

## Standard error estimates for pump example

	$m = 1000$	$m = 100$	Independent
$\lambda_1$	0.000752	0.000710	0.000075
$\lambda_2$	0.003992	0.002769	0.000205
$\lambda_3$	0.000885	0.001063	0.000111
$\lambda_4$	0.000446	0.000555	0.000094
$\lambda_5$	0.012051	0.011193	0.001009
$\lambda_6$	0.002258	0.002373	0.000439
$\lambda_7$	0.060813	0.040679	0.002970
$\lambda_8$	0.048219	0.037656	0.002807
$\lambda_9$	0.033030	0.028835	0.002945
$\lambda_{10}$	0.007568	0.007264	0.001428
$\mu$	0.019808	0.013840	0.005729
$\sigma^2$	0.139560	0.099674	0.056767

Monte Carlo and Optimization

Monte Carlo EM (MCEM) / Monte Carlo MLE

Simulated Annealing

Monte Carlo EM

Wei and Tanner (1990), Guo and Thompson (1992), Levine and Casella (2001)

As we have seen before, a function we wish to optimize might be an integral.

For example, in the EM setup, the likelihood to be optimized is of the form

$$g(X|\theta) = \int f(X, Y|\theta) dY$$

where  $x$  is the observed data and  $y$  is the “missing” data.

In EM, optimization of  $g(x|\theta)$  is performed by the optimization of the sequence of

$$Q(\theta|\theta_n) = E \left[ \log f(X, Y|\theta) | X, \theta_n \right]$$

where  $Y$  is chosen so that  $Q(\theta|\theta_n)$  is easy to determine and optimize.

In particular,  $Y$  is chosen so that the integral

$$\int \log f(X, Y|\theta) f(Y|X, \theta_n) dY$$

is easy to evaluate.

One idea is to replace the direct integration with Monte Carlo Integration.

Monte Carlo EM (Wei and Tanner version)

MC E-step

Sample  $y_1, \dots, y_m \sim f(Y|X, \theta_n)$

Let  $\hat{Q}(\theta|\theta_n) = \frac{1}{m} \sum_{i=1}^m \log f(X, y_i|\theta)$

$\hat{Q}(\theta|\theta_n)$  is a MC estimate of  $Q(\theta|\theta_n)$

MC M-step

Set

$$\theta_{n+1} = \arg \sup \hat{Q}(\theta | \theta_n)$$

Example: Genetic Linkage

$$(Y_1, Y_2, Y_3, Y_4)$$

$$\sim \text{Multi} \left( 197, \left( \frac{\lambda}{4}, \frac{1-\lambda}{4}, \frac{1-\lambda}{4}, \frac{2+\lambda}{4} \right) \right)$$

$$(X_1, X_2, X_3, X_4, X_5)$$

$$\sim \text{Multi} \left( 197, \left( \frac{\lambda}{4}, \frac{1-\lambda}{4}, \frac{1-\lambda}{4}, \frac{\lambda}{4}, \frac{1}{2} \right) \right)$$

MC E-step:

$$Q(\lambda | \lambda_n)$$

$$= E \left[ (X_1 + X_4) \log \lambda + (X_2 + X_3) \log (1 - \lambda) \mid Y, \lambda_n \right]$$

$$Q(\lambda | \lambda_n) = Y_1 \log \lambda + (Y_2 + Y_3) \log (1 - \lambda)$$

$$+ E \left[ X_4 \log \lambda \mid Y, \lambda_n \right]$$

$$= Y_1 \log \lambda + (Y_2 + Y_3) \log (1 - \lambda)$$

$$+ E \left[ X_4 \log \lambda \mid Y_4, \lambda_n \right]$$

$$= (Y_1 + \hat{X}_4) \log \lambda + (Y_2 + Y_3) \log (1 - \lambda)$$

where  $\hat{X}_4 = E[X_4 | Y_4, \lambda_n]$ .

