

Mean-Shift Tracker

16-385 Computer Vision

Mean Shift Algorithm

A 'mode seeking' algorithm

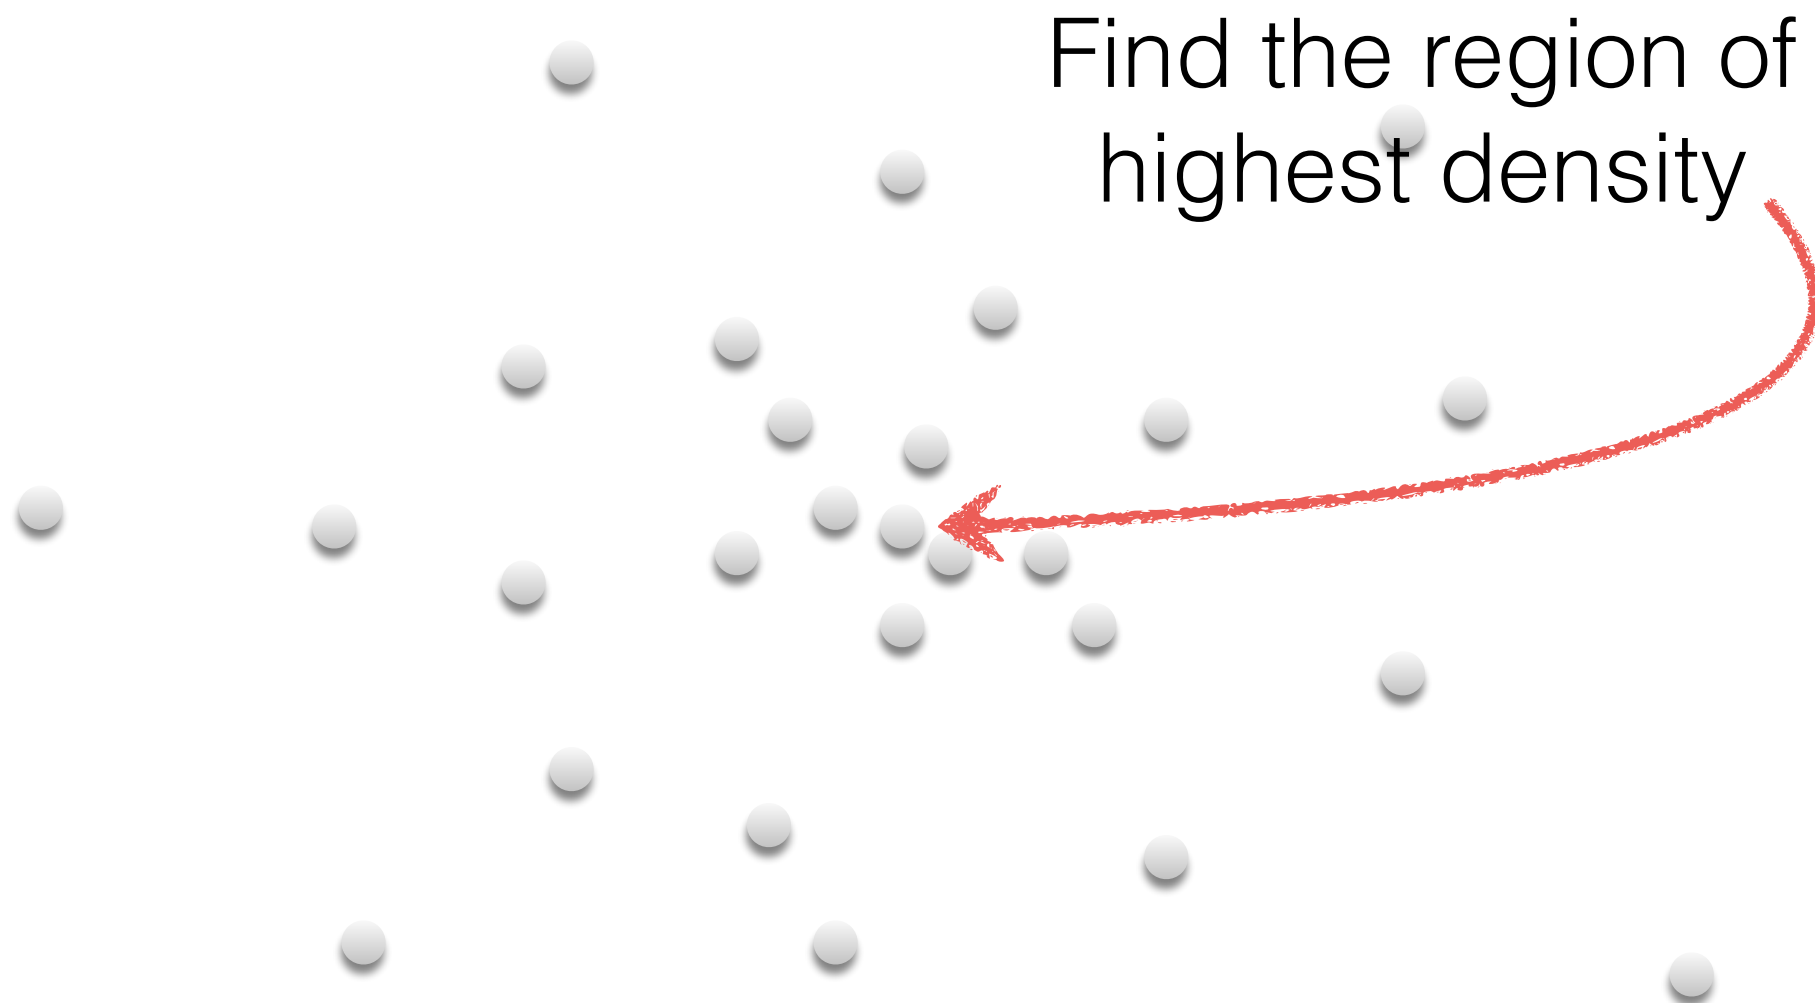
Fukunaga & Hostetler (1975)



Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)



Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

Pick a point

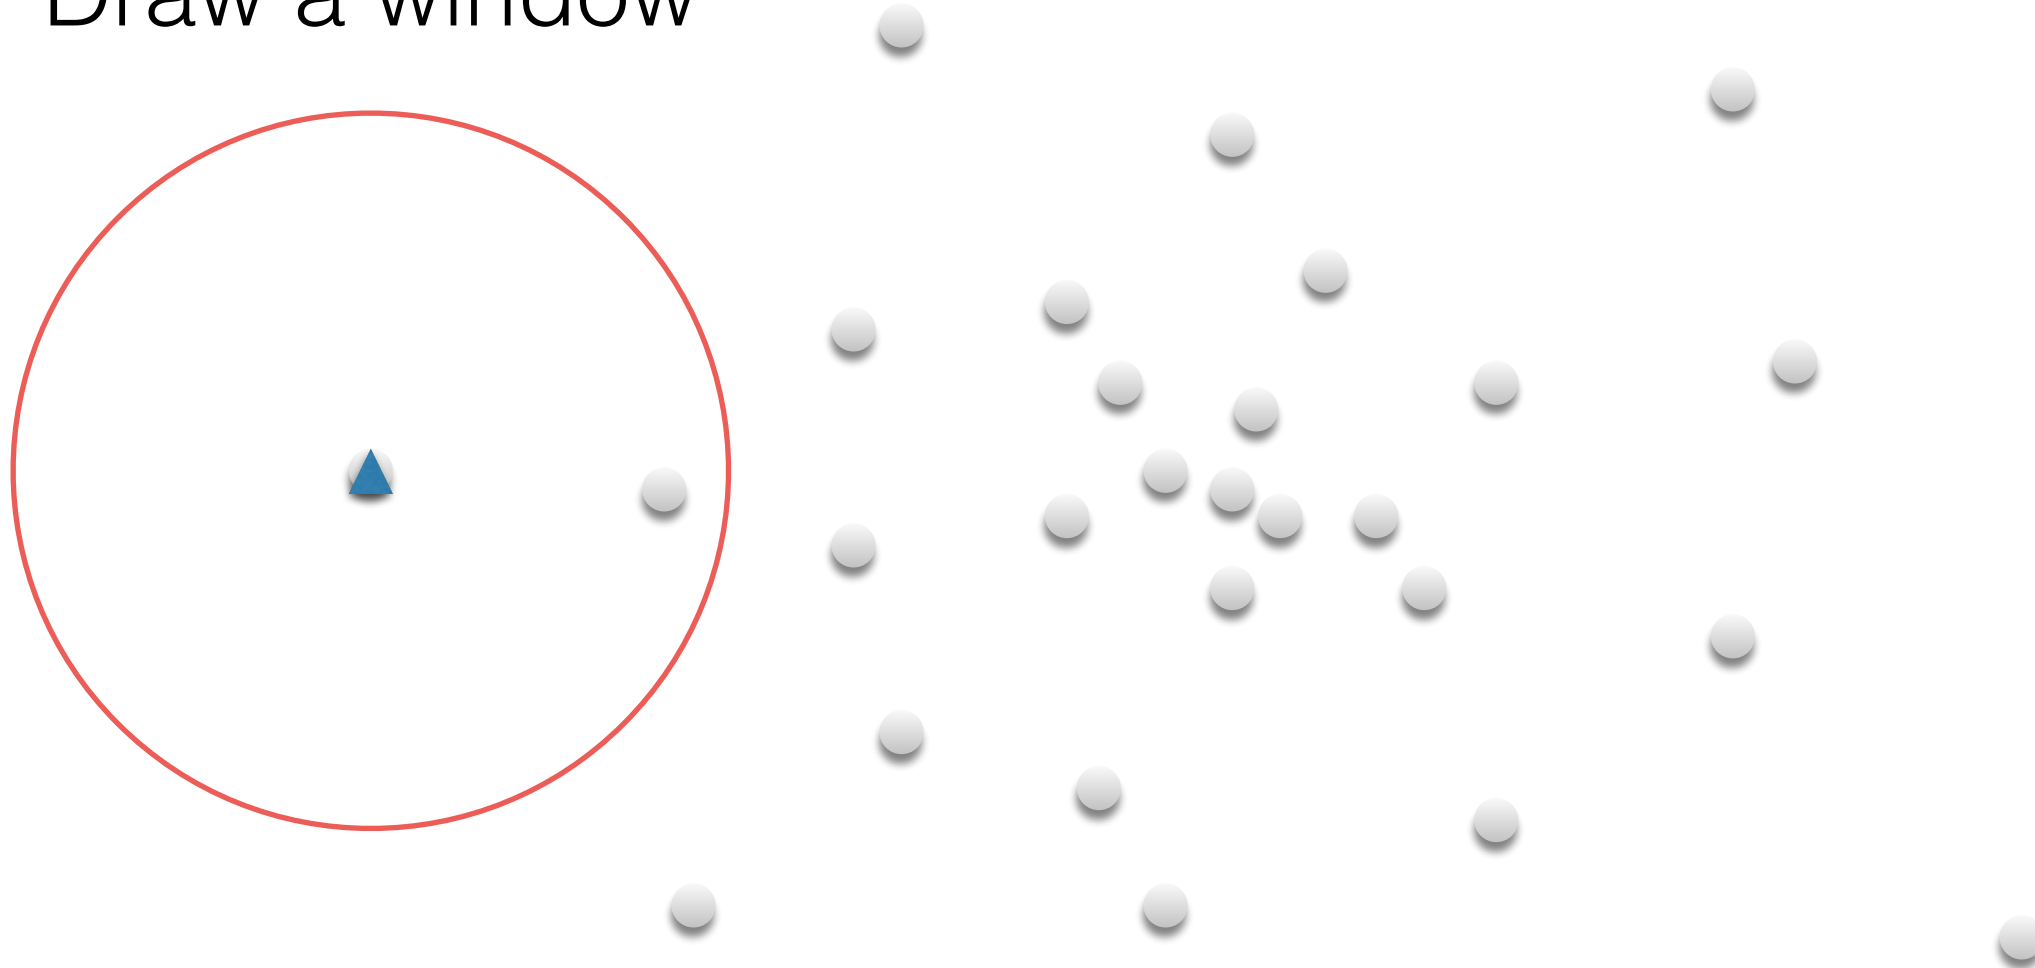


