

SPUDD: Stochastic Planning using Decision Diagrams

Kushagra Chandak

Paper by: Hoey et al, UAI 1999

30th Nov, 2020

Introduction: SPUDD

- ▶ Large state space: Curse of dimensionality.

Introduction: SPUDD

- ▶ Large state space: Curse of dimensionality.
- ▶ Abstraction/Aggregation techniques to obviate state enumeration.

Introduction: SPUDD

- ▶ Large state space: Curse of dimensionality.
- ▶ Abstraction/Aggregation techniques to obviate state enumeration.
- ▶ SPUDD: VI for MDPs and POMDPs using ADDs to represent value functions and policies.

Introduction: SPUDD

- ▶ Large state space: Curse of dimensionality.
- ▶ Abstraction/Aggregation techniques to obviate state enumeration.
- ▶ SPUDD: VI for MDPs and POMDPs using ADDs to represent value functions and policies.
- ▶ ADDs: Generalization of BDDs.

Introduction: SPUDD

- ▶ Large state space: Curse of dimensionality.
- ▶ Abstraction/Aggregation techniques to obviate state enumeration.
- ▶ SPUDD: VI for MDPs and POMDPs using ADDs to represent value functions and policies.
- ▶ ADDs: Generalization of BDDs.
- ▶ Derives from SPI which uses decision trees (unscalable) to represent π and V .

Introduction: SPUDD

- ▶ Large state space: Curse of dimensionality.
- ▶ Abstraction/Aggregation techniques to obviate state enumeration.
- ▶ SPUDD: VI for MDPs and POMDPs using ADDs to represent value functions and policies.
- ▶ ADDs: Generalization of BDDs.
- ▶ Derives from SPI which uses decision trees (unscalable) to represent π and V .
- ▶ Disjunctive structure in probability exploited by decision graphs.

Value Iteration (VI) Recap

- ▶ Bellman Equation: $V_{\pi}(s) = R(s) + \gamma \sum_{t \in \mathcal{S}} Pr(s, \pi(s), t) \cdot V_{\pi}(t)$

Value Iteration (VI) Recap

- ▶ Bellman Equation: $V_{\pi}(s) = R(s) + \gamma \sum_{t \in \mathcal{S}} Pr(s, \pi(s), t) \cdot V_{\pi}(t)$
- ▶ VI eqn: $V^{n+1}(s) = R(s) + \max_{a \in \mathcal{A}} \left\{ \gamma \sum_{t \in \mathcal{S}} Pr(s, a, t) \cdot V^n(t) \right\}$

Value Iteration (VI) Recap

- ▶ Bellman Equation: $V_{\pi}(s) = R(s) + \gamma \sum_{t \in \mathcal{S}} Pr(s, \pi(s), t) \cdot V_{\pi}(t)$
- ▶ VI eqn: $V^{n+1}(s) = R(s) + \max_{a \in \mathcal{A}} \left\{ \gamma \sum_{t \in \mathcal{S}} Pr(s, a, t) \cdot V^n(t) \right\}$
- ▶ For some finite n , a 's that maximize VI eqn form an opt π and V^n approximates its value.

Value Iteration (VI) Recap

- ▶ Bellman Equation: $V_{\pi}(s) = R(s) + \gamma \sum_{t \in \mathcal{S}} Pr(s, \pi(s), t) \cdot V_{\pi}(t)$
- ▶ VI eqn: $V^{n+1}(s) = R(s) + \max_{a \in \mathcal{A}} \left\{ \gamma \sum_{t \in \mathcal{S}} Pr(s, a, t) \cdot V^n(t) \right\}$
- ▶ For some finite n , a 's that maximize VI eqn form an opt π and V^n approximates its value.
- ▶ Stopping criterion: $\|V^{n+1} - V^n\| < \frac{\epsilon(1-\gamma)}{2\gamma}$
 $\|X\| = \max\{|x| : x \in X\}$. Resulting π is ϵ -opt and V^{n+1} is within $\epsilon/2$ of V^*

ADD

- ▶ BDD: A function, $f : \mathcal{B}^n \rightarrow \mathcal{B}$
- ▶ ADD: Generalize BDD, $f : \mathcal{B}^n \rightarrow \mathcal{R}$.
 - ▶ Terminal node: $f(.) = c$
 - ▶ Non-terminal node: $f(x_1 \dots x_n) = x_1 \cdot f_{then}(x_2 \dots x_n) + \overline{x_1} \cdot f_{else}(x_2 \dots x_n)$




Figure 1: Binary Decision Diagram

ADD Representation of MDP

- ▶ State space: $X = \{X_1 \dots X_n\}$. Can be extended to multi-valued variables.

