

18734 Homework 2

Released: Sep 25, 2017

Due: 12 noon Eastern, 9am Pacific, on Oct 4, 2017

1 Policy Language (15 marks)

This question asks you to explore the design space for privacy policy languages and understand the trade-offs offered by three languages. Specifically, we ask you to compare P3P, XACML, and the privacy logic described in lectures of Sept 6, 11 along these axes.

- Q1 Expressiveness: What privacy concepts can each language represent and what concepts can it not represent? If you think one is more expressive than another, give an example of a concept that can be expressed in one language but not the other.
- Q2 Enforceability: What privacy concepts can each automatically enforce and what concepts can it not enforce? Provide pointers to the algorithms used for enforcement of concepts that can be enforced.
- Q3 Usability: What is the expected class of users of this language? What background is expected of the users to be able to use these languages effectively?

Hint: You might find it useful to revisit the readings for the lectures of Sept 7, 12.

2 Logic (35 marks)

Q1 (10 marks)

- 1a: (3 marks) Assume we are working with first order logic, and the domain is the space of positive integers. Consider the formula $\exists x, y, z. \psi(x, y, z)$, where $\psi(x, y, z)$ is $x^{10} + y^{10} = z^{10}$. Is this true? (No proof is required, just say true or false)
- 1b: (7 marks) For finite domains, say x_1, \dots, x_n , $\exists x. \phi(x)$ can be written as $\phi(x_1) \vee \dots \vee \phi(x_n)$. A naive attempt at proving $\exists x, y, z. x^{10} + y^{10} = z^{10}$ would be to substitute every possible value of x, y, z and check the equality. What is the problem with such an approach?

Q2 (10 marks)

For all subparts of this question, $\psi(x, y, z)$ is $x^{10} + y^{10} = z^{10}$ and the domain is the space of positive integers. Which of the following formulas can be checked using the naive approach above:

2a: $\exists x, y, z. (x \leq 10) \wedge \psi(x, y, z)$

2b: $\exists x, y, z. (x \leq 10) \wedge (y \leq 10) \wedge \psi(x, y, z)$

2c: $\exists x, y, z. (x \leq 10) \wedge (y \leq 10) \wedge (z \leq 10) \wedge \psi(x, y, z)$

Q3 (15 marks)

In our translations of HIPAA clauses, we have often seen clauses of the form

$$\begin{aligned} \forall [p_1, p_2, m, q, t, u] \\ & (send(p_1, p_2, m) \\ & \wedge tagged(m, q, t, u) \\ & \wedge attr_in(t, phi)) \\ & \implies maysend(p_1, p_2, m) \end{aligned}$$

Sometimes, we may need to move one (or more) predicates across the \implies connective. Using the rules of inference from the first recitation slides (dated: Sep 11) and the Wikipedia article ¹, prove that the following expression is equivalent to the above statement. [You may not use the Law of Exportation. The objective of this homework is to prove the Law of Exportation.]

$$\begin{aligned} \forall [p_1, p_2, m, q, t, u] \\ & (send(p_1, p_2, m) \\ & \wedge tagged(m, q, t, u)) \\ & \implies (attr_in(t, phi)) \\ & \implies maysend(p_1, p_2, m) \end{aligned}$$

3 Case Study: MDPs (50 marks)

MDP is a tuple $\langle S, A, P(., .), R(., .) \rangle$, where:

- S is a finite set of states, represented as a 1-dimensional matrix (vector).
- A is the finite action set, represented as a 1-dimensional matrix (vector).
- $P_a(s, s') = \Pr(s_{t+1} = s' \mid s_t = s, a_t = a)$ is the probability that action a in state s at time t will lead to state s' at time $t + 1$. The transition function is represented as a 3-dimensional matrix.
- $R_a(s, s')$ is the immediate reward (or expected immediate reward) received after transition to state s' from state s with action a . The reward function is represented as a 3-dimensional matrix.

Read more on MDPs from the lecture slides and Wikipedia.

¹http://en.wikipedia.org/wiki/List_of_rules_of_inference

Q1: AdX Revenue

An online tracking company AdX is deciding whether or not to track the website `personalhealth.com`. It wants to serve a health related ad ad_h which has the potential to generate a lot of ad-revenue. Given that AdX is tracking the website, an impression of ad_h has a probability of 0.5 of getting clicked. If there is no tracking, this probability reduces to a mere 0.01. Figure 1 shows the MDP for the purpose: *maximize-revenue*. For any action a , if there is no edge from s to s' in the Figure, it means that $P_a(s, s') = 0$. For example, $P_{start-tracking}(not-tracking, ad-clicked) = 0$. The starting state is *not-tracking*.

Note that this is a much simpler setting than real-world ad serving companies. We assume that the states *ad-not-clicked* and *ad-clicked* are possible final states, i.e. these states do not allow any further actions.

