1 HMMs: Warmup

1. What are the three components of a hidden markov model? What makes it "hidden"?

   - Initial distribution: $P(X_0)$
   - Transition model: $P(X_t|X_{t-1})$
   - Sensor model: $P(E_t|X_t)$

   The hidden part of hidden markov models comes from the fact that we do not observe the state variables X_i directly, rather we observe the evidence variables E_i and must make conclusions about the underlying true state.

2. Write an expression for the joint distribution of a hidden markov model consisting of states $X_0, \ldots, X_n$ and evidence variables $E_1, \ldots, E_N$. How does the expression reflect the underlying structure of the model?

   $P(X_0, \ldots, X_N, E_1, \ldots, E_N) = P(X_0) \prod_{t=1}^{N} P(X_t|X_{t-1})P(E_t|X_t)$

   This expression reflects that the a state is only directly influenced by its previous state, and that the evidence is independent of everything else given the corresponding state.

3. For each of the following descriptions in English of an inference task, write the corresponding probability expression:

   - Draw conclusions about our current underlying state given evidence up to the current time step
   - Draw conclusions about our future underlying state given evidence up to the current time step
   - Draw conclusions about a past underlying state given evidence up to the current time step
   - Draw conclusions about the sequence of underlying states given evidence up to the current time step
   - Draw conclusions about the most likely sequence of underlying states given evidence up to the current time step

   - Filtering: $P(X_t|E_{1:t})$
   - Prediction: $P(X_{t+k}|E_{1:t}), k > 0$
   - Smoothing: $P(X_k|E_{1:t}), 1 \leq k < t$
   - Explanation: $P(X_{1:t}|E_{1:t})$
   - Most likely explanation: $\text{argmax}_{X_{1:t}} P(X_{1:t}|E_{1:t})$

4. Hidden Markov Models can be extended in a number of ways to incorporate additional relations. Since the independence assumptions are different in these extended Hidden Markov Models, the forward algorithm updates will also be different. What is the forward algorithm updates for the extended Hidden Markov Models specified by the following Bayes net?

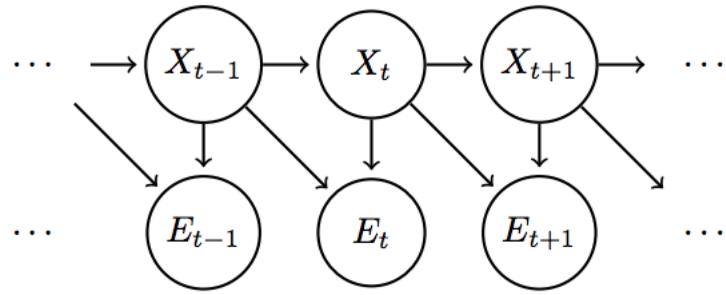

$P(X_t|e_{1:t}) = \alpha \sum_{x_{t-1}} P(e_t|x_t, x_{t-1})P(x_t|x_{t-1})P(x_{t-1}|e_{1:t-1})$

2 HMMs: Tracking a Jabberwock

You have been put in charge of a Jabberwock for your friend Lewis. The Jabberwock is kept in a large tugley wood which is conveniently divided into an $N \times N$ grid. It wanders freely around the N^2 possible cells. At each time step $t = 1, 2, 3, \ldots$, the Jabberwock is in some cell $X_t \in \{1, \ldots, N\}^2$, and it moves to cell X_{t+1} randomly as follows: with probability $1 - \epsilon$, it chooses one of the (up to 4) valid neighboring cells uniformly at random; with probability ϵ, it uses its magical powers to teleport to a random cell uniformly at random among the N^2 possibilities (it might teleport to the same cell). Suppose $\epsilon = \frac{1}{2}$, $N = 10$ and that the Jabberwock always starts in $X_1 = (1, 1)$.

(a) Compute the probability that the Jabberwock will be in $X_2 = (2, 1)$ at time step 2. What about $P(X_2 = (4, 4))$?

$P(X_2 = (2, 1)) = 1/2 \cdot 1/2 + 1/2 \cdot 1/100 = 0.255$
$P(X_2 = (4, 4)) = 1/2 \cdot 1/100 = 0.005$

At each time step t, you don't see X_t but see E_t, which is the row that the Jabberwock is in; that is, if $X_t = (r, c)$, then $E_t = r$. You still know that $X_1 = (1, 1)$.

(b) Suppose we see that $E_1 = 1$, $E_2 = 2$. Fill in the following table with the distribution over X_t after each time step, taking into consideration the evidence. Your answer should be concise. Hint: you should not need to do any heavy calculations.