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

Draw a window

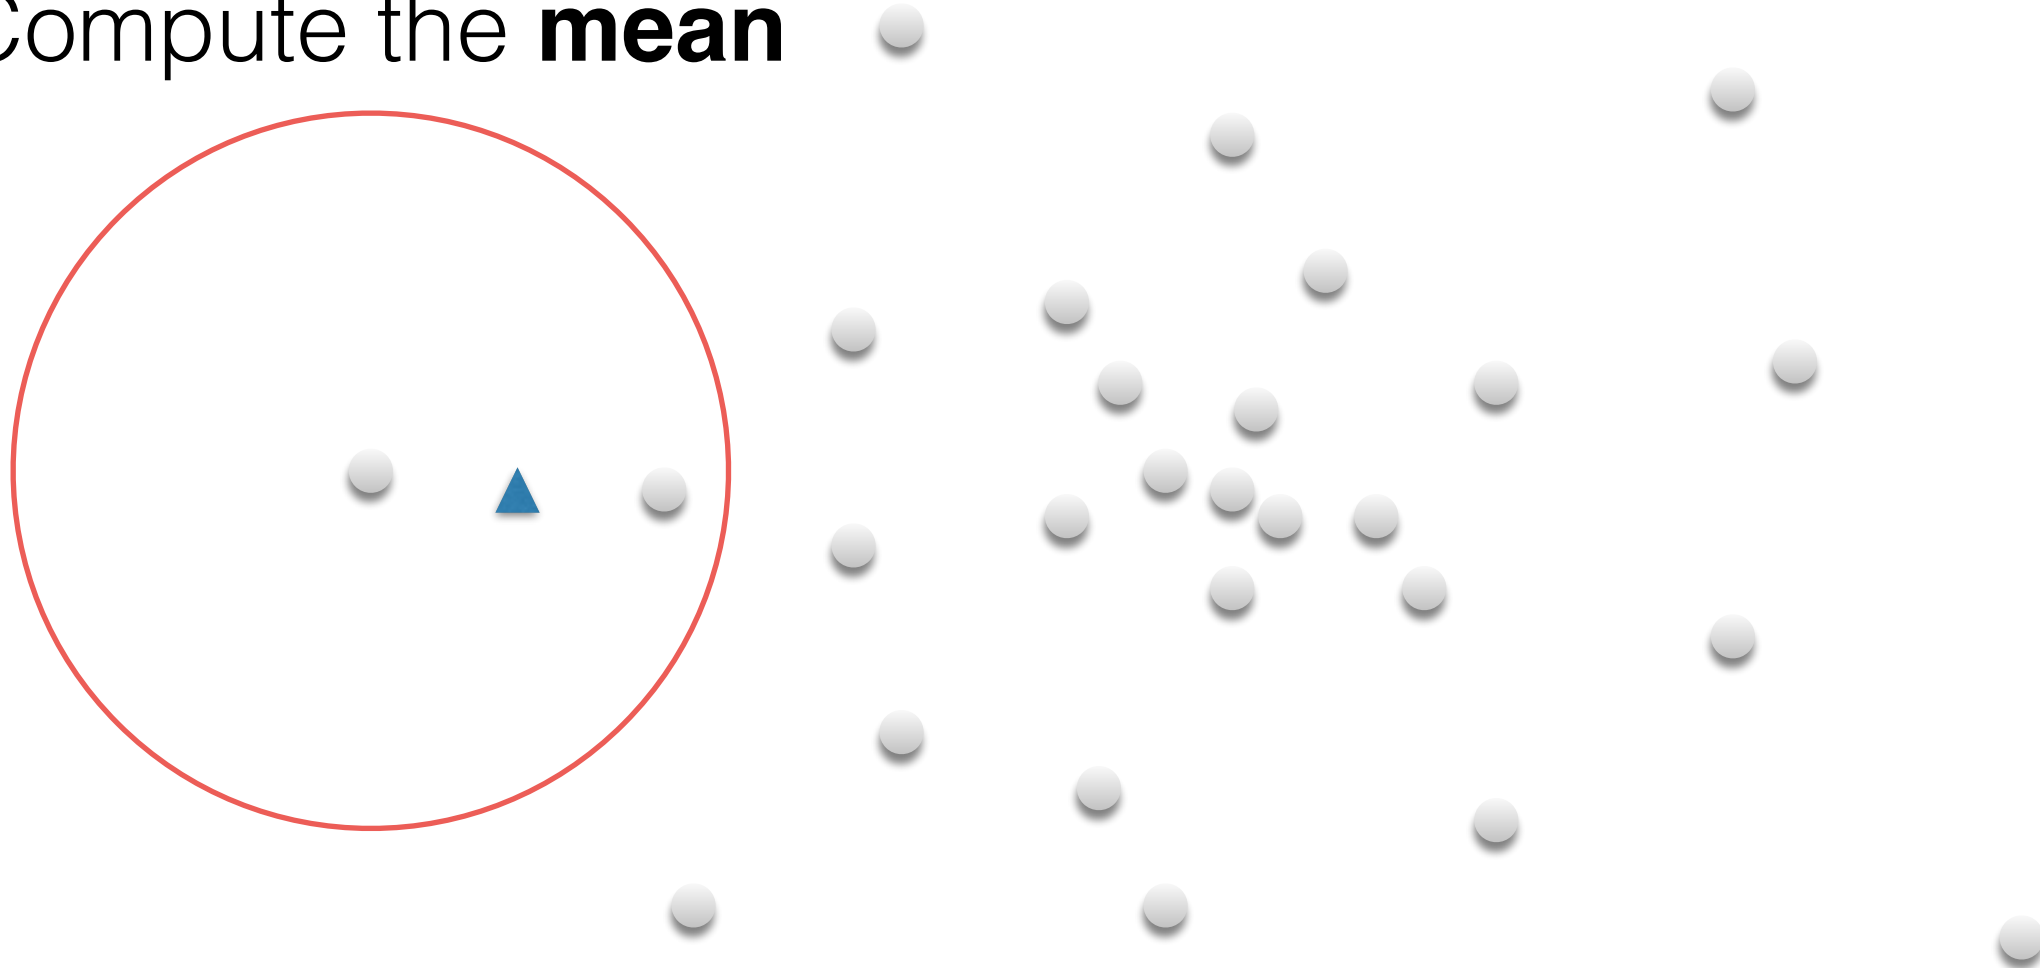


