

CS 152 Programming Languages

Probabilistic Programming & Probabilistic Programming Languages

Yizhou Zhang

University of Waterloo

What is a Probabilistic Program?

What is a Probabilistic Program?

Drawing samples

What is a Probabilistic Program?

Drawing samples

Conditioning

specifies samples that are **good**

What is a Probabilistic Program?


Drawing samples

→ Describes a distribution

Conditioning

specifies samples that are **good**


→ Describes a conditional distribution



What is a Probabilistic Program?



$$p(Z = z | X = x)$$



What is a Probabilistic Program?

Drawing samples

➔ Describes a distribution


Conditioning

specifies samples that are **good**

➔ Describes a conditional distribution

$$p(\mathbf{Z} = z \mid \mathbf{X} = x)$$

latent observed




What is a Probabilistic Program?



Bayes' Theorem

$$p(\mathbf{Z} = z \mid \mathbf{X} = x) = \frac{p(\mathbf{Z} = z, \mathbf{X} = x)}{p(\mathbf{X} = x)}$$

latent observed



What is a Probabilistic Program?

Drawing samples

➔ Describes a distribution


Conditioning

specifies samples that are **good**

➔ Describes a conditional distribution

Bayes' Theorem

$$\underbrace{p(z | x)}_{\text{posterior}} = \frac{\overbrace{p(z, x)}^{\text{joint}}}{\underbrace{p(x)}_{\text{evidence}}}$$



What is a Probabilistic Program?

Drawing samples

➔ Describes a distribution


Conditioning

specifies samples that are **good**

➔ Describes a conditional distribution

Bayes' Theorem

$$\begin{array}{c} p(z|x) \\ \text{posterior} \end{array} = \frac{\begin{array}{c} \text{joint} \\ p(z, x) \end{array}}{\begin{array}{c} p(x) \\ \text{evidence} \end{array}} = \frac{\begin{array}{c} \text{likelihood} \\ p(x|z) \end{array} \begin{array}{c} \text{prior} \\ p(z) \end{array}}{p(x)}$$



What is a Probabilistic Program?

Drawing samples

➔ Describes a distribution


Conditioning

specifies samples that are **good**

➔ Describes a conditional distribution


Bayes' Theorem

$$\underbrace{p(z|x)}_{\text{posterior}} = \frac{\underbrace{p(z,x)}_{\text{joint}}}{\underbrace{p(x)}_{\text{evidence}}} = \frac{\underbrace{p(x|z)}_{\text{likelihood}} \underbrace{p(z)}_{\text{prior}}}{\underbrace{\int p(z,x) dz}_{\text{marginal likelihood}}}$$




Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example



Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example



Aside on notation:

$$p(W = T | S = T, R = F)$$

probability mass

$$p(W | S = T, R = F)$$


distribution

$$p(W | S, R)$$

family of distributions

Example: Rain-Sprinkler-Grass

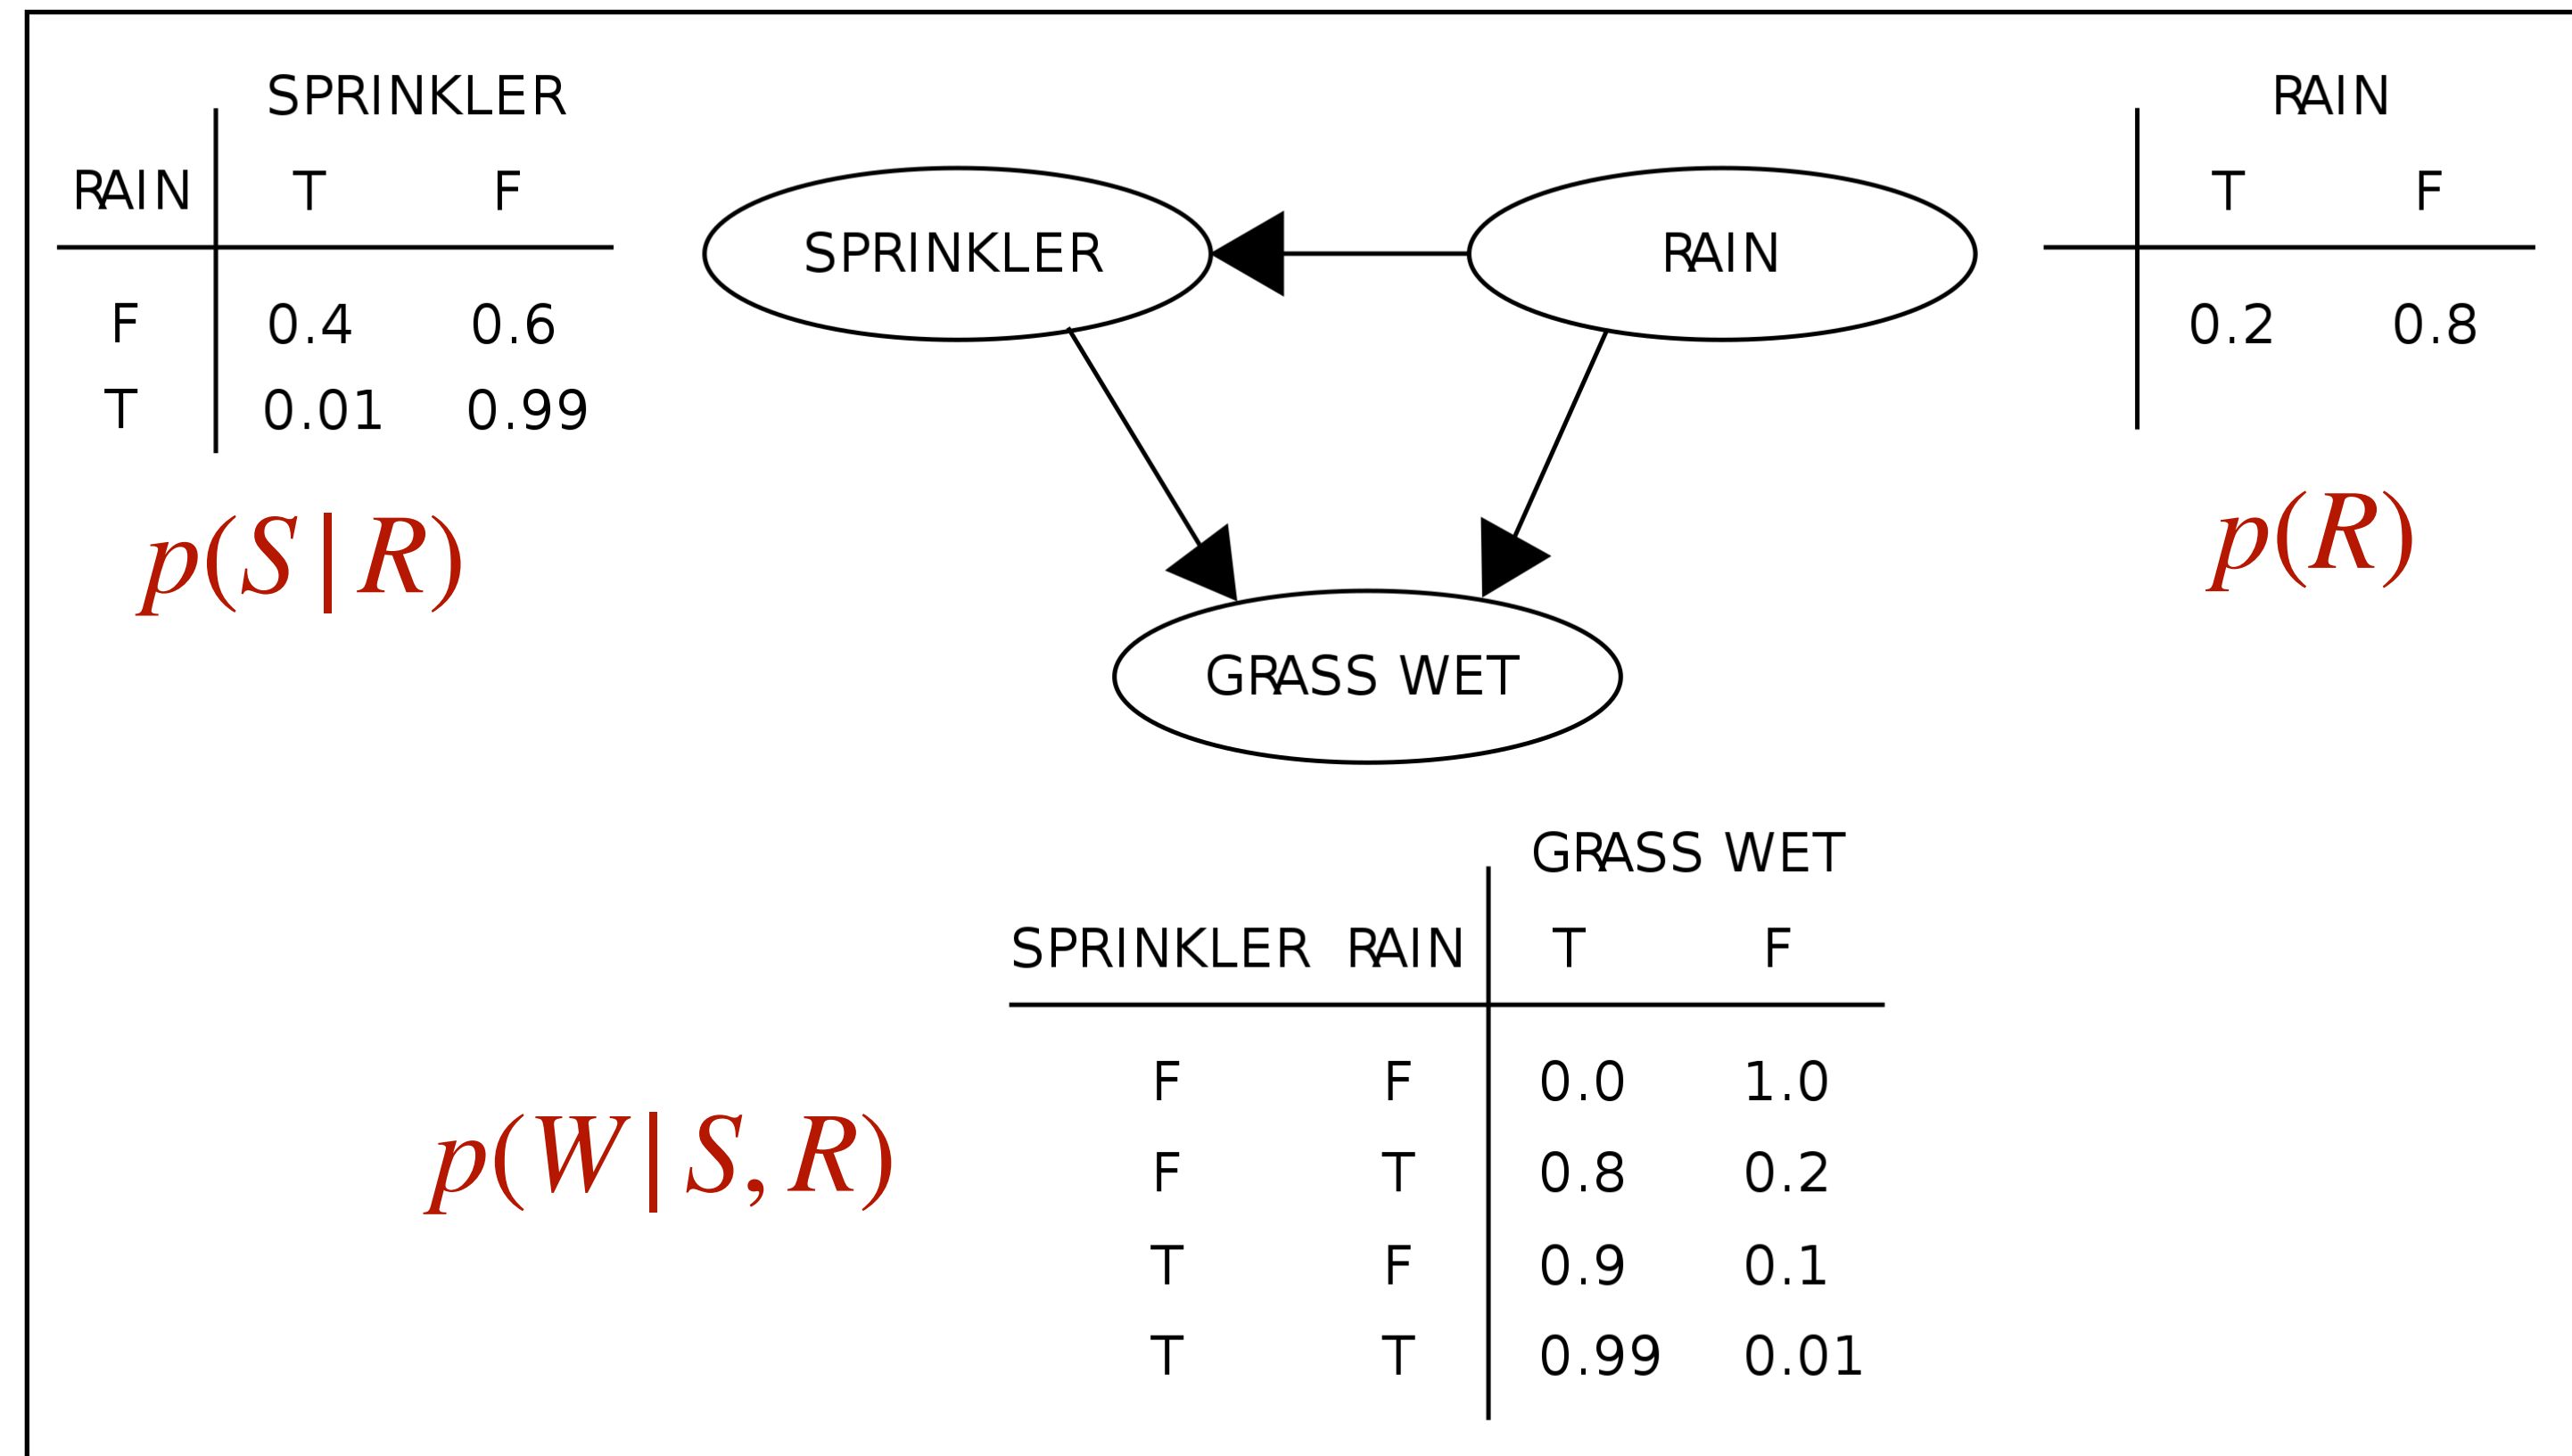
https://en.wikipedia.org/wiki/Bayesian_network#Example



Q1: Given that it rained, how likely is that the sprinkler was active?

Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example




Q1: Given that it rained, how likely is that the sprinkler was active?

Q2: Given that it rained, how likely is that the grass is wet?

$$\begin{aligned}
 & p(w | R=T) \\
 &= \sum_s p(w, s | R=T) \\
 &= \sum_s p(w | S=s, R=T) p(s | R=T) \\
 &= 0.99 \times 0.01 + 0.8 \times 0.99 \\
 &= 0.8019
 \end{aligned}$$

Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example



Q1: **Given that it rained**, how likely is that the sprinkler was active?

Q2: **Given that it rained**, how likely is that the grass is wet?

Q3: **Given that grass is wet**, how likely is that it rained?

Example: Rain-Sprinkler-Grass

```
var model = function() {
  // Pr(R)
  var r = sample(Bernoulli({p : 0.2}))

  // Pr(S|R=r)
  var s = sample(Bernoulli({p : r ? 0.01 : 0.4}))

  // Pr(W|R=r,S=s)
  var w = sample(Bernoulli({p :
    r ? (s ? 0.99 : 0.8) : (s ? 0.9 : 0.00)
  }))

  // condition model on W being true
  condition(w == true);

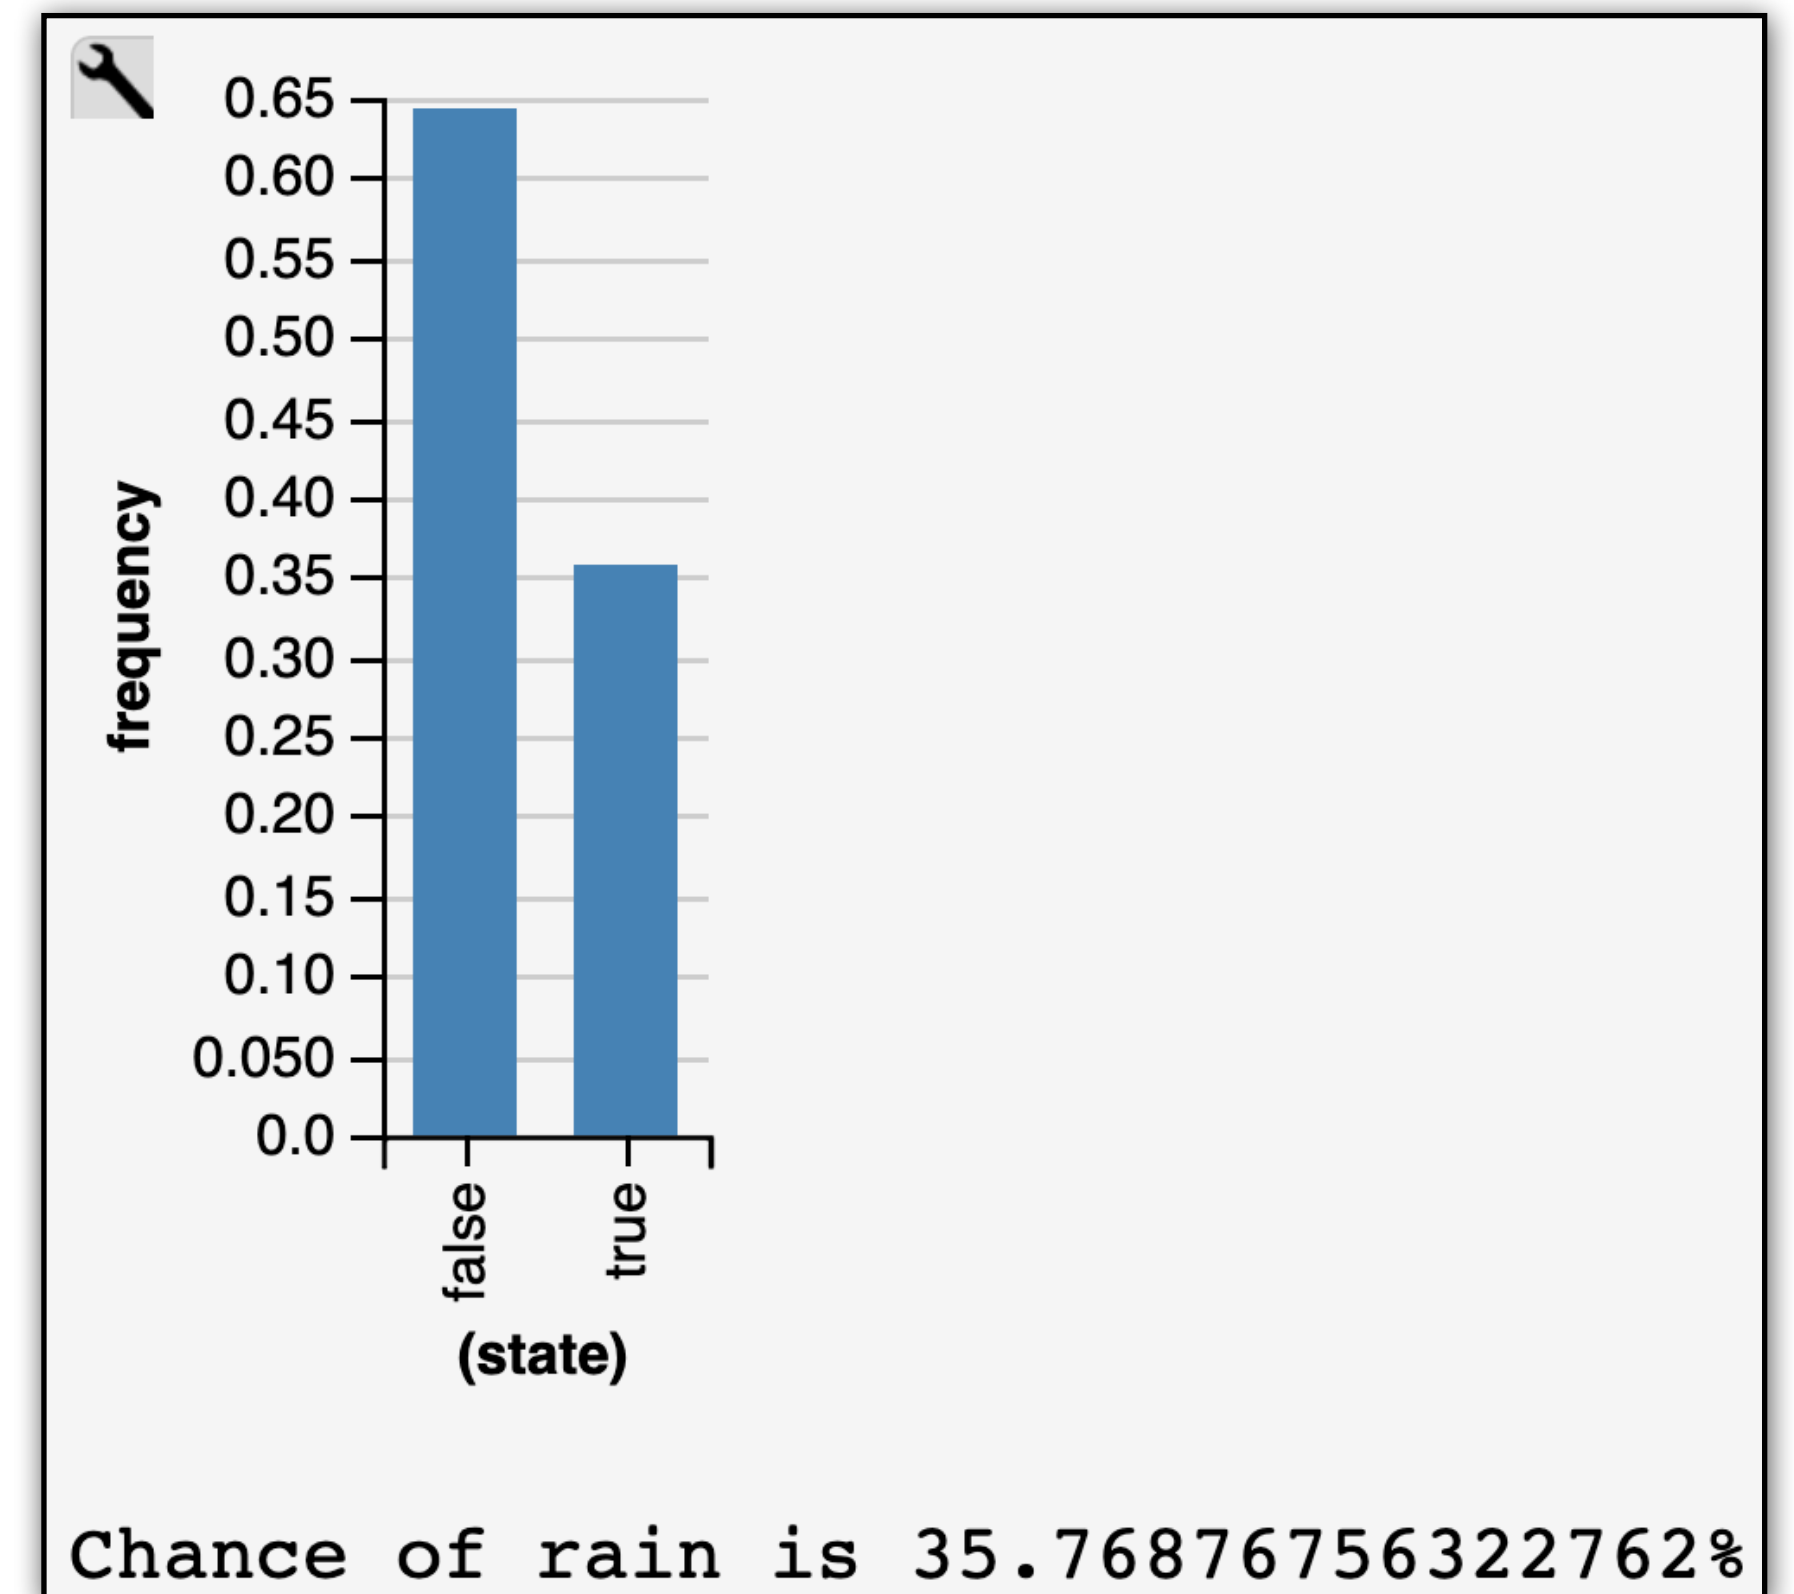
  return {R: r, S: s, W: w}
}

// apply Bayesian inference
var R_dist = Infer({
  method: 'enumerate',
  model: function() {
    var result = model()
    return result.R
  }
})
```

