



```
1 from pylab import *
```



Assume we have a state space consisting of  $N$  states  $0 \dots N - 1$  .

A *Markov chain* is an infinite sequence of random states  $X_i$  such that

$$P(X_{n+1} = x | X_1 = x_1 \dots X_n = x_n) = P(X_{n+1} = x | X_n = x_n)$$

That is, the probability of each subsequent state only depends on the previous state.

There are many variations:

- ▶ usually, the probability is independent of  $n$  (*time homogeneous*)
- ▶ the probability can depend on a bounded number of previous states giving a higher order Markov chain
- ▶ the set of states can be countably infinite
- ▶ the state index  $n$  can be continuous (continuous time Markov process)

We represent time homogeneous Markov chains with a finite state space as a matrix of transition probabilities  $M_{ij}$  .

Here,  $P(j \rightarrow i) = M_{ij}$  and  $\sum_i M_{ij} = 1$  .

```
1 M = rand(5,5)
```

```
1 sum(M,axis=0)
```

```
array([ 1.85666173,  0.66606344,  2.39313239,  2.00524954,  3.69764657])
```



```
1 M /= sum(M,axis=0)[newaxis,:]
2 sum(M,axis=0)
```

```
array([ 1.,  1.,  1.,  1.,  1.])
```

We call  $M$  a *transition matrix* or a *stochastic matrix*.

(Note that the matrix is often written the other way around in the literature.)

We often draw a transition diagram; this shows the state and the probabilities of transitions between them:

![[transitions]](

[http://upload.wikimedia.org/wikipedia/commons/thumb/2/2b/Markovkate\\_01.svg/220px-Markovkate\\_01.svg.png](http://upload.wikimedia.org/wikipedia/commons/thumb/2/2b/Markovkate_01.svg/220px-Markovkate_01.svg.png))

![[transitions]](

<http://upload.wikimedia.org/wikipedia/commons/thumb/4/47/MarkovChain1.png/200px-MarkovChain1.png>)

We can compute the probability of being in the next state given that we are in state 0 with a matrix multiplication.

```
1 state = zeros(5)
2 state[0] = 1
3 dot (M,state)
```

```
array([ 0.01215787,  0.40654973,  0.16527515,  0.36084789,  0.05516936])
```



```
1 v = add.accumulate(M[:,0])
```

```
1 def binsearch(a,x,lo,hi):
2     while lo<hi:
3         mid = (lo+hi)//2
4         v = a[mid]
5         if v<x: lo = mid+1
6         elif v>x: hi = mid
7         else: return mid
8     assert lo==hi
9     return lo
```



```
1 binsearch(v,0.9,0,len(M))
```

3

```
1 def rsample(dist):
2     assert amin(dist)>=0.0 and amax(dist)<=1.0
3     v = add.accumulate(dist)
4     assert abs(v[-1]-1.0)<1e-6
5     val = rand()
6     result = binsearch(v, val, 0, len(v))
7     return result
```

```
1 def sample_chain(M, state, N):
2     result = [state]
3     for i in range(N):
4         state = rsample(M[:, state])
5         result.append(state)
6     return result
```

```
1 sample_chain(M,0,10)
```

```
[0, 4, 3, 4, 3, 2, 0, 2, 2, 4, 1]
```

1 `M.shape`

`(5, 5)`

An  $n$ -gram is a sequence of  $n$  letters, syllables, words or other linguistic entity (usually words).

An  $n$ -gram model is a model of the probability that  $n$  words occur.

Usually,  $n$ -gram models are expressed as a conditional probability, given a very long sequence of words  $w_1 \dots w_L$  :

$$P(w_i | w_{i-1} \dots w_{i-n})$$

Sometimes, it might also be expressed as a joint probability:

$$P(w_i \dots w_{i-n})$$

(Boundary conditions are handled by adding special symbols to “fill things up”.)

An  $n$ -gram model defines a Markov chain.

Google has released a large  $n$ -gram model for English and other languages, and you can experiment with it here:

```
1 s = """Humpty Dumpty sat on a wall,  
2 Humpty Dumpty had a great fall;  
3 All the King's horses and all the King's men  
4 Couldn't put Humpty together again""".lower()
```

```
1 import re
2 words = re.split(r'\W+',s)
3 words[:10]
```

```
['humpty', 'dumpty', 'sat', 'on', 'a', 'wall', 'humpty', 'dumpty', 'had', 'a']
```

```
1 from collections import Counter
2 bigrams = Counter()
3 for i in range(len(words)-1):
4     bigrams[tuple(words[i:i+2])] += 1
```



```
1 bigrams
```

```
Counter({'the', 'king'): 2, ('king', 's'): 2, ('all', 'the'): 2, ('humpty', 'dumpt
```

```
1 sorted(bigrams.items())
```

```
[(('a', 'great'), 1),  
 (('a', 'wall'), 1),  
 (('all', 'the'), 2),  
 (('and', 'all'), 1),  
 (('couldn', 't'), 1),  
 ...  
 (('sat', 'on'), 1),  
 (('t', 'put'), 1),  
 (('the', 'king'), 2),  
 (('together', 'again'), 1),  
 (('wall', 'humpty'), 1)]
```

So, we see that:

$$P(\text{great}|\text{a}) = 0.5 \quad (1)$$

$$P(\text{wall}|\text{a}) = 0.5 \quad (2)$$

$$P(\text{anything else}|\text{a}) = 0.0 \quad (3)$$

Note that there are many pairs of words (even in this restricted vocabulary) that we don't see, even though their probability isn't zero.