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

Compute the **mean**

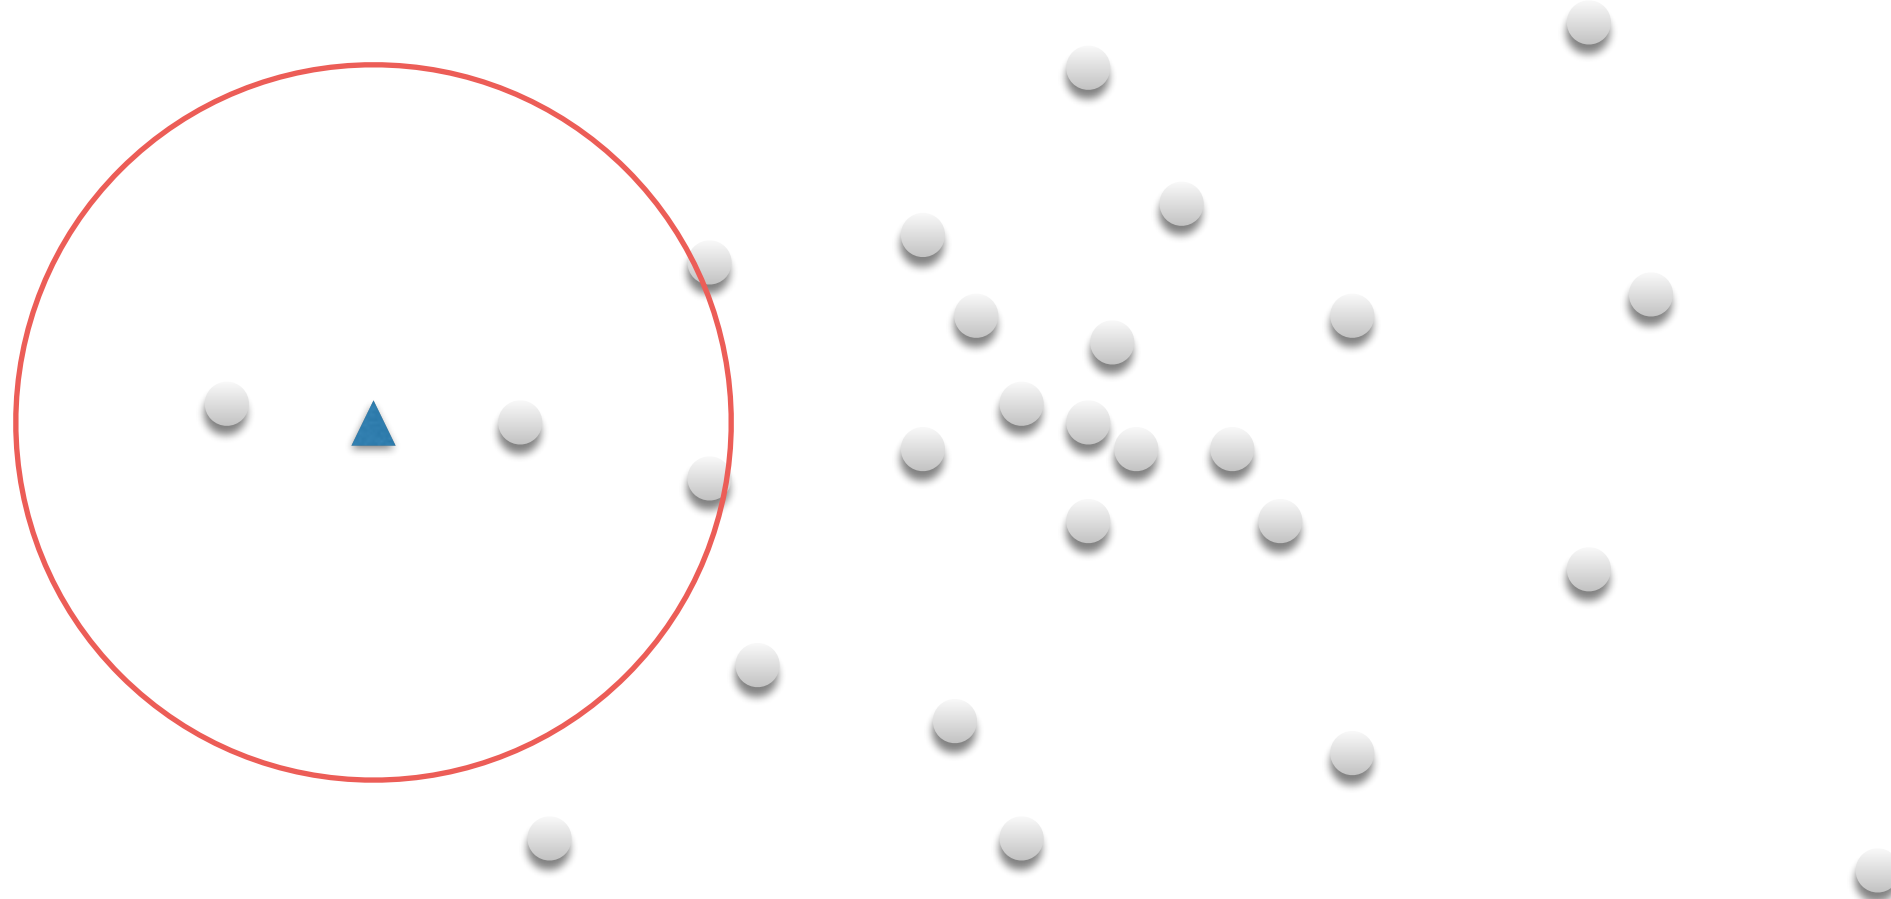