Q3: Given that grass is wet,
how likely is that it rained?

Example: Rain-Sprinkler-Grass


```
var model = function() {  
  // Pr(R)  
  var r = sample(Bernoulli({p : 0.2}))  
  
  // Pr(S|R=r)  
  var s = sample(Bernoulli({p : r ? 0.01 : 0.4}))  
  
  // Pr(W|R=r,S=s)  
  var w = sample(Bernoulli({p :  
    r ? (s ? 0.99 : 0.8) : (s ? 0.9 : 0.00)  
  })))  
  
  // condition model on W being true  
  condition(w == true);  
  
  return {R: r, S: s, W: w}  
}  
  
// apply Bayesian inference  
var R_dist = Infer({  
  method: 'enumerate',  
  model: function() {  
    var result = model()  
    return result.R  
  }  
})
```



Q3: Given that grass is wet, how likely is that it rained?

Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example



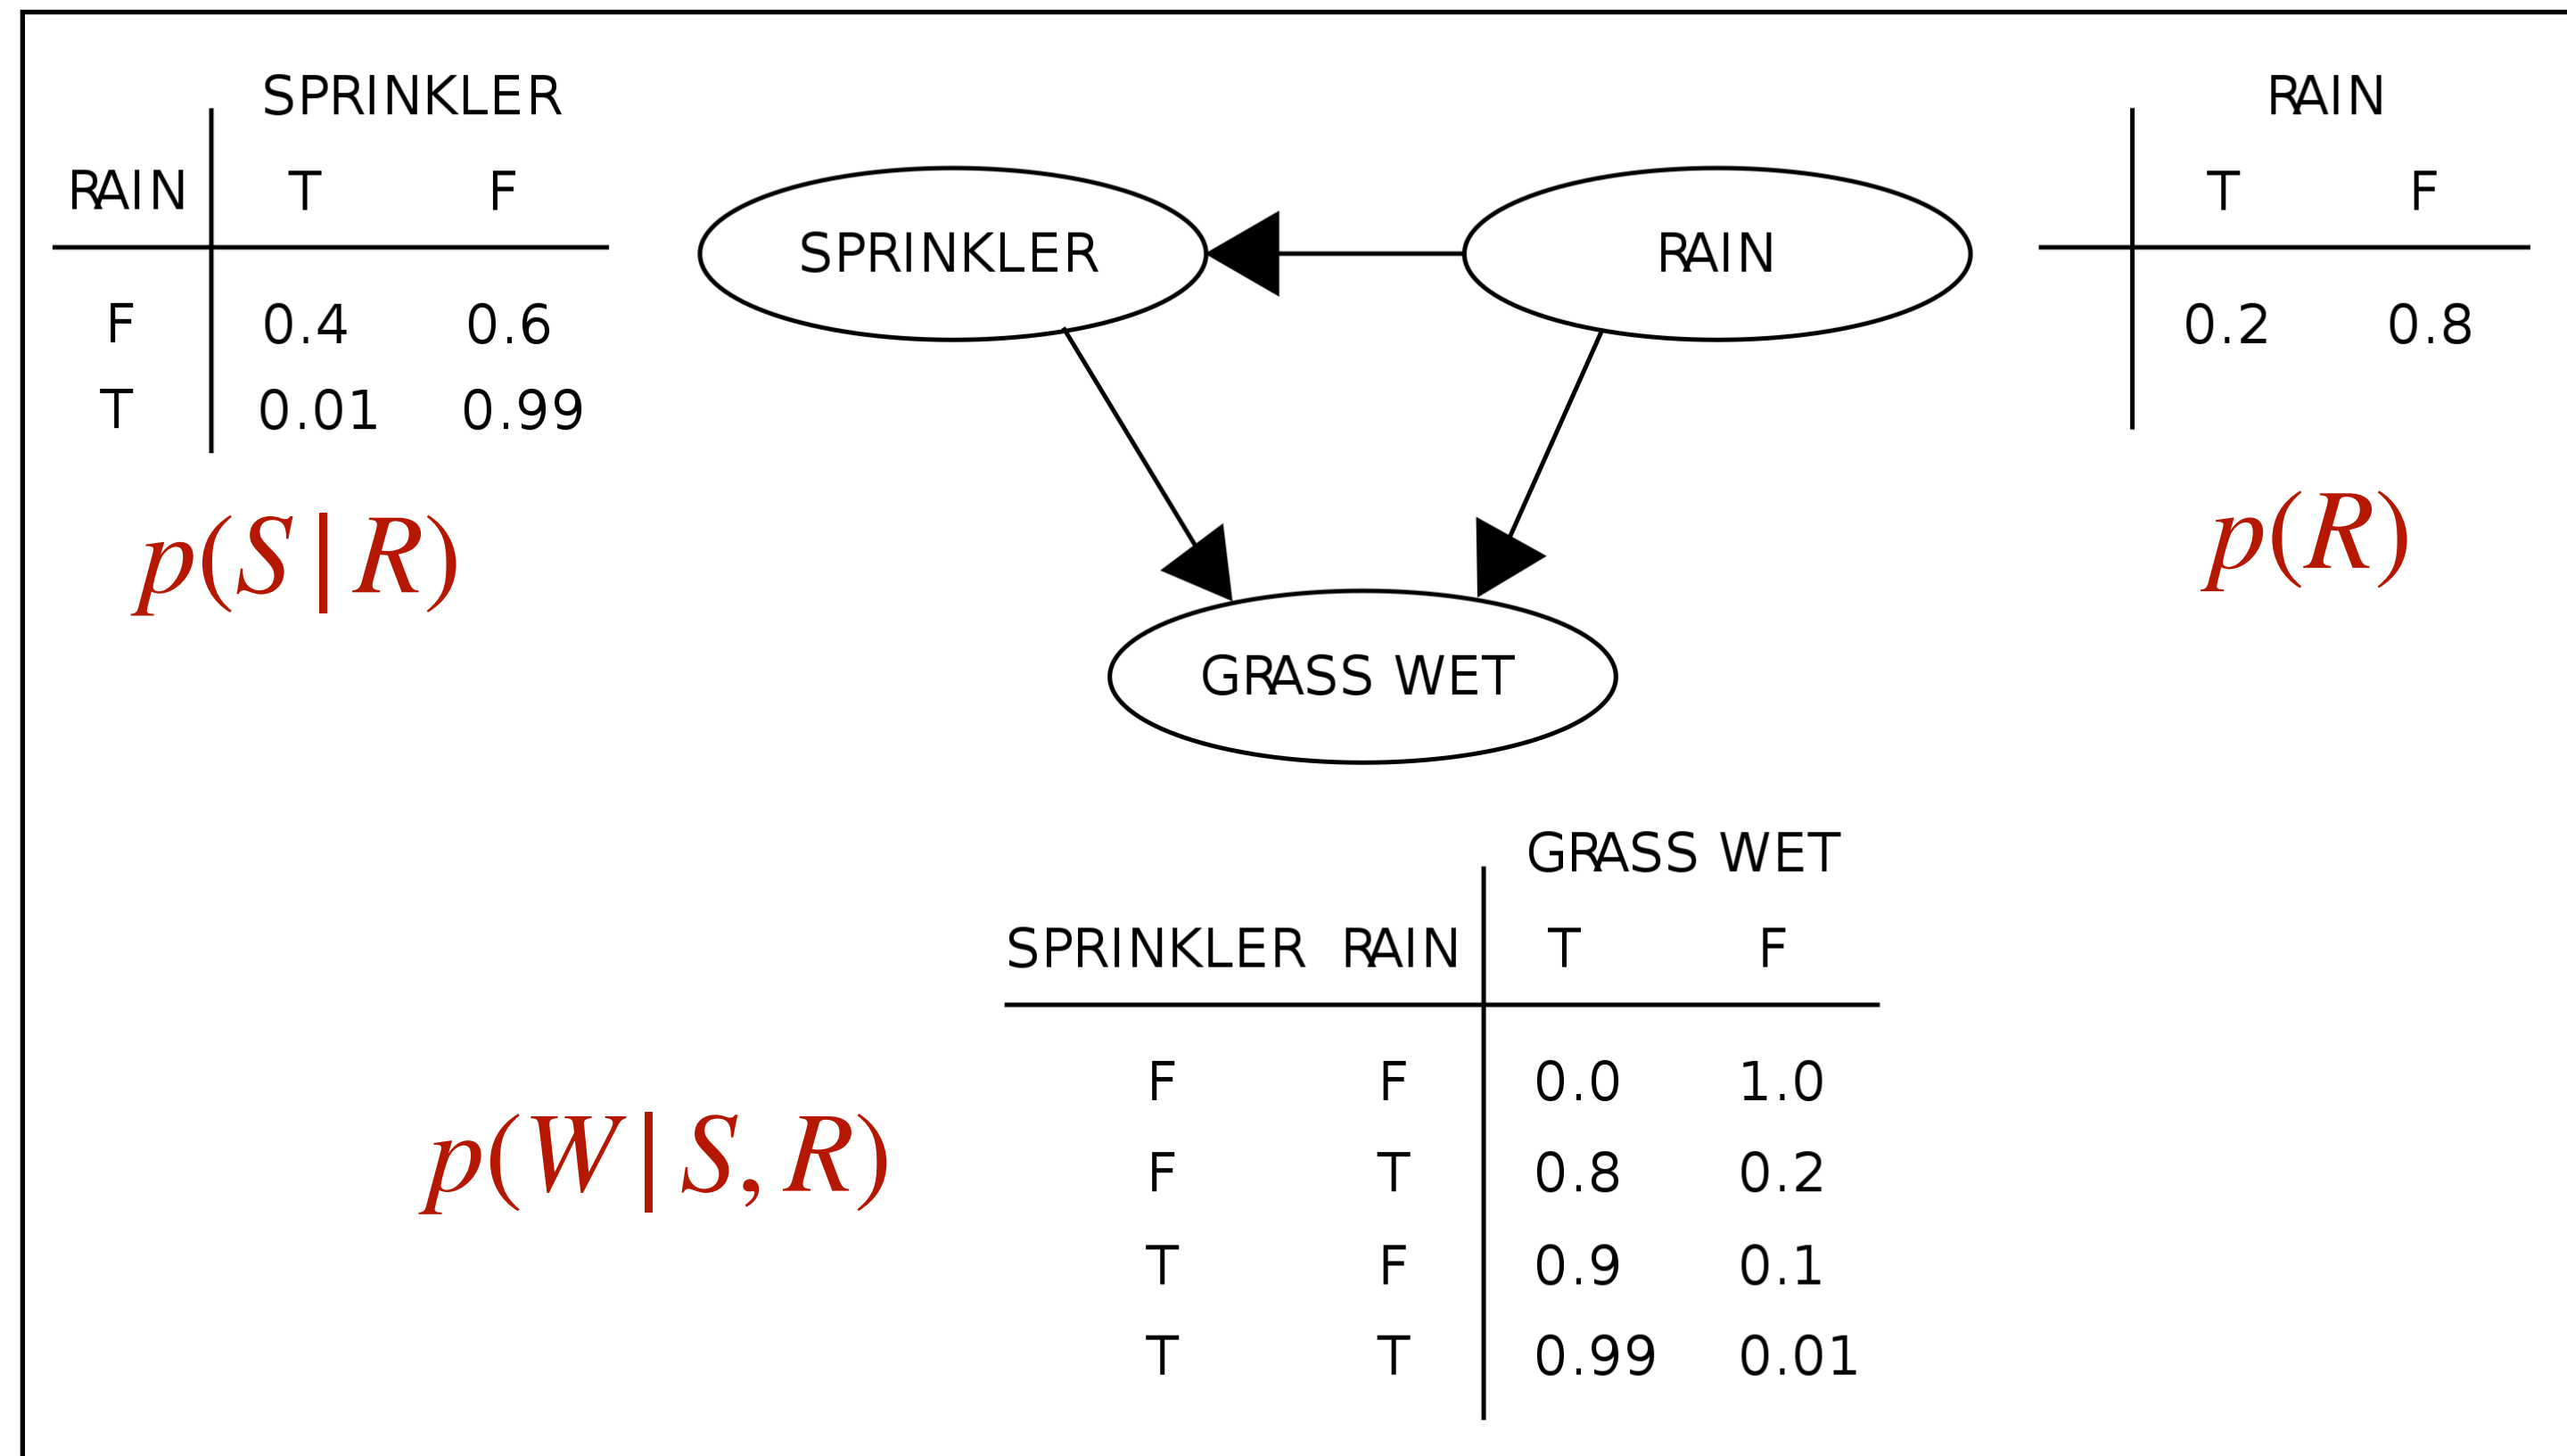
Q3: Given that grass is wet, how likely is that it rained?

associational

Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example

Ladder of Causation



Q3: Given that grass is wet, how likely is that it rained?


associational



Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example

Ladder of Causation



Q4: If we were to turn on the sprinkler, how likely would the grass be wet?

interventional

Q3: Given that grass is wet, how likely is that it rained?


associational



Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example

Ladder of Causation



Q5: Given that the sprinkler is active, had we turned off the sprinkler, how likely would the grass still be wet?
counterfactual

Q4: If we were to turn on the sprinkler, how likely would the grass be wet?
interventional