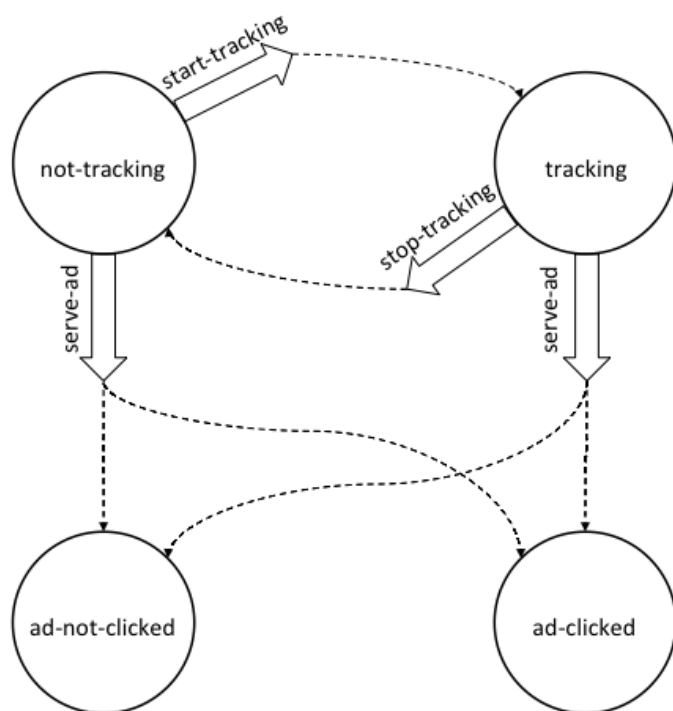


Figure 1: Markov Decision Process for maximizing revenue for AdX

1a: States and Actions (1 mark)

What is the set of states (S) for the MDP? What is the set of actions (A) for the MDP?

1b: Transition function (5 marks)

Assign every dotted edge in the figure the corresponding transition probability. Represent $P(.,.)$ as a 3-dimensional matrix, and write it out as you would assign a 3D matrix in a programming language (like C/C++/Java/Python).

1c: Reward function (5 marks)

AdX receives revenue $R = \$10$ whenever ad_h is clicked. When the ad is not clicked, the revenue is \$0. AdX does not receive any revenue for tracking websites. The reward for an action a taken in state s to reach s' is the expected reward for that action to reach the state, i.e. $R_a(s, s') = R \times P_a(s, s')$. Using this information compute the full reward function for every pair for states and every action. Assign every dotted edge in the figure the corresponding reward. Write it out as you would assign a 3D matrix in a programming language (like C/C++/Java/Python).

1d: Purpose (4 marks)

Is the action *start-tracking* for the purpose *maximize-revenue*? Explain your answer based on the MDP for maximizing revenue in Figure 1.

Q2: AdX Inference

AdX also wants to draw inferences about the gender of online consumers. In Figure 2, find the MDP for the purpose: *draw-inference*. The p values on the dotted edges represent the corresponding transition probabilities. It turns out that by tracking `personalhealth.com`, AdX does not gain any additional information that would help them infer the gender of users. Hence the probability of making a correct inference after tracking (0.5) is the same as the probability of making a correct inference without tracking (0.5). The same holds for incorrect inferences. The reward for a correct inference is $R_i = \$10$, whereas the reward for an incorrect inference is $R_i = \$5$. The starting state is *not-tracking*.

2a: Formalize MDP (5 marks)

Follow the same exercise as Q1 and formally represent the MDP as the tuple $\langle S, A, P(.,.), R(.,.) \rangle$. The reward for an action a taken in state s to reach s' is the expected reward for that action to reach the state, i.e. $R_a(s, s') = R_i \times P_a(s, s')$. The non-zero transition probabilities are provided in Figure 2.

2b: Policy Adherence (5 marks)

AdX states in its privacy policy that

Users will not be tracked on `personalhealth.com` for the purpose of drawing inferences about their gender.

If they perform the sequence of actions $\{start-tracking, infer-gender\}$, does AdX violate their policy? How about the sequence $\{infer-gender\}$? Explain your answer based on the MDP for drawing inference in Figure 2.




Figure 2: Markov Decision Process for drawing inference for AdX

Q3: Two sites and two ads

Given, there are two sites `pregnancy.com` and `sports.com`, and two ads ad_p and ad_s . It is known that if ad_p is served while `pregnancy.com` is tracked, then the probability that ad_p is clicked is 0.5. Also, given that `sports.com` is tracked, the probability that ad_s is clicked is 0.7. If the corresponding sites are not tracked, both these probabilities fall to 0.01. AdX receives revenue $R_p = \$10$ whenever the ad_p is clicked, and $R_s = \$5$ whenever ad_s is clicked. For this question, you have to model this decision process as an MDP for the purpose *maximize-revenue*.

REMEMBER: Although AdX has two ads in its repository, it can only serve only one ad (either ad_p or ad_s). So, actions like *serve-both-ads* or a sequence of actions where both ads are served are NOT allowed. However, AdX can track none, one, or both websites. Secondly, the final state of the MDP is when one of the ads is clicked/not clicked, i.e. there can be no actions from any state representing that an ad is clicked/not clicked. There cannot be further actions after an action that serves an ad.

This question does not have just one correct answer. You should argue why your design works with all the constraints.

3a: States (3 marks)

What is the set of states (S) in your MDP? For each state, explain why it is necessary for the MDP.

3b: Actions (2 marks)

What is the set of actions (A) in your MDP? For each action, explain why it is necessary for the MDP.

3c: Transition and Reward Functions (10 marks)

Draw a neat figure of the MDP and label it with all transition probabilities and rewards. If the figure gets too messy, you may use two figures: one for the transition probabilities, one for the rewards. If there is a missing edge in the figure, I will assume the transition probability and reward are 0 for that edge. Once again, the reward for an action a taken in state s to reach s' is the expected reward for that action to reach the state, i.e. $R_a(s, s') = R_x \times P_a(s, s')$, where R_x is R_p or R_s depending on which ad is clicked.

3d: Optimal plan (10 marks)

Based on your MDP, what would you advise AdX do if their purpose is to maximize revenue? Specifically, mention the sequence of actions that maximizes AdX's revenue.

Submission

Submit your answers in a pdf file with the name <your_andrew_id>_HW2.pdf through Canvas. The submission deadline is 12 noon Eastern/ 9am Pacific, Oct 4, 2017.