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

Shift the window

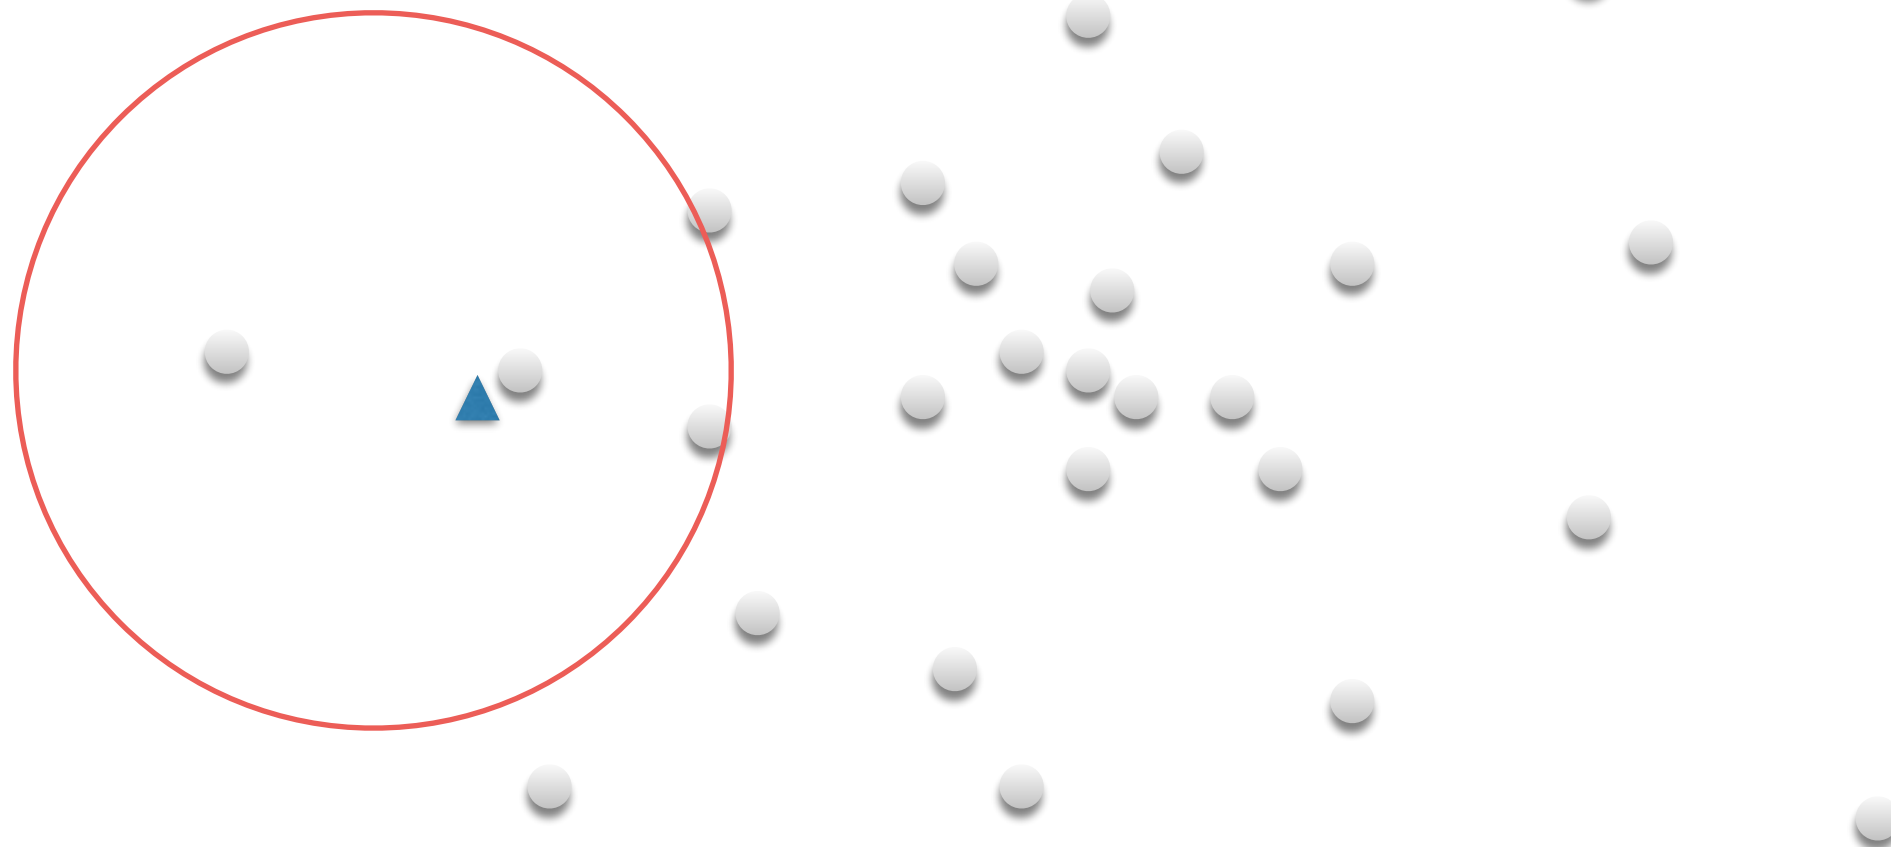


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

