Algorithm 11.7 The Mean-Field approximation algorithm

Procedure Mean-Field (
  Φ, // factors that define P_Φ
  Q_0 // Initial choice of Q
)
1 $Q \leftarrow Q_0$
2 $\mathit{Unprocessed} \leftarrow \mathcal{X}$
3 **while** $\mathit{Unprocessed} \neq \emptyset$
4 Choose X_i from *Unprocessed*
5 $Q_{old}(X_i) \leftarrow Q(X_i)$
6 **for** $x_i \in \mathit{Val}(X_i)$ **do**
7 $Q(x_i) \leftarrow \exp\left\{\sum_{\phi: X_i \in \mathit{Scope}[\phi]} \boldsymbol{E}_{(\boldsymbol{U}_\phi - \{X_i\}) \sim Q}[\ln \phi[\boldsymbol{U}_\phi, x_i]]\right\}$
8 Normalize $Q(X_i)$ to sum to one
9 **if** $Q_{old}(X_i) \neq Q(X_i)$ **then**
10 $\mathit{Unprocessed} \leftarrow \mathit{Unprocessed} \cup \left(\cup_{\phi: X_i \in \mathit{Scope}[\phi]} \mathit{Scope}[\phi]\right)$
11 $\mathit{Unprocessed} \leftarrow \mathit{Unprocessed} - \{X_i\}$
12 **return** Q

to another marginal $Q(X_j)$ may result in a different optimal parameterization for $Q(X_i)$. Thus, the algorithm repeats these steps until convergence. Note that, in practice, we do not test for equality in line 9, but rather for equality up to some fixed small-error tolerance.

A key property of the coordinate ascent procedure is that each step leads to an increase in the energy functional. **Thus, each iteration of** Mean-Field **results in a better approximation Q to the target density P_Φ, guaranteeing convergence.**

Theorem 11.10 *The* Mean-Field *iterations are guaranteed to converge. Moreover, the distribution Q^* returned by* Mean-Field *is a stationary point of $F[\tilde{P}_\Phi, Q]$, subject to the constraint that $Q(\mathcal{X}) = \prod_i Q(X_i)$ is a distribution.*

Proof We showed earlier that each iteration of Mean-Field is monotonically nondecreasing in $F[\tilde{P}_\Phi, Q]$. Because the energy functional is bounded, the sequence of distributions represented by successive iterations of Mean-Field must converge. At the convergence point the fixed-point equations of theorem 11.9 hold for all the variables in the domain. As a consequence, the convergence point is a stationary point of the energy functional. ∎

As we discussed, the distribution Q^* returned by Mean-Field is not necessarily a local optimum of the algorithm. However, local minima and saddle points are not stable convergence points of the algorithm, in the sense that a small perturbation of Q followed by optimization will lead to a better convergence point. Because the algorithm is unlikely to accidentally land precisely on the unstable point and get stuck there, in practice, the convergence points of the algorithm are local maxima.

In general, however, the result of the mean field approximation is a local maximum, and not necessarily a global one.