

Purpose Restrictions on Information Use^{*}

Michael Carl Tschantz^{1, **}, Anupam Datta², and Jeannette M. Wing^{3, **}

¹ University of California, Berkeley, mct@berkeley.edu

² Carnegie Mellon University, danupam@cmu.edu

³ Microsoft Research, wing@microsoft.com

Abstract. Privacy policies in sectors as diverse as Web services, finance and healthcare often place restrictions on the purposes for which a governed entity may use personal information. Thus, automated methods for enforcing privacy policies require a semantics of *purpose restrictions* to determine whether a governed agent *used information* for a purpose. We provide such a semantics using a formalism based on planning. We model planning using Partially Observable Markov Decision Processes (POMDPs), which supports an explicit model of information. We argue that information use is for a purpose if and only if the information is used while planning to optimize the satisfaction of that purpose under the POMDP model. We determine information use by simulating ignorance of the information prohibited by the purpose restriction, which we relate to noninterference. We use this semantics to develop a sound audit algorithm to automate the enforcement of purpose restrictions.

1 Introduction

Purpose is a key concept for privacy policies. Some policies limit the use of certain information to an explicit list of purposes. The privacy policy of The Bank of America states, “Employees are authorized to access Customer Information for business purposes only.” [1]. The HIPAA Privacy Rule requires that healthcare providers in the U.S. use protected health information about a patient with that patient’s authorization or only for a fixed list of allowed purposes, such as treatment and billing [2]. Other policies prohibit using certain information for a purpose. For example, Yahoo!’s privacy policy states “Yahoo!’s practice on Yahoo! Mail Classic is not to use the content of messages stored in your Yahoo! Mail account for marketing purposes.” [3].

Each of these examples presents a constraint on the purposes for which the organization may use information. We call these constraints *purpose restrictions*.

^{*} This research was supported by the U.S. Army Research Office grants DAAD19-02-1-0389 and W911NF-09-1-0273 to CyLab, by the National Science Foundation (NSF) grants CCF0424422 and CNS1064688, and by the U.S. Department of Health and Human Services grant HHS 90TR0003/01. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

^{**} The authors conducted most of this work while at Carnegie Mellon University.

Let us consider a purpose restriction in detail. As a simplification of the Yahoo! example, consider an advertising network attempting to determine which advertisement to show for marketing to a visitor of a website (such as an email website). To improve its public image and to satisfy government regulations, the network adopts a privacy policy containing a restriction prohibiting the use of the visitor's gender for the purpose of marketing.

The network has access to a database of information about potential visitors, which includes their gender. Since some advertisements are more effective, on average, for some demographics than others, using this information is in the network's interest. However, the purpose restriction prohibits the use of gender for selecting advertisements since it is a form of marketing. Since tension exists between selecting the most effective ad and obeying the purpose restriction, internal compliance officers and government regulators should audit the network to determine whether it has complied with the privacy policy.

However, the auditors may find manually auditing the network difficult and error prone leading them to desire automated tools to aid them. Indeed, the difficulty of manually auditing purpose restrictions has led to commercial software for this task (e.g., [4]). However, their approaches have been ad hoc.

Our goal is to place purpose restrictions governing information use on a formal footing and to automate their enforcement. In the above example, intuitively, the auditor must determine what information the network used while planning which ads to show to a user. In general, determining whether the purpose restriction was obeyed involves determining facts about how the audited agent (a person, organization, or computer system) planned its actions. In particular, philosophical inquiry [5] and an empirical study [6] show that the behavior of an audited agent is for a purpose when the agent chooses that behavior while planning to satisfy the purpose. Our prior work has used a formal model of planning to automate the auditing of purpose restrictions that limit visible actions to certain purposes [6].

We build upon that work to provide formal semantics and algorithms for purpose restrictions limiting *information uses*, whose occurrence the auditor cannot directly observe. For example, while the ad network is prohibited from using the visitor's gender, it may access the database to use other information even if the database returns the gender as part of a larger record. Thus, our model must elucidate whether the network *used* the gender component of the accessed information.

To provide auditing algorithms, we need a formal model of planning. Fortunately, research in artificial intelligence has provided a variety of formal models of planning. To select an appropriate model for auditing, we examine the key features of our motivating example of the ad network. First, it shows that purposes are not just goals to be achieved since the purpose of marketing is quantitative: marketing can be satisfied to varying degrees and more can always be done. Second, the example shows that outcomes can be probabilistic since the network does not know what ad will be best for each visitor but does have statistical

information about various demographics. Lastly, the policy is governing the use of information. Thus, our model needs an explicit model of information.

The first two features suggest using Markov Decision Processes (MDPs), which we have successfully used in an auditing algorithm for purpose restrictions on observable actions [6]. However, needing an explicit model of information requires us to use an extension of MDPs, Partially Observable Markov Decision Processes (POMDPs), which make the ability of the planning agent to observe its environment and collect information explicit. We use a POMDP to model the agent’s environment where the purpose in question defines the reward function of the POMDP. The explicitness of observations (inputs) in the POMDP model allows us to go beyond standard research on planning to provide a semantics of *information use* by considering how the agent would plan if some observations were conflated to ignore information of interest.

In more detail, we quotient the POMDP’s space of observations to express information use. Intuitively, to use information is to see a distinction, and to not use information corresponds to ignoring this distinction. Thus, we quotient by an equivalence relation that treats two observations as indistinguishable if they differ only by information whose use is prohibited by a purpose restriction. For example, the ad network promising not to use gender should quotient its observations by an equivalence relation that treats the genders as equivalent. By conflating observations that differ only by gender, the network will ignore gender, simulating ignorance of it. Such quotienting is defined for POMDPs since observations probabilistically constrain the space of possible current states of the agent’s environment, and quotienting just decreases the constraint’s accuracy.

We use our quotienting operation to provide two different definitions of what it means for an agent to obey a purpose restriction involving information use. The first requires that the agent uses the quotiented POMDP to select its behavior. We call this definition *cognitive* since it refers to the agent’s cognitive process of selecting behavior. Since the auditor cannot examine the agent’s cognitive processes and might only care about their external consequences, we offer a second weaker definition that depends upon the agent’s observable behavior. The *behaviorist* definition only requires that the agent’s behaviors be consistent with using the quotiented POMDP. It does not depend upon whether the agent actually used that POMDP or a different process to select its behavior.

We use the behaviorist definition as the basis of an auditing algorithm that compares the behaviors of an agent to each of the behaviors that is acceptable under our notion of simulated ignorance. Despite comparing to multiple behaviors, our algorithm only needs to optimize the quotiented POMDP once. For the behaviorist definition, we prove that the algorithm is sound (Theorem 1) and is complete when the POMDP can be optimized exactly (Theorem 2).