<table>
<tr><th>t</th><th>$P(X_t \mid e_{1:t-1}, X_1 = (1,1))$</th><th>$P(X_t \mid e_{1:t}, X_1 = (1,1))$</th></tr>
<tr><td>1</td><td><table><tr><th>X_1</th><th>$P(X_1)$</th></tr><tr><td>(1, 1)</td><td></td></tr><tr><td>all other values</td><td></td></tr></table></td><td><table><tr><th>X_1</th><th>$P(X_1)$</th></tr><tr><td>(1, 1)</td><td></td></tr><tr><td>all other values</td><td></td></tr></table></td></tr>
<tr><td>2</td><td><table><tr><th>X_2</th><th>$P(X_2 \mid e_1, X_1 = (1,1))$</th></tr><tr><td>(1, 2)</td><td></td></tr><tr><td>(2, 1)</td><td></td></tr><tr><td>all other values</td><td></td></tr></table></td><td><table><tr><th>X_2</th><th>$P(X_2 \mid e_{1:2}, X_1 = (1,1))$</th></tr><tr><td>(2, 1)</td><td></td></tr><tr><td>$(2, a)$ $(\forall a, a > 1)$</td><td></td></tr><tr><td>all other values</td><td></td></tr></table></td></tr>
</table>

<table>
<tr><th>t</th><th>$P(X_t \mid e_{1:t-1}, X_1 = (1,1))$</th><th>$P(X_t \mid e_{1:t}, X_1 = (1,1))$</th></tr>
<tr><td>1</td><td><table><tr><th>X_1</th><th>$P(X_1)$</th></tr><tr><td>(1, 1)</td><td>1</td></tr><tr><td>all other values</td><td>0</td></tr></table></td><td><table><tr><th>X_1</th><th>$P(X_1)$</th></tr><tr><td>(1, 1)</td><td>1</td></tr><tr><td>all other values</td><td>0</td></tr></table></td></tr>
<tr><td>2</td><td><table><tr><th>X_2</th><th>$P(X_2 \mid e_1, X_1 = (1,1))$</th></tr><tr><td>(1, 2)</td><td>51/200</td></tr><tr><td>(2, 1)</td><td>51/200</td></tr><tr><td>all other values</td><td>1/200</td></tr></table></td><td><table><tr><th>X_2</th><th>$P(X_2 \mid e_{1:2}, X_1 = (1,1))$</th></tr><tr><td>(2, 1)</td><td>51/60</td></tr><tr><td>$(2, a)$ $(\forall a, a > 1)$</td><td>1/60</td></tr><tr><td>all other values</td><td>0</td></tr></table></td></tr>
</table>

You are a bit unsatisfied that you can't pinpoint the Jabberwock exactly. But then you remembered Lewis told you that the Jabberwock teleports only because it is frumious on that time step, and it becomes frumious independently of anything else. Let us introduce a variable $F_t \in \{0, 1\}$ to denote whether it will teleport at time t. We want to to add these frumious variables to the HMM. Consider the two candidates:

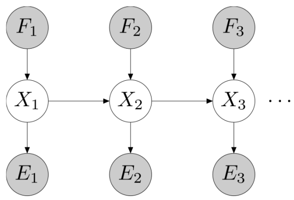

(A)

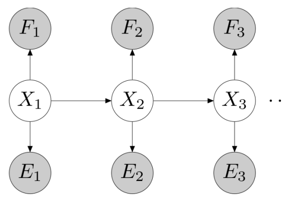

(B)

(A)	(B)
$X_1 \perp\!\!\!\perp X_3 \mid X_2$	$X_1 \perp\!\!\!\perp X_3 \mid X_2$
$X_1 \perp\!\!\!\perp E_2 \mid X_2$	$X_1 \perp\!\!\!\perp E_2 \mid X_2$
$X_1 \perp\!\!\!\perp F_2 \mid X_2$	$X_1 \perp\!\!\!\perp F_2 \mid X_2$
$X_1 \perp\!\!\!\perp E_4 \mid X_2$	$X_1 \perp\!\!\!\perp E_4 \mid X_2$
$X_1 \perp\!\!\!\perp F_4 \mid X_2$	$X_1 \perp\!\!\!\perp F_4 \mid X_2$
$E_3 \perp\!\!\!\perp F_3 \mid X_3$	$E_3 \perp\!\!\!\perp F_3 \mid X_3$
$E_1 \perp\!\!\!\perp F_2 \mid X_2$	$E_1 \perp\!\!\!\perp F_2 \mid X_2$
$E_1 \perp\!\!\!\perp F_2 \mid E_2$	$E_1 \perp\!\!\!\perp F_2 \mid E_2$

(c) For each model, circle the conditional independence assumptions above which are true in that model.

