individuals who are insensitive to anxiety, as a personal trait or contextual contingency (e.g. high self-confidence, routine situation) are more likely to display goal-oriented behavior. This property can be verified in Figure 3 with $W = 0$. The agent behaves as a an anxiety-free reward-maximizing agent.

Mild Sensitivity: Impatience. Individuals with mild sensitivity to anxiety tend to prefer an early resolution of uncertainty, even at the expense of some objective utility or taking short-term risks. This property can be verified with $W = 0.1$. Figure 3 highlights that the agent crosses through the right passage, which is slightly longer and more risky, but reduces the duration during which the agent is exposed to the uncertainty of crossing the obstacle field.

Moderate Sensitivity: Hypervigilance and Dysfunctional Performance. At higher sensitivity levels, individuals tend to take irrational precautions for minimizing risks, even at the expanse of highly ineffective behaviours. This property is verified with $W = 0.5$, as displayed in Figure 3. In this figure, the trajectory ensures a zero-risk of hitting any obstacles, despite leading to a path that is irrationally long.

Severe Sensitivity: Task Avoidance and Self-Harm. The highest degrees of anxiety tend to trigger (pathological) behaviors, including task avoidance (i.e. refusing to attempt activities for avoiding subsequent uncertainty even when outcomes are exclusively positive) and self-harm (i.e. actively seeking an unnecessarily negative outcome for the sake of suppressing uncertainty). This property is observed with $W = 1$, Figure 3. In this case, the agent demonstrates task-avoidance in the pure-blue tiles by refusing to engage in the task, as moving towards the goal creates some uncertainty on whether the goal will be reached and how long it will take (Maner and Schmidt 2006). The area marked with a red circle indicates self-destructive behaviors, as the agent prefers to reach the goal by going through the wall once rather than taking the risk of hitting an obstacle while crossing the obstacle field in the future (Taylor et al. 2011).

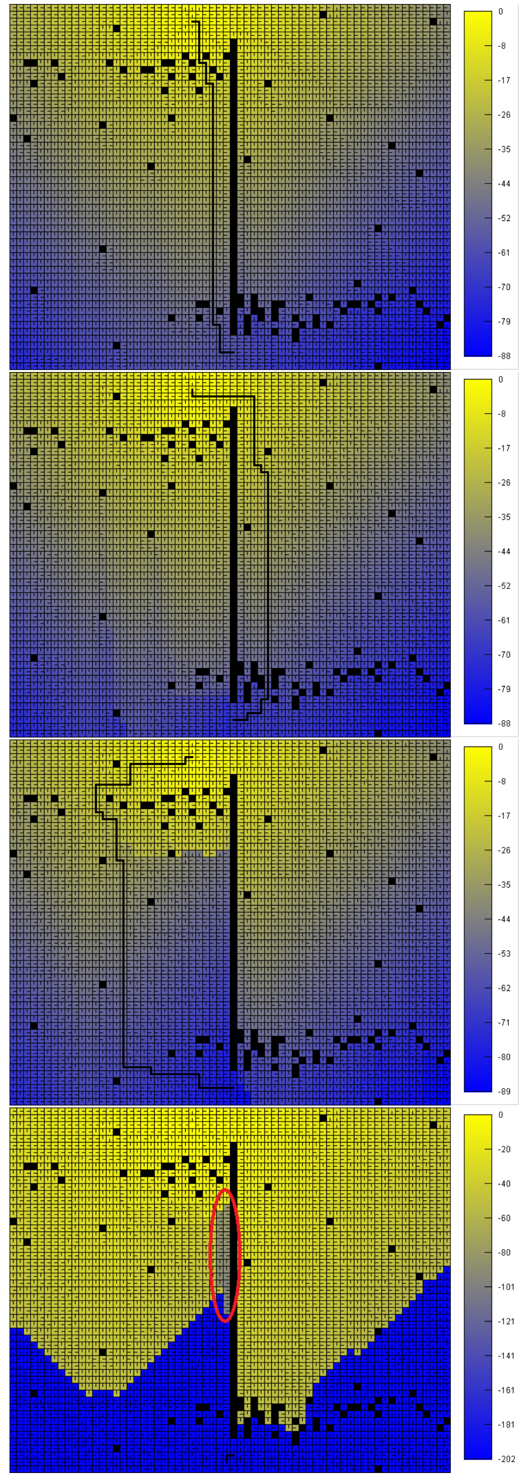

Figure 3: Expected value depending on W, from top to bottom: $W = 0$, $W = 0.1$, $W = 0.5$, and $W = 1$