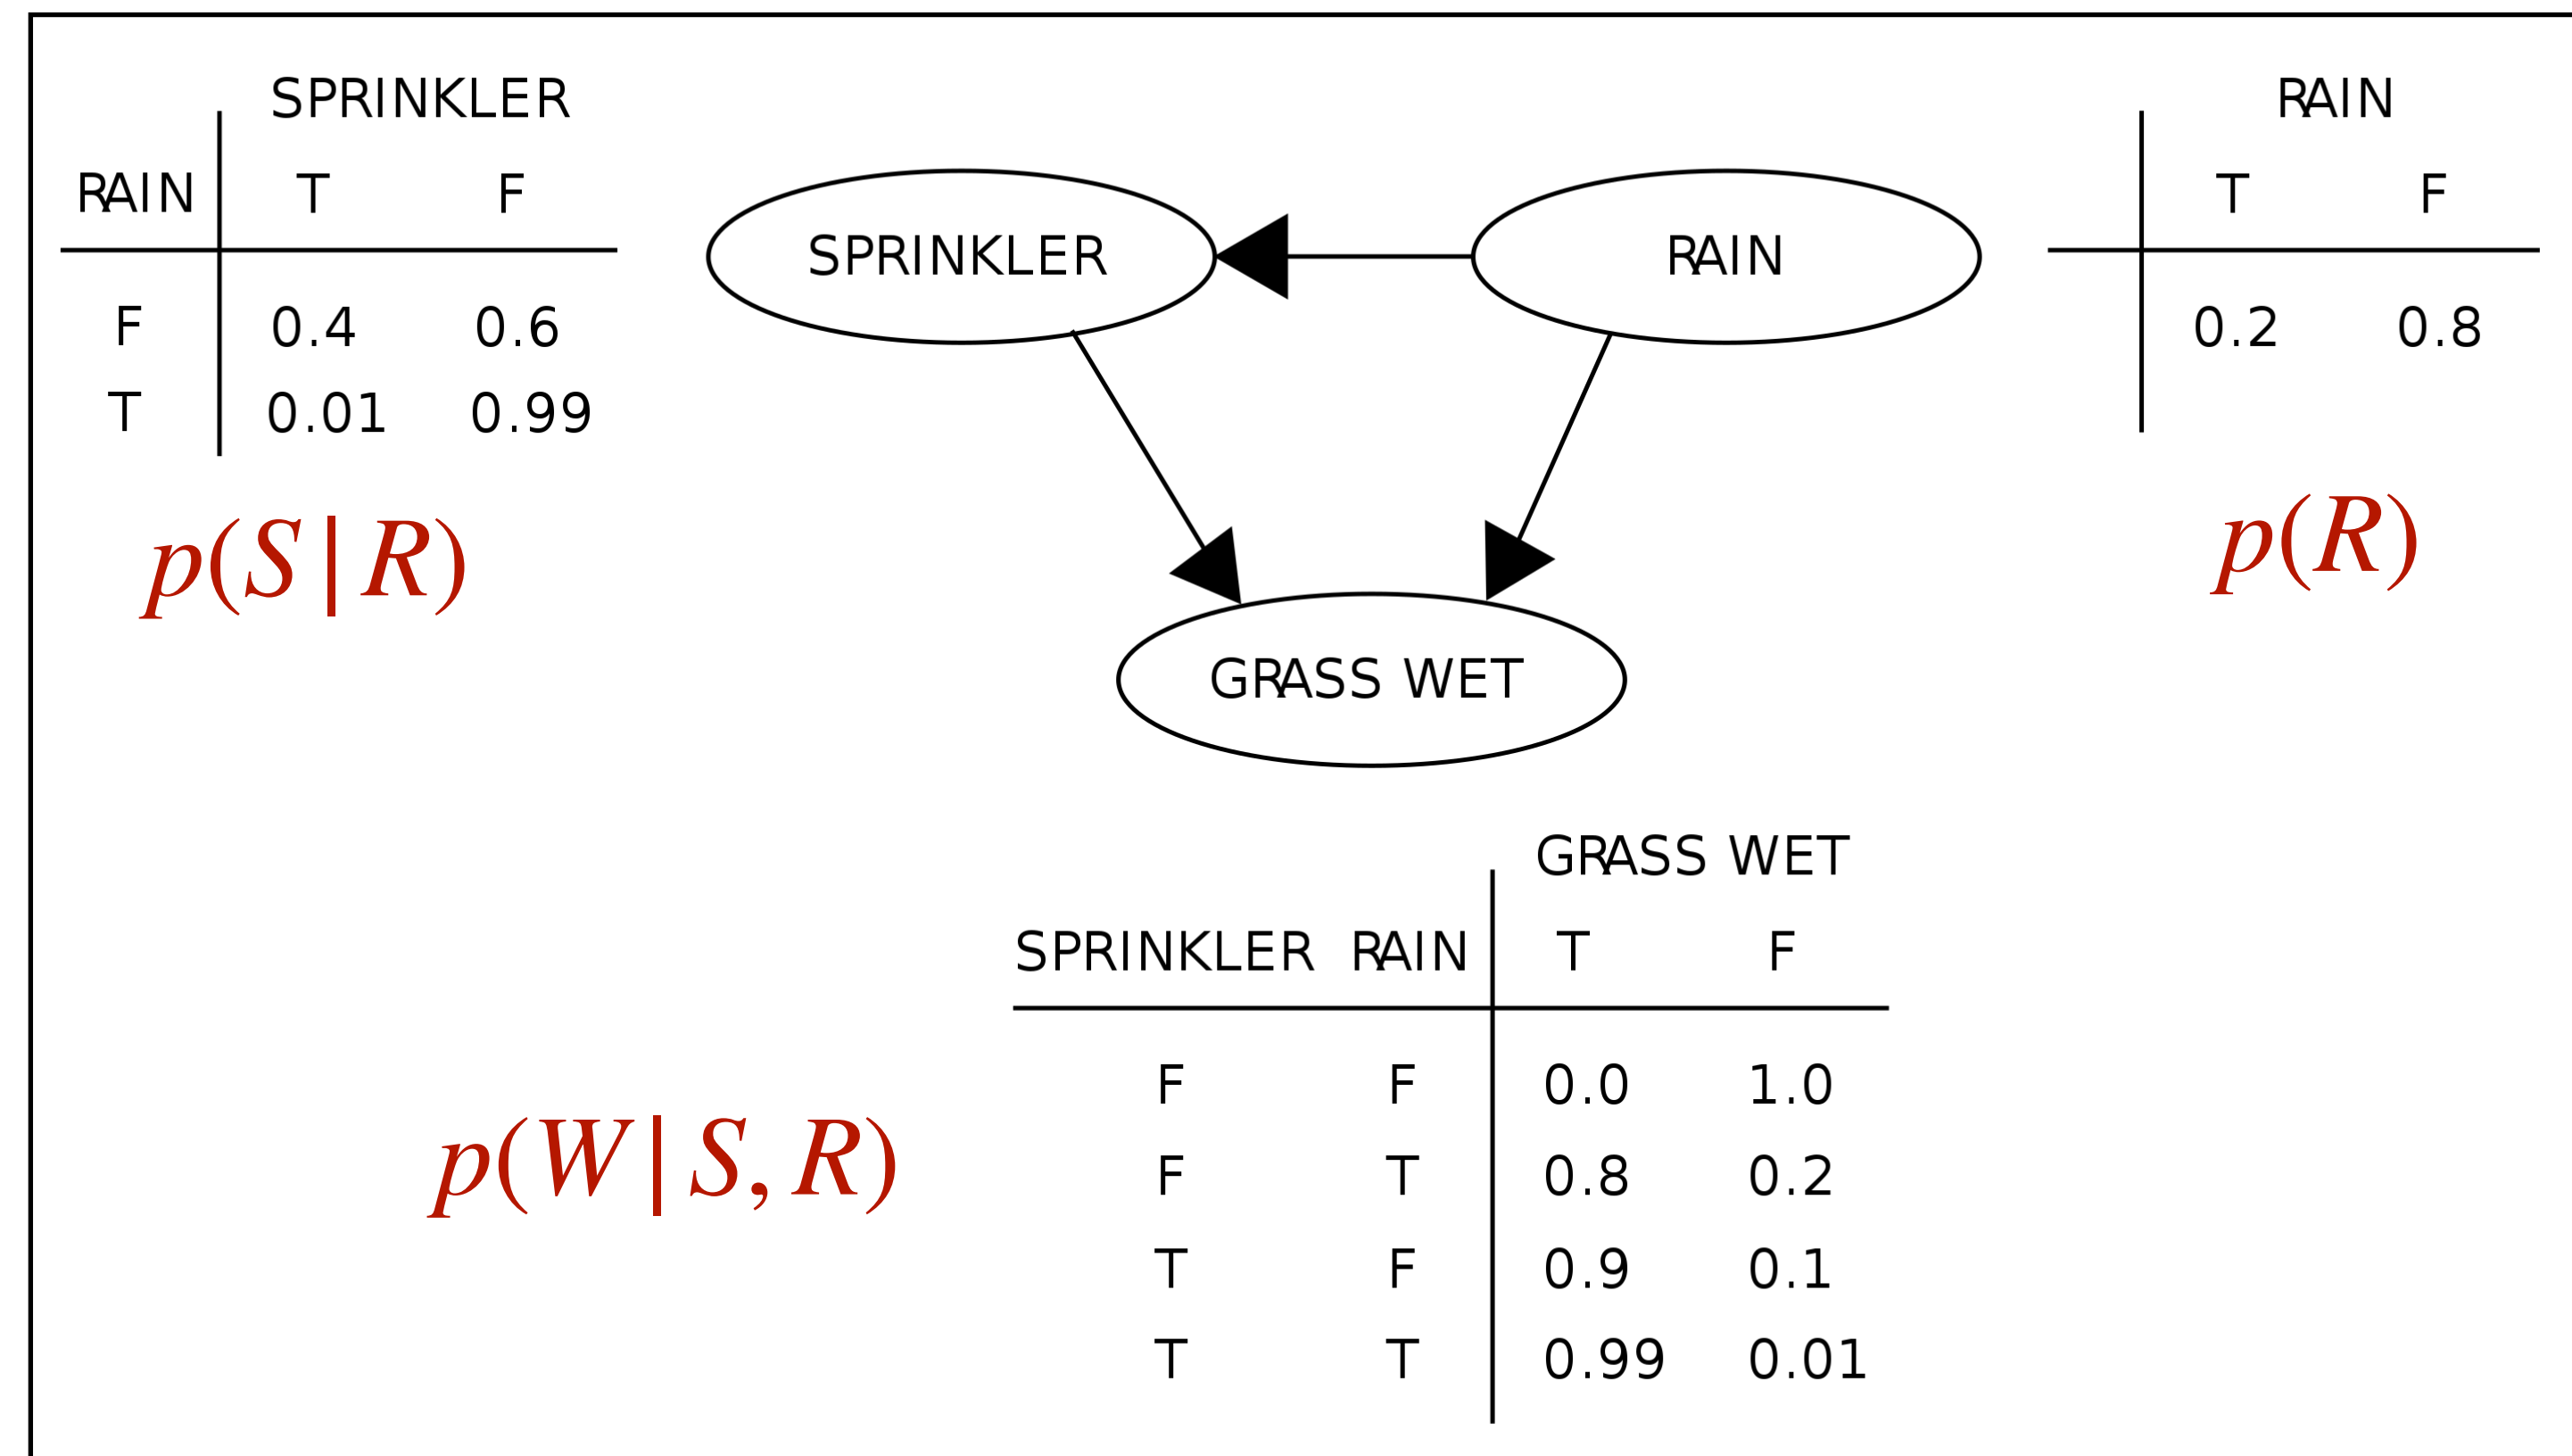
Q3: Given that grass is wet, how likely is that it rained?
associational



Example: Rain-Sprinkler-Grass

https://en.wikipedia.org/wiki/Bayesian_network#Example

Ladder of Causation




Q5: Given that the sprinkler is active, had we turned off the sprinkler, how likely would the grass still be wet?
counterfactual

Q4: If we were to turn on the sprinkler, how likely would the grass be wet?
interventional