To show that our semantics is strong enough, we compare it to *noninterference*, a prior formalization of information use for automata [7]. This definition examines how an input to an automaton affects the automaton’s output. Our approach is similar but uses POMDPs instead of automata. We relate the two models by defining how an automaton can implement a strategy for a quotiented

POMDP, which allows us to prove that the cognitive definition implies a form of noninterference (Theorem 3). On the other hand, we show that an agent can obey the behaviorist definition while still exhibiting interference. However, interestingly, such interference cannot further the restricted purpose showing that the behaviorist definition is strong enough to prevent interference *for that purpose*.

Since an action’s purpose can depend upon how it fits into a chain of actions, we focus on post-hoc auditing. Nevertheless, other enforcement mechanisms can employ our semantics. Despite focusing on privacy policies, our semantics and algorithm may aid the enforcement of other policies restricting the use of information to only certain purposes, such as those governing intellectual property.

Contributions and Outline. We start by reviewing related work and POMDPs (Sections 2 and 3). Our first contribution is definitional: we use our quotienting characterization of information use to provide both the cognitive and behaviorist definitions of complying with a purpose restriction on information use (Section 4). Our second contribution is our auditing algorithm accompanied by theorems of soundness and a qualified form of completeness (Section 5). Our final contribution is relating our formalization to noninterference with a theorem showing that the cognitive definition implies noninterference (Sections 6). We end with conclusions (Sections 7). All proofs are in a related technical report [8].

2 Prior Work

Information Flow Analysis. Research on information flow analysis led to noninterference [7], a formalization of information flow, or use. However, prior methods of detecting noninterference have typically required access to the program running the system in question. These analyses either used the program for directly analyzing its code (see [9] for a survey), for running an instrumented version of the system (e.g., [10–13]), or for simulating multiple executions of the system (e.g., [14–16]). Traditionally, the requirement of access to the program has not been problematic since the analysis has been motivated as a tool for software engineers securing a program that they have designed.

However, in our setting of enforcing purpose restrictions, such access is not always possible since the analyzed system can be a person who could be adversarial and whose behavior the auditor can only observe. On the other hand, the auditor has information about the purposes that the system should be pursuing. Since the system is a purpose-driven agent, the auditor can understand its behavior in terms of a POMDP model of its environment. Thus, while prior work provides a definition of information use, it does not provide appropriate models or methods for determining whether it occurs in our setting.

Enforcing Purpose Restrictions. Most prior work on using formal methods for enforcing purpose restrictions has focused on when observable actions achieve a purpose [17–24]. That is, they define an action as being for a purpose if that action (possibly as part of a chain of actions) results in that purpose being achieved. Our work differs from these works in two ways.

First, we define an action as being for a purpose when that action is part of a plan for maximizing the satisfaction of that purpose. Our definition differs by treating purposes as rewards that can be satisfied to varying degrees and by focusing on the plans rather than outcomes, which allows an action to be for a purpose even if it probabilistically fails to improve it. The semantics of purpose we use follows from informal philosophical inquiry [5] and our prior work using Markov Decision Processes to formalize purpose restrictions for actions [6]. Jafari et al. offer an alternative view of planning and purposes in which a purpose is high-level action related to low-level actions by a plan [25]. Our views are complementary in that theirs picks up where ours leaves off: Our model of planning can justify the plans that their model accepts as given while their model allows for reasoning about the relationships among purposes with a logic.

Second, we consider information use. While the aforementioned works address restrictions on information *access*, they do not have a model of information *use*, such as noninterference [7]. In particular, we extend our prior work [6] to model how the agent uses information while selecting actions with a POMDP. We show that if the agent does not use the information under our model, then noninterference holds. While Martinelli et al. have used POMDPs for enforcing quantitative access controls, we differ by using POMDPs to model information use itself instead of for modeling policies governing information use treated as observable actions [26]. Hayati and Abadi provide a type system for tracking information flow in programs with purpose restrictions in mind [27]. However, their work presupposes that the programmer can determine the purpose of a function and provides no formal guidance for making this determination.

Minimal disclosure requires that the amount of information granted should be as little as possible while still achieving the purpose behind the grant. This differs from purpose restrictions, which do not require the amount of information used to be minimal and often involve purposes that are never fully achieved (e.g., more marketing is always possible). Thus, unlike works on minimal disclosure [28, 29], we model purposes as being satisfied to varying degrees. Furthermore, we model probabilistic failures of the agent’s plan, which allows us to identify when information use is for a purpose despite not increasing the purpose’s satisfaction.

Planning. Since our formal definition is in terms of planning, automating auditing depends upon automated *plan recognition* [30]. We build upon works that use models of planning to recognize plans (e.g., [31–34]). The most related work has provided methods of determining when a sequence of actions are for a purpose (or “goal” in their nomenclature) given a POMDP model of the environment [34]. Our algorithm for auditing is similar to their algorithm. However, whereas their algorithm attempts to determine the probability that a sequence of actions are for a purpose, we are concerned with whether a use of information could be for a purpose. Thus, we must first develop a formalism for information use. We must also concern ourselves with the soundness of our algorithm rather than its accuracy in terms of a predicted probability. Additionally, we use traditional POMDPs to model purposes that are never fully satisfied instead of the goal POMDPs used in their work.

3 Modeling Purpose-Driven Agents

We review the Partially Observable Markov Decision Process (POMDP) model and then show how to model the above motivating example as one. We start with an agent, such as a person, organization, or artificially intelligent computer, that attempts to maximize the satisfaction of a purpose. The agent uses a POMDP to plan its actions. The POMDP models the agent’s environment and how its actions affects the environment’s state and the satisfaction of the purpose. The agent selects a plan that optimizes the expected total discounted reward (degree of purpose satisfaction) under the POMDP. This plan corresponds to the program running the audited system.

POMDPs. To define POMDPs, let $\text{Dist}(X)$ denote the space of all distributions over the set X and let \mathbb{R} be the set of real numbers. A POMDP is a tuple $\langle \mathcal{Q}, \mathcal{A}, \tau, \rho, \mathcal{O}, \nu, \gamma \rangle$ where

- \mathcal{Q} is a finite state space representing the states of the agent’s environment;
- \mathcal{A} , a finite set of actions;
- $\tau : \mathcal{Q} \times \mathcal{A} \rightarrow \text{Dist}(\mathcal{Q})$, a transition function from a state and an action to a distribution over states representing the possible outcomes of the action;
- $\rho : \mathcal{Q} \times \mathcal{A} \rightarrow \mathbb{R}$, a reward function measuring the immediate impact on the satisfaction of the purpose when the agent takes the given action in the given state;
- \mathcal{O} , a finite observation space containing any observations the agent may perceive while performing actions;
- $\nu : \mathcal{A} \times \mathcal{Q} \rightarrow \text{Dist}(\mathcal{O})$, a distribution over observations given an action and the state resulting from performing that action; and
- γ , a discount factor such that $0 \leq \gamma < 1$.