ADD Representation of MDP

- ▶ State space: $X = \{X_1 \dots X_n\}$. Can be extended to multi-valued variables.
- ▶ Action space: DBN. $X = \{X_1, \dots, X_n\} \xrightarrow{a} X' = \{X'_1 \dots X'_n\}$

ADD Representation of MDP

- ▶ State space: $X = \{X_1 \dots X_n\}$. Can be extended to multi-valued variables.
- ▶ Action space: DBN. $X = \{X_1, \dots, X_n\} \xrightarrow{a} X' = \{X'_1 \dots X'_n\}$
- ▶ Directed arcs from variables in X to variables in X' denote direct causal relationship.

ADD Representation of MDP

- ▶ State space: $X = \{X_1 \dots X_n\}$. Can be extended to multi-valued variables.
- ▶ Action space: DBN. $X = \{X_1, \dots, X_n\} \xrightarrow{a} X' = \{X'_1 \dots X'_n\}$
- ▶ Directed arcs from variables in X to variables in X' denote direct causal relationship.
- ▶ CPT for each post-action variable X'_i defines a conditional distribution $P_{X'_i}^a$ over X'_i : $P_{X'_i}^a(X_1 \dots X_n)$.

Example

- ▶ Process planning problem: A factory agent is tasked to connect 2 objects A and B.
- ▶ One way the agent can connect is take take action *bolt*.
- ▶ State C (objects connected) is independent of variable P (objects painted).
- ▶ If obj A is punched (APU) after bolting depends only on whether it was punched before bolting.
- ▶ Use ADDs to represent the functions $P_{X_i}^a$ (to capture regularities in the CPTs)
- ▶ ADDs also exploit context-specific independence in the distributions.

Example

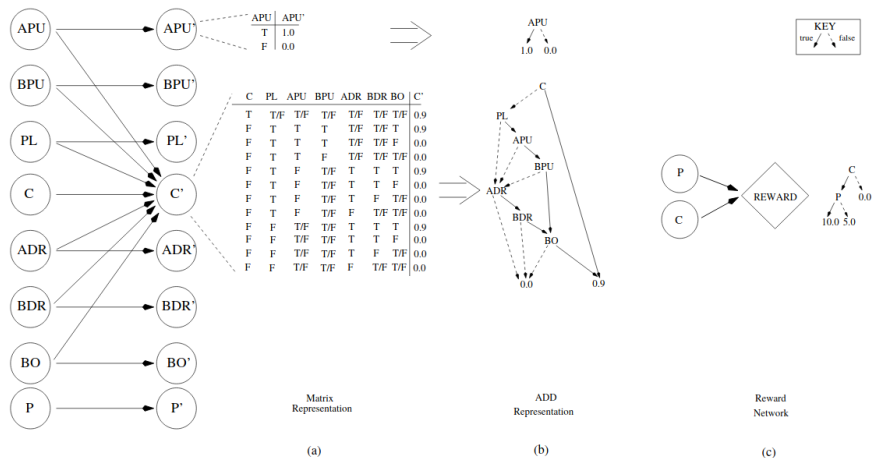


Figure 2: Small FACTORY example: (a) action network for action *bolt*; (b) ADD representation of CPTs (action diagrams); and (c) immediate reward network and ADD representation of the reward table.

Example

- ▶ Regularity in CPT: $Pr_{C'}^{bolt}(C, PL, APU, BPU, ADR, BDR, BO) = [C + \bar{C}[(PL \cdot \overline{APU} \cdot \overline{PL}) \cdot ADR \cdot BDR + PL \cdot APU \cdot BPU] \cdot BO] \cdot 0.9$
- ▶ Reward function as ADD: $R(C, P) = C \cdot P \cdot 10 + C \cdot \bar{P} \cdot 5$
- ▶ Disjunctive structure exploited by ADD. Eg., CPT for C' : Variety of distinct conditions each give give rise to successfully connecting the 2 parts. (Similar to paths)
- ▶ ADDs more compact than trees (and tables): 7 internal nodes and 2 leaves vs 11 internal nodes and 12 leaves. Std matrix: 128 parameters.
- ▶ ADDs more compact than trees *most* times but not always.

SPUDD Algorithm: Overview

- ▶ Avoids explicit enumeration of the state space. (Similar to SDD)

SPUDD Algorithm: Overview

- ▶ Avoids explicit enumeration of the state space. (Similar to SDD)
- ▶ Classical VI but uses ADDs to represent V s and CPTs.

SPUDD Algorithm: Overview

- ▶ Avoids explicit enumeration of the state space. (Similar to SDD)
- ▶ Classical VI but uses ADDs to represent V s and CPTs.
- ▶ Savings both in space and computational time.