Q3: Given that grass is wet, how likely is that it rained?
associational



Programming




Probabilistic Programming




$$p(z | x) = \frac{p(x | z) p(z)}{p(x)}$$

Example: Bouncing Balls

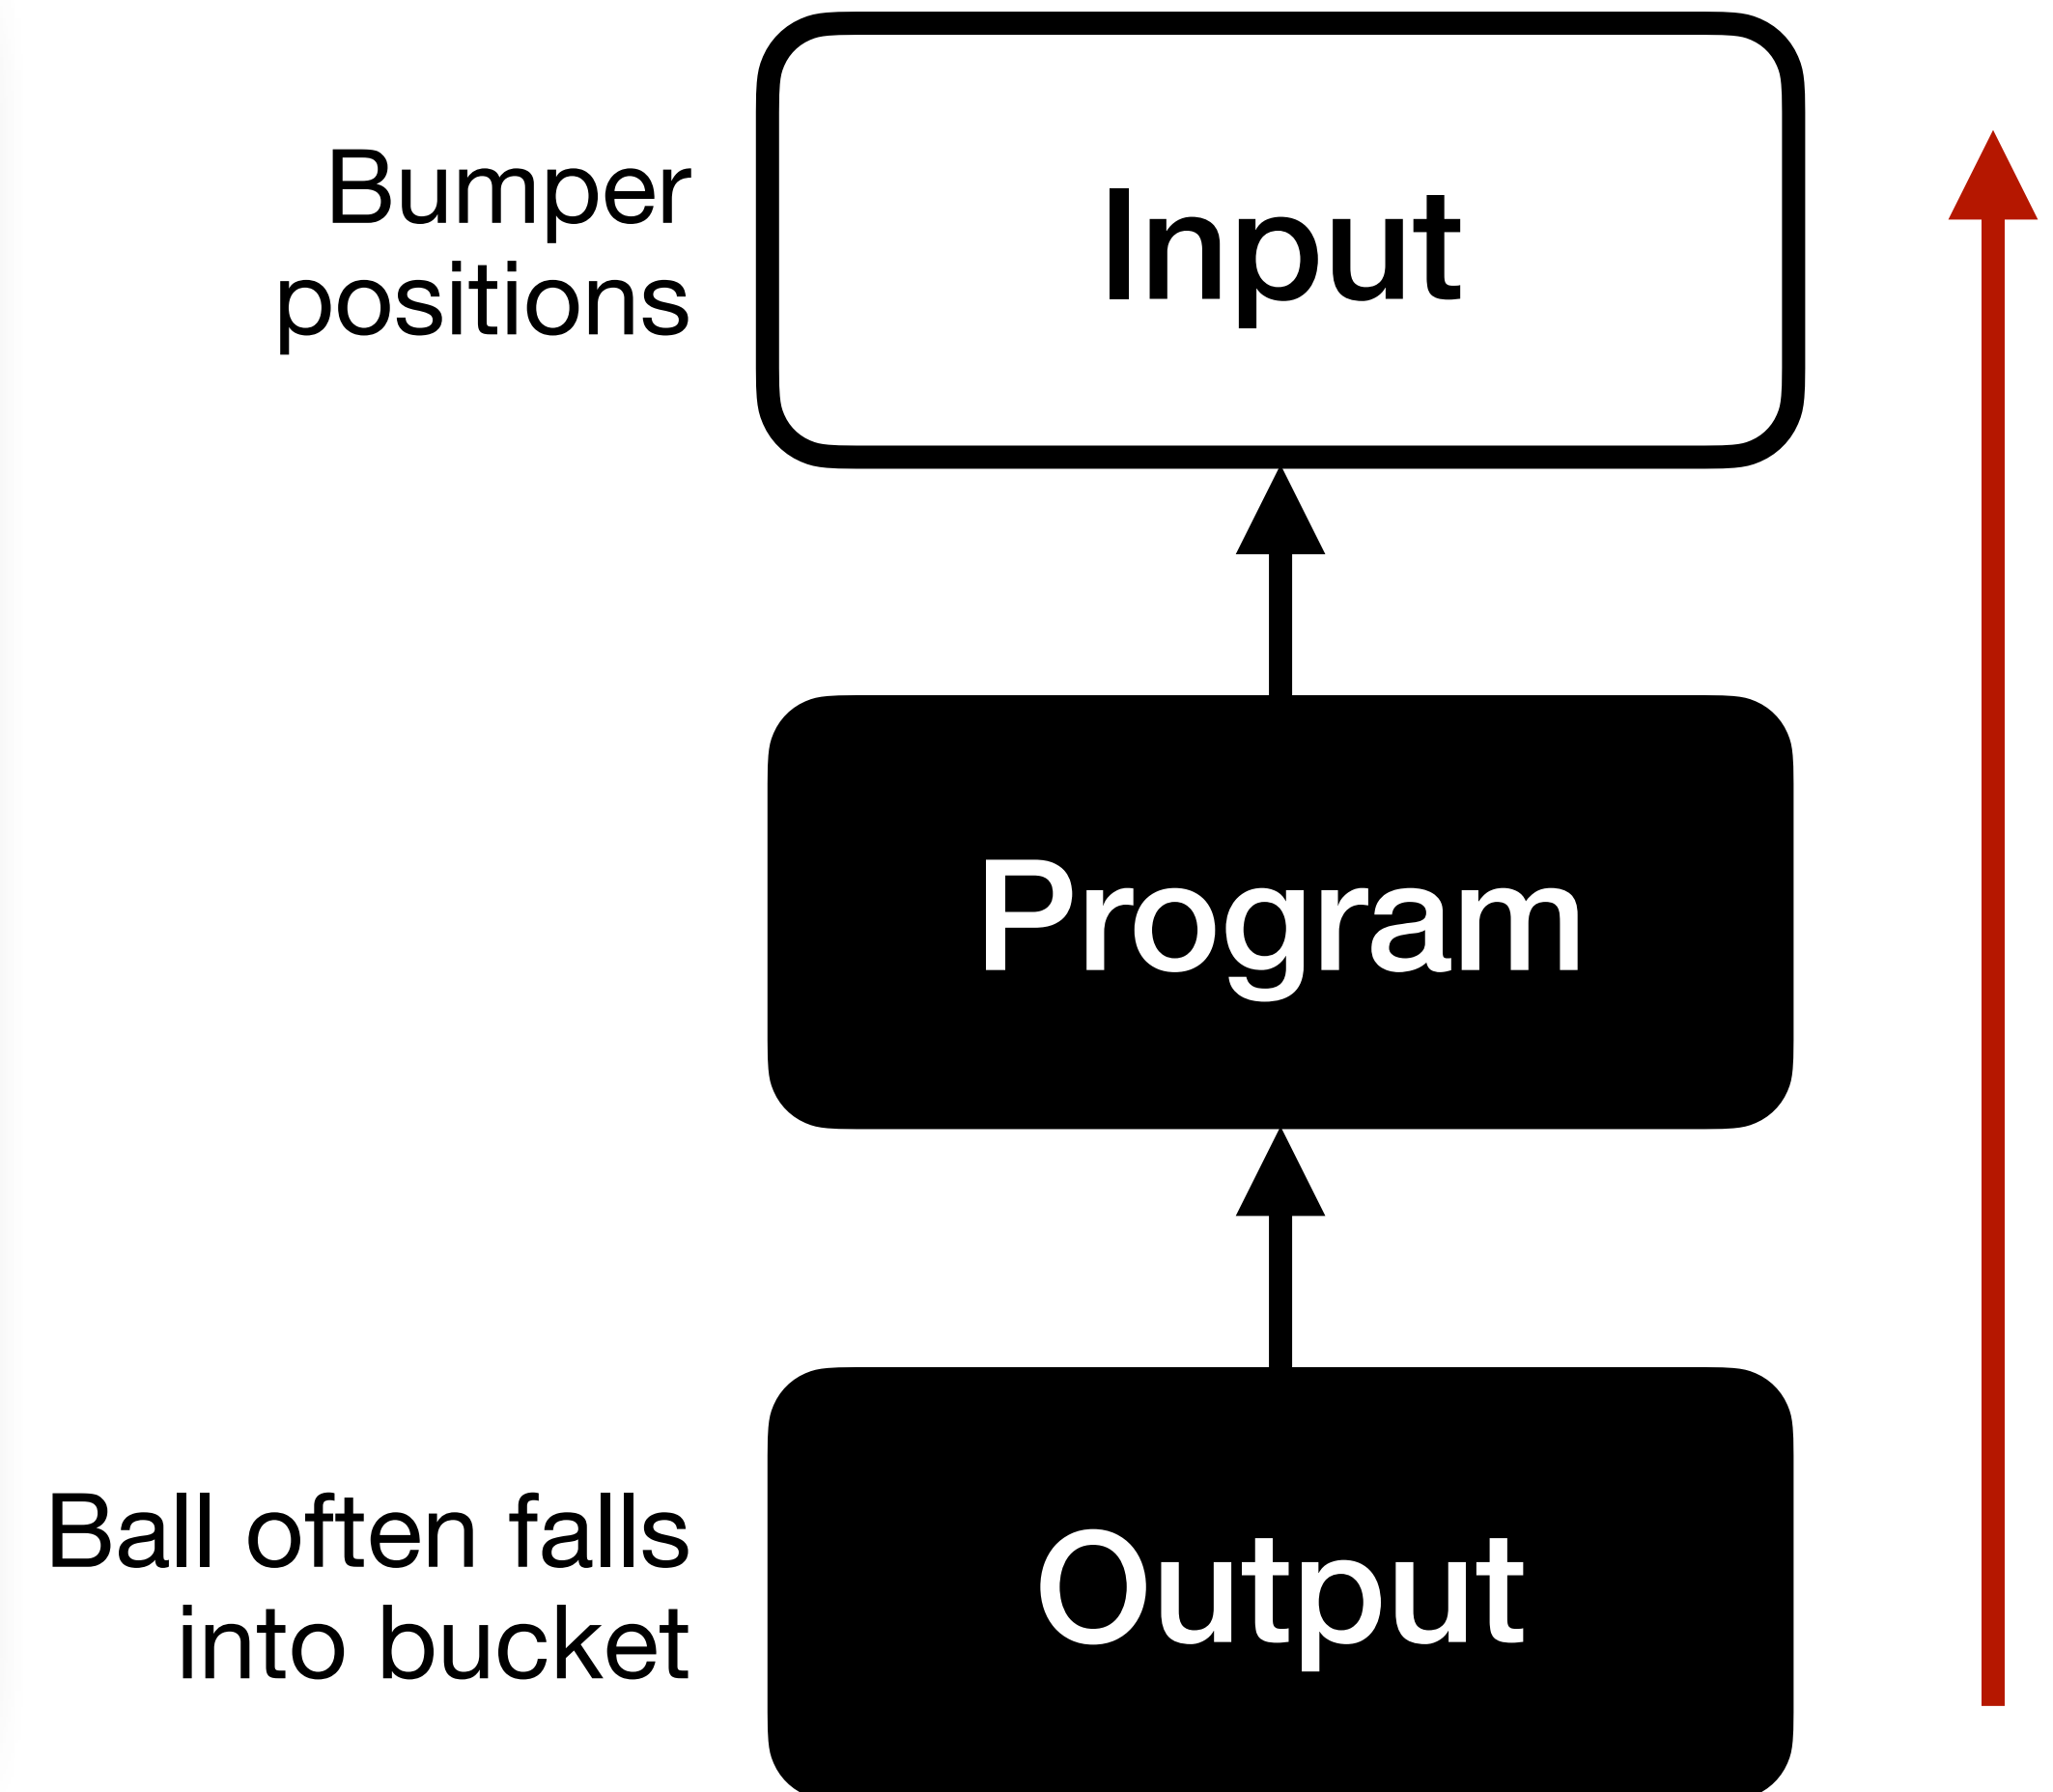



Ball rarely falls into bucket

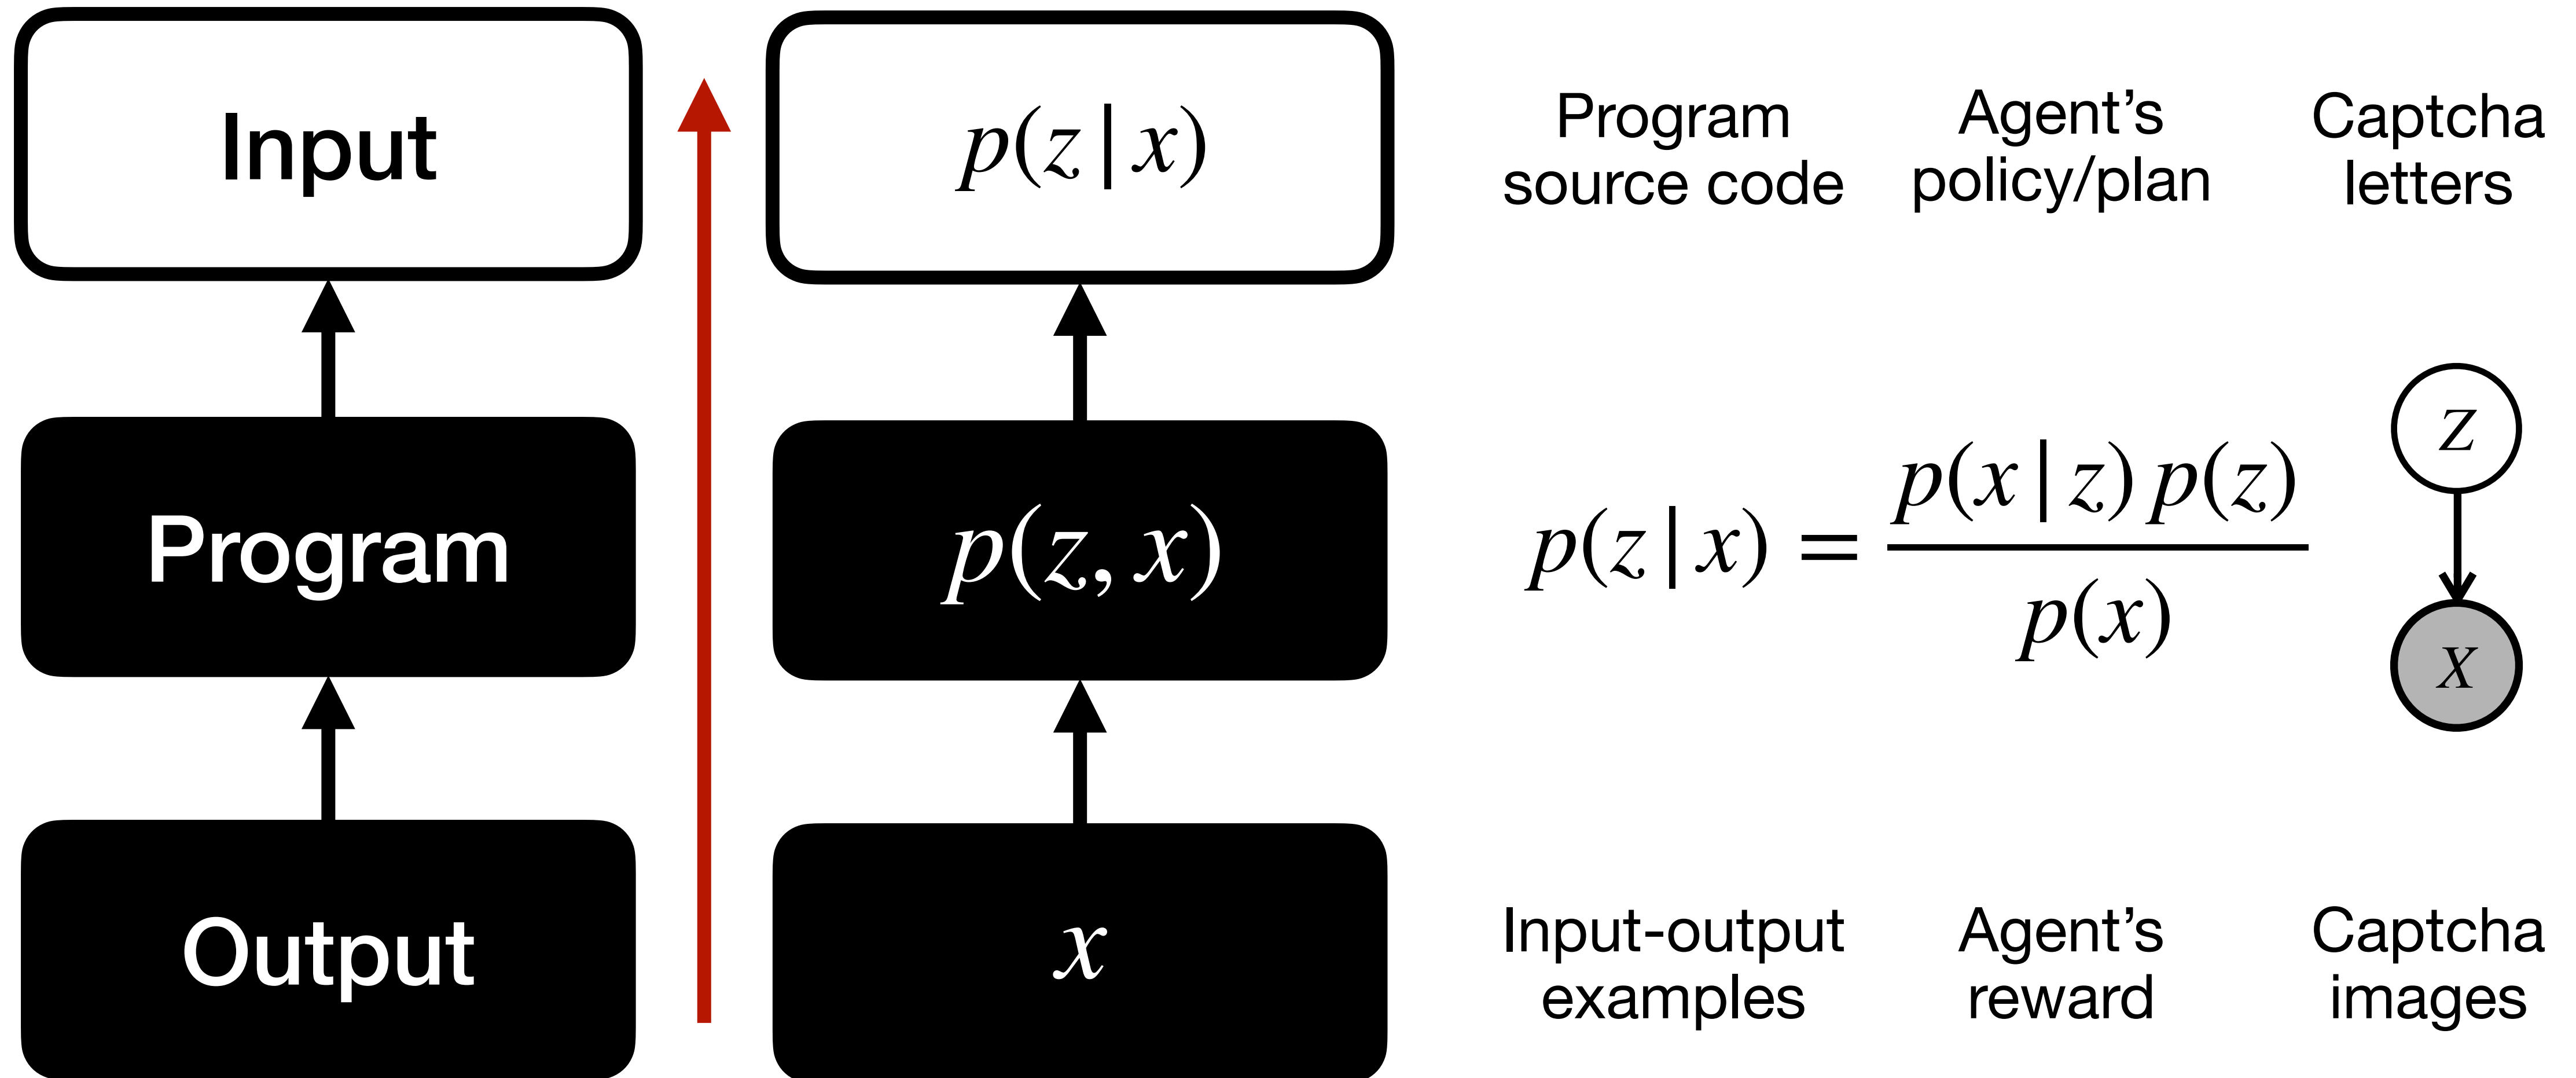
Bumper positions



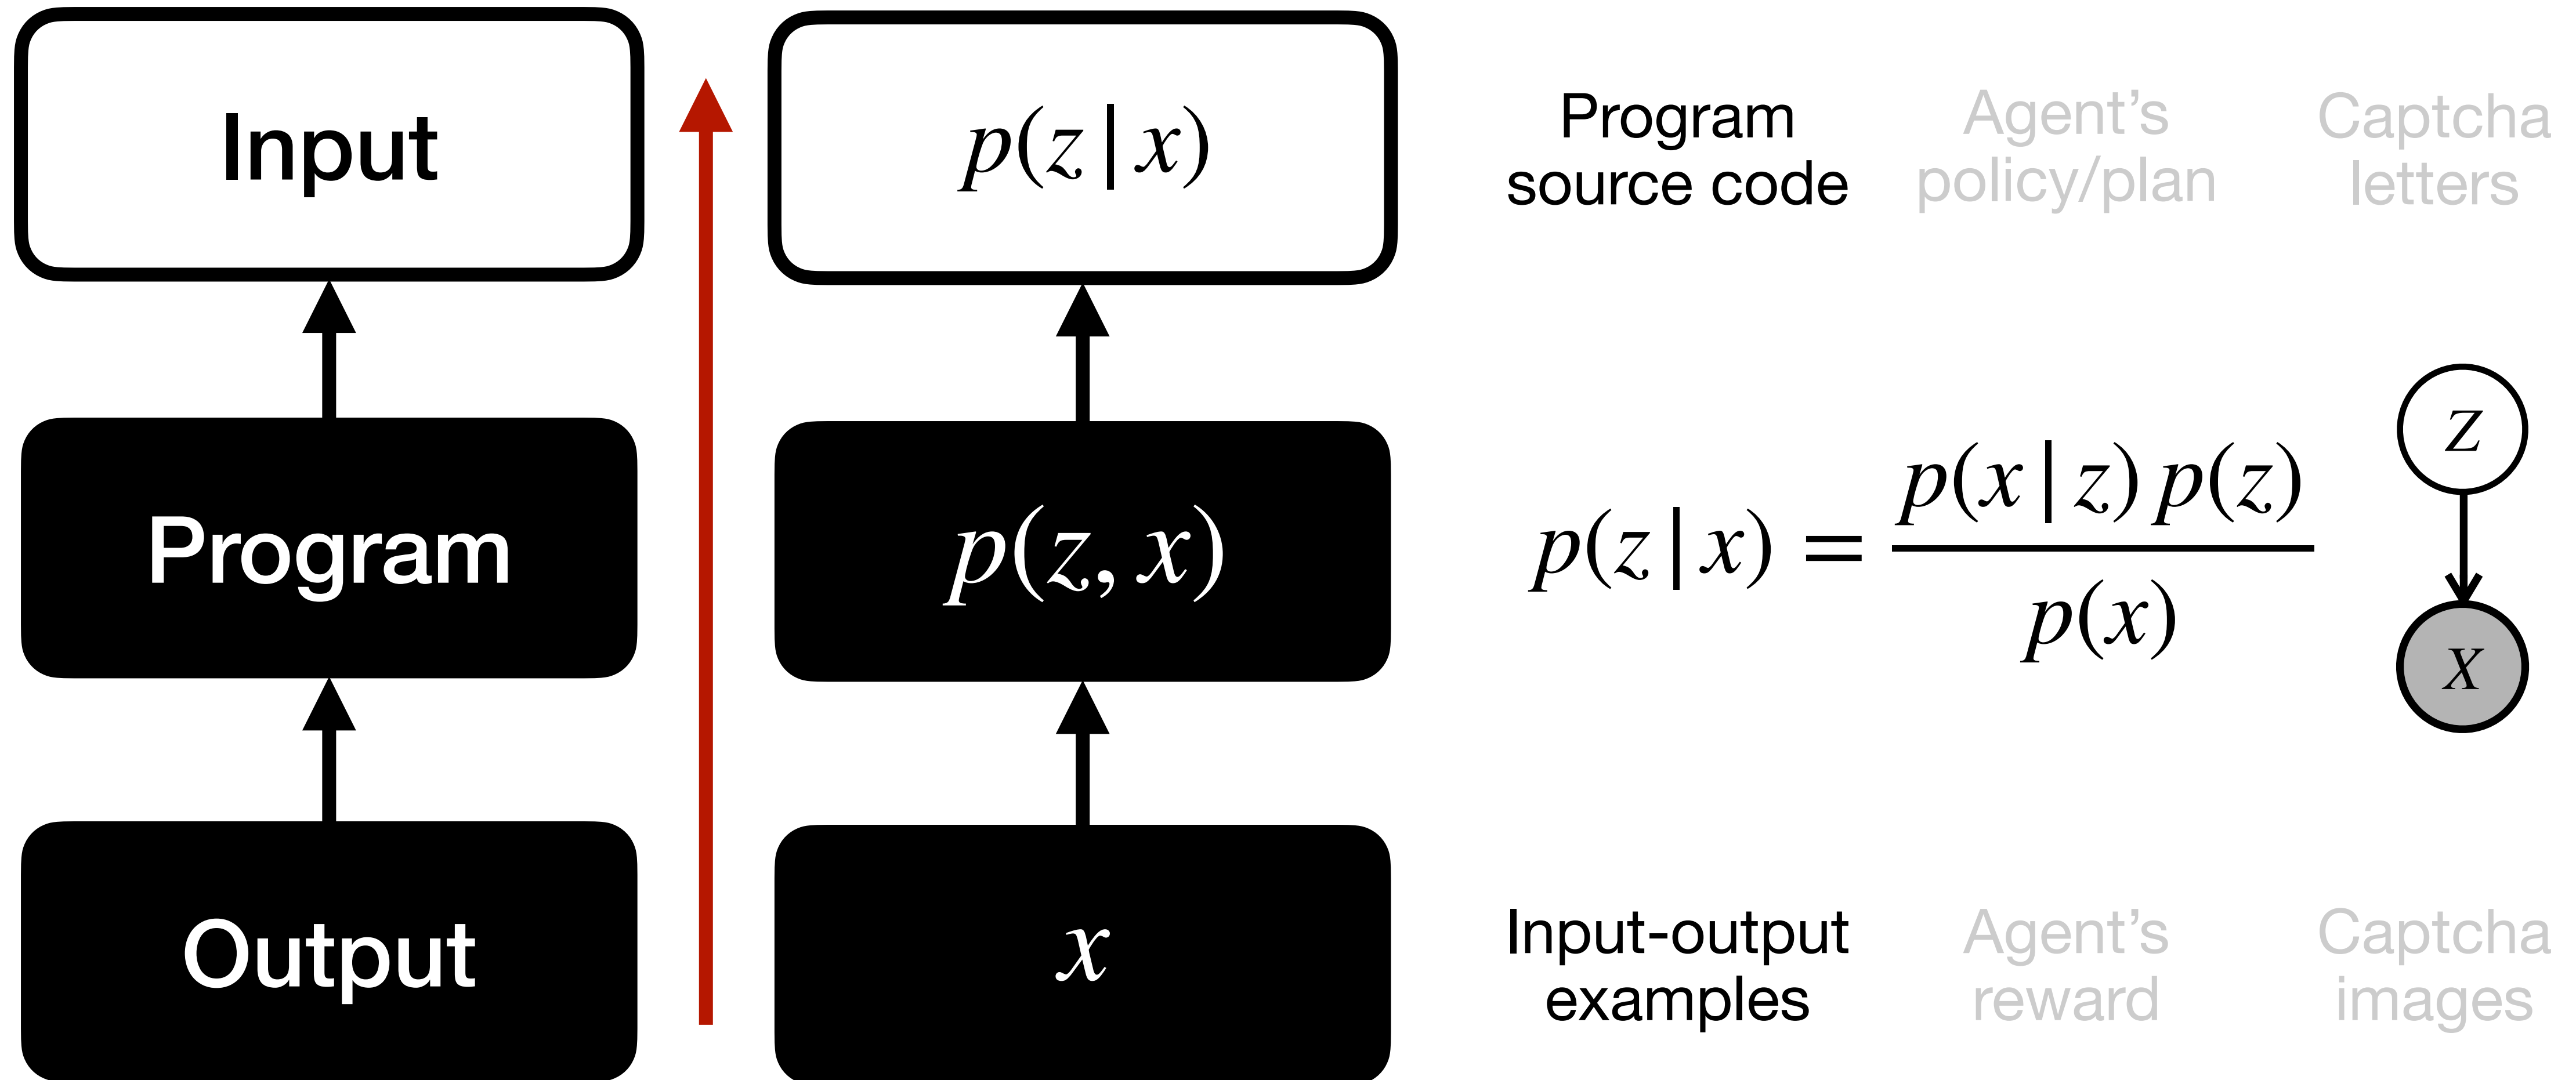
Example: Bouncing Balls into Bucket



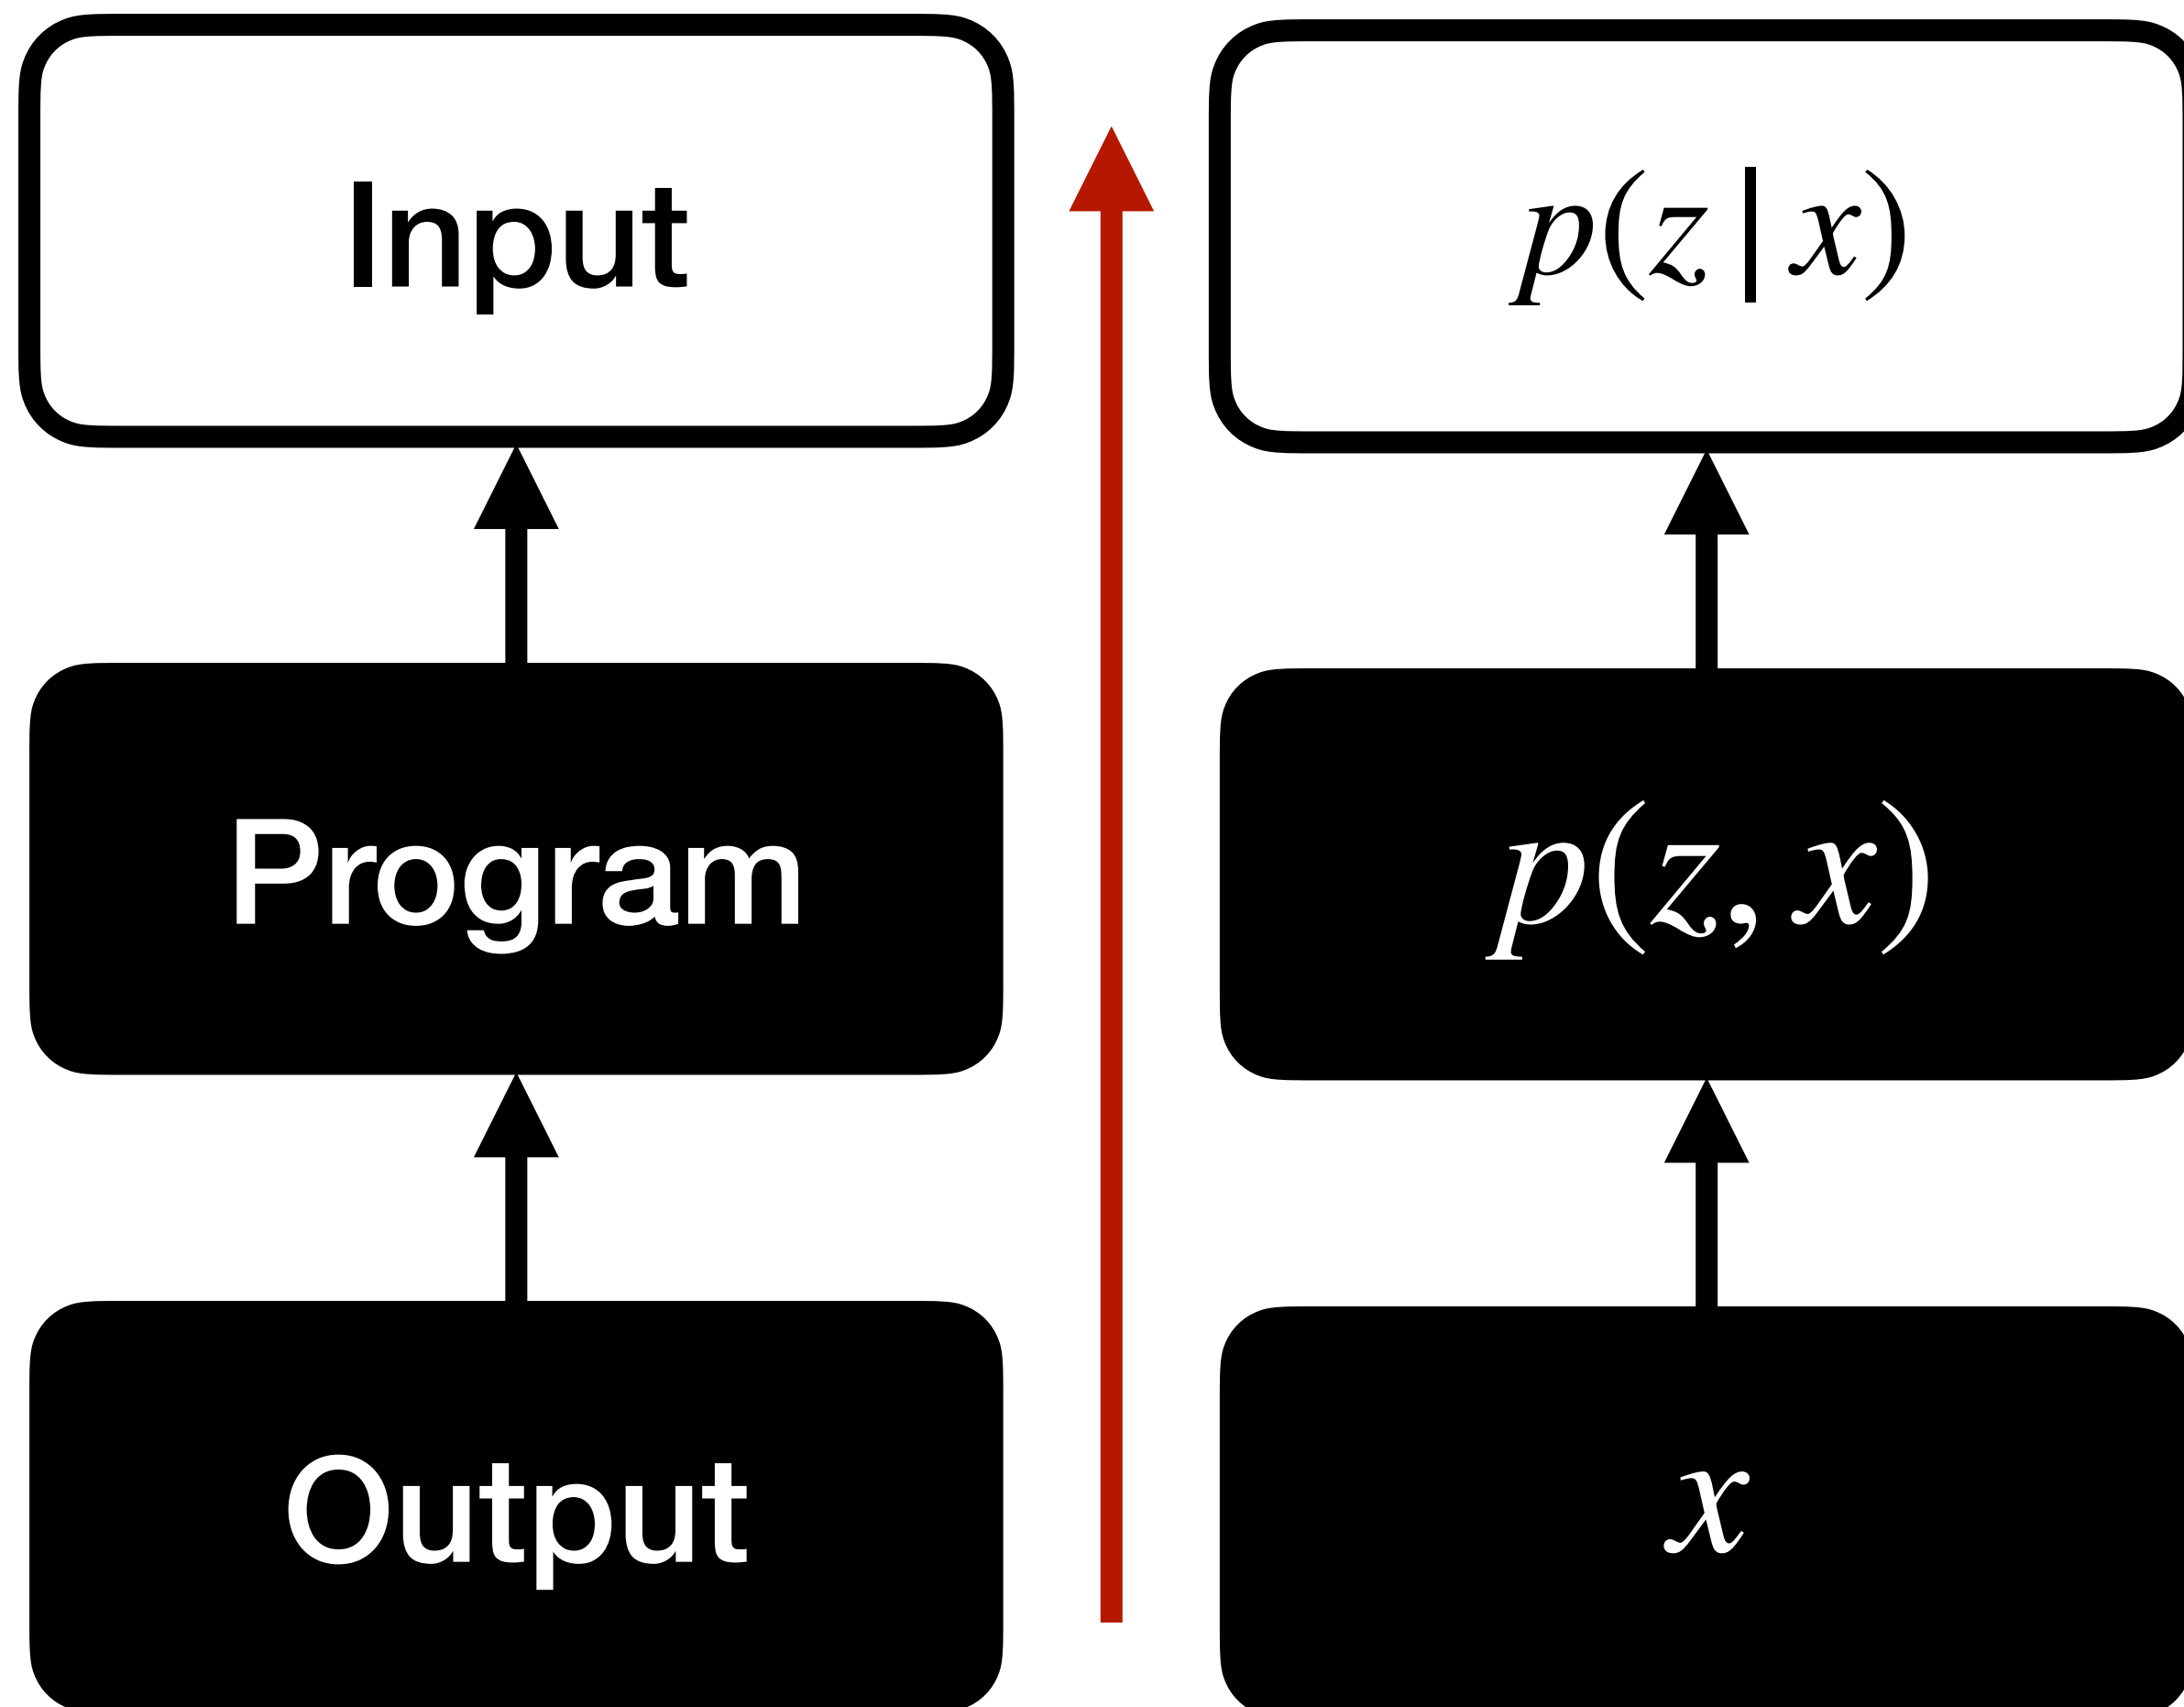

Applications of Probabilistic Programming



Applications of Probabilistic Programming



Applications of Probabilistic Programs




Program
source code

Agent's
policy/plan

Captcha
letters

$$p(z|x) = \frac{p(x|z)p(z)}{p(x)}$$

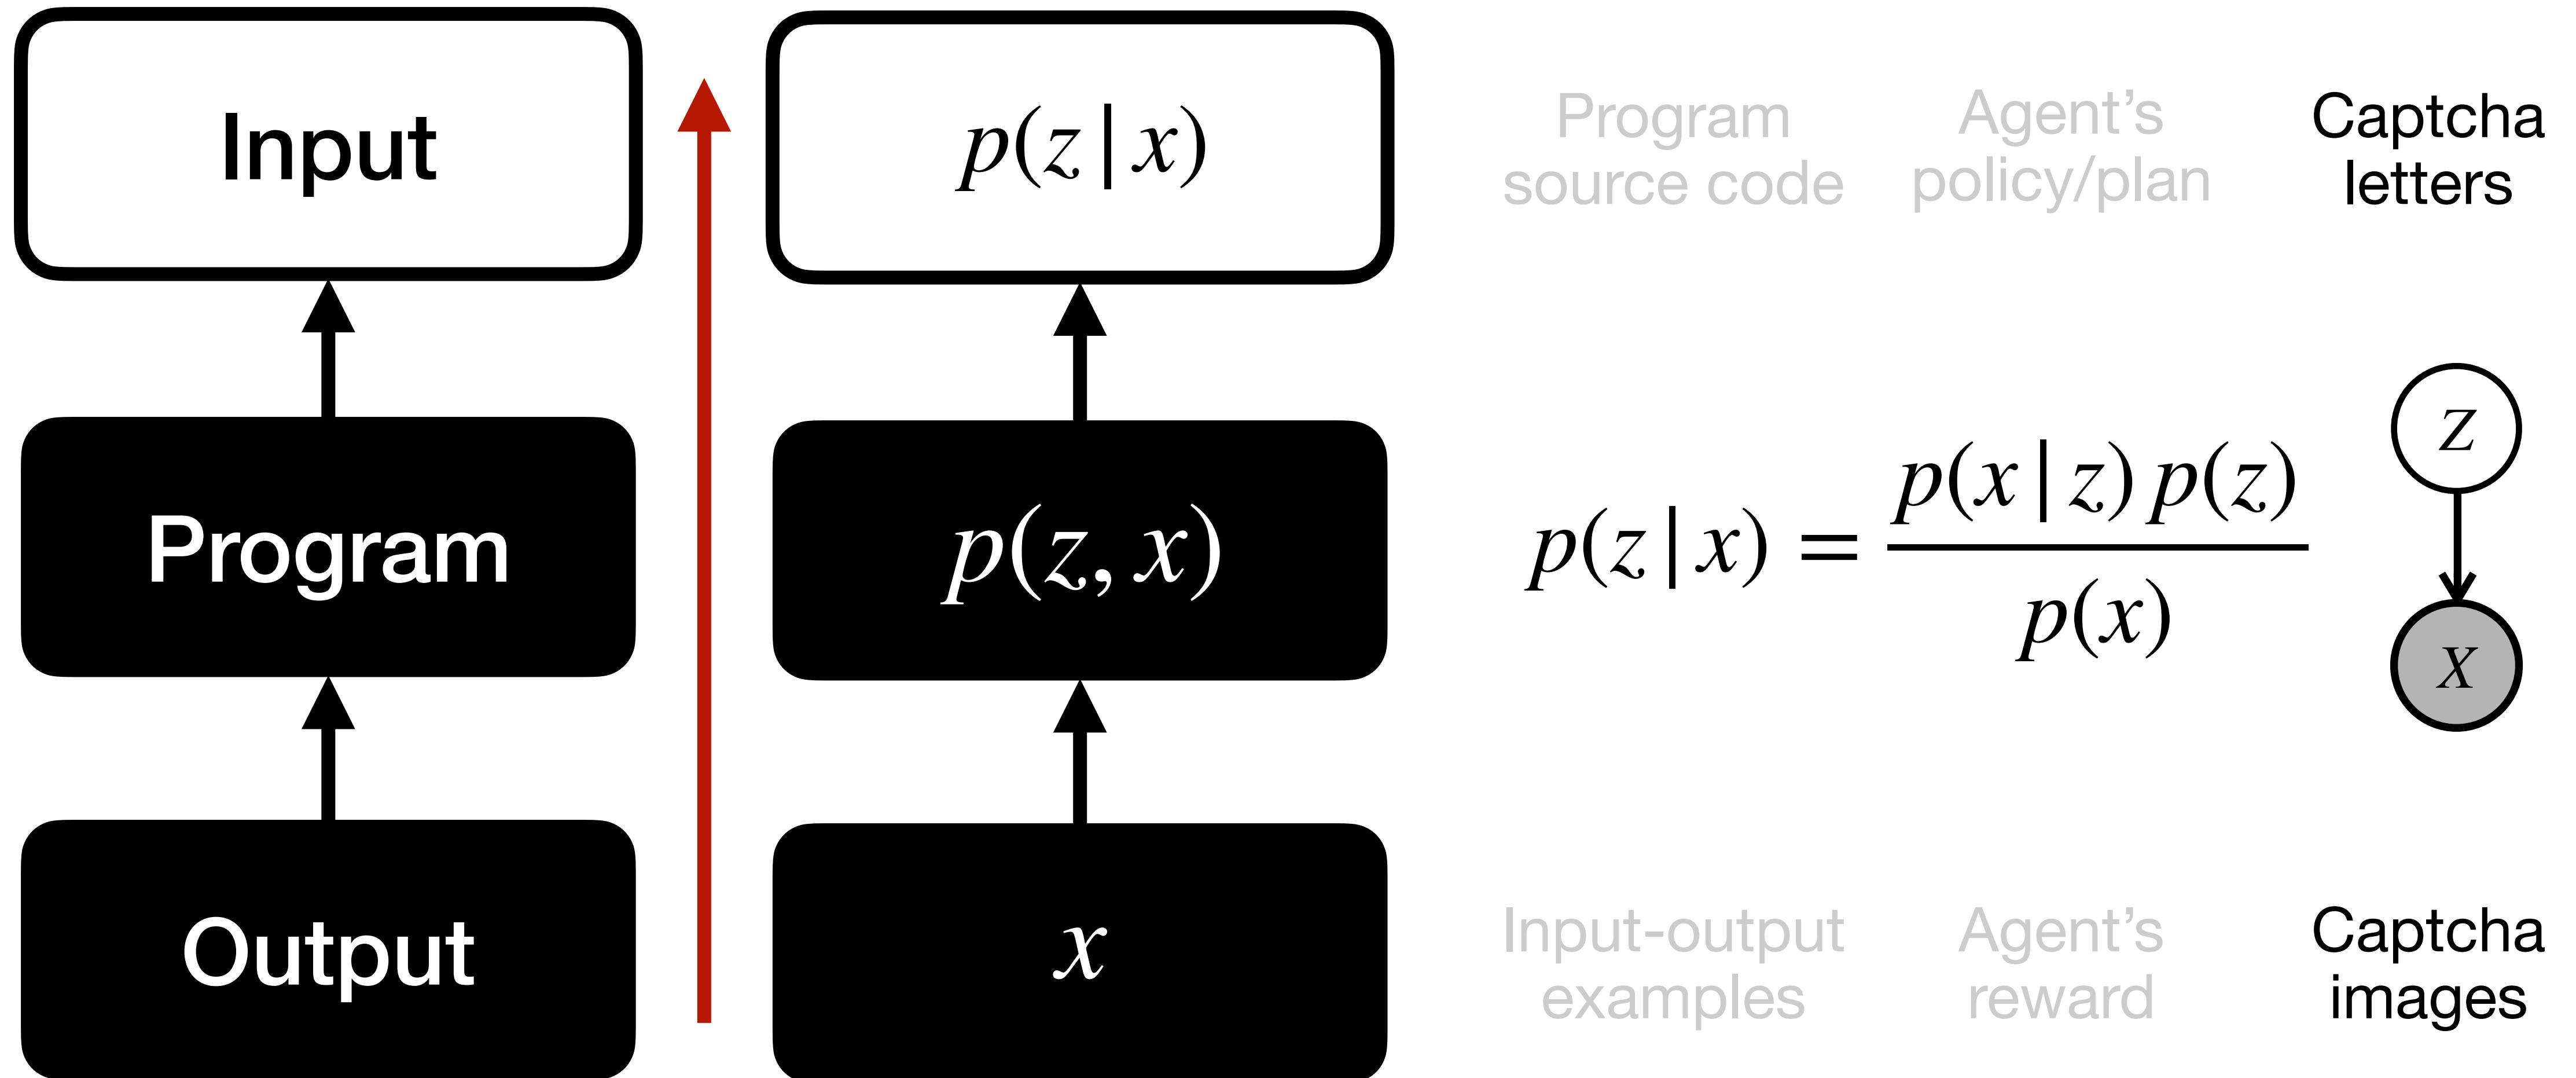


Input-output
examples

Agent's
reward

Captcha
images

Applications of Probabilistic Programming



PL as an Abstraction Layer

Programs



PL

Interpreter / Compiler

PL as an Abstraction Layer


Programs

expression

PL

Interpreter / Compiler

solving



PPL as an Abstraction Layer

Probabilistic Programs

expression

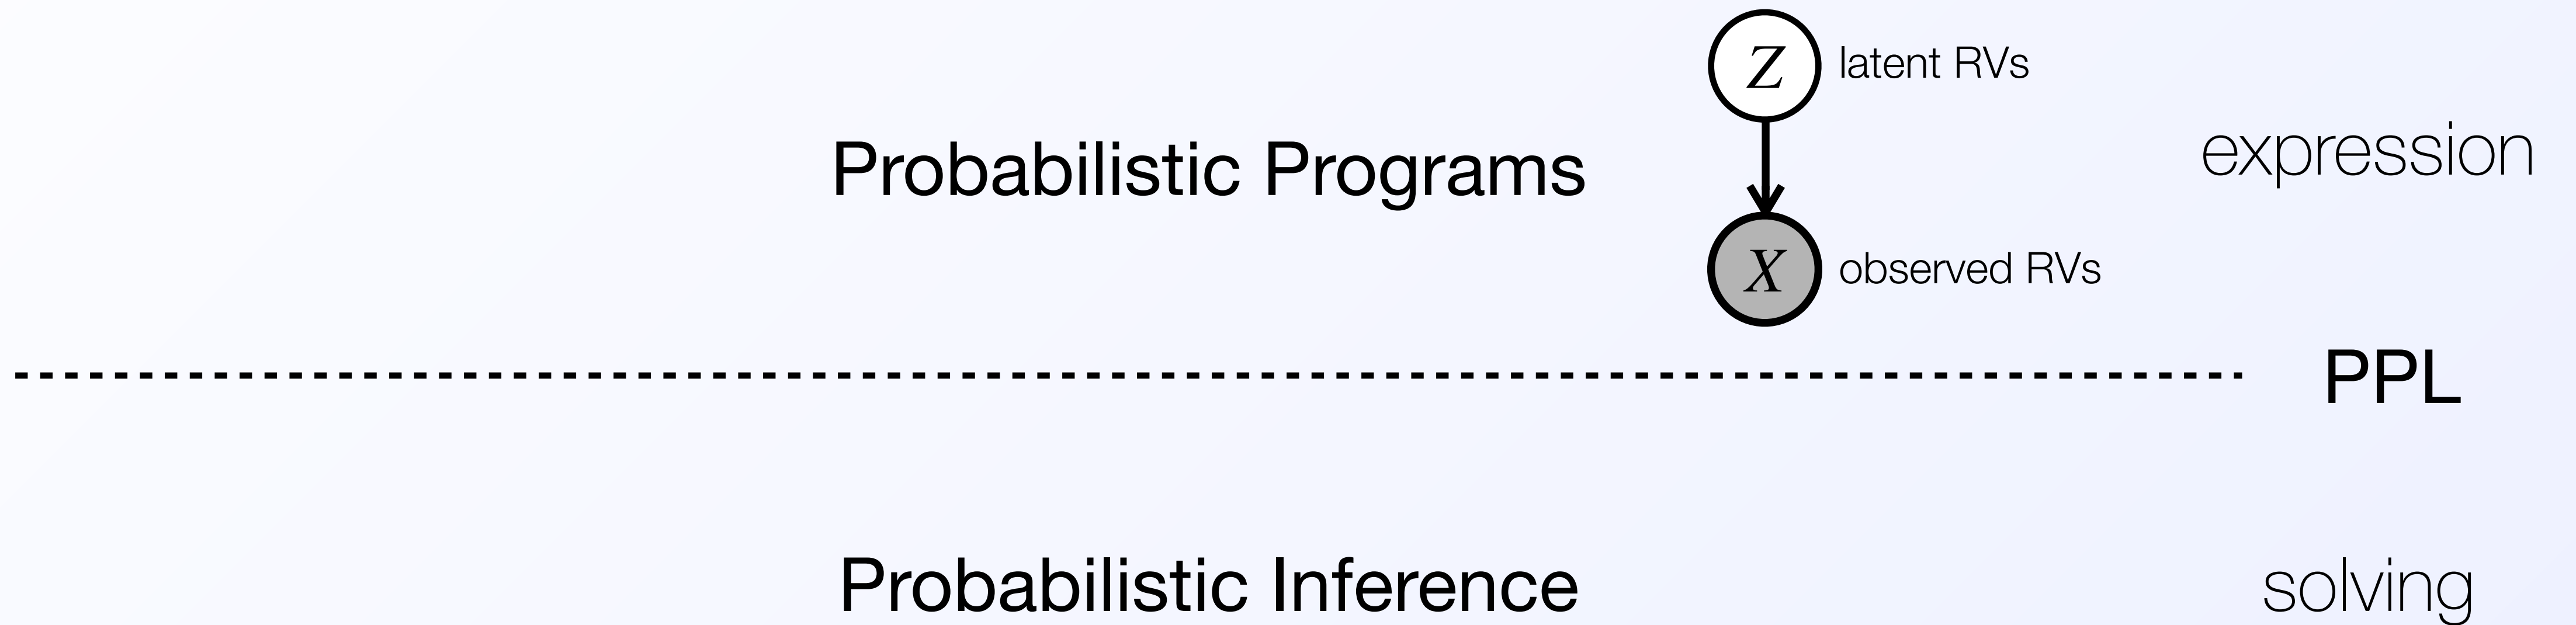


PPL

Probabilistic Inference

solving


PPL as an Abstraction Layer



PPL as an Abstraction Layer

data-generation process
generative model
stochastic simulation
decoders
inductive bias

Probabilistic Programs



expression

PPL

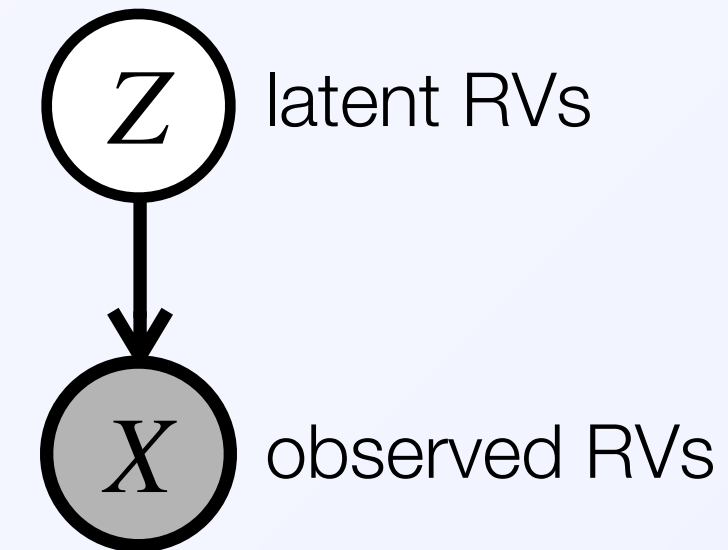
Probabilistic Inference

solving

PPL as an Abstraction Layer

data-generation process
generative model
stochastic simulation
decoders
inductive bias

Probabilistic Programs



expression

PPL

Probabilistic Inference

Interpreters & Compilers


solving

$$p(z | x) = \frac{p(x | z) p(z)}{p(x)}$$

What is hard about Bayesian inference?