We say that a POMDP *models a purpose* if ρ measures the degree to which the purpose is satisfied. To select actions for that purpose, the agent should select those that maximizes its expected total discounted reward, $\mathbb{E} [\sum_{i=0}^{\infty} \gamma^i u_i]$ where i represents time and u_i , the reward from the agent’s i th action.

This goal is complicated by the agent not knowing *a priori* which of the possible states of the POMDP is the current state of its environment. Rather it holds beliefs about which state is the current state. In particular, the agent assigns a probability to each state q according to how likely the agent believes that the current state is the state q . A *belief state* β captures these beliefs as a distribution over states of \mathcal{Q} (i.e., $\beta \in \text{Dist}(\mathcal{Q})$). An agent updates its belief state as it performs actions and makes observations. When an agent takes the action a and makes the observation o starting with the beliefs β , the agent develops the new beliefs β' where $\beta'(q')$ is the probability that q' is the next state.

We define $\text{up}_m(\beta, a, o)$ to equal the updated beliefs β' . β' assigns to the state q' the probability $\beta'(q') = \Pr[Q'=q'|O=o, A=a, B=\beta]$ where Q' is a random variable over next states, $B=\beta$ identifies the agent’s current belief state as β , $A=a$ identifies the agent’s current action as a , and $O=o$ identifies the observation

the agent makes while performing action a as o . We may reduce $\text{up}_m(\beta, a, o)$ to the following formula in terms of the POMDP model:

$$\text{up}_m(\beta, a, o)(q') = \frac{\nu(a, q')(o) \sum_{q \in \mathcal{Q}} \beta(q) * \tau(q, a)(q')}{\sum_{q' \in \mathcal{Q}} \nu(a, q')(o) \sum_{q \in \mathcal{Q}} \beta(q) * \tau(q, a)(q')}$$

To maximize its expected total discounted reward, the agent does not need to track its history of actions and observations independently of its beliefs as such beliefs are a sufficient statistic. Thus, the agent need only consider for each possible belief β it can have, what action it would perform. That is, the agent can plan by selecting a *strategy*: a function from the space of beliefs $\text{Dist}(\mathcal{Q})$ to the space of actions \mathcal{A} . (We use the word “strategy” instead of the more common “policy” to avoid confusion with privacy policies.)

The goal of the agent is find the optimal strategy. By the Bellman equation [35], the expected value of a belief state β under a strategy σ is

$$V_m(\sigma, \beta) = R_m(\beta, \sigma(\beta)) + \gamma \sum_{o \in \mathcal{O}} N_m(\beta, \sigma(\beta))(o) * V_m(\sigma, \text{up}_m(\beta, \sigma(\beta), o)) \quad (1)$$

where R_m and N_m are ρ and ν raised to work over beliefs: $R_m(\beta, a) = \sum_{q \in \mathcal{Q}} \beta(q) * \rho(q, a)$ and $N_m(\beta, a)(o) = \sum_{q, q' \in \mathcal{Q}} \beta(q) * \tau(q, a)(q') * \nu(a, q')(o)$. A strategy σ is optimal if it maximizes V_m for all belief states, that is, if for all β , $V_m(\sigma, \beta)$ is equal to $V_m^*(\beta) = \max_{\sigma'} V_m(\sigma', \beta)$. Prior work has provided algorithms for finding optimal strategies by reducing the problem to one of finding an optimal strategy for a related Markov Decision Process (MDP) that uses these belief states as its state space (e.g., [36]). (For a survey, see [37].)

Example. We can formalize the motivating example provided in Section 1 as a POMDP m_{ex} . Here, we provide an overview that is sufficient for understanding the rest of the paper; [8] provides additional details.

For simplicity, we assume that the only information relevant to advertising is the gender of the visitor. Thus, the state space \mathcal{Q} is determined by three factors: the visitor’s gender, the gender (if any) recorded in the database, and what advertisement (if any) the network has shown to the visitor.

Also for simplicity, we assume that the network is choosing among three advertisements. We use the action space $\mathcal{A} = \{\text{lookup}, \text{ad}_1, \text{ad}_2, \text{ad}_3\}$. The actions ad_1 , ad_2 , and ad_3 correspond to the network showing the visitor one of the three possible advertisements while lookup corresponds to the network looking up information on the visitor. We presume ad_1 is the best for females and the worst for males, ad_3 is the best for males and the worst for females, and ad_2 strikes a middle ground. In particular, we use $\rho(q, \text{ad}_1) = 9$ for a state q in which the visitor is a female and has not yet seen an ad. The reward 9 could refer to a measure of the click through rate or the average preference assigned to the ad by females during market research. If the visitor were instead a male, the reward would be 3. For ad_3 , the rewards are reversed with 3 for females and 9 for males. For ad_2 , the reward is 7 for both genders. The action lookup or showing a second ad produces reward of zero. We use a discounting factor of $\gamma = 0.9$.

The function τ shows how actions change the environment’s state while ν shows how observations accompany these actions. τ enforces that showing an ad changes the state into one in which showing a second ad produces no further rewards. It also specifies that performing `lookup` does not change the state of the environment. On the other hand, ν shows that `lookup` can change the state of the agent’s knowledge. In particular, it shows that performing `lookup` produces an observation $\langle d, \alpha \rangle$. The observation reveals that the database holds data d about the visitor’s gender and α about what if any ad the visitor has seen. Thus, the observation space is $\mathcal{O} = \{\mathbf{f}, \mathbf{m}, \perp\} \times \{\mathbf{ad}_1, \mathbf{ad}_2, \mathbf{ad}_3, \emptyset\}$ with \mathbf{f} for the database showing a female, \mathbf{m} for a male, \perp for no gender entry, \mathbf{ad}_i for the visitor having seen \mathbf{ad}_i , and \emptyset for the visitor having not seen an ad.

How the network will behave depends upon the network’s initial beliefs β_{ex1} . We presume that the network believes its database’s entries to be correct, that it has not shown an advertisement to the visitor yet, and that visitors are equally likely to be female or male. Under these assumptions, the optimal plan for the network is to first check whether the database contains information about the visitor. If the database records that the visitor is a female, then the network shows her \mathbf{ad}_1 . If it records a male, the network shows \mathbf{ad}_3 . If the database does not contain the visitor’s gender (holds \perp), then the network shows \mathbf{ad}_2 . The optimal plan is not constrained as to what the agent does after showing the advertisement as it does not affect the reward. (In [8], we discuss using the idea of *non-redundancy* to eliminate this artifact.)