Application: Anxiety-Aware Pathfinding

This section is dedicated to demonstrating that 1) AA-MDPs allow augmenting any existing MDP-based solution with anxiety-avoidance with minimal redesign effort; and 2) the

City	$\|S\|$	Dist.	V_{π^*}	$V_{\pi^{AA}}$	AV_{π^*}	$AV_{\pi^{AA}}$
Boston	1056376	35664	6546	7659	1048	10
Berlin	442523	23072	12158	12938	741	6
Johannesb.	290153	43933	7160	10019	2043	139
Kiev	120348	17645	5424	6611	913	64
London	445513	11893	3320	4162	731	0
Madrid	436073	17417	4890	6191	1099	15
Paris	697704	26299	3795	4484	790	3
Rome	155866	13653	3466	4068	769	6
Seoul	507341	69662	11092	14104	3327	11
Stockholm	146626	13468	5081	5492	565	12
Yokohama	656239	14161	4059	5237	1455	118

Table 2: Number of states, map width (in meters), expected value V and anxiety AV of anxiety-blind π^* and anxiety-aware π^{AA} policies, depending on the city (V and AV are negative values), $R_\triangle = R^\sigma_{\boldsymbol{\Delta}}$

suitability of AA-MDPs for lowering human anxiety when compared with pute reward maximization.

Itinerary-planning was selected as an example of how to include anxiety-sensitivity in an ubiquitous AI-based model. Whereas a car navigation map has been used for the convenience of display, the application directly expands to public transportation in general (Carrel and Walker 2017; Cheng 2010; Gobind 2018). The MDP model consists of a classic navigation in a network: states S are locations in a city, actions A are moves towards other locations, transitions T represent the location reached when deciding for a move moving, and rewards R represents the travel time for completing the desired move. The value V of a policy thus represents the total expected travel time. Delays (e.g. traffic jams, delayed transports) are a prevalent source of uncertainty when planning itineraries as they cause variability on the duration of move actions (formally, S is augmented with a marker indicating what delay occurred when playing the last action, thus increasing the cost and T has a small probability of reaching a delayed state). For simplicity and readability, a small probability of a long delay was associated to moving through highways –which can of course be replaced by specific traffic data if available. For demonstrating the realism of the approach, the navigation graphs used for the experiments were used based on OpenStreetMaps (OSM) data of capital cities around the world. Formally, OSM places labelled as small streets are associated with low speed and high fluidity (crossed at a speed of $10m/s$ with a probability of 1); and OSM places labelled as highways are associated with faster speed but possible congestion (crossed $30m/s$ with a probability of 0.8 and a speed of $3m/s$; with a probability of 0.2). Simple graphical user interfaces can be set up for every user to tailor the plan generation to their own needs, such as a slider for setting W (i.e. tradeoff between performance and anxiety). W was set to 0.1 as to demonstrate how much basic care for anxiety can lead to dramatic anxiety reduction for moderate performance costs.

The experiment shown in Table 2 compares anxiety-blind (π^*) and anxiety-sensitive (π^{AA}) policies for 100 trials for each city (random start and end destinations). As a remarkable result, anxiety-avoiding policies 97.1% of the average AV cost at the expanse of 17.2% increase in travel time. Figure 4 depicts navigation trees to the destination, marked by a star. Anxiety-aware policies tend to avoid highways and prefer smaller, more certain streets, unless using highways offer a significant gain.