Bayes' Theorem

$$p(z | x) \underset{\text{posterior}}{=} \frac{\overset{\text{joint}}{p(z, x)}}{\underset{\text{evidence}}{p(x)}} = \frac{\overset{\text{likelihood}}{p(x | z)} \overset{\text{prior}}{p(z)}}{\int p(z, x) dz \underset{\text{marginal likelihood}}{}}$$




Integration of joint over **all** execution traces \Rightarrow Enumerating all traces is unrealistic


What is hard about Bayesian inference?

Bayes' Theorem

$$p(z | x) \underset{\text{posterior}}{=} \underset{\text{joint}}{\frac{p(z, x)}{p(x)}} \underset{\text{evidence}}{=} \frac{\underset{\text{likelihood}}{p(x | z)} \underset{\text{prior}}{p(z)}}{\int p(z, x) dz \underset{\text{marginal likelihood}}{}}$$




Integration of joint over **all** execution traces



Enumerating all traces is unrealistic

Joint is defined by a **program**




Integration rarely has analytical solutions


Have to Approximate or Limit Expressivity

Bayes' Theorem

$$p(z | x) \underset{\text{posterior}}{=} \frac{\overset{\text{joint}}{p(z, x)}}{\underset{\text{evidence}}{p(x)}} = \frac{\overset{\text{likelihood}}{p(x | z)} \overset{\text{prior}}{p(z)}}{\int p(z, x) dz \underset{\text{marginal likelihood}}{}}$$




Integration of joint over **all** execution traces



Enumerating all traces is unrealistic

Joint is defined by a **program**



Integration rarely has analytical solutions

Have to Approximate or Limit Expressivity

Approximate

Rejection Sampling

Likelihood Weighting

Importance Sampling

Sequential Monte Carlo (SMC)

Markov Chain Monte Carlo (MCMC)

Variational Inference

Have to Approximate or Limit Expressivity

Approximate

Rejection Sampling

Likelihood Weighting

Importance Sampling

