Deterministic Approximations 1

Motivation:
Representing distributions more compactly and often more quickly than a bag of samples from MCMC

Iain Murray
http://iainmurray.net/

Example: inferring dark matter

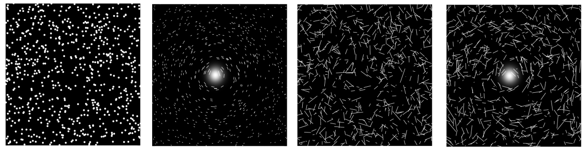

A. Distant circular galaxies (or dots in this case) are randomly distributed in the sky. Each galaxy has an (x,y) coordinate corresponding to the position in the sky from 0:4200

B. By placing a Dark Matter halo in the middle of the sky between us and the background galaxies, they are altered such that they become elliptical. The lines show the orientation and size of the major axis of the galaxy.

C. However unfortunately galaxies are NOT circular and infact they are inherently elliptical. This property is random, however since the Universe has no preferred ellipticity this averages out to zero in the case of no other influence.

D. Therefore if we placed a Dark Matter halo into a field of randomly elliptical galaxies we would get a field that does not average out to zero. If we can use the fact that Dark Matter makes the pattern seen in B, we should be able to detect the position of the central halo.

http://www.kaggle.com/c/DarkWorlds
http://homepages.inf.ed.ac.uk/imurray2/pub/12kaggle_dark/

These slides are for motivation. Details about dark matter non-examinable!

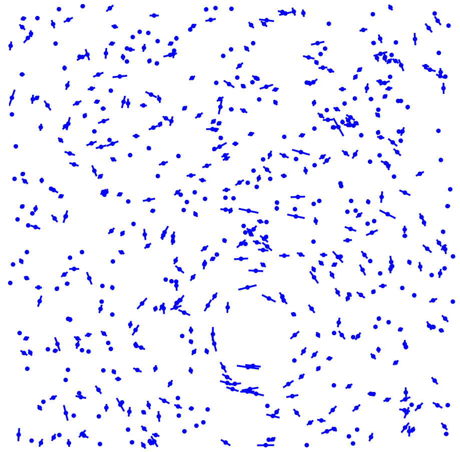

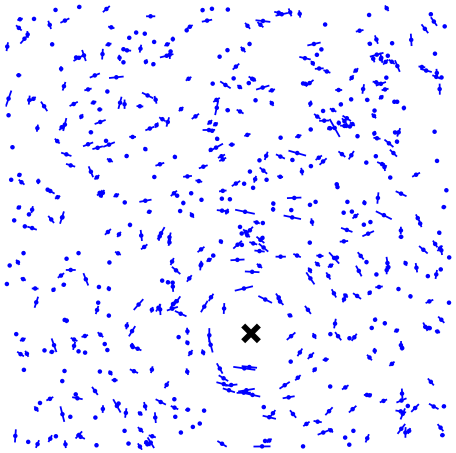

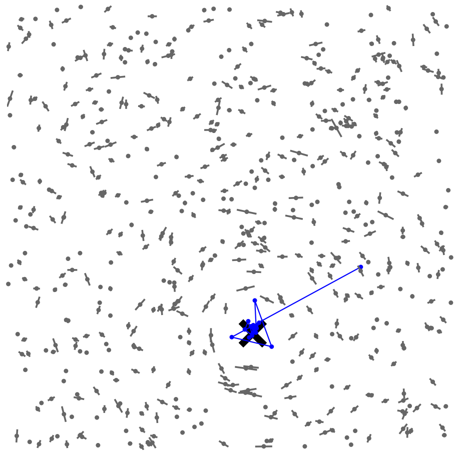

Answer is obvious

Could optimize likelihood of dark matter position

Faster than MCMC!

Want some way to report error bars though.

Usually dark matter locations *not* obvious. . .