An analytic cross-comparison (Greenberg and Buxton 2008) between AA-MDP policies and human anxiety-avoiding strategies as depicted by psychology studies has been put in place for validating the suitability of AA-MDPs for reducing the anxiety of human users[4]. Commuting anxiety, notably caused by the irregularity of transportation, is a well-recognized phenomenon and individuals are willing to compromise a longer average travel time for more certain travel duration (Carrel and Walker 2017; Cheng 2010; Gobind 2018). Accounting for and alleviating transport anxiety as done by anxiety-sensitive itinerary planner has thus the prospect of reaching out and reducing the exposition of anxiety to billions of daily commuters. Of course, a (large-scale) deployment of such a system is likely to entail side-effects, for which in-depth social analysis is required. Such analysis opens for future work for studying how to model the anxiety a system exposes third parties and societies to (e.g. simulation models featuring anxiety).

Related Work

To our knowledge, this paper pioneers in anxiety-aware planning, despite extensive former research being dedicated to aspects tightly related to anxiety: trust and deception, explainability, predictability, privacy, transparency, robustness/risk sensitivity, recognition (Bäuerle and Rieder 2013; Castelfranchi and Tan 2001; Chakraborti et al. 2019; Cowie et al. 2001; Franzoni, Milani, and Vallverdú 2017; Hou, Yeoh, and Varakantham 2014; Mueller et al. 2018). To a certain extent, anxiety can be seen as a unifying factor for the aforementioned aspects and may play a role of a "cognitive currency" as to help the compromise between the various aspects. However, despite some commonalities, each of these aspects involves fundamentally different conceptual models and deliberation frameworks. For example, trust and risk can be tied, but trust-building and risk-avoidance can involve in practice largely differing perspectives and courses of action.

Anxiety-avoidance partly relates to risk-avoidance in that

[4]Human Computer Interface research provides an array of validation methods for such an artefact. Human experiments are one instance which we left out of this article due to its well-recognized limitations (Greenberg and Buxton 2008; Lieberman 2003): rather than detailed a single, local-scale experimentation, thus case-specific and sensitive to statistical variability, our validation approach allows crossing AA-MDP policies to more extensive psychology findings that underwent significant empirical validation. From a methodological standpoint, tools such as the Visual Analogue Scale for Anxiety (VAS-A) (Abend et al. 2014; Facco et al. 2011), the Spielberger's State-Trait Anxiety Inventory (STAI) (Spielberger et al. 1999), and the Multiple Affect Adjective Check-List Revised (MAACL-R) (Zuckerman, Lubin, and Rinck 1985) can be used for assessing whether a certain system impacts the experienced anxiety.

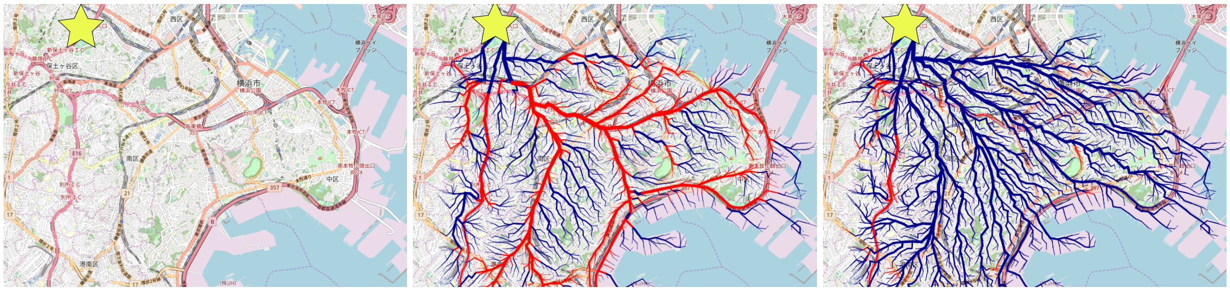

(a) Raw map (b) Anxiety-blind navigation tree (c) Anxiety-aware navigation tree

Figure 4: Navigation tree depicting all routes leading to a destination (star). In Figure 4a, red and orange routes are highways. Figure 4b shows trajectories generated by an anxiety-blind policy. Figure 4c shows the trajectories generated by an anxiety-aware policy. Red lines are highways and blue lines are regular streets. Thicker lines mean that more paths (i.e. from more starting points) reach the destination using that route.