This optimal plan characterizes the form of the set of optimal strategies. The set contains multiple optimal strategies since the network is unconstrained in the actions it performs after showing the advertisement. The optimal strategies must also specify how the network would behave under other possible beliefs it could have had. For example, if the network believed that all visitors are females regardless of what its database records, then it would always show \mathbf{ad}_1 without first checking its database.

Intuitively, using any of these optimal strategies would violate the privacy policy prohibiting using gender for marketing. The reason is that the network selected which advertisement to show using the database’s information about the visitor’s gender.

We expect the network constrained to obeying the policy will show \mathbf{ad}_2 to all visitors (presuming approximately equal numbers of female and male visitors). Our reasoning is that the network must plan as though it does not know and cannot learn the visitor’s gender. In this state of simulated ignorance, the best plan the network can select is the middle ground of \mathbf{ad}_2 . The next section formalizes this planning under simulated ignorance.

4 Constraining POMDPs for Information Use

We now provide a formal characterization of how an agent pursuing a purpose should behave when prohibited from using a class of information. Recall the intuition that using information is using a distinction and that not using it cor-

responds to ignoring the distinction. We use this idea to model sensitive information with an equivalence relation \equiv . We set $o_1 \equiv o_2$ for any two observations o_1 and o_2 that differ only by sensitive information.

From \equiv and a POMDP m , we construct a POMDP m/\equiv that ignores the prohibited information. For each equivalence class of \equiv , m/\equiv will conflate its members by treating every observation in it as indistinguishable from one another. To ignore these distinctions, on observing o , the agent updates its belief state as though it has seen some element of $\equiv[o]$ but is unsure of which one where $\equiv[o]$ is the equivalence class that holds the observation o .

To make this formal, we define a quotient POMDP m/\equiv that uses a quotiented space of observations. Let \mathcal{O}/\equiv be the set of equivalence classes of \mathcal{O} under \equiv . Let ν/\equiv give the probability of seeing any observation of an equivalence class: $\nu/\equiv(a, q')(O) = \sum_{o \in O} \nu(a, q')(o)$ where O is an equivalence class in \mathcal{O}/\equiv . Given $m = \langle \mathcal{Q}, \mathcal{A}, \tau, \rho, \mathcal{O}, \nu, \gamma \rangle$, let m/\equiv be $\langle \mathcal{Q}, \mathcal{A}, \tau, \rho, \mathcal{O}/\equiv, \nu/\equiv, \gamma \rangle$.

Proposition 1. *For all POMDPs m and equivalences \equiv , m/\equiv is a POMDP.*

Example. Returning to the example POMDP of Section 3, the policy governing the network states that the network will not use the database’s entry about the visitor’s gender for determining the advertisement to show the visitor. The auditor must decide how to formally model this restriction. One way would be to define \equiv_{ex} such that for all g and g' in $\{\text{f}, \text{m}, \perp\}$, and α in $\{\text{ad}_1, \text{ad}_2, \text{ad}_3, \emptyset\}$, $\langle g, \alpha \rangle \equiv_{\text{ex}} \langle g', \alpha \rangle$, conflating the gender for all observations. Under this requirement, $m_{\text{ex}}/\equiv_{\text{ex}}$ will be such that the optimal strategy will be determined solely by the network’s initial beliefs and performing the action `lookup` will be of no benefit. Any optimal strategy for $m_{\text{ex}}/\equiv_{\text{ex}}$ will call for performing `ad2` from the initial beliefs $\beta_{\text{ex}1}$ discussed above.

Alternatively, the auditor might conclude that the policy only forces the network to ignore whether the database records the visitor as a female or male and not whether the database contains this information. In this case, the auditor would use a different equivalence \equiv'_{ex} such that $\langle \text{f}, \alpha \rangle \equiv'_{\text{ex}} \langle \text{m}, \alpha \rangle$ but $\langle \text{f}, \alpha \rangle \not\equiv'_{\text{ex}} \langle \perp, \alpha \rangle \not\equiv'_{\text{ex}} \langle \text{m}, \alpha \rangle$ for all α . Under the initial beliefs $\beta_{\text{ex}1}$, the network would behave identically under \equiv_{ex} and \equiv'_{ex} . However, if the network’s beliefs were such that it is much more likely to not know a female’s gender than a male’s, then it might choose to show `ad1` instead of `ad2` in the case of observing $\langle \perp, \emptyset \rangle$.

The next proposition proves that we constructed the POMDP m/\equiv so that beliefs are updated as if the agent only learns that some element of an equivalence class of observations was observed but not which one. That is, we prove that the updated belief $\text{up}_{m/\equiv}(\beta, a, \equiv[o])(q')$ is equal to the probability that the next environmental state is q' given the distribution β over possible last states, that the last action was a , and that the observation was a member of $\equiv[o]$. Recall that Q' is a random variable over the next state while O , A , and B identify the last observation, action, and belief state, respectively.

Proposition 2. *For all POMDPs m , equivalences \equiv , beliefs β , actions a , observations o , and states q' , $\text{up}_{m/\equiv}(\beta, a, \equiv[o])(q') = \Pr[Q'=q' \mid O \in \equiv[o], A=a, B=\beta]$.*

Propositions 1 and 2 show that m/\equiv is a POMDP that ignores the distinctions among observations that only differ by sensitive information. They justify the following definition, which explains how a purpose-driven agent should act when prohibited from using certain information. They show that it correctly prevents the use of the prohibited information. The definition’s appeal to optimizing a POMDP is justified by our prior work showing that an action is for a purpose when that action is selected as part of a plan optimizing the satisfaction of that purpose [6]. We extend this result to information by concluding that information used to select an action is used for that action’s purpose.

Definition 1 (Cognitive). *An agent obeys the purpose restriction to perform actions for the purpose modeled by the POMDP m without using the information modeled by \equiv iff the agent selects an strategy by optimizing m/\equiv .*

We call the above definition *cognitive* since it refers to the strategy selected by the agent as part of a cognitive process that the auditor cannot measure. Rather, the auditor can only view the agent’s external behavior and visible aspects of the environment. That is, the auditor can only view the agent’s actions and observations, which we refer to collectively as the agent’s *execution*.

We can formalize the agent’s execution using a function `exe`. Even when the agent uses the POMDP m/\equiv with observation space \mathcal{O}/\equiv to select a strategy, the actual observations the agent makes lie in \mathcal{O} , complicating `exe`. We recursively define `exe($m, \equiv, \sigma, \beta_1, \mathbf{o}$)` to be the agent’s execution that arises from it employing a strategy σ observing a sequence of observations $\mathbf{o} = [o_1, \dots, o_n]$ in \mathcal{O}^* starting with beliefs β_1 for a POMDP m/\equiv . For the empty sequence $[]$ of observations, `exe($m, \equiv, \sigma, \beta, []$)` = $[\sigma(\beta)]$ since the agent can only make one action before needing to wait for the next observation and updating its beliefs. For non-empty sequences $\mathbf{o}:\mathbf{o}$, it is equal to $\sigma(\beta):o:\text{exe}(m, \equiv, \sigma, \text{up}_{m/\equiv}(\beta, \sigma(\beta), \equiv[o]), \mathbf{o})$ where $x:y$ denotes prepending element x to the sequence y .

