

# Lecture 2: Gaussian Mixture Models; Expectation-Maximization


Lester Mackey

April 2, 2014


# Blackboard discussion

- See lecture notes

# GMM EM Example




# GMM EM Example







# GMM EM Example




# GMM EM Example



# GMM EM Example




# GMM EM Example




# *k*-means vs. GMM EM


Credit:  
Michael I. Jordan




(a)




(b)




(c)



(d)




(e)




(f)

## *k*-means iterates


- EM often converges more slowly than *k*-means
- Common to initialize EM with *k*-means solution




(a)




(b)




(c)




(d)




(e)




(f)



(g)



(h)



(i)

## GMM EM iterates 9