both are sensitive to worse outcomes. However, anxiety fundamentally differs from risk in that and anxiety are fundamentally different in that 1) anxiety-avoidance concerns uncertainty in general, which covers both positive and negative events, whereas risk-avoidance only considers worst-cases; and 2) anxiety-avoidance is sensitive to the temporal component, i.e. how to minimize the degree of uncertainty and time during which this uncertainty is met, whereas risk-avoidance only considers the outcome –regardless of how long is the exposition to this uncertainty. For example, engaging early to slightly greater risks (e.g. engaging in conflicts, overspeeded) is a recognized anxiety-avoiding strategy (Hofstede, Hofstede, and Minkov 2010, p.197-198) that counteracts with risk-avoidance. Former models of risk-avoidance, include Conditional Value-at-Risk (CVaR) and bounded risks of reaching pit-states (Bisi et al. 2020; Chow et al. 2015; Hou, Yeoh, and Varakantham 2014; Tian, Sun, and Tomizuka 2021). Some recent models alleviate the focus on purely negative outcomes by seeking to reduce variance, which is also a limited form of anxiety (Whiteson, Zhang, and Liu 2021). Related emotion-motivated models include regret i.e. minimizing worse-case outcome given partly unknown reward and transition functions (Rigter, Lacerda, and Hawes 2021), which differs from anxiety in that it only focuses on negative outcomes and may engage in high-risk policies; fear and hope (avoiding worse and seeking best-case outcomes) (Moerland, Broekens, and Jonker 2016), which are confined forms of anxiety[5]. Last, stability-seeking models focus on the stability of sequences of rewards over time (e.g. preferring “0,0” over “+10,-10”) (White 1988). Both anxiety and stability are often tied to greater control over rewards, which lower uncertainty, but cover different phenomena, in that perfectly-controlled high variations may raise little anxiety (e.g. double pay every two months) and little variation can cause cumulative anxiety (e.g. being exposed to a single high-risk situation in a far future).

[5]Models of fear and hope detailed in the annexes as pessimistic and optimistic Practical AA-MDP reward functions

Conclusion

This paper introduces and formalizes the first anxiety-aware planning formalism, which we call Anxiety-Aware Markov Decision Processes (AA-MDPs), built on fundamental psychology research (Barlow 2002a; Miceli and Castelfranchi 2005; Rapee et al. 1996). AA-MDPs allow generating policies that balance value-maximization and expected anxiety, where this anxiety is modelled as the distribution of sequences of instant anxiety over time and instant anxiety is modelled as the distribution of values expected to be obtained by following a given policy π. This article proves that finding optimal AA-MDP policies is NP-complete, while also introducing an approximate polynomial algorithm for computing realistic anxiety-avoiding policies. The soundness of AA-MDPs has been demonstrated by showing that it reproduces central properties of anxiety-coping behaviors as described in fundamental psychology research. The direct applicability and relevance of AA-MDPs has been shown by demonstrating that any existing MDP-based solution can be expanded with anxiety-sensitivity and that the policies generated by AA-MDP expansions align with lowered human anxiety on, at least, pathplanning applications, which may reach billions of users.

The array of possible applications is extremely broad. Consider the example of managing the COVID-19 given in the introduction, for which multiple advanced MDP-based models have been proposed (Capobianco et al. 2021; Just and Echaust 2020). AA-MDPs would allow providing very concrete answers on questions such as: “how do policies for lockdowns impact the experienced anxiety due to infections over future infections?”. Lockdowns reduce immediate anxiety by reducing immediate risks but increase exposure on the long run. Intermittent lockdown policies in reaction to semi-random factors actually cause even further anxiety. Such anxiety-sensitive AI tools would have been very handy for anticipating and even mitigating the anxiety-induced ailments caused by the pandemic, which are now endemically observed. AA-MDPs allow to perform such assessment for minimal cost.