Now  $X_4 | Y_4 \sim \text{Bin}\left(Y_4, \frac{\lambda}{\lambda + 2}\right)$  so

sample  $x_4^{(1)}, \dots, x_4^{(m)} \sim \text{Bin}(Y_4, \lambda_n)$

$$\begin{aligned}\hat{X}_4 &= \frac{1}{m} \sum_{i=1}^m x_4^{(i)} \\ &= \hat{E}[X_4 | Y_4, \lambda_n]\end{aligned}$$

M-step:

$$\lambda_{n+1} = \frac{Y_1 + \hat{X}_4}{Y_1 + Y_2 + Y_3 + \hat{X}_4}$$

Sample Run

$$m = 1000$$

$$\text{TOL} = 1 \times 10^{-6}$$


$$n_{\max} = 20$$

$$\lambda_0 = 0.5$$


Iteration	$\lambda_n$
0	0.5000000
1	0.6086751
2	0.6232812
3	0.6258296
4	0.6262932
5	0.6273085
6	0.6271988
7	0.6259806
8	0.6262454
9	0.6256674
10	0.6263483
11	0.6265943
12	0.6274692

and so up to  $n_{\max} = 20$

**Basic Monte Carlo EM**



**Basic Monte Carlo EM**



Notes:

- 1) The number of imputations at each step can vary. Often you want to start with  $m_n$  relatively small and increase it towards the end.
- 2) Any sampling scheme can be used here, independent, importance, or MCMC.
- 3) Can be combined with exact calculation. Split the missing data into two pieces. Simulate one piece and calculate conditional expectations of second piece based on simulated data. (Hybrid MCEM)
- 3) In the M-step, a maximizer is not needed, just a value that increases  $\hat{Q}(\theta|\theta_n)$ .
- 4) Unlike normal EM, this procedure does not converge to a stationary point of the likelihood surface, but to a stochastic process centered at a stationary point of the likelihood surface.
- 5) This method is similar to the Stochastic EM method of Diebolt and Ip (1996). In there approach  $m = 1$  but you keep the old samples in calculating  $\hat{Q}(\theta|\theta_n)$ .



Reweighted Monte Carlo EM (Irwin, 1995, Levine and Casella, 2001)

Sampling step

$$\text{Sample } y_1, \dots, y_m \sim f(Y | X, \theta^*)$$

MC E-step

Update weights

$$w_n(y_i) = \frac{f(X, y_i | \theta_n)}{f(X, y_i | \theta^*)}$$

$$\tilde{w}_n(y_i) = \frac{w_n(y_i)}{\sum w_n(y_j)}$$

$$\text{Let } \hat{Q}(\theta | \theta_n) = \sum_{i=1}^m \tilde{w}_n(y_i) \log f(X, y_i | \theta)$$

$\hat{Q}(\theta | \theta_n)$  is a MC estimate of  $Q(\theta | \theta_n)$

MC M-step

Set

$$\theta_{n+1} = \arg \sup \hat{Q}(\theta | \theta_n)$$

## Sample Run


$$m_1 = m_2 = 1000, m_3 = 10000$$

$$\text{TOL} = 1 \times 10^{-6}$$


$$n_{\max} = 20$$

$$\text{Pass 1: } \lambda_0 = \lambda^* = 0.5$$


Pass 2 & 3:  $\lambda_0 = \lambda^*$  = estimate from previous pass




**Reweighted MCEM - Pass 1**



**Reweighted MCEM - Pass 2**





### Notes:

- 1) Usually you want to run this algorithm at least twice. Once to get in a region of the optimum, and the second one to tighten up the estimate.
- 2) The number of imputations at each step can vary. Often you want to start with  $m_1$  to be moderate in size and increase it towards the end.
- 3) Particularly useful when simulation is time consuming. Reweighting is often much quicker and easier than sampling.