A single execution e can be consistent with both an optimal strategy for m/\equiv and a strategy that is not optimal for m/\equiv . Consider for example, the execution $e = [\text{ad}_2] = \text{exe}(m_{\text{ex}}, \equiv_{\text{ex}}, \sigma, \beta_{\text{ex}}, [])$ that arises from an optimal strategy σ for $m_{\text{ex}}/\equiv_{\text{ex}}$. This execution can also arise from the agent planning for a different purpose, such as maximizing kickbacks for showing certain ads, provided that ad_2 also just so happens to maximize that purpose. Since the auditor only observes the execution e and not the cognitive process that selected the action ad_2 , the auditor cannot know by which process the agent selected the ad. Thus, the auditor cannot determine from an execution that an agent obeyed a purpose restriction under Definition 1.

Some auditors may find this fundamental limitation immaterial since such an agent’s actions are still consistent with an allowed strategy. Since the actual reasons behind the agent selecting those actions do not affect the environment, an auditor might not find concerning an agent doing the right actions for the wrong reasons. To capture this more consequentialist view of compliance, we provide a weaker definition that focuses on only the agent’s execution.

Definition 2 (Behaviorist). *An agent performing execution \mathbf{e} obeys the purpose restriction to perform actions for the purpose modeled by the POMDP m and initial beliefs β_1 without using the information modeled by the equivalence relation \equiv given the observations \mathbf{o} iff $\mathbf{e} = \text{exe}(m, \equiv, \sigma, \beta_1, \mathbf{o})$ for some σ that is an optimal strategy of m/\equiv .*

5 Auditing Algorithm

Under the behaviorist definition, to determine whether an agent obeyed a prohibition against using certain information for a purpose pursued by the agent, the auditor can compare the agent’s behaviors to the appropriate strategies. The auditor records the agent’s execution in a log ℓ that shows the actions and observations of the agent. For example, databases for electronic medical records log many of the actions and observations of healthcare providers. The auditor may then compare the recorded behavior to that dictated by Definition 2, i.e., to the optimal strategies for the quotient POMDP modeling the purpose while ignoring disallowed information.

Given our formal model, we can automate the comparison of the agent’s behavior to the allowable behavior. We use an algorithm AUDIT that takes as inputs a POMDP m , an equivalence relation \equiv , and a log $\ell = [a_1, o_1, a_2, o_2, \dots, a_n, o_n]$ such that the audited agent is operating in the environment m under a policy prohibiting information as described by \equiv and took action a_i followed by observation o_i for all $i \leq n$. For simplicity, we assume that ℓ records all relevant actions and observations. AUDIT returns whether the agent’s behavior, as recorded in ℓ , is inconsistent with optimizing the POMDP m/\equiv .

AUDIT operates by first constructing the quotient POMDP m/\equiv from m and \equiv . Next, similar to a prior algorithm [34], for each i , AUDIT checks whether performing the recorded action a_i in the current belief state β_i is optimal under m/\equiv . The algorithm constructs these belief states from the observations and initial belief state β_1 . Due to the complexity of solving POMDPs [38], we use an approximation algorithm to solve for the value of performing a_i in β_i (denoted $Q_{m/\equiv}^*(\beta_i, a_i)$) and the optimal value $V_{m/\equiv}^*(\beta_i)$. Unlike prior work, for soundness, we require an approximation algorithm SOLVEPOMDP that produces both lower bounds V_{low}^* and upper bounds V_{up}^* on $V_{m/\equiv}^*(\beta_i)$. Many such algorithms exist (e.g., [39–42]). For each β_i and a_i in ℓ , AUDIT checks whether these bounds show that $Q_{m/\equiv}^*(\beta_i, a_i)$ is strictly less than $V_{m/\equiv}^*(\beta_i)$. If so, then the action a_i is sub-optimal for β_i and AUDIT returns true. Pseudo-code for AUDIT follows:

```

AUDIT( $\langle \mathcal{Q}, \mathcal{A}, \tau, \rho, \mathcal{O}, \nu, \gamma \rangle, \equiv, \beta_1, [a_1, o_1, a_2, o_2, \dots, a_n, o_n]$ ):
01   $m' = \langle \mathcal{Q}, \mathcal{A}, \tau, \rho, \mathcal{O}/\equiv, \nu/\equiv, \gamma \rangle$ 
02   $\langle V_{\text{low}}^*, V_{\text{up}}^* \rangle := \text{SOLVEPOMDP}(m')$ 
03  for ( $i := 1; i \leq n; i++$ ):
04      if ( $Q_{\text{up}}^*(V_{\text{up}}^*, \beta_i, a_i) < V_{\text{low}}^*(\beta_i)$ ):
05          return true
06       $\beta_{i+1} := \text{up}_{m/\equiv}(\beta_i, a_i, \equiv[o_i])$ ;
07  return false

```

where $Q_{\text{up}}^*(V_{\text{up}}^*, \beta, a)$ is a function that uses V_{up}^* to return an upper bound on $Q_{m/\equiv}^*(\beta, a)$. $Q_{\text{up}}^*(V_{\text{up}}^*, \beta, a)$ equals:

$$R_m(\beta, a) + \gamma \sum_{O \in \mathcal{O}/\equiv} N_m(\beta, a)(O) * V_{\text{up}}^*(\text{up}_{m'}(\beta, \sigma(\beta), O))$$

Theorem 1 (Soundness). *If AUDIT returns true, then the agent did not follow an optimal strategy for m/\equiv , violating both Definitions 1 and 2.*

Thus, if AUDIT returns true, either the agent optimized some other purpose, used information it should not have, used a different POMDP model of its environment, or failed to correctly optimize the POMDP. Each of these possibilities should concern the auditor and is worthy of further investigation.

If the algorithm returns false, then the auditor cannot find the agent’s behavior inconsistent with an optimal strategy and should spend his time auditing other agents. However, AUDIT is incomplete and such a finding does not mean that the agent surely performed its actions for the purpose without using the prohibited information. For the cognitive definition, incompleteness is unavoidable since the definition depends upon cognitive constructs that the auditor cannot measure. For example, recall that the network could display the execution $e = [\text{ad}_2]$ either from performing the allowed optimization or by performing some disallowed optimization that also results in the action ad_2 being optimal.