This paper is intended as a foundational article dedicated to opening new branches of innovative high-impact research and future applications, by demonstrating that anxiety sensitivity and avoidance can be modelled and achieved in practice. First, Practical AA-MDPs, by directly overlaying MDPs, enable to turn any already-built MDP-based system into an anxiety-sensitive system for minimal cost, as illustrated by our itinerary-planning application. AA-MDPs also directly apply for developing anxiety-sensitive agents, such as virtual characters for serious games and social simulations. Second, AA-MDPs lay the structures for including additional forms of anxiety, beyond outcome-based, covering nearly all forms of uncertainty described by (Bach and Dolan 2012): uncertain state and partial observability (state anxiety) can be covered by expanding AA-MDPs to Partially Observable MDPs models (Bernstein et al. 2002) and partial knowledge of the environment (transition uncertainty (Bach and Dolan 2012)) can be covered by expanding AA-MDPs with Reinforcement Learning models (Sutton and Barto 1998). Third, the conceptual foundations provided by this paper open for proactive user management through representations of user mental models (Vanhée, Jeanpierre, and Mouaddib 2021), including self-escalating anxiety loops (e.g. a robot slowing down when being watched, a bus system guaranteeing connections when buses are infrequent) and long-term repercussions of anxiety (e.g. recommendations on strategies when taking a loan). Such models are directly suited for generating tailored optimized therapeutic strategies: e.g. how to best cut maladaptive anxiety loops, given the dynamics of the specific traumatic experience, anxiety-induced disability, patient engagement, and therapeutic resources (Wells 2011).

Anxiety is a two-sided blade for human intelligence (Barlow 2002b): anxiety can both cause crippling avoidance and risky behavior as it fuels the inner drive for proactively solving problems, knowing more and becoming better. Anxiety-sensitive AI has not only a potential for reducing tremendous health cost but also for supporting decisions that help society and individuals to flourish out of the shackles of uncertainty.

Acknowledgements

The first author, Loïs Vanhée, acknowledges the support of the Knut and Alice Wallenberg Foundation (project number 570080103), and of the project "GEDAI: Growing Ethical Designers of Artificial Intelligence" at Umeå University, Sweden. The first author also wishes to thank Melania Borit for her support proofreading the text and structuring the arguments, Christophe Chesneau for his support identifying the presence of mixture distributions in our model, and Arvid Horned for his support reading and commenting over the text.

Ethical Impact

AA-MDPs allow generating plans that balance effectiveness and exposition to outcome uncertainty for plan users. While the stage of research is too early to anticipate all consequences, AA-MDPs reasonably tie to humans as a system for generating 1) anxiety-sensitive plans to human users; 2) plans for an artefact that avoids anxiety "for itself", notably suiting 2.a) mimicking human-like intelligence, e.g. virtual characters for the game, social simulations; and 2.b) systems being observed by a user who can experience anxiety by proxy (e.g. if user's utility is tied to the artefact performance, like for a trading agent). For all these applications, the user can be reasonably expected to retain the freedom not use the system, or else specific ethical inquiry is required.

The HLEG guidelines (on AI 2019) bring forward four core ethical principles to be strived towards: respect for human autonomy, prevention of harm, fairness and explicability and prescribes designers to find compromises that best cover these four principles. AA-MDPs have a very clear positive valence regarding human autonomy and prevention of harm as they reduce the risk of anxiety-induced ailments that damage both health and autonomy. AA-MDPs have a neutral to positive valence regarding fairness in that anxiety supports subgroups most exposed to anxiety and uncertainty and groups that are usually discriminated against (gender, age, disability). However, AA-MDPs fundamentally offering alternative tradeoffs, require a scrutiny of the fairness provided by the new tradeoff: one's preference for less anxiety can cause someone else's loss (e.g. more accurate buses schedules can incur greater costs). AA-MDPs have a positive valence regarding explainability as they are built upon the same explicit models as used in planning expanded with explicit markers of anxiety measurement, allowing for straightforward high-level explanations.

Besides a generally positive valence, prospective risk areas are to be considered as to anticipate possible harm. First, a major reduction to the exposure to anxiety may lead to unexpected psychological impact in the long run (e.g. reduced general ability to cope with anxiety), thus calling for dedicated psychology and sociology studies in case of a large-scale prolonged uptake of the technology. Second, intentional misuses of AA-MDPs may lead to lead to causing greater anxiety as they can be used for building social simulations that seek to optimize anxiety-inducing activities (e.g. terrorism), though the benefit with regard to existing simulation tools remains open to debate. An accidental misuse of AA-MDPs is likely to be detected early as it would lead to the generation of clearly suboptimal plans.