- 4) Any sampling scheme can be used here, independent, importance sampling, or MCMC. Hybrid MCEM also works here and is preferable if possible.
- 5) In the M-step, a maximizer is not needed, just a value that increases  $\hat{Q}(\theta|\theta_n)$
- 6) The starting point in the first E-step  $\theta_0$  need not be the simulation point  $\theta^*$ . In fact using different  $\theta_0$  allows for easy searching for multiple modes.
- 7) Weight properties

$$w_\theta(y) = \frac{p_\theta(y)}{p_{\theta^*}(y)}$$

$$E[w_\theta(Y)|X, \theta^*] = \frac{g(X|\theta)}{g(X|\theta^*)} = l(\theta, \theta^*)$$

$$\hat{l}(\theta, \theta^*) = \frac{1}{m} \sum_{i=1}^m w_\theta(y_i)$$

As seen before  $\hat{l}(\theta, \theta^*)$  is an unbiased estimator of  $l(\theta, \theta^*)$ .

8) Convergence properties

i)  $\hat{l}(\theta_{n+1}, \theta^*) \geq \hat{l}(\theta_n, \theta^*)$ .

ii) procedure converges to stationary point of  $\hat{l}(\theta, \theta^*)$ .

9) Properties of MC MLE:

MLE:  $\hat{\theta}$  maximizer of  $g(X|\theta)$  (or  $l(\theta, \theta^*)$ ).

MC MLE:  $\tilde{\theta}$  maximizer of  $\hat{l}(\theta, \theta^*)$ .

i)  $\tilde{\theta}$  is a consistent estimator of  $\hat{\theta}$ .

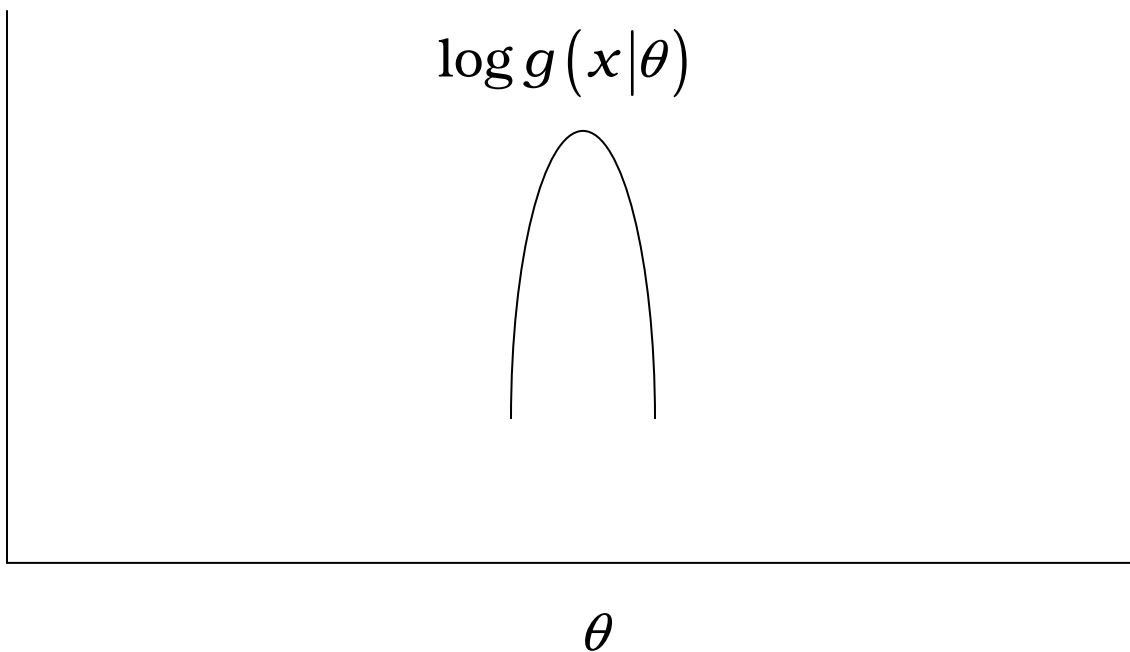

ii) Asymptotically normal

iii)  $\text{Var}(\tilde{\theta}) = \frac{1}{m} I^{-1}(\hat{\theta}) V(\hat{\theta}) I^{-1}(\hat{\theta})$

where

$$V(\theta) = \text{Var} \left( \frac{f(y|x, \theta)}{f(y|x, \theta^*)} \frac{D \log f(y|x, \theta)}{D\theta} \Big|_{x, \theta^*} \right)$$

$$I(\theta) = \frac{D^2 \log f(x|\theta)}{D^2 \theta}$$



This variance formula may only hold when the sampler generates independent realizations. I'm not sure about MCMC.