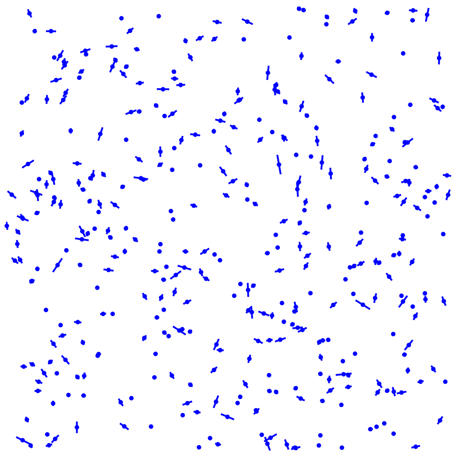

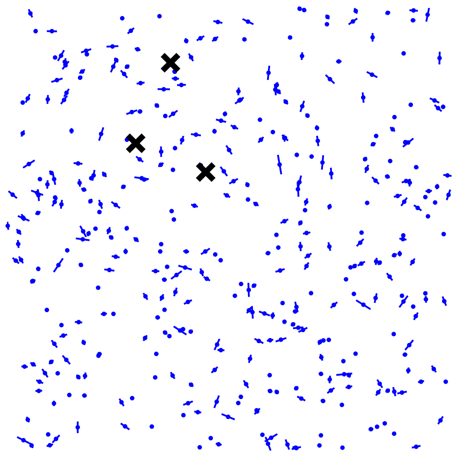

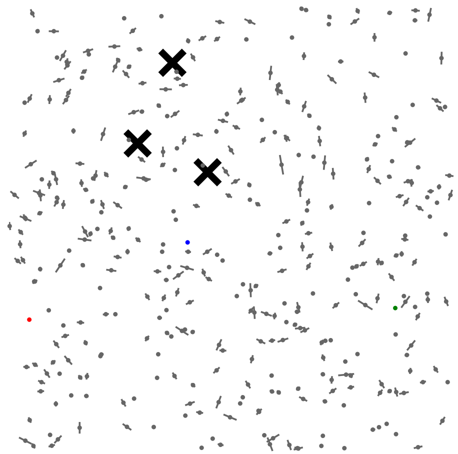

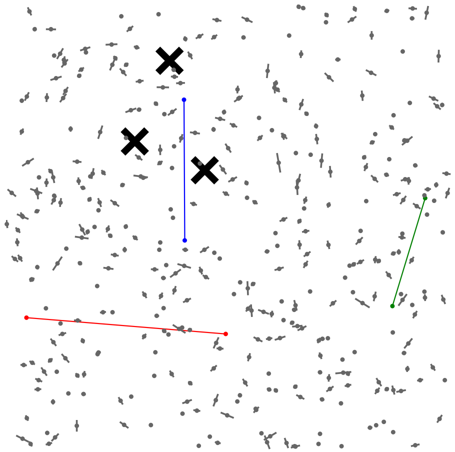

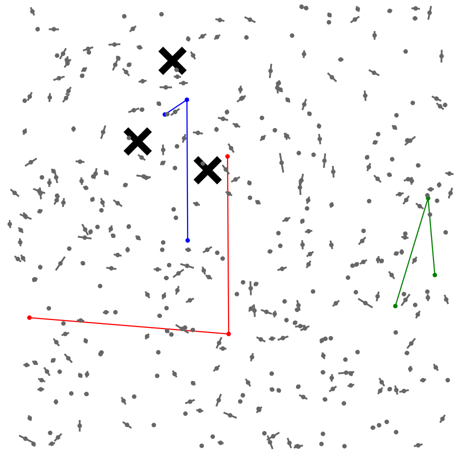

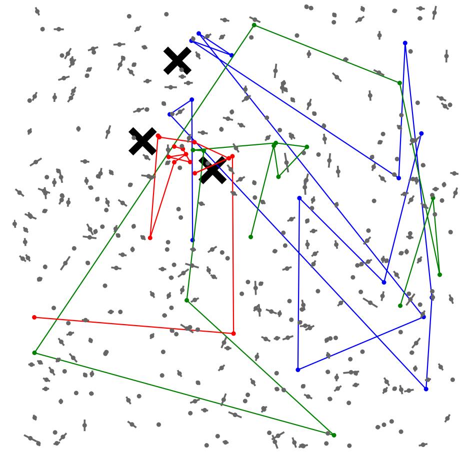

. . .

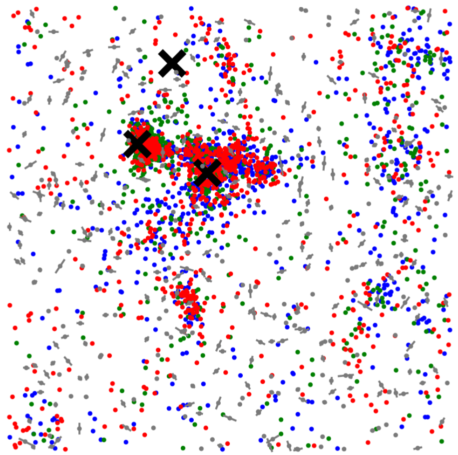

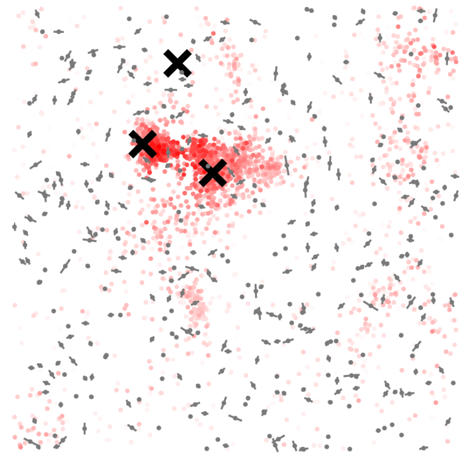

Summarizing beliefs?

— Average/mean sample?

— Most probable sample?

— Cluster?

I have several answers. But still a research question. For this course:

Some complicated distributions most easily represented by samples

Then predict under each possible world

Lower dimensional example

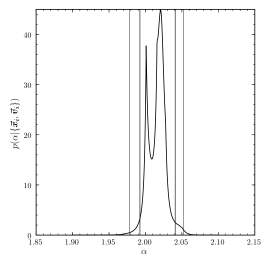

A posterior over some quantity α from http://iopscience.iop.org/0004-637X/711/2/1157/

Might summarize with the vertical credible intervals containing 95% and 99% of probability mass.

Mean, mode, median?

Weirder example

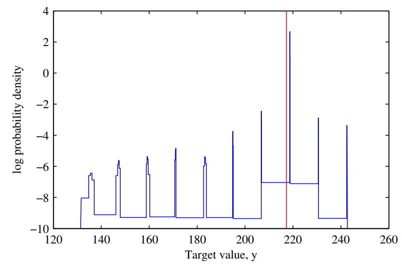

Red vertical bar is the mean (not a probable point, note log scale)

Median? Mode?

from http://link.springer.com/chapter/10.1007%2F11736790_3

Gaussian approximations

Finite parameter vector θ

$P(\theta\,|\,lots\ of\ data)$ often nearly Gaussian around the mode

Need to identify which Gaussian it is: mean, covariance