References

Abend, R.; Dan, O.; Maoz, K.; Raz, S.; and Bar-Haim, Y. 2014. Reliability, validity and sensitivity of a computerized visual analog scale measuring state anxiety. *Journal of Behavior Therapy and Experimental Psychiatry*, 45(4): 447–453.

Alonso, J.; Liu, Z.; Evans-Lacko, S.; Sadikova, E.; Sampson, N.; and Others. 2018. Treatment gap for anxiety disorders is global: Results of the World Mental Health Surveys in 21 countries. *Depression and anxiety*, 35(3): 195–208.

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

Andrews, G.; Anderson, T. M.; Slade, T.; and Sunderland, M. 2008. Classification of anxiety and depressive disorders: problems and solutions. *Depression and anxiety*, 25(4): 274–281.

Avery, S. N.; Clauss, J. A.; and Blackford, J. U. 2016. The human BNST: functional role in anxiety and addiction. *Neuropsychopharmacology*, 41(1): 126–141.

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.

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

Barlow, D. H. 2002a. Fear, anxiety, and theories of emotion. *Anxiety and its disorders: The nature and treatment of anxiety and panic*, 37–63.

Barlow, D. H. 2002b. The experience of anxiety: Shadow of intelligence or specter of death. *Anxiety and its Disorders: The Nature and Treatment of Anxiety and Panic, 2nd edn. Guilford Press, New York*, 1–36.

Bäuerle, N.; and Rieder, U. 2013. More risk-sensitive Markov decision processes. *Mathematics of Operations Research*, 39(1): 105–120.

Beck, A. T.; and Clark, D. A. 1997. An information processing model of anxiety: Automatic and strategic processes. *Behaviour research and therapy*, 35(1): 49–58.

Belik, S.-L.; Sareen, J.; and Stein, M. B. 2009. Anxiety disorders and physical comorbidity. In *Oxford handbook of anxiety and related disorders*.

Bellman, R. 1956. Dynamic Programming and the Smoothing Problem. *Management Science*, 3(1): 111–113.

Bernstein, D.; Givan, R.; Immerman, N.; and Zilberstein, S. 2002. The complexity of decentralized control of Markov decision processes. *Mathematics of Operations Research*, 27(4): 819–840.

Bisi, L.; Sabbioni, L.; Vittori, E.; Papini, M.; and Restelli, M. 2020. Risk-averse trust region optimization for reward-volatility reduction. *Proceedings of the International Joint Conference on Artificial Intelligence*.

Brandes, M.; and Bienvenu, O. J. 2006. Personality and anxiety disorders. *Current psychiatry reports*, 8(4): 263–269.

Briscese, G.; Lacetera, N.; Macis, M.; and Tonin, M. 2020. Compliance with covid-19 social-distancing measures in italy: the role of expectations and duration. Technical report, National Bureau of Economic Research.

Capobianco, R.; Kompella, V.; Ault, J.; Sharon, G.; Jong, S.; Fox, S.; Meyers, L.; Wurman, P. R.; and Stone, P. 2021. Agent-based markov modeling for improved COVID-19 mitigation policies. *Journal of Artificial Intelligence Research*, 71: 953–992.

Carleton, R. N. 2016. Into the unknown: A review and synthesis of contemporary models involving uncertainty. *Journal of anxiety disorders*, 39: 30–43.

Carrel, A.; and Walker, J. L. 2017. Understanding future mode choice intentions of transit riders as a function of past experiences with travel quality. *European Journal of Transport and Infrastructure Research*, 17(3).

Castelfranchi, C.; and Tan, Y.-H. 2001. *Trust and deception in virtual societies*. Springer.

Chakraborti, T.; Kulkarni, A.; Sreedharan, S.; Smith, D. E.; and Kambhampati, S. 2019. Explicability? legibility? predictability? transparency? privacy? security? The emerging landscape of interpretable agent behavior. In *International conference on automated planning and scheduling*.