Sequential Monte Carlo (SMC)

Markov Chain Monte Carlo (MCMC)

Variational Inference

Limit Expressivity

Reduced expressive power

➡ Improved run-time efficiency

Have to Approximate or Limit Expressivity

Approximate

Rejection Sampling

Likelihood Weighting

Importance Sampling

Sequential Monte Carlo (SMC)

Markov Chain Monte Carlo (MCMC)

Variational Inference

Limit Expressivity

Reduced expressive power

➡ Improved run-time efficiency

Most common restriction to impose:

Ban recursion/unbounded loops
(think of finite graphical models)

Examples: Stan, Infer.NET, ...

PPL as an Abstraction Layer

Probabilistic Programs

expression





PPL



Probabilistic Inference

solving

Example: Reinforcement Learning



Example: Reinforcement Learning




Environment


```
def reward(state) // immediate reward  
def transition(state, action) // step the environment
```

Agent

```
def MDP(state): // recursive MDP description
```



Example: Reinforcement Learning




Environment

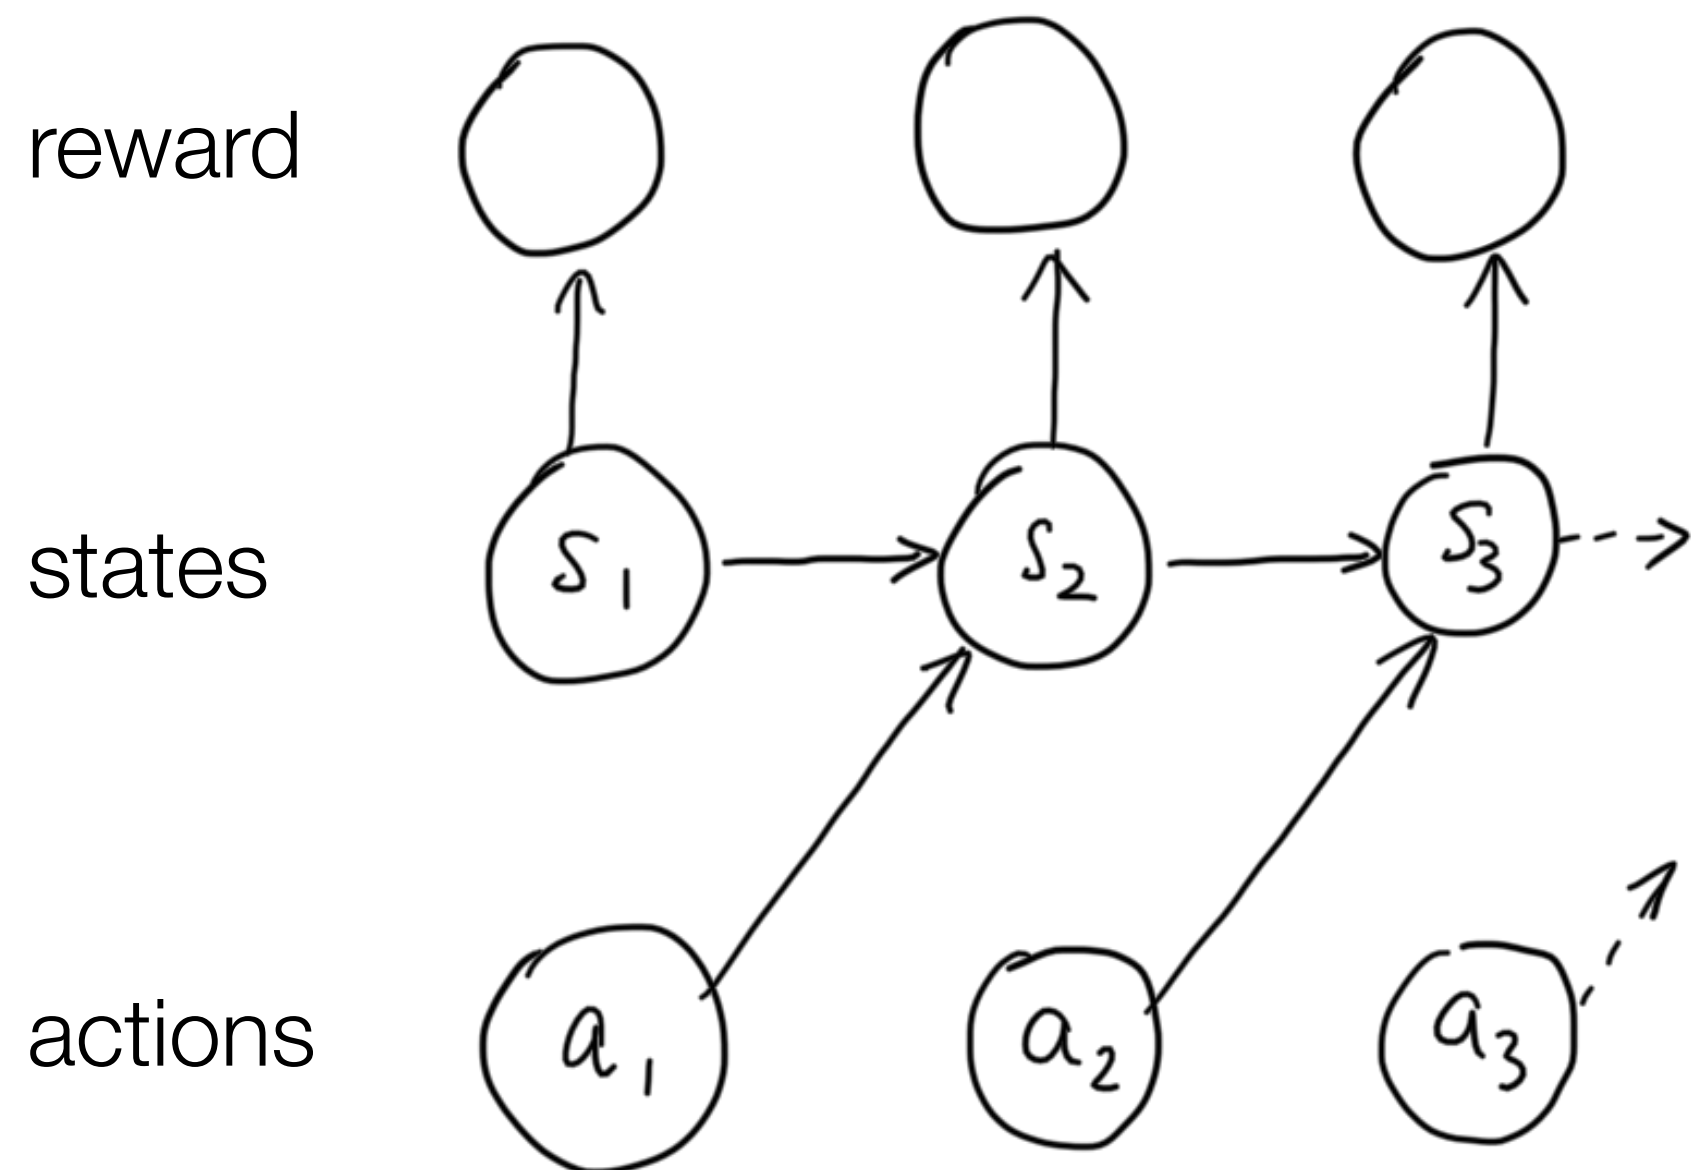

```
def reward(state) // immediate reward  
def transition(state, action) // step the environment
```

Agent

```
def MDP(state): // recursive MDP description  
  if (terminal(state))  
    return  
  action = sample(...) // sample action from prior  
  nextState = transition(state, action)  
  
  MDP(nextState) // recurse
```



Example: Reinforcement Learning



Goal of inference:

a policy function $\pi : \text{State} \rightarrow \text{Action}$


Environment

```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```

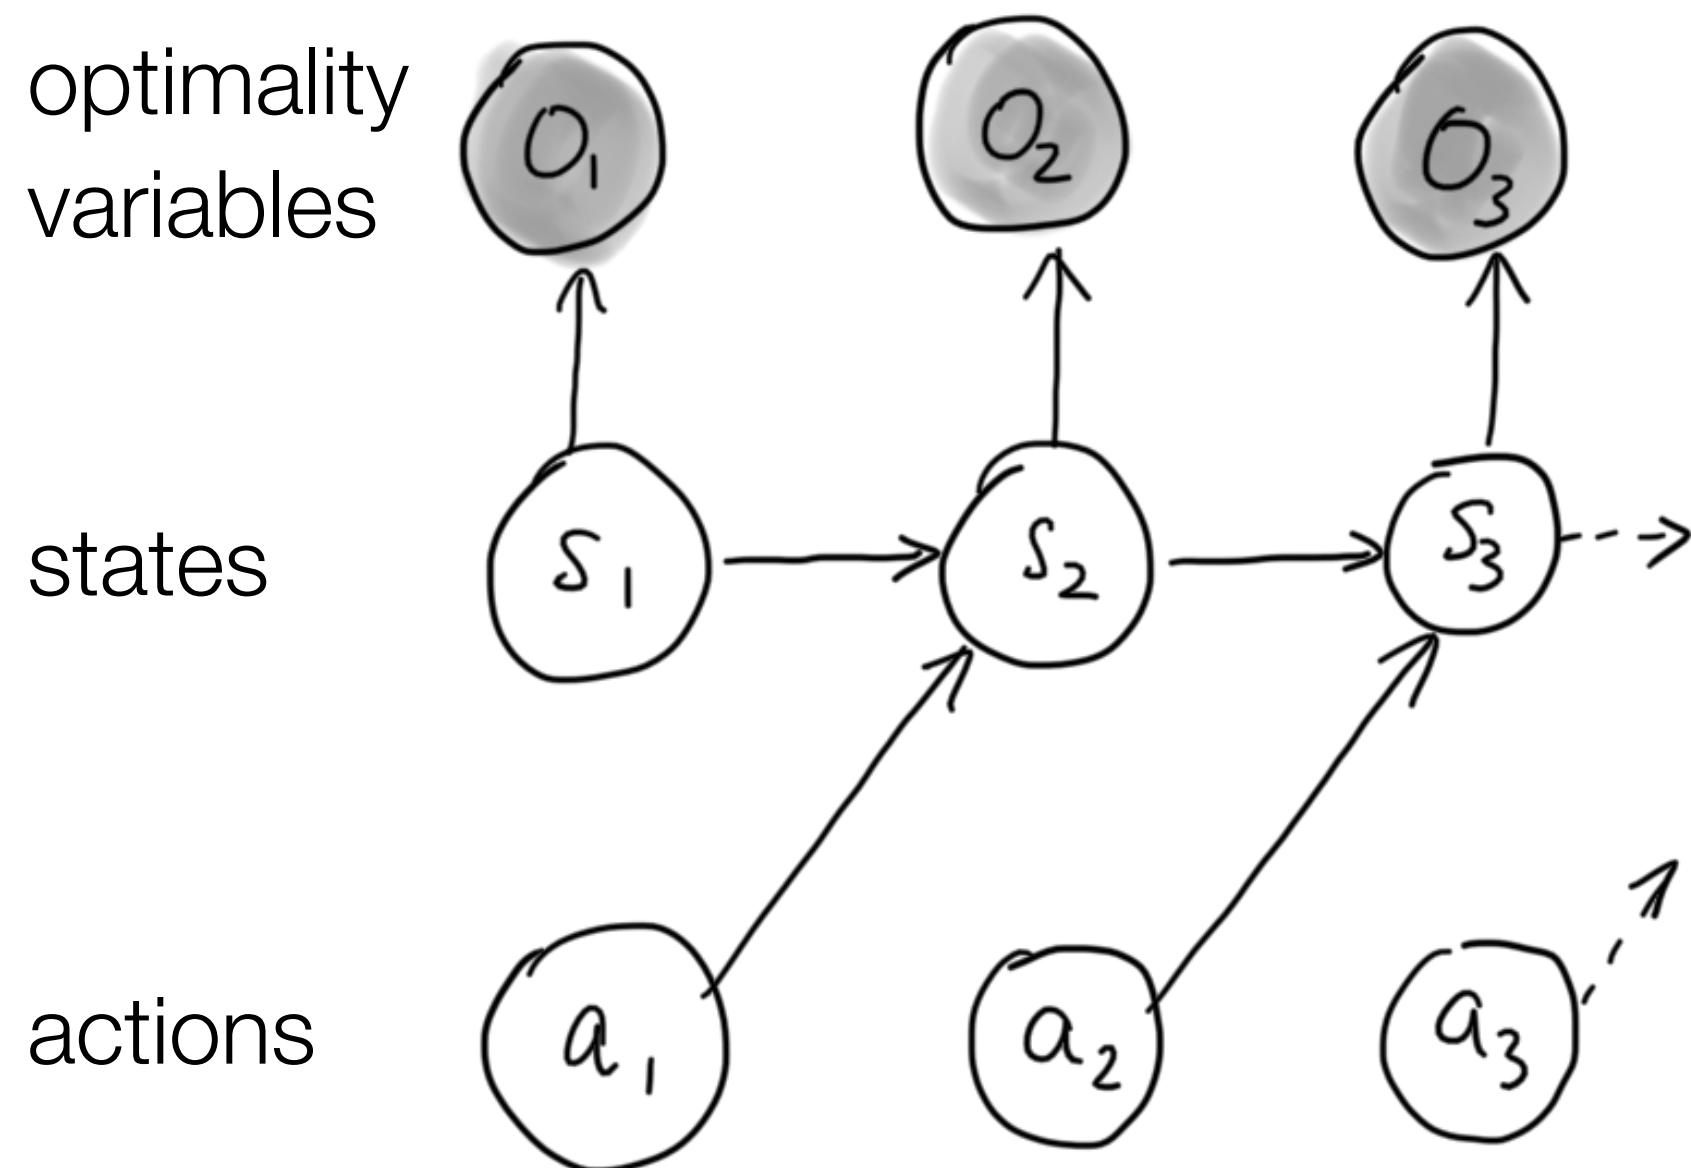

Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action)

  MDP(nextState) // recurse
```