For the behaviorist definition, incompleteness results since a better approximation might actually show that $Q_{m/\equiv}^*(\beta_i, a_i) < V_{m/\equiv}^*(\beta_i)$ for some i . In principle this source is avoidable by using an exact POMDP solver instead of an approximate one. However, the exact solution to some POMDPs is undecidable [43]. Nevertheless, we can prove that this inability is the only source of incompleteness.

Theorem 2 (Qualified Completeness). *If Audit using an oracle to exactly solve POMDPs returns false, then the agent obeyed the purpose restriction according to the behaviorist definition (Definition 2).*

Other Purpose Restrictions. AUDIT is specialized for determining whether or not the audited agent performed its actions for a purpose without using some prohibited information. While such a question is relevant to an internal compliance officer auditing employees, it does not correspond to the purpose restrictions found in outward-facing privacy policies.

One type of restriction found in such policies is the *not-for* restriction prohibiting information from being used for a purpose. For example, Yahoo! promised to *not* use contents of emails *for* marketing. This restriction is similar to the condition checked by AUDIT, but is weaker in that audited agent may obey it either (1) by performing actions for that purpose without using that information (which AUDIT checks) or (2) by not performing actions for that purpose.

A second type is the *only-for* restriction, which limits the agent to using a class of information only for a purpose. For example, HIPAA requires that medical records are used *only for* certain purposes such as treatment. It is also

weak in that the agent can obey it either (1) by performing actions for the purpose (which AUDIT checks using equality for \equiv to allow the agent to use the information) or (2) by not using the information in question while performing actions for some other purpose.

For both of these types, our algorithm can handle the first option (1) for compliance. However, for both these types, the second option (2) for compliance involves an open-ended space of possible alternative purposes that could have motivated the agent’s actions. In some cases (e.g., healthcare), this space may be small enough to check each alternative (e.g., treatment, billing, research, training) with AUDIT. In other cases, the auditor might have the authority to compel the agent to explain what its purpose was. In either of these cases, the auditor could use AUDIT to explore these alternative purposes.

Modeling. AUDIT requires a POMDP that models how various actions affect the purpose in question. Future work could ease the process of model construction using techniques from reinforcement learning, such as SARSA [44], that automatically construct models from observing the behavior of multiple agents.

In some cases, the auditor might be able to compel the agent to provide the POMDP used. In this case, AUDIT would check whether the agent’s story is consistent with its actions.

6 Relationship with Noninterference

We have provided a definition of information use in terms of a POMDP. Prior work provides the *noninterference* definition of information use for automata [7]. In this section, we show that our definition implies a form of noninterference. In particular, we show that agents using strategies optimizing m/\equiv has noninterference for \equiv , which suggests that our definition is sufficiently strong to rule out information use. We start by reviewing automata and noninterference.

Automaton Model of Systems. The agent using the POMDP to select a strategy can implement that strategy as a *control system* or *controller* (e.g., [45]). We follow Goguen and Meseguer’s work and model systems as deterministic automata [7]. However, since we do not analyze the internal structure of systems (it is unavailable to the auditor), our approach can be applied to other models. We limit our discussion to deterministic systems since there are many competing generalizations of noninterference to the nondeterministic setting (e.g., [46–48]), but the main competitors collapse into standard noninterference in the deterministic case [49].

A system automaton $s = \langle t, r \rangle$ consists of a labeled transition system (LTS) t and a current state r . An LTS $t = \langle \mathcal{R}, \mathcal{O}, \mathcal{A}, \text{next}, \text{act} \rangle$ describes the automaton’s behavior where \mathcal{R} is a set of states; \mathcal{O} , a set of observations (inputs); \mathcal{A} , a set of actions (outputs); $\text{next} : \mathcal{R} \times \mathcal{O} \rightarrow \mathcal{R}$ is a transition function; and $\text{act} : \mathcal{R} \rightarrow \mathcal{A}$ is a function identifying the action that the automation selects given its current state. The current state $r \in \mathcal{R}$ changes as the system makes observations and takes actions.

As with POMDPs, an execution of a system s modeled as an automaton corresponds to an interleaving of observations from the environment and actions taken by the system. Let $\text{exe}(s, \mathbf{o})$ denote the execution of s on a sequence \mathbf{o} of observations. As for POMDPs, we define exe for systems recursively: $\text{exe}(\langle t, r \rangle, []) = [\text{act}(r)]$ and $\text{exe}(\langle t, r \rangle, \mathbf{o}:\mathbf{o}) = \text{act}(r):o:\text{exe}(\langle r, \text{next}(r, o) \rangle, \mathbf{o})$ where $t = \langle \mathcal{R}, \mathcal{O}, \mathcal{A}, \text{next}, \text{act} \rangle$.

Noninterference. Recall that we set $o_1 \equiv o_2$ for any two observations o_1 and o_2 that differ only by sensitive information. To not use the sensitive information, the system s should treat such related observations identically.

To formalize this notion, we raise \equiv to work over sequences of observations and actions (i.e., executions and sequences of observations). For such sequences \mathbf{x} and \mathbf{y} in $(\mathcal{O} \cup \mathcal{A})^*$, $\mathbf{x} \equiv \mathbf{y}$ iff they are of the same length and for each pair of elements x and y at the same position in \mathbf{x} and \mathbf{y} , respectively, $x \equiv y$ where \equiv is treated as equality when comparing actions.

Definition 3. *A system s has noninterference for \equiv iff for all observation sequences \mathbf{o}_1 and \mathbf{o}_2 in \mathcal{O}^* , $\mathbf{o}_1 \equiv \mathbf{o}_2$ implies that $\text{exe}(s, \mathbf{o}_1) \equiv \text{exe}(s, \mathbf{o}_2)$.*

Our definition corresponds to the form of noninterference enforced by most type systems for information flow. (See [9] for a survey.) Unlike Goguen and Meseguer’s definition, ours does not require the system’s behavior to remain unchanged regardless of whether or not it receives sensitive information. Rather, the system’s behavior may change upon receiving sensitive information, but this change must be the same regardless of the value of the sensitive information. (See [50] for a discussion.)

Relationship. We now characterize the relationship between our quotienting definition of information use and noninterference. We do so by considering a control system s operating in an environment modeled by a POMDP m . We require that s and m share the same sets of actions \mathcal{A} and observations \mathcal{O} . However, the state spaces \mathcal{R} of s and \mathcal{Q} of m differ with \mathcal{R} representing the internal states of the system and \mathcal{Q} representing the external states of the environment.

We relate systems and strategies by saying that a system s *implements* a strategy σ for m/\equiv and beliefs β_1 iff for all \mathbf{o} in \mathcal{O}^* , $\text{exe}(s, \mathbf{o}) = \text{exe}(m, \equiv, \sigma, \beta_1, \mathbf{o})$. We denote the set of such implementing systems as $\text{Imp}(m, \equiv, \sigma, \beta_1)$. This definition allows us to formalize the intuition that agents using strategies optimizing m/\equiv has noninterference for \equiv . In fact, systems implementing any strategy for m/\equiv has noninterference since any such implementation respects \equiv .