## Standard Errors

Louis' formula can easily be approximated by Monte Carlo. For Reweighted MCEM, the estimator of the information matrix is

$$\begin{aligned}\hat{I}(\theta) = & -\sum \tilde{w}(y_i) \frac{D^2 \log f(x, y_i | \theta)}{D^2 \theta} \\ & - \sum \tilde{w}(y_i) \left( \frac{D \log f(x, y_i | \theta)}{D \theta} \right)^2 \\ & + \left( \sum \tilde{w}(y_i) \frac{D \log f(x, y_i | \theta)}{D \theta} \right)^2\end{aligned}$$

## Simulating Annealing

Another Monte Carlo approach for optimizing a function.

Approach for determining global optima, not local.

Useful for high dimensional, multimodal, complicated functions.



## Examples where useful

- 1) Phylogentic tree reconstruction (Salter, 2000)

Want to infer the evolutionary history based  $n$  observations (morphology, genetic, etc data) and describe it with a binary tree.

Number of possible unrooted binary trees is

$$\prod_{i=1}^{n-1} (2i - 3)$$

For  $n = 30$ , this is  $8.69 \times 10^{36}$ .

- 2) Traveling Salesman Problem

A salesman has  $n$  towns to be visited. Wants to minimize the traveling distance and finish at his starting point

This is known to be an NP complete problem

The name comes from condensed matter physics. Annealing is a thermal process for obtaining low-energy states of a solid in a heat path. The procedure has two steps

- 1) Raise the temperature of the heat bath enough for the solid (metal) to melt.
- 2) Decrease the temperature slowly to near zero so that the particles in the system can arrange themselves in the ground state of the solid (i.e. crystallize).

Problem: Want to minimize a function  $h(x)$  or maximize a function  $-h(x)$ . This is equivalent to maximizing the function  $\exp(-h(x)/T)$  for any given temperature  $T$ .

Method:

- 1) Let  $T_1 > T_2 > T_3 > \dots$  be a sequence of temperatures in which  $T_1$  is relatively large and  $\lim_{k \rightarrow \infty} T_k = 0$ .

2) At each temperature  $T_k$  run  $N_k$  steps of a M-H sampler with stationary distribution

$$\pi_k(x) \propto \exp(-h(x)/T_k)$$

Pass the final configuration of  $x$  to the next iteration and set  $k$  to  $k + 1$ .

Why does SA work:

For any system in which

$$\int \exp(-h(x)/T) dx < \infty$$


for all  $T$ , distribution  $\pi_k$ , as  $k$  increases puts more and more of its probability mass (converging to 1) into the vicinity of the global maximum of  $-h(x)$ .





Hence a sample from  $\pi_k$  will almost certainly be in the vicinity of the global optimum of  $h(x)$  when  $T_k$  is close to 0.

And as  $T \rightarrow \infty$ ,  $\exp(-h(x)/T)$  goes to a Uniform distribution

## Example function

$$-h(x) = \cos(14.5x - 0.3) + \exp(-(x-1)^2)$$





In addition, as we have seen (sort of) M-H algorithms will tend to accept more likely proposals. They definitely will with symmetric proposal distributions.

So SA will tend to move to points when  $-h(x)$  increases, but can move to points where  $-h(x)$  is lower so it can move out of a local, but lower, mode, into a better mode.

It can be shown that the global optimum can be reached with SA with probability 1 if the temperature variable  $T_k$  decreases sufficiently slowly enough ( $O(\log L_k^{-1})$ ) where

$$L_k = N_1 + \dots + N_k.$$

This slow temperature decrease is required to make sure the chain has run long enough to get into the dominant mode.

The large  $T_k$  allow for the space to be completely sampled and the small  $T_k$  focus on the global optimum.

In practice cooling schedules like these aren't used. Linear or exponential is more common as these give more manageable computational burden.

However you are no longer guaranteed to find the global optima. However you usually find a reasonable answer.