Cheng, Y.-H. 2010. Exploring passenger anxiety associated with train travel. *Transportation*, 37(6): 875–896.

Chow, Y.; Tamar, A.; Mannor, S.; and Pavone, M. 2015. Risk-sensitive and robust decision-making: a cvar optimization approach. In *Advances in Neural Information Processing Systems*, 1522–1530.

Cowie, R.; Douglas-Cowie, E.; Tsapatsoulis, N.; Votsis, G.; Kollias, S.; Fellenz, W.; and Taylor, J. G. 2001. Emotion recognition in human-computer interaction. *IEEE Signal processing magazine*, 18(1): 32–80.

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.

Ehlers, A.; and Clark, D. M. 2000. A cognitive model of posttraumatic stress disorder. *Behaviour research and therapy*, 38(4): 319–345.

Everitt, B. S.; and Hand, D. J. 1981. Mixtures of discrete distributions. In *Finite Mixture Distributions*, 89–105. Dordrecht: Springer Netherlands.

Facco, E.; Zanette, G.; Favero, L.; Bacci, C.; Sivolella, S.; Cavallin, F.; and Manani, G. 2011. Toward the validation of visual analogue scale for anxiety. *Anesthesia progress*, 58(1): 8–13.

Firth, J.; Siddiqi, N.; Koyanagi, A.; Siskind, D.; Rosenbaum, S.; Galletly, C.; Allan, S.; Caneo, C.; Carney, R.; Carvalho, A. F.; and Others. 2019. The Lancet Psychiatry Commission: a blueprint for protecting physical health in people with mental illness. *The Lancet Psychiatry*, 6(8): 675–712.

Franzoni, V.; Milani, A.; and Vallverdú, J. 2017. Emotional affordances in human-machine interactive planning and negotiation. In *Proceedings of the International Conference on Web Intelligence*, 924–930. ACM.

Gobind, J. 2018. Transport anxiety and work performance. *SA Journal of Human Resource Management*, 16(1): 1–7.

Greenberg, S.; and Buxton, B. 2008. Usability evaluation considered harmful (some of the time). In *Proceedings of the SIGCHI conference on Human factors in computing systems*, 111–120.

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.

Hartley, S.; Barrowclough, C.; and Haddock, G. 2013. Anxiety and depression in psychosis: a systematic review of associations with positive psychotic symptoms. *Acta Psychiatrica Scandinavica*, 128(5): 327–346.

Hofstede, G.; Hofstede, G.-J.; and Minkov, M. 2010. *Cultures and Organizations: Software of the Mind, Third Edition*. McGraw-Hill Professional.

Hou, P.; Yeoh, W.; and Varakantham, P. 2014. Revisiting risk-sensitive MDPs: New algorithms and results. In *Conference on Automated Planning and Scheduling*.

Huppert, J. D. 2009. Anxiety disorders and depression comorbidity. *Anxiety*, 35(27): 65.

IEEE. 2017. Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems. Technical report, IEEE.

Just, M.; and Echaust, K. 2020. Stock market returns, volatility, correlation and liquidity during the COVID-19 crisis: Evidence from the Markov switching approach. *Finance Research Letters*, 37: 101775.

Kahneman, D. 2011. *Thinking, fast and slow*, volume 1. Farrar, Straus and Giroux New York.

Kulynych, B.; Overdorf, R.; Troncoso, C.; and Gürses, S. 2020. POTs: protective optimization technologies. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 177–188.

Levy, I.; and Schiller, D. 2020. Neural computations of threat. *Trends in cognitive sciences*.

Lieberman, H. 2003. The tyranny of evaluation. *CHI Fringe*.

Loizou, M.; Hartley, T.; Slater, S.; Newman, R.; and Pannese, L. 2012. Emotions for intelligent agents in crisis management simulations: A survey. In *International Conference on Computer Games*, 213–219. IEEE.

Maner, J. K.; and Schmidt, N. B. 2006. The role of risk avoidance in anxiety. *Behavior Therapy*, 37(2): 181–189.