Theorem 3. *For all systems q , POMDPs m , initial beliefs β_1 , strategies σ , and equivalences \equiv , if s is in $\text{Imp}(m, \equiv, \sigma, \beta_1)$, then s has noninterference for \equiv .*

Agents obeying a purpose restriction under the cognitive definition (Definition 1) will employ a system in $\text{Imp}(m, \equiv, \sigma, \beta_1)$. Thus, Theorem 3 shows that the cognitive definition is sufficiently strong to rule out information use.

Information Use for Other Purposes. The situation is subtler for the weaker behaviorist definition (Definition 2) and the algorithm AUDIT based upon it. Systems exist that will pass AUDIT and satisfy the behaviorist definition despite having interference by using the protected information for some purpose other than the restricted one. The key is that there could be more than one optimal strategy for a POMDP and that the agent may use the choice among optimal strategies to communicate information. The behavior of such a system will be consistent with whichever optimal strategy it selects, satisfying the behaviorist definition and AUDIT. However, such a system will not actually implement any strategy for the quotiented POMDP m/\equiv since it distinguishes between observations conflated by \equiv .

For example, consider modifying the motivating example found in Section 3 in two ways to make the POMDP m'_{ex} . First, let ad_2 come in two versions, ad_2^a and ad_2^b , which are otherwise the same as the original ad_2 . Second, change the POMDP so that the network must perform the action `lookup` before showing any ads. The agent can optimize m'_{ex}/\equiv by either using a strategy σ^a or σ^b . Under σ^a , starting from the initial beliefs $\beta_{\text{ex}1}$ discussed above, the network will first perform `lookup` and then show ad_2^a . Under σ^b , it will show ad_2^b after `lookup`.

The network’s ability to choose between σ^a and σ^b can result in interference. In particular, the network might not implement either of them and instead delay the choice between ad_2^a and ad_2^b until after the observation from `lookup` informs it of the visitor’s gender. The network could then use ad_2^a for a female and ad_2^b for a male. While such a system would use the information and have interference, it obeys the behaviorist definition with its actions consistent with either σ^a or σ^b .

Since such systems use the prohibited information to choose between *optimal* strategies, doing so does not actually increase its satisfaction of the purpose. Thus, this information use is not intuitively for that purpose and the agent must be motivated by some other purpose. Thus, the behaviorist definition does not allow the agent to use the information for the purpose prohibited by the restriction, but rather allows the agent to use the information for other purposes.

The auditor might want to prevent such interference since it violates the cognitive definition. The modifications to the example illustrate two ways that the auditor can do so if he has sufficient control over the agent’s environment. The first is to ensure that only a single strategy is optimal. The second is to make sure that the agent can avoid learning the protected information (such as by performing the action `lookup`) and that learning it incurs a cost. When learning information is optional and costly, the agent will only be able to learn it if doing so increases its total reward, and not just to select among optimal strategies that do not depend upon using that information. A third possible modification is to require the agent to perform an action committing it to a single strategy before it can learn the protected information.

In some cases an auditor can detect such information flows without modifying the POMDP. For example, intuitively, we would expect the ad network to handle more than one visitor. The auditor could compare the network’s behavior when

given a female to that when given a male. A difference in treatment indicates that the network is not consistently implementing either of the optimal strategies.

7 Conclusion

We use planning to create the first formal semantics for determining when information is used for a purpose. We have provided an auditing algorithm based on our formalism. We have discussed applying our algorithm to the problem of enforcing purpose restrictions found in privacy policies.

Our methods have applications beyond enforcing purpose restrictions. For example, due to privacy concerns, much interest exists in determining how third-party data collection agencies use the information they collect. (See [51] for a survey.) Despite being a question of information flow, program analyses are inapplicable since the programs are unavailable, as in our setting. Unlike our setting, these agencies typically do not subject themselves to purpose restrictions. Nevertheless, their desire for profit implicitly restrains their behavior in a manner similar to a purpose restriction. Thus, our semantics and algorithm provide a starting point for investigating such agencies.

Acknowledgments. We appreciate discussions with Lorrie Faith Cranor, Joseph Y. Halpern, and Manuela M. Veloso on this work. We thank Amit Datta, Dilsun Kaynar, and Divya Sharma for many helpful comments on this paper.

References

1. Bank of America Corp.: Bank of America privacy policy for consumers (2005)
2. Office for Civil Rights: Summary of the HIPAA privacy rule. OCR Privacy Brief, U.S. Department of Health and Human Services (2003)
3. Yahoo!: Privacy policy: Yahoo Mail (2013)
4. FairWarning: Privacy breach detection for healthcare. White Paper (2010)
5. Taylor, R.: Action and Purpose. Prentice-Hall (1966)
6. Tschantz, M.C., Datta, A., Wing, J.M.: Formalizing and enforcing purpose restrictions in privacy policies. In: 2012 IEEE Symp. on Security and Privacy. (2012) 176–190
7. Goguen, J.A., Meseguer, J.: Security policies and security models. In: IEEE Symp. on Security and Privacy. (1982) 11–20
8. Tschantz, M.C., Datta, A., Wing, J.M.: Purpose restrictions on information use. Technical Report CMU-CyLab-13-005 and CMU-CS-13-116, Carnegie Mellon University (June 2013)
9. Sabelfeld, A., Myers, A.C.: Language-based information-flow security. IEEE Journal on Selected Areas in Communications **21**(1) (2003) 5–19
10. Vachharajani, N., Bridges, M.J., Chang, J., Rangan, R., Ottoni, G., Blome, J.A., Reis, G.A., Vachharajani, M., August, D.I.: RIFLE: An architectural framework for user-centric information-flow security. In: 37th Annual IEEE/ACM Intl. Symp. on Microarchitecture. (2004) 243–254