SPUDD Algorithm: Overview

- ▶ Avoids explicit enumeration of the state space. (Similar to SDD)
- ▶ Classical VI but uses ADDs to represent V s and CPTs.
- ▶ Savings both in space and computational time.
- ▶ V at each step is represented as an ADD. ($V^0 = R$)

SPUDD Algorithm: Overview

- ▶ Avoids explicit enumeration of the state space. (Similar to SDD)
- ▶ Classical VI but uses ADDs to represent V s and CPTs.
- ▶ Savings both in space and computational time.
- ▶ V at each step is represented as an ADD. ($V^0 = R$)
- ▶ Exploit ADD structure of V^i and MDP representation to get ADD structure for V^{i+1} .

SPUDD Algorithm

- ▶ Variables in V^i are replaced by their primed counterparts (post-action vars).

SPUDD Algorithm

- ▶ Variables in V^i are replaced by their primed counterparts (post-action vars).
- ▶ Goal: For each a , compute ADD for V_a^{i+1} : Exp value of performing action a .

SPUDD Algorithm

- ▶ Variables in V^i are replaced by their primed counterparts (post-action vars).
- ▶ Goal: For each a , compute ADD for V_a^{i+1} : Exp value of performing action a .
 - ▶ Negative action diagrams: $\overline{P_{X'_i}^a}(X_1 \dots X_n) = 1 - P_{X'_i}^a(X_1 \dots X_n)$:
Probability that a will make X'_i false.

SPUDD Algorithm

- ▶ Variables in V^i are replaced by their primed counterparts (post-action vars).
- ▶ Goal: For each a , compute ADD for V_a^{i+1} : Exp value of performing action a .
 - ▶ Negative action diagrams: $\overline{P_{X'_i}^a}(X_1 \dots X_n) = 1 - P_{X'_i}^a(X_1 \dots X_n)$:
Probability that a will make X'_i false.
 - ▶ Dual action diagrams:
 $Q_{X'_i}^a(X'_i; X_1 \dots X_n) = X'_i \cdot P_{X'_i}^a(X_1 \dots X_n) + \overline{X'_i} \cdot \overline{P_{X'_i}^a}(X_1 \dots X_n)$

SPUDD Algorithm

- ▶ Variables in V^i are replaced by their primed counterparts (post-action vars).
- ▶ Goal: For each a , compute ADD for V_a^{i+1} : Exp value of performing action a .
 - ▶ Negative action diagrams: $\overline{P_{X'_i}^a}(X_1 \dots X_n) = 1 - P_{X'_i}^a(X_1 \dots X_n)$:
Probability that a will make X'_i false.
 - ▶ Dual action diagrams:
 $Q_{X'_i}^a(X'_i; X_1 \dots X_n) = X'_i \cdot P_{X'_i}^a(X_1 \dots X_n) + \overline{X'_i} \cdot \overline{P_{X'_i}^a}(X_1 \dots X_n)$
- ▶ Intuitively, Q denotes $P(X'_i = x'_i | X_1 = x_1 \dots X_n = x_n)$ (under action a)

SPUDD Algorithm

- ▶ Variables in V^i are replaced by their primed counterparts (post-action vars).
- ▶ Goal: For each a , compute ADD for V_a^{i+1} : Exp value of performing action a .
 - ▶ Negative action diagrams: $\overline{P_{X'_i}^a}(X_1 \dots X_n) = 1 - P_{X'_i}^a(X_1 \dots X_n)$:
Probability that a will make X'_i false.
 - ▶ Dual action diagrams:
 $Q_{X'_i}^a(X'_i; X_1 \dots X_n) = X'_i \cdot P_{X'_i}^a(X_1 \dots X_n) + \overline{X'_i} \cdot \overline{P_{X'_i}^a}(X_1 \dots X_n)$
- ▶ Intuitively, Q denotes $P(X'_i = x'_i | X_1 = x_1 \dots X_n = x_n)$ (under action a)
- ▶ To generate $V_a^{i+1}(s)$: Combine $V_a^i(t)$ with probability of reaching t from s .

SPUDD Algorithm

- ▶ To get V_a^{i+1} : Multiply dual action diagrams X'_j by V''^i and then eliminate X'_j .

SPUDD Algorithm

- ▶ To get V_a^{i+1} : Multiply dual action diagrams X'_j by V^{ii} and then eliminate X'_j .
- ▶ $Q_{X'_j}^a \cdot V^{ii} = f(X'_1 \dots X'_n, X_1 \dots X_n)$

SPUDD Algorithm

- ▶ To get V_a^{i+1} : Multiply dual action diagrams X'_j by V^{ii} and then eliminate X'_j .
- ▶ $Q_{X'_j}^a \cdot V^{ii} = f(X'_1 \dots X'_n, X_1 \dots X_n)$
- ▶ $f(x'_1, \dots, x'_n, x_1 \dots x_n) = V^{ii}(x'_1 \dots x'_n)P(x'_j|x_1 \dots x_n)$