The process of fixing this is called *smoothing*.

A simple mechanism is *pseudocounts*: here, we assume that anything that hasn't occurred has actually occurred some small number of times (1, 0.5, or smaller).

Alternatively, we can compute  $n$ -grams for different  $n$  and linearly interpolate.

Note that  $n$ -grams for  $n > 2$  give rise to higher order Markov chains.

We can transform higher order Markov chains into first order Markov chains by associating states with each context of  $n - 1$  words. However, then the emitted symbol differs from the state label.



Assume we are given a vector of probabilities  $p$  of states.

The probability distribution after one state transition is just the product  $M \cdot p$ .

```
1 p = zeros(5)
2 p[0] = 1.0
```



```
1 for i in range(10):  
2     print p  
3     p = dot(M,p)
```

```
[ 1.  0.  0.  0.  0.]  
[ 0.01215787  0.40654973  0.16527515  0.36084789  0.05516936]  
[ 0.06709196  0.0425824   0.21430012  0.36536528  0.31066025]  
[ 0.11259592  0.11789572  0.24105906  0.240486    0.28796329]  
[ 0.10613046  0.13175719  0.23142557  0.25488324  0.27580354]  
[ 0.10373109  0.12629671  0.23029253  0.26201374  0.27766592]  
[ 0.10424015  0.12572506  0.23096495  0.26074477  0.27832507]  
[ 0.1043648   0.12608833  0.23100254  0.26036384  0.27818048]  
[ 0.1043276   0.12610693  0.23095915  0.26046004  0.27814629]  
[ 0.10432159  0.12608359  0.23095882  0.26047919  0.27815681]
```

This should look familiar: after a lot of applications of the matrix, it converges to a *steady state* distribution.

This steady state distribution is an *eigenvector* of the matrix.

There are a number of concepts about Markov chains that are important:

- ▶ a state  $j$  is *accessible* from a state  $i$  if there is some path connecting the two states with non-zero probability
- ▶ two states *communicate* if they are accessible from each other
- ▶ a Markov chain is *irreducible* if all states are accessible to each other
- ▶ a state is *transient* if there is a non-zero probability that we never return to it
- ▶ a state is *recurrent* if it is not transient
- ▶ a state is *positive recurrent* if the expected recurrence time is finite
- ▶ a state is *periodic* with period  $k$  if we are guaranteed to return to it exactly every  $k$  steps
- ▶ a state is *absorbing* if there are no transitions out of it
- ▶ a Markov chain is *absorbing* if every state can reach an absorbing state
- ▶ a state is *ergodic* if it is positive recurrent and not periodic

Can we construct matrices illustrating these concepts?

A Markov chain is *reversible* if there is some probability distribution  $p$  such that

$$p \cdot M = M \cdot p$$

Note that  $p \cdot M$  is the Markov chain running in reverse.

This means that if  $p$  is the *steady state distribution*, running the Markov chain forwards or backwards gives the same result.

Reversible Markov chains are very important for modern pattern recognition, statistics, physics, and simulations. They are used in Gibbs samplers and the Metropolis algorithm for probabilistic decision making.

In Markov chains,

- ▶ the state sequence is what we are directly interested in
- ▶ we can “observe” the state sequence directly

In *Hidden Markov Models*...

- ▶ states are associated with actions or observations
- ▶ often we can only observe the actions or observations and want to infer the state

That is, each state  $x$  *emits* one symbol  $y$  from a set of  $k$  symbols.

Extremely widely used:

- ▶ speech recognition: hidden state is part-of-phoneme, observation is acoustic signal
- ▶ handwriting recognition: hidden state is part-of-handwriting, observation is pen position
- ▶ bioinformatics: hidden state is part-of-gene, observation is DNA “letter”
- ▶ natural language processing: hidden state is part-of-speech, observation is word

To specify an HMM, we need:

- ▶ sets of states, a symbol alphabet
- ▶ a matrix of transition probabilities between states
- ▶ a matrix of emission probabilities, one for each state

Usually, we are only given a training set of observations (outputs)  $Y_1 \dots Y_M$ , but may also be given a model (state transition matrix, emission probabilities).

There are several different things we may want to compute:

- ▶ given a model  $M$  and a sequence of observations  $Y$
- ▶ determine  $P_M(Y)$
- ▶ determine the most likely state sequence  $X$  (most likely explanation)
- ▶ determine the probability distribution of the next state (filtering)
- ▶ determine the probability distribution for state  $X_k$  (smoothing)
- ▶ given a set of training observations  $Y$  and constraints on the model
- ▶ find a model that assigns the overall highest likelihood to the training data

Why do we need constraints? Is there a trivial example?