11. Newsome, J., Song, D.X.: Dynamic taint analysis for automatic detection, analysis, and signature generation of exploits on commodity software. In: Network and Distributed System Security Symp., The Internet Society (2005)
12. Venkatakrisnan, V.N., Xu, W., DuVarney, D.C., Sekar, R.: Provably correct runtime enforcement of non-interference properties. In: 8th Intl. Conf. on Information and Communications Security, Springer-Verlag (2006) 332–351
13. McCamant, S., Ernst, M.D.: A simulation-based proof technique for dynamic information flow. In: 2007 Wksp. on Programming Languages and Analysis for Security, ACM (2007) 41–46
14. Yumerefendi, A.R., Mickle, B., Cox, L.P.: Tightlip: keeping applications from spilling the beans. In: 4th USENIX Conf. on Networked Systems Design and Implementation. (2007) 12–12
15. Capizzi, R., Longo, A., Venkatakrisnan, V.N., Sistla, A.P.: Preventing information leaks through shadow executions. In: 2008 Annual Computer Security Applications Conf., IEEE Computer Society (2008) 322–331
16. Devriese, D., Piessens, F.: Noninterference through secure multi-execution. In: 2010 IEEE Symp. on Security and Privacy. (2010) 109–124
17. Agrawal, R., Kiernan, J., Srikant, R., Xu, Y.: Hippocratic databases. In: VLDB '02: 28th Intl. Conf. on Very Large Data Bases, VLDB Endowment (2002) 143–154
18. Byun, J.W., Bertino, E., Li, N.: Purpose based access control of complex data for privacy protection. In: SACMAT '05: Tenth ACM Symp. on Access Control Models and Technologies. (2005) 102–110
19. Al-Fedaghi, S.S.: Beyond purpose-based privacy access control. In: Eighteenth Australasian Database Conf., Australian Computer Society, Inc. (2007) 23–32
20. Byun, J.W., Li, N.: Purpose based access control for privacy protection in relational database systems. The VLDB Journal **17**(4) (2008) 603–619
21. Peng, H., Gu, J., Ye, X.: Dynamic purpose-based access control. In: Intl. Symp. on Parallel and Distributed Processing with Applications, IEEE Computer Society (2008) 695–700
22. Jafari, M., Safavi-Naini, R., Sheppard, N.P.: Enforcing purpose of use via workflows. In: WPES '09: 8th ACM Wksp. on Privacy in the Electronic Society. (2009) 113–116
23. Ni, Q., Bertino, E., Lobo, J., Brodie, C., Karat, C.M., Karat, J., Trombetta, A.: Privacy-aware role-based access control. ACM Trans. Inf. Syst. Secur. **13** (2010) 24:1–24:31
24. Enamul Kabir, M., Wang, H., Bertino, E.: A conditional purpose-based access control model with dynamic roles. Expert Syst. Appl. **38** (2011) 1482–1489
25. Jafari, M., Fong, P.W., Safavi-Naini, R., Barker, K., Sheppard, N.P.: Towards defining semantic foundations for purpose-based privacy policies. In: First ACM Conf. on Data and Application Security and Privacy. (2011) 213–224
26. Martinelli, F., Morisset, C.: Quantitative access control with partially-observable markov decision processes. In: Second ACM Conf. on Data and Application Security and Privacy. (2012) 169–180
27. Hayati, K., Abadi, M.: Language-based enforcement of privacy policies. In: PET 2004: Wksp. on Privacy Enhancing Technologies, Springer-Verlag (2005) 302–313
28. Massacci, F., Mylopoulos, J., Zannone, N.: Hierarchical Hippocratic databases with minimal disclosure for virtual organizations. The VLDB Journal **15**(4) (2006) 370–387
29. Barth, A., Mitchell, J., Datta, A., Sundaram, S.: Privacy and utility in business processes. In: CSF '07: 20th IEEE Computer Security Foundations Symp. (2007) 279–294

30. Schmidt, C., Sridharan, N., Goodson, J.: The plan recognition problem: An intersection of psychology and artificial intelligence. *Artificial Intelligence* **11**(1-2) (1978) 45 – 83
31. Baker, C.L., Tenenbaum, J.B., Saxe, R.R.: Bayesian models of human action understanding. In: *Advances in Neural Information Processing Systems 18*, MIT Press (2006) 99–106
32. Baker, C.L., Saxe, R., Tenenbaum, J.B.: Action understanding as inverse planning. *Cognition* **113**(3) (2009) 329–349
33. Ramírez, M., Geffner, H.: Plan recognition as planning. In: *21st International Joint Conf. on Artificial Intelligence*. (2009) 1778–1783
34. Ramírez, M., Geffner, H.: Goal recognition over POMDPs: Inferring the intention of a POMDP agent. In: *22nd International Joint Conf. on Artificial Intelligence, IJCAI/AAAI* (2011) 2009–2014
35. Bellman, R.: On the theory of dynamic programming. *National Academy of Sciences* **38** (1952) 716–719
36. Sondik, E.J.: The optimal control of partially observable Markov processes. PhD thesis, Stanford University (1971)
37. Monahan, G.E.: A survey of partially observable Markov decision processes: Theory, models, and algorithms. *Management Science* **28**(1) (1982) 1–16
38. Papadimitriou, C., Tsitsiklis, J.N.: The complexity of Markov decision processes. *Math. Oper. Res.* **12** (1987) 441–450
39. Zhou, R., Hansen, E.A.: An improved grid-based approximation algorithm for POMDPs. In: *17th International Joint Conf. on Artificial Intelligence*. Volume 1., Morgan Kaufmann (2001) 707–714
40. Smith, T., Simmons, R.: Point-based POMDP algorithms: Improved analysis and implementation. In: *Conf. on Uncertainty in Artificial Intelligence*. (July 2005)
41. Kurniawati, H., Hsu, D., Lee, W.S.: SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In: *Proc. Robotics: Science and Systems*. (2008)
42. Poupart, P., Kim, K.E., Kim, D.: Closing the gap: Improved bounds on optimal POMDP solutions. In: *Intl. Conf. on Automated Planning and Scheduling, AAAI* (2011)
43. Madani, O.: Complexity Results for Infinite-Horizon Markov Decision Processes. PhD thesis, University of Washington (2000)
44. Rummery, G.A., Niranjan, M.: On-line Q-learning using connectionist systems. Technical Report CUEF/F-INFENG/TR 166, Cambridge University Engineering Department (1994)
45. Kaelbling, L.P., Littman, M.L., Cassandra, A.R.: Planning and acting in partially observable stochastic domains. *Artif. Intell.* **101** (1998) 99–134
46. McCullough, D.: Noninterference and the composability of security properties. In: *IEEE Symp. on Security and Privacy*. (1988) 177–186
47. Wittbold, J.T., Johnson, D.M.: Information flow in nondeterministic systems. In: *IEEE Symp. on Security and Privacy*. (1990) 144–161
48. McLean, J.: A general theory of composition for trace sets closed under selective interleaving functions. In: *1994 IEEE Symp. on Security and Privacy*. (1994) 79
49. Clark, D., Hunt, S.: Non-interference for deterministic interactive programs. In: *Formal Aspects in Security and Trust*, Springer-Verlag (2009) 50–66
50. Tschantz, M.C., Wing, J.M.: Extracting conditional confidentiality policies. In: *Sixth IEEE Intl. Conferences on Software Engineering and Formal Methods*. (2008)
51. Mayer, J.R., Mitchell, J.C.: Third-party web tracking: Policy and technology. In: *IEEE Symp. on Security and Privacy*. (2012) 413–427