Example: Reinforcement Learning



Goal of inference:

a policy function $\pi : \text{State} \rightarrow \text{Action}$

Environment



```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```

Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action)
  MDP(nextState) // recurse
```

$$p(O_t = 1 | s_t) \stackrel{\text{def}}{=} ?$$

Example: Reinforcement Learning



Goal of inference:

a policy function $\pi : \text{State} \rightarrow \text{Action}$

Environment


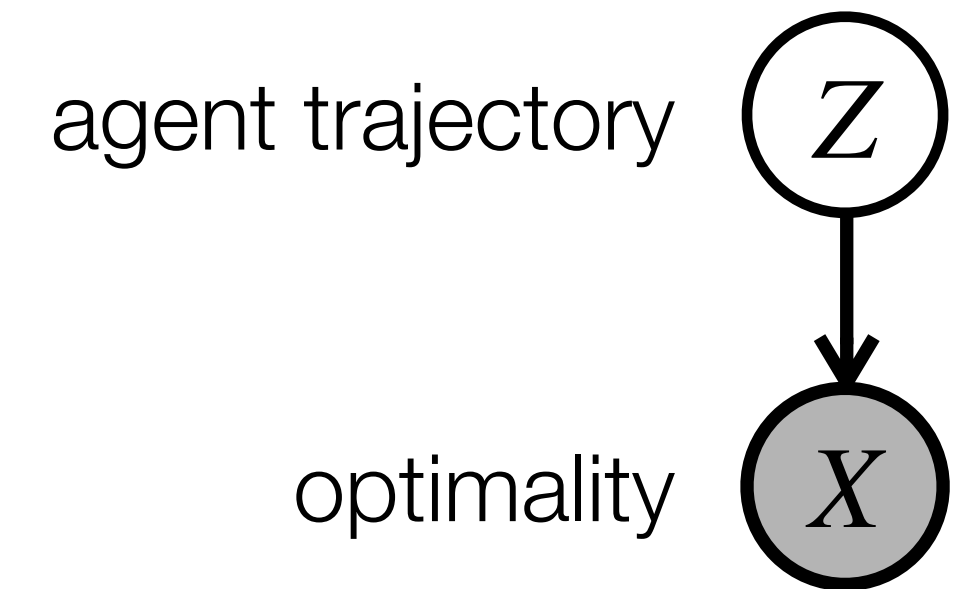
```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```

Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action) // condition on optimality
  MDP(nextState) // recurse
```

$$p(O_t = 1 | s_t) \stackrel{\text{def}}{=} \exp(r(s_t))$$

Example: Reinforcement Learning



Goal of inference:

a policy function $\pi : \text{State} \rightarrow \text{Action}$

Environment

```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```



Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action)
  factor(reward(nextState)) // condition on optimality
  MDP(nextState) // recurse
```

$$p(O_t = 1 | s_t) \stackrel{\text{def}}{=} \exp(r(s_t))$$



Example: Reinforcement Learning



Goal of inference:

a policy function $\pi : \text{State} \rightarrow \text{Action}$ leading to optimal trajectory, rather than a trajectory per se

Environment

```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```



Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action)
  factor(reward(nextState)) // condition on optimality
  MDP(nextState) // recurse
```

$$p(O_t = 1 | s_t) \stackrel{\text{def}}{=} \exp(r(s_t))$$



Example: Reinforcement Learning



Goal of inference:

a policy function $\pi : \text{State} \rightarrow \text{Dist}[\text{Action}]$ leading to optimal trajectory, rather than a trajectory per se

Environment

```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```



Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action)
  factor(reward(nextState)) // condition on optimality
  MDP(nextState) // recurse
```

$$p(O_t = 1 | s_t) \stackrel{def}{=} \exp(r(s_t))$$



Example: Reinforcement Learning



Goal of inference: $\pi(s) = q(a | s; \phi)$

a policy function $\pi : \text{State} \rightarrow \text{Dist}[\text{Action}]$ leading to optimal trajectory, rather than a trajectory per se

Environment

```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```



Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action)
  factor(reward(nextState)) // condition on optimality
  MDP(nextState) // recurse
```

$$p(O_t = 1 | s_t) \stackrel{\text{def}}{=} \exp(r(s_t))$$



Example: Reinforcement Learning



Goal of inference: $\pi(s) = q(a | s; \phi)$

a policy function $\pi : \text{State} \rightarrow \text{Dist}[\text{Action}]$ leading to optimal trajectory, rather than a trajectory per se

Environment



```
def reward(state) // immediate reward
def transition(state, action) // step the environment
```

Agent

```
def MDP(state): // recursive MDP description
  if (terminal(state))
    return
  action = sample(...) // sample action from prior
  nextState = transition(state, action)
  factor(reward(nextState)) // condition on optimality
  MDP(nextState) // recurse
```

$$p(a_t) \leftarrow$$
$$p(s_{t+1} | s_t, a_t) \leftarrow$$
$$p(O_t = 1 | s_t) \stackrel{def}{=} \exp(r(s_t)) \leftarrow$$

Example: Reinforcement Learning



Goal of inference:

$$\pi(s) = q(a | s; \phi)$$

a policy function $\pi : \text{State} \rightarrow \text{Dist}[\text{Action}]$ leading to optimal trajectory, rather than a trajectory per se

$$p(s_1, a_1, \dots | \text{optimality}) \propto p(s_1) \prod_t p(a_t) p(s_{t+1} | s_t, a_t) p(O_t = 1 | s_t)$$

Agent

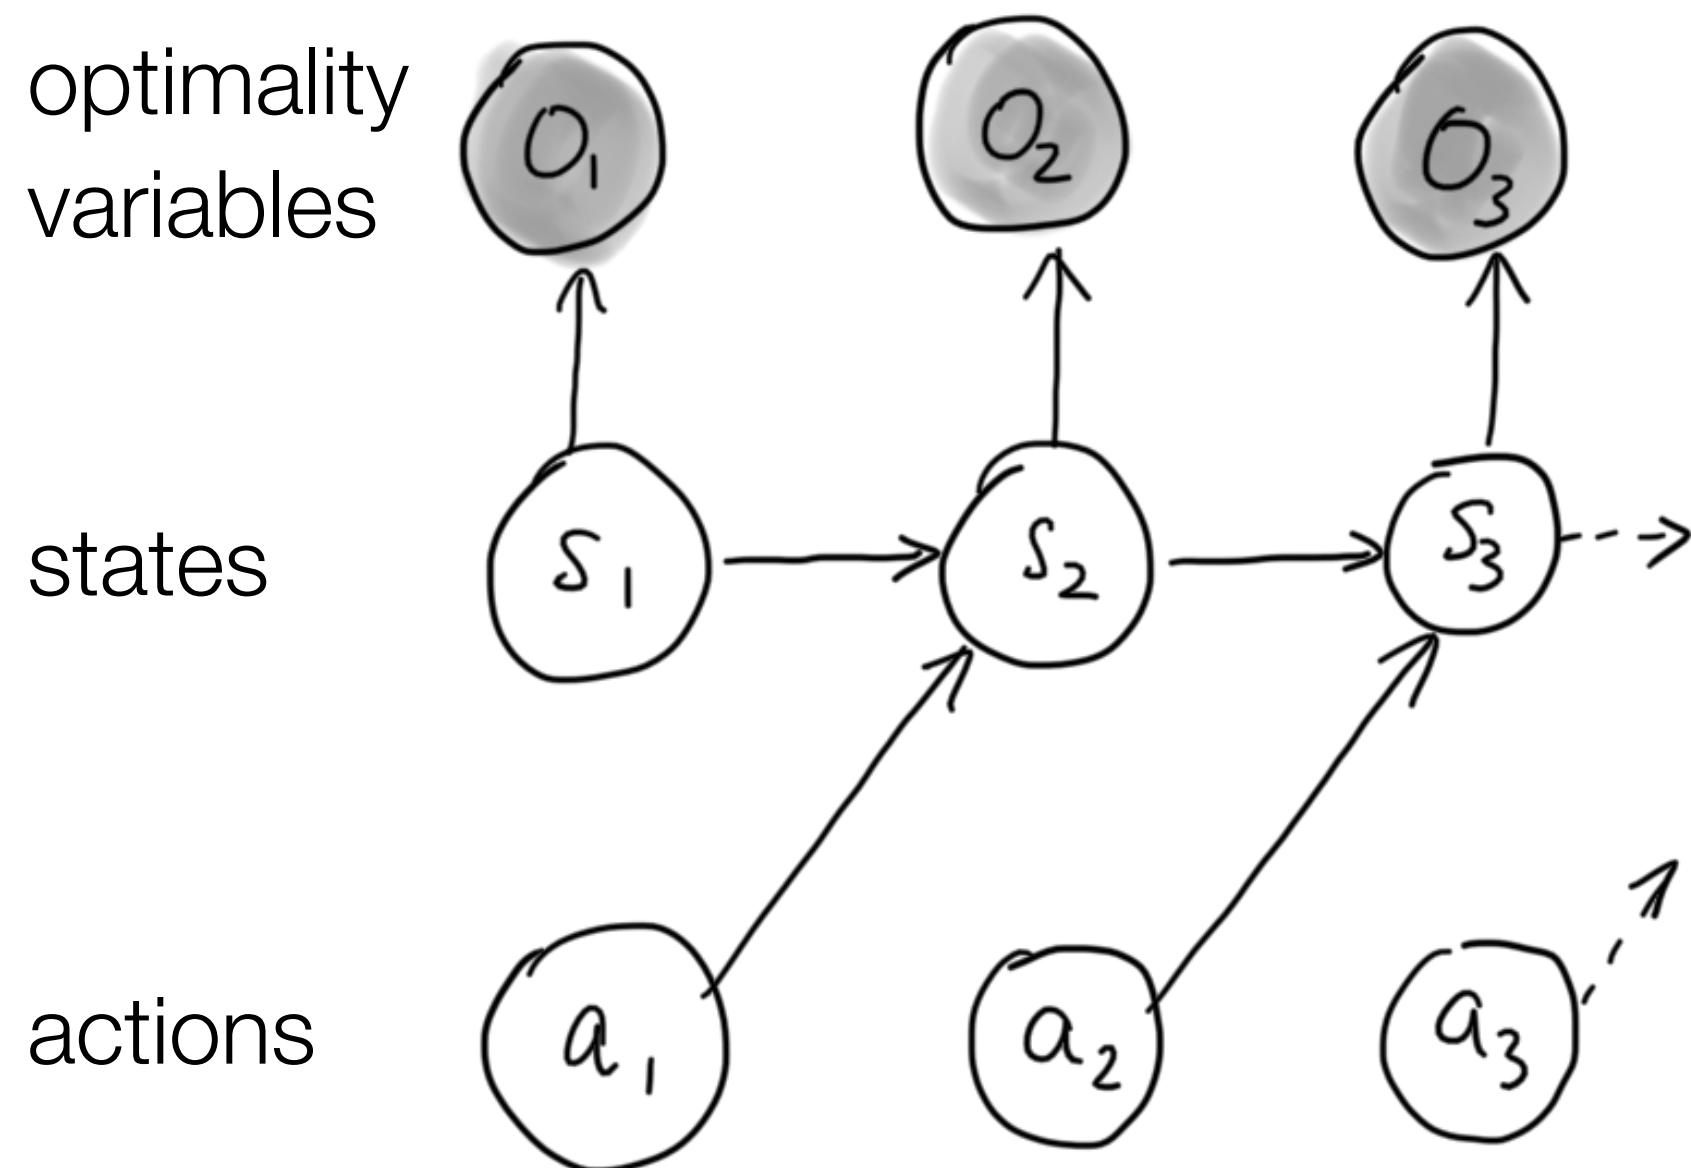

```
def MDP(state): // recursive MDP description
    if (terminal(state))
        return
    action = sample(...) // sample action from prior
    nextState = transition(state, action)
    factor(reward(nextState)) // condition on optimality
    MDP(nextState) // recurse
```

$p(a_t)$ ←

$p(s_{t+1} | s_t, a_t)$ ←

$p(O_t = 1 | s_t) \stackrel{\text{def}}{=} \exp(r(s_t))$ ←

Example: Reinforcement Learning



Goal of inference: $\pi(s) = q(a | s; \phi)$

a policy function $\pi : \text{State} \rightarrow \text{Dist}[\text{Action}]$ leading to optimal trajectory, rather than a trajectory per se

$$p(s_1, a_1, \dots | \text{optimality}) \propto p(s_1) \prod_t p(a_t) p(s_{t+1} | s_t, a_t) p(O_t = 1 | s_t)$$

$$q(s_1, a_1, \dots, s_t, a_t; \phi) = p(s_1) \prod_t q(a_t | s_t; \phi) p(s_{t+1} | s_t, a_t)$$

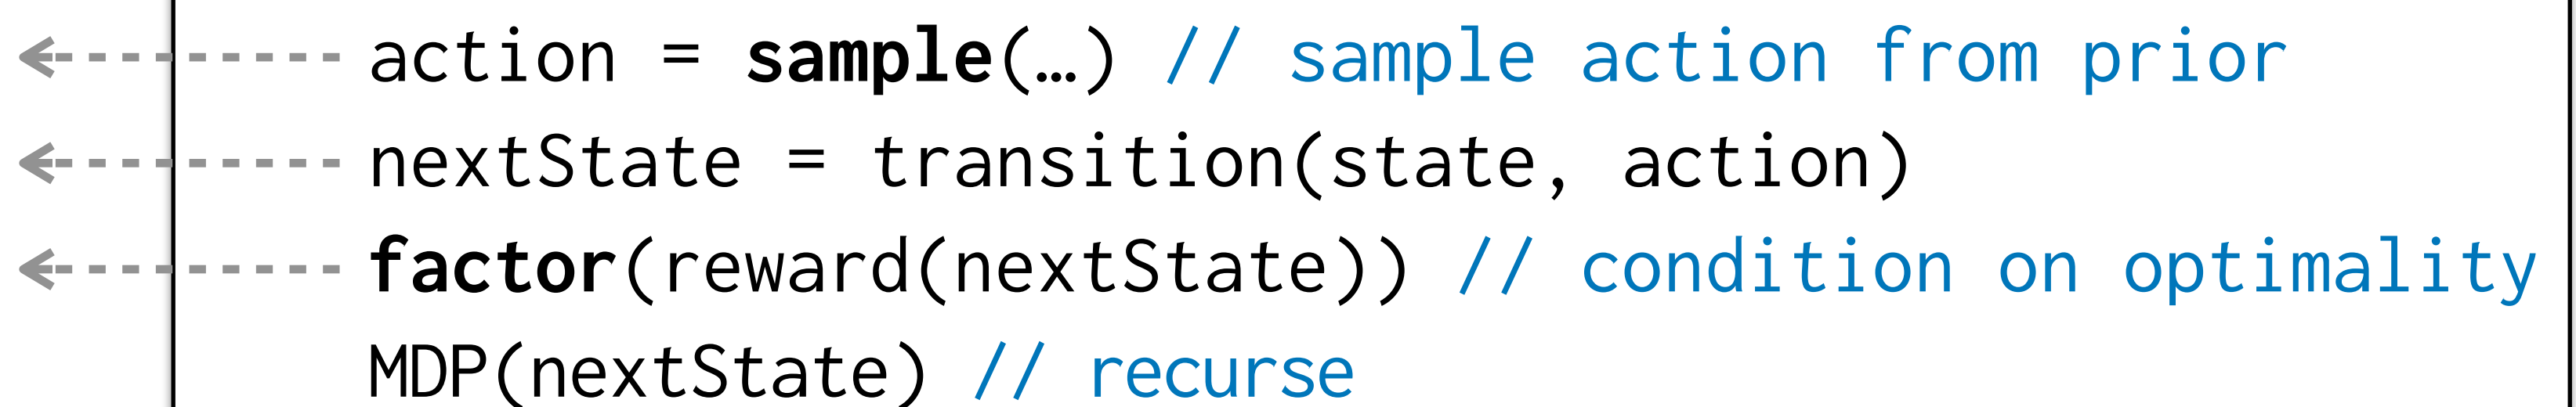
Agent

```
def MDP(state): // recursive MDP description
    if (terminal(state))
        return
    action = sample(...) // sample action from prior
    nextState = transition(state, action)
    factor(reward(nextState)) // condition on optimality
    MDP(nextState) // recurse
```

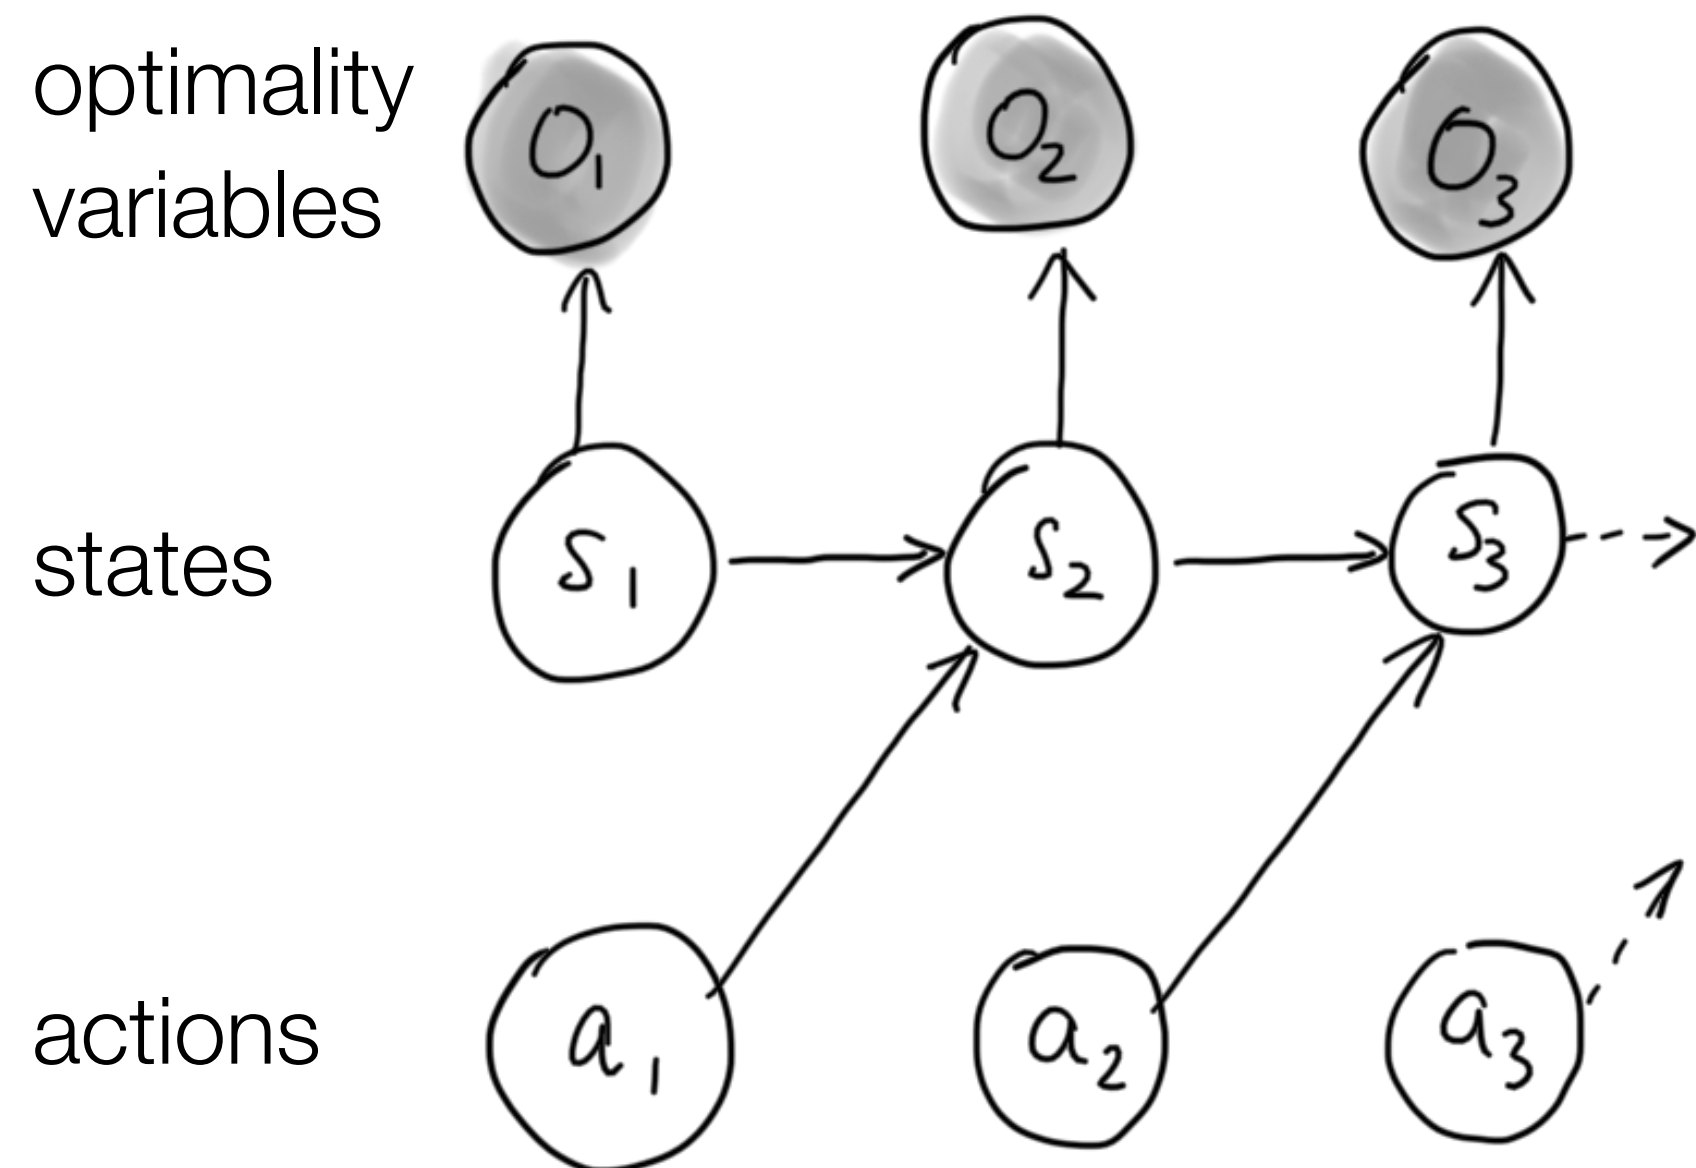

$p(a_t)$

$p(s_{t+1} | s_t, a_t)$

$p(O_t = 1 | s_t) \stackrel{\text{def}}{=} \exp(r(s_t))$



Example: Reinforcement Learning




Goal of inference: $\pi(s) = q(a | s; \phi)$
a policy function $\pi : \text{State} \rightarrow \text{Dist}[\text{Action}]$ leading to optimal trajectory, rather than a trajectory per se

$$p(s_1, a_1, \dots | \text{optimality}) \propto p(s_1) \prod_t p(a_t) p(s_{t+1} | s_t, a_t) p(O_t = 1 | s_t)$$
$$q(s_1, a_1, \dots, s_t, a_t; \phi) = p(s_1) \prod_t q(a_t | s_t; \phi) p(s_{t+1} | s_t, a_t)$$


$$\min_{\phi} D_{\text{KL}} \left(q(s_1, a_1, \dots, s_t, a_t; \phi) || p(s_1, a_1, \dots | \text{optimality}) \right)$$

Example: Two-Player Game




Example: Two-Player Game

```
var act = function(state, player) {  
  var move = sample(movePrior(state));  
  var u = utility(state, move, player);  
  factor(alpha * u);  
  return move;  
};
```



Example: Two-Player Game

```
var act = function(state, player) {  
  var move = sample(movePrior(state));  
  var u = utility(state, move, player);  
  factor(alpha * u);  
  return move;  
};  
  
var utility = function(state, move, player) {  
  var outcome = simulate(state, move, player);  
  if (hasWon(state, player)) { return 10; }  
  else if (isDraw(state)) { return 0; }  
  else { return -10; }  
}
```




Example: Two-Player Game

```
var act = function(state, player) {
  var move = sample(movePrior(state));
  var u = utility(state, move, player);
  factor(alpha * u);
  return move;
};

var utility = function(state, move, player) {
  var outcome = simulate(state, move, player);
  if (hasWon(state, player)) { return 10; }
  else if (isDraw(state)) { return 0; }
  else { return -10; }
}

var simulate = function(state, action, player) {
  var nextState = transition(state, action, player);
  if (isTerminal(nextState)) {
    return nextState;
  } else {
    var nextPlayer = otherPlayer(player);
    return sample(Infer({ model() {
      var nextMove = act(nextState, nextPlayer);
      return simulate(nextState, nextMove, nextPlayer);
    }}}));
  }
};
```




Example: Two-Player Game

```
var act = function(state, player) {
  var move = sample(movePrior(state));
  var u = utility(state, move, player);
  factor(alpha * u);
  return move;
};

var utility = function(state, move, player) {
  var outcome = simulate(state, move, player);
  if (hasWon(state, player)) { return 10; }
  else if (isDraw(state)) { return 0; }
  else { return -10; }
}

var simulate = function(state, action, player) {
  var nextState = transition(state, action, player);
  if (isTerminal(nextState)) {
    return nextState;
  } else {
    var nextPlayer = otherPlayer(player);
    return sample(Infer({ model() {
      var nextMove = act(nextState, nextPlayer);
      return simulate(nextState, nextMove, nextPlayer);
    }}}));
  }
};
```




Example: Two-Player Game

```
var act = function(state, player) {
  var move = sample(movePrior(state));
  var u = utility(state, move, player);
  factor(alpha * u);
  return move;
};

var utility = function(state, move, player) {
  var outcome = simulate(state, move, player);
  if (hasWon(state, player)) { return 10; }
  else if (isDraw(state)) { return 0; }
  else { return -10; }
}

var simulate = function(state, action, player) {
  var nextState = transition(state, action, player);
  if (isTerminal(nextState)) {
    return nextState;
  } else {
    var nextPlayer = otherPlayer(player);
    return sample(Infer({ model() {
      var nextMove = act(nextState, nextPlayer);
      return simulate(nextState, nextMove, nextPlayer);
    }}}));
  }
};
```



Inference inside inference => Thinking about thinking


Example: Two-Player Game

```
var act = function(state, player) {
  var move = sample(movePrior(state));
  var u = utility(state, move, player);
  factor(alpha * u);
  return move;
};

var utility = function(state, move, player) {
  var outcome = simulate(state, move, player);
  if (hasWon(state, player)) { return 10; }
  else if (isDraw(state)) { return 0; }
  else { return -10; }
}

var simulate = function(state, action, player) {
  var nextState = transition(state, action, player);
  if (isTerminal(nextState)) {
    return nextState;
  } else {
    var nextPlayer = otherPlayer(player);
    return sample(Infer({ model() {
      var nextMove = act(nextState, nextPlayer);
      return simulate(nextState, nextMove, nextPlayer);
    }}}));
  }
};
```

Amortizing nested inference




Example: Two-Player Game

```
var act = function(state, player) {
  var move = sample(movePrior(state));
  var u = utility(state, move, player);
  factor(alpha * u);
  return move;
};

var utility = function(state, move, player) {
  var outcome = simulate(state, move, player);
  if (hasWon(state, player)) { return 10; }
  else if (isDraw(state)) { return 0; }
  else { return -10; }
}

var simulate = function(state, action, player) {
  var nextState = transition(state, action, player);
  if (isTerminal(nextState)) {
    return nextState;
  } else {
    var nextPlayer = otherPlayer(player);
    return sample(Infer({ model() {
      var nextMove = act(nextState, nextPlayer);
      return simulate(nextState, nextMove, nextPlayer);
    }}}));
  }
};
```

Amortizing nested inference



A Glimpse into Formal Semantics

terms	$t ::=$	x	variable
		$\lambda x. t$	abstraction
		$t_1 t_2$	application
		r	real number
		$\text{op}_n(t_1, \dots, t_n)$	n -ary operation invocation
		sample	sampling
		factor t	conditioning

$$\rho_n(\ell, t, V) \stackrel{\text{def}}{=} \begin{cases} r & \langle \ell \mid t \rangle \longrightarrow_n \langle \epsilon \mid v \rangle \bullet r \text{ and } v \in V \\ 0 & \text{otherwise} \end{cases}$$

$$\mu_n(t, V) \stackrel{\text{def}}{=} \int \rho_n(\ell, t, V) d\ell$$

$$\mu(t, V) \stackrel{\text{def}}{=} \lim_{n \rightarrow \infty} \mu_n(t, V)$$

$$\boxed{t_1 \longrightarrow t_2}$$

$$\text{[KTX]} \quad \frac{t_1 \longrightarrow t_2}{K[t_1] \longrightarrow K[t_2]}$$

$$\text{[BETA]} \quad (\lambda x. t) v \longrightarrow t \{v/x\}$$

$$\text{[OP]} \quad \frac{\llbracket \text{op}_n \rrbracket(r_1, \dots, r_n) = r}{\text{op}_n(r_1, \dots, r_n) \longrightarrow r}$$

$$\boxed{\langle \ell_1 \mid t_1 \rangle \longrightarrow \langle \ell_2 \mid t_2 \rangle \bullet r}$$

$$\text{[PURE]} \quad \frac{t_1 \longrightarrow t_2}{\langle \ell \mid t_1 \rangle \longrightarrow \langle \ell \mid t_2 \rangle \bullet 1}$$

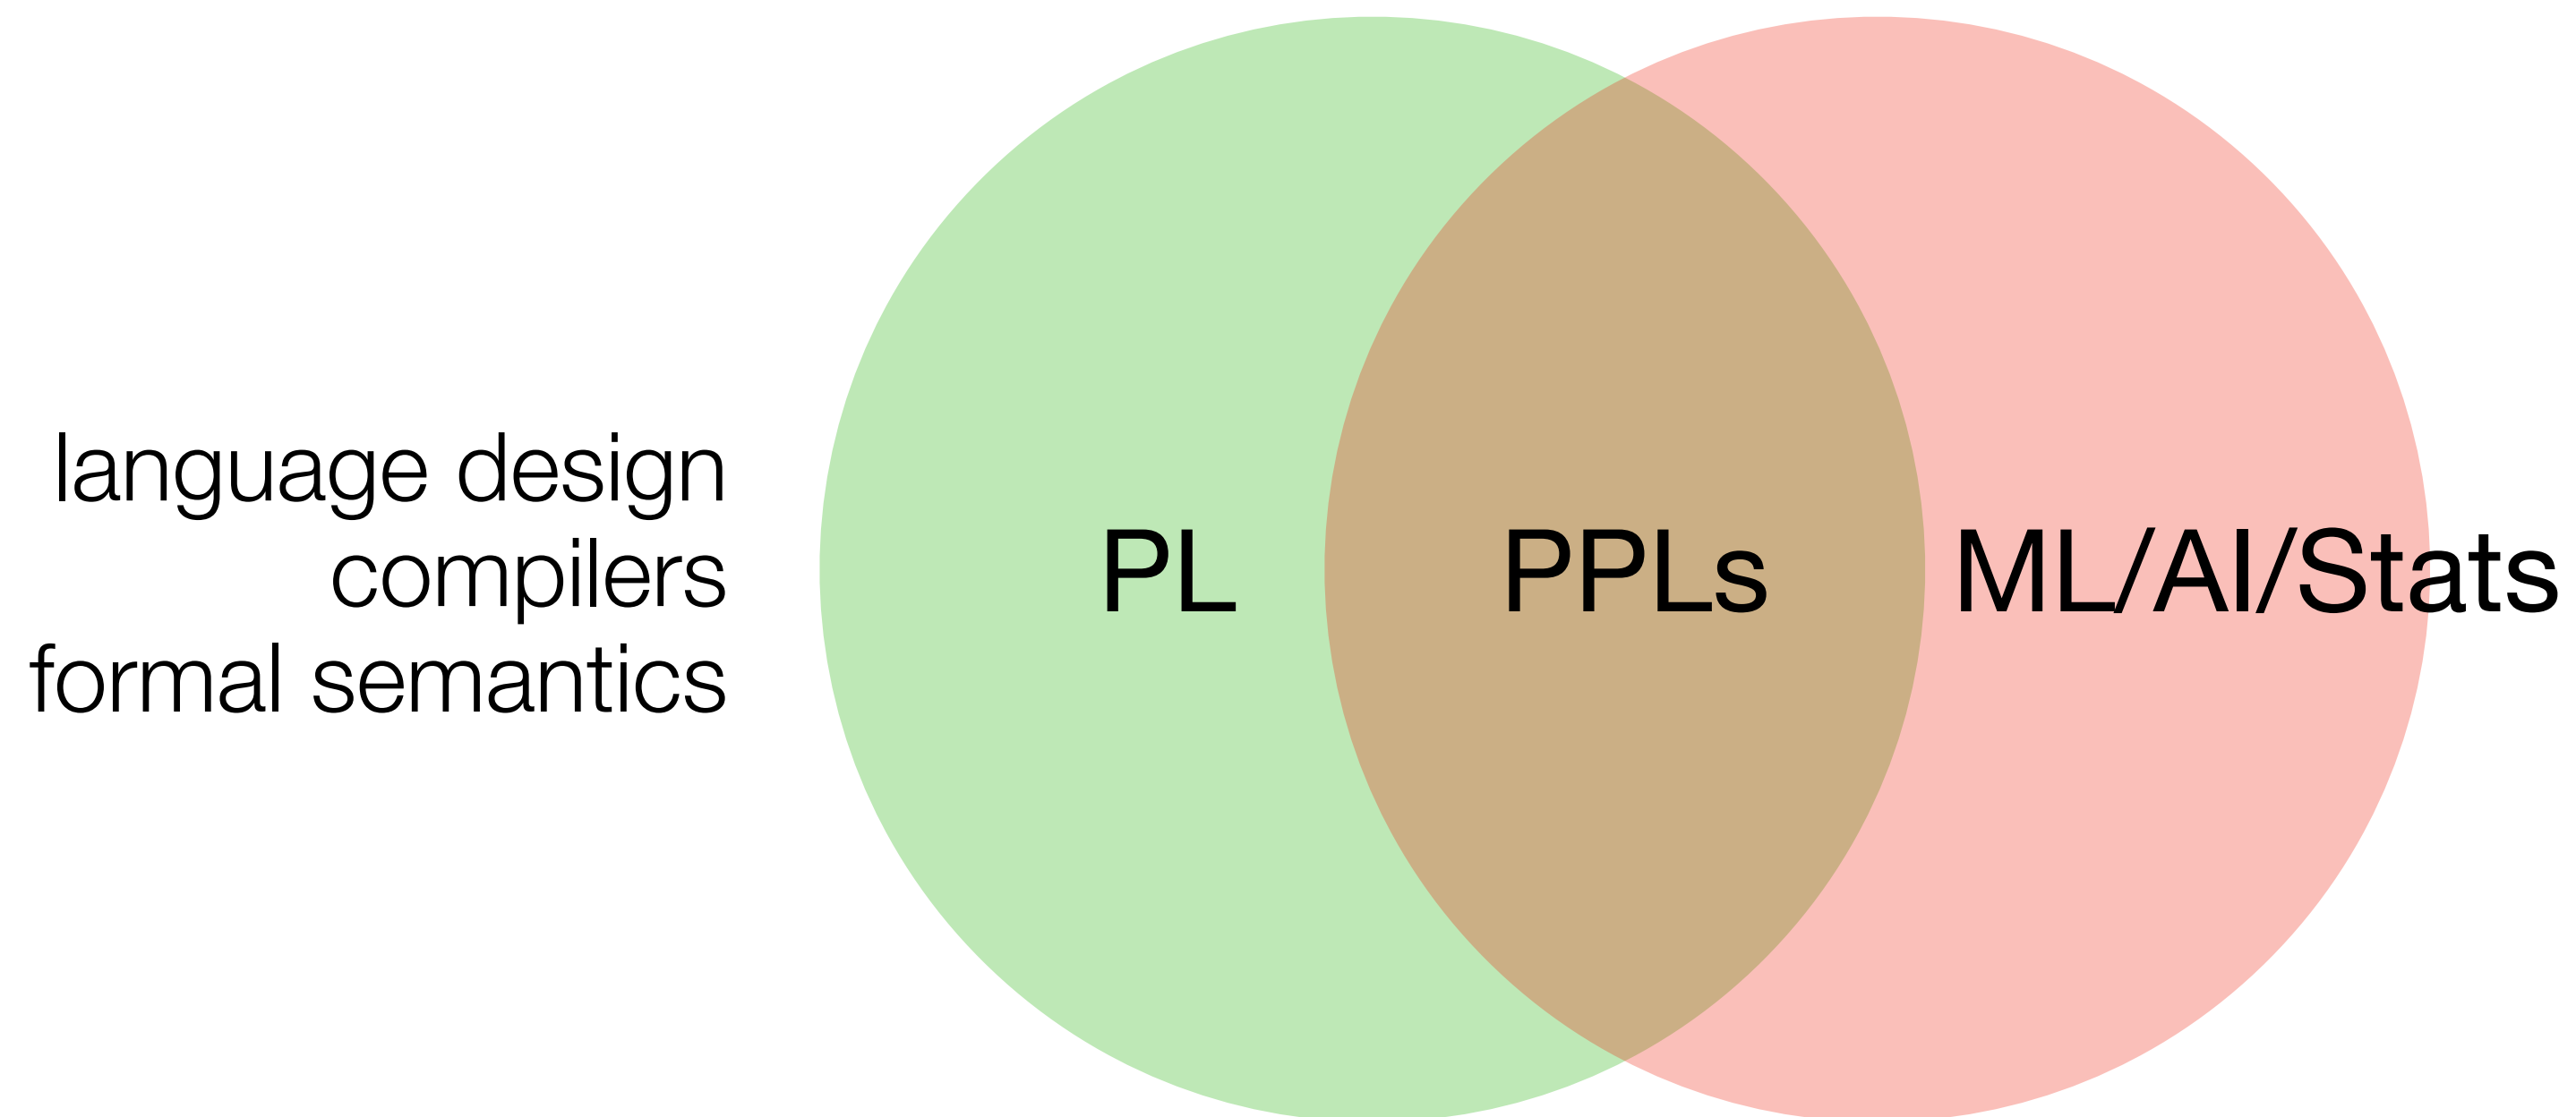
$$\text{[SAMPLE]} \quad \frac{0 \leq r \leq 1}{\langle r, \ell \mid \text{sample} \rangle \longrightarrow \langle \ell \mid r \rangle \bullet 1}$$

$$\text{[FACTOR]} \quad \frac{0 < r}{\langle \ell \mid \text{factor } r \rangle \longrightarrow \langle \ell \mid 0 \rangle \bullet r}$$

Takeaway messages

PPLs are powerful tools for probabilistic modeling and inference

Exciting area of ongoing research



CS 152 Programming Languages

Probabilistic Programming & Probabilistic Programming Languages

Yizhou Zhang

University of Waterloo