SPUDD Algorithm

- ▶ To get V_a^{i+1} : Multiply dual action diagrams X_j' by V^{ii} and then eliminate X_j' .
- ▶ $Q_{X_j'}^a \cdot V^{ii} = f(X_1' \dots X_n', X_1 \dots X_n)$
- ▶ $f(x_1', \dots, x_n', x_1 \dots x_n) = V^{ii}(x_1' \dots x_n')P(x_j'|x_1 \dots x_n)$
- ▶ Elimination of X_j' (Summing over left and right subgraphs of the ADD for f)

SPUDD Algorithm

- ▶ To get V_a^{i+1} : Multiply dual action diagrams X'_j by V^{ii} and then eliminate X'_j .
- ▶ $Q_{X'_j}^a \cdot V^{ii} = f(X'_1 \dots X'_n, X_1 \dots X_n)$
- ▶ $f(x'_1, \dots, x'_n, x_1 \dots x_n) = V^{ii}(x'_1 \dots x'_n)P(x'_j | x_1 \dots x_n)$
- ▶ Elimination of X'_j (Summing over left and right subgraphs of the ADD for f)
- ▶ $g(X'_1 \dots X'_{j-1}, X'_{j+1} \dots X'_n, X_1 \dots X_n) = \sum_{x'_j} V^{ii}(X'_1 \dots x'_j \dots X'_n)P(x'_j | X_1 \dots X_n)$

SPUDD Algorithm

- ▶ Repeat the previous step for each post-action variable X'_j that occurs in ADD for V^i : Multiply by Q and eliminate the prime variable.

SPUDD Algorithm

- ▶ Repeat the previous step for each post-action variable X'_j that occurs in ADD for V^i : Multiply by Q and eliminate the prime variable.
- ▶ After eliminating all prime variables, we get $h(X_1 \dots X_n) = \sum_{x'_1 \dots x'_n} V^i(x'_1 \dots x'_n) P(x'_1 | X_1 \dots X_n) \dots P(x'_n | X_1 \dots X_n)$

SPUDD Algorithm

- ▶ Repeat the previous step for each post-action variable X'_j that occurs in ADD for V^i : Multiply by Q and eliminate the prime variable.
- ▶ After eliminating all prime variables, we get $h(X_1 \dots X_n) = \sum_{x'_1 \dots x'_n} V^i(x'_1 \dots x'_n) P(x'_1 | X_1 \dots X_n) \dots P(x'_n | X_1 \dots X_n)$
- ▶ $R + h$: ADD representation of V_a^{i+1}

SPUDD Algorithm

- ▶ Repeat the previous step for each post-action variable X'_j that occurs in ADD for V^i : Multiply by Q and eliminate the prime variable.
- ▶ After eliminating all prime variables, we get $h(X_1 \dots X_n) = \sum_{x'_1 \dots x'_n} V^i(x'_1 \dots x'_n) P(x'_1 | X_1 \dots X_n) \dots P(x'_n | X_1 \dots X_n)$
- ▶ $R + h$: ADD representation of V_a^{i+1}
- ▶ ADD for $V^{i+1} = \max_{a \in \mathcal{A}} V_a^{i+1}$

Summary and Conclusion

- ▶ SPUDD: VI for solving MDPs using ADDs.

Summary and Conclusion

- ▶ SPUDD: VI for solving MDPs using ADDs.
- ▶ ADDs capture regularities in the system dynamics, reward and value: Compact representation of the problem. (vs explicit matrix and decision tree methods)

Summary and Conclusion

- ▶ SPUDD: VI for solving MDPs using ADDs.
- ▶ ADDs capture regularities in the system dynamics, reward and value: Compact representation of the problem. (vs explicit matrix and decision tree methods)
- ▶ Drawback: Boolean variables only. (Multi-valued variables can be split into Boolean variables)

Summary and Conclusion

- ▶ SPUDD: VI for solving MDPs using ADDs.
- ▶ ADDs capture regularities in the system dynamics, reward and value: Compact representation of the problem. (vs explicit matrix and decision tree methods)
- ▶ Drawback: Boolean variables only. (Multi-valued variables can be split into Boolean variables)
- ▶ Ordering of variables fixed; dynamic ordering could reduce size. (Done by SDD)

Summary and Conclusion

- ▶ SPUDD: VI for solving MDPs using ADDs.
- ▶ ADDs capture regularities in the system dynamics, reward and value: Compact representation of the problem. (vs explicit matrix and decision tree methods)
- ▶ Drawback: Boolean variables only. (Multi-valued variables can be split into Boolean variables)
- ▶ Ordering of variables fixed; dynamic ordering could reduce size. (Done by SDD)
- ▶ Extensions to other dynamic programming algorithms.