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

Moerland, T. M.; Broekens, J.; and Jonker, C. M. 2016. Fear and Hope Emerge from Anticipation in Model-Based Reinforcement Learning. In *International Joint Conference on Artificial Intelligence*, 848–854.

Mueller, S. T.; Hoffman, R. R.; Clancey, W.; Emrey, A.; and Klein, G. 2018. Explanation in human-AI systems: A literature meta-review, synopsis of key ideas and publications, and bibliography for explainable AI. *Institute for Human and Machine Cognition*.

Neumann, I. D.; Veenema, A. H.; and Beiderbeck, D. I. 2010. Aggression and anxiety: social context and neurobiological links. *Frontiers in behavioral neuroscience*, 4: 12.

on AI, H.-L. E. G. 2019. Ethics guidelines for trustworthy AI. Technical report.

Patel, V.; Saxena, S.; Lund, C.; Thornicroft, G.; Baingana, F.; Bolton, P.; Chisholm, D.; Collins, P. Y.; Cooper, J. L.; Eaton, J.; and Others. 2018. The Lancet Commission on global mental health and sustainable development. *The Lancet*, 392(10157): 1553–1598.

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.

Restauri, N.; and Sheridan, A. D. 2020. Burnout and post-traumatic stress disorder in the COVID-19 pandemic: intersection, impact, and interventions. *Journal of the American College of Radiology*, 17(7): 921–926.

Rigter, M.; Lacerda, B.; and Hawes, N. 2021. Minimax regret optimisation for robust planning in uncertain markov decision processes. *AAAI Conference on Artificial Intelligence*.

Roijers, D.; Vamplew, P.; and Whiteson, S. 2013. A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research*, 48: 67–113.

Salari, N.; Hosseinian-Far, A.; Jalali, R.; Vaisi-Raygani, A.; and Others. 2020. Prevalence of stress, anxiety, depression among the general population during the COVID-19 pandemic: a systematic review and meta-analysis. *Globalization and health*, 16(1): 1–11.

Schuller, D.; and Schuller, B. W. 2018. The age of artificial emotional intelligence. *Computer*, 51(9): 38–46.

Sigaud, O.; and Buffet, O. 2010. *Markov decision processes in artificial intelligence*. John Wiley & Sons.

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 Anger Expression Inventory (STAXI).* Lawrence Erlbaum Associates.

Sutton, R. S.; and Barto, A. G. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.

Taylor, D. J.; Lichstein, K. L.; Durrence, H. H.; Reidel, B. W.; and Bush, A. J. 2005. Epidemiology of insomnia, depression, and anxiety. *Sleep*, 28(11): 1457–1464.

Taylor, P. J.; Gooding, P.; Wood, A. M.; and Tarrier, N. 2011. The role of defeat and entrapment in depression, anxiety, and suicide. *Psychological bulletin*, 137(3): 391.

Tian, R.; Sun, L.; and Tomizuka, M. 2021. Bounded Risk-Sensitive Markov Games: Forward Policy Design and Inverse Reward Learning with Iterative Reasoning and Cumulative Prospect Theory. *AAAI Conference on Artificial Intelligence*, 35(7): 6011–6020.

Vanhée, L.; Jeanpierre, L.; and Mouaddib, A.-I. 2021. Optimizing Requests for Support in Context-Restricted Autonomy. In *International Conference on Intelligent Robots and Systems*.

Wells, A. 2011. *Metacognitive therapy for anxiety and depression*. Guilford press.

White, D. J. 1988. Mean, variance, and probabilistic criteria in finite Markov decision processes: a review. *Journal of Optimization Theory and Applications*, 56(1): 1–29.

Whiteson, S.; Zhang, S.; and Liu, B. 2021. Mean-variance policy iteration for risk-averse reinforcement learning. In *AAAI Conference on Artificial Intelligence*.

Wittchen, H.-U. 2002. Generalized anxiety disorder: prevalence, burden, and cost to society. *Depression and anxiety*, 16(4): 162–171.

Zuckerman, M.; Lubin, B.; and Rinck, C. M. 1985. Multiple Affect Adjective Check List–Revised. *Journal of Behavioral Assessment*.