(A)	(B)
$X_1 \perp\!\!\!\perp X_3 \mid X_2$ ✓	$X_1 \perp\!\!\!\perp X_3 \mid X_2$ ✓
$X_1 \perp\!\!\!\perp E_2 \mid X_2$ ✓	$X_1 \perp\!\!\!\perp E_2 \mid X_2$ ✓
$X_1 \perp\!\!\!\perp F_2 \mid X_2$	$X_1 \perp\!\!\!\perp F_2 \mid X_2$ ✓
$X_1 \perp\!\!\!\perp E_4 \mid X_2$ ✓	$X_1 \perp\!\!\!\perp E_4 \mid X_2$ ✓
$X_1 \perp\!\!\!\perp F_4 \mid X_2$ ✓	$X_1 \perp\!\!\!\perp F_4 \mid X_2$ ✓
$E_3 \perp\!\!\!\perp F_3 \mid X_3$ ✓	$E_3 \perp\!\!\!\perp F_3 \mid X_3$ ✓
$E_1 \perp\!\!\!\perp F_2 \mid X_2$	$E_1 \perp\!\!\!\perp F_2 \mid X_2$ ✓
$E_1 \perp\!\!\!\perp F_2 \mid E_2$	$E_1 \perp\!\!\!\perp F_2 \mid E_2$

(d) Which Bayes net is more appropriate for the problem domain here, (A) or (B)? Justify your answer.

(A) because the choice of X depends on F in the problem description.

For the following questions, your answers should be fully general for models of the structure shown above, not specific to the teleporting Jabberwock.

(e) For (A), express $P(X_{t+1}, e_{1:t+1}, f_{1:t+1})$ in terms of $P(X_t, e_{1:t}, f_{1:t})$ and the conditional probability tables used to define the network. Assume the E and F nodes are all observed.

$$P(x_{t+1}, e_{1:t+1}, f_{1:t+1}) = P(e_{t+1}|x_{t+1})P(f_{t+1}) \sum_{x_t} P(x_{t+1}|x_t, f_{t+1})P(x_t, e_{1:t}, f_{1:t}).$$

We're already provided with $P(x_t, e_{1:t}, f_{1:t})$. To get $P(x_t + 1, e_{1:t}, f_{1:t})$, we can sum over all x_t and multiply by $P(x_{t+1} \mid x_t, f_{t+1})$, the conditional probability table of x_{t+1}.
Then, to get the joint probability $P(x_t + 1, e_{1:t+1}, f_{1:t+1})$, we multiply the above quantity with the emission probability $(P(e_{t+1} \mid x_{t+1}))$ and $P(f_{t+1})$, the CPT of $P(f_{t+1})$.

(f) For (B), express $P(X_{t+1}, e_{1:t+1}, f_{1:t+1})$ in terms of $P(X_t, e_{1:t}, f_{1:t})$ and the CPTs used to define the network. Assume the E and F nodes are all observed.

$$P(x_{t+1}, e_{1:t+1}, f_{1:t+1}) = P(e_{t+1}|x_{t+1})P(f_{t+1}|x_{t+1}) \sum_{x_t} P(x_{t+1}|x_t)P(x_t, e_{1:t}, f_{1:t}).$$

Similar idea as above, except this time we multiply the joint probability by $P(x_{t+1}|x_t)$, since x_{t+1} now no longer depends on f_{t+1}).

Suppose that we don't actually observe the F_ts.

(g) For (A), express $P(X_{t+1}, e_{1:t+1})$ in terms of $P(X_t, e_{1:t})$ and the CPTs used to define the network.

$$P(x_{t+1}, e_{1:t+1}) = P(e_{t+1}|x_{t+1}) \sum_{f_{t+1}} P(f_{t+1}) \sum_{x_t} P(x_{t+1}|x_t, f_{t+1})P(x_t, e_{1:t}).$$

(h) For (B), express $P(X_{t+1}, e_{1:t+1})$ in terms of $P(X_t, e_{1:t})$ and the CPTs used to define the network.

$$P(x_{t+1}, e_{1:t+1}) = P(e_{t+1}|x_{t+1}) \sum_{x_t} P(x_{t+1}|x_t)P(x_t, e_{1:t}).$$

For (g) and (h), we essentially use the same logic as (e) and (f). However, we no longer need the F_ts in the joint probability - so for any probability values that are conditioned on an f_t, we multiply by $P(f_t)$ and sum over all possible f_t values. If not (i.e., for graph (B)), we simply drop that term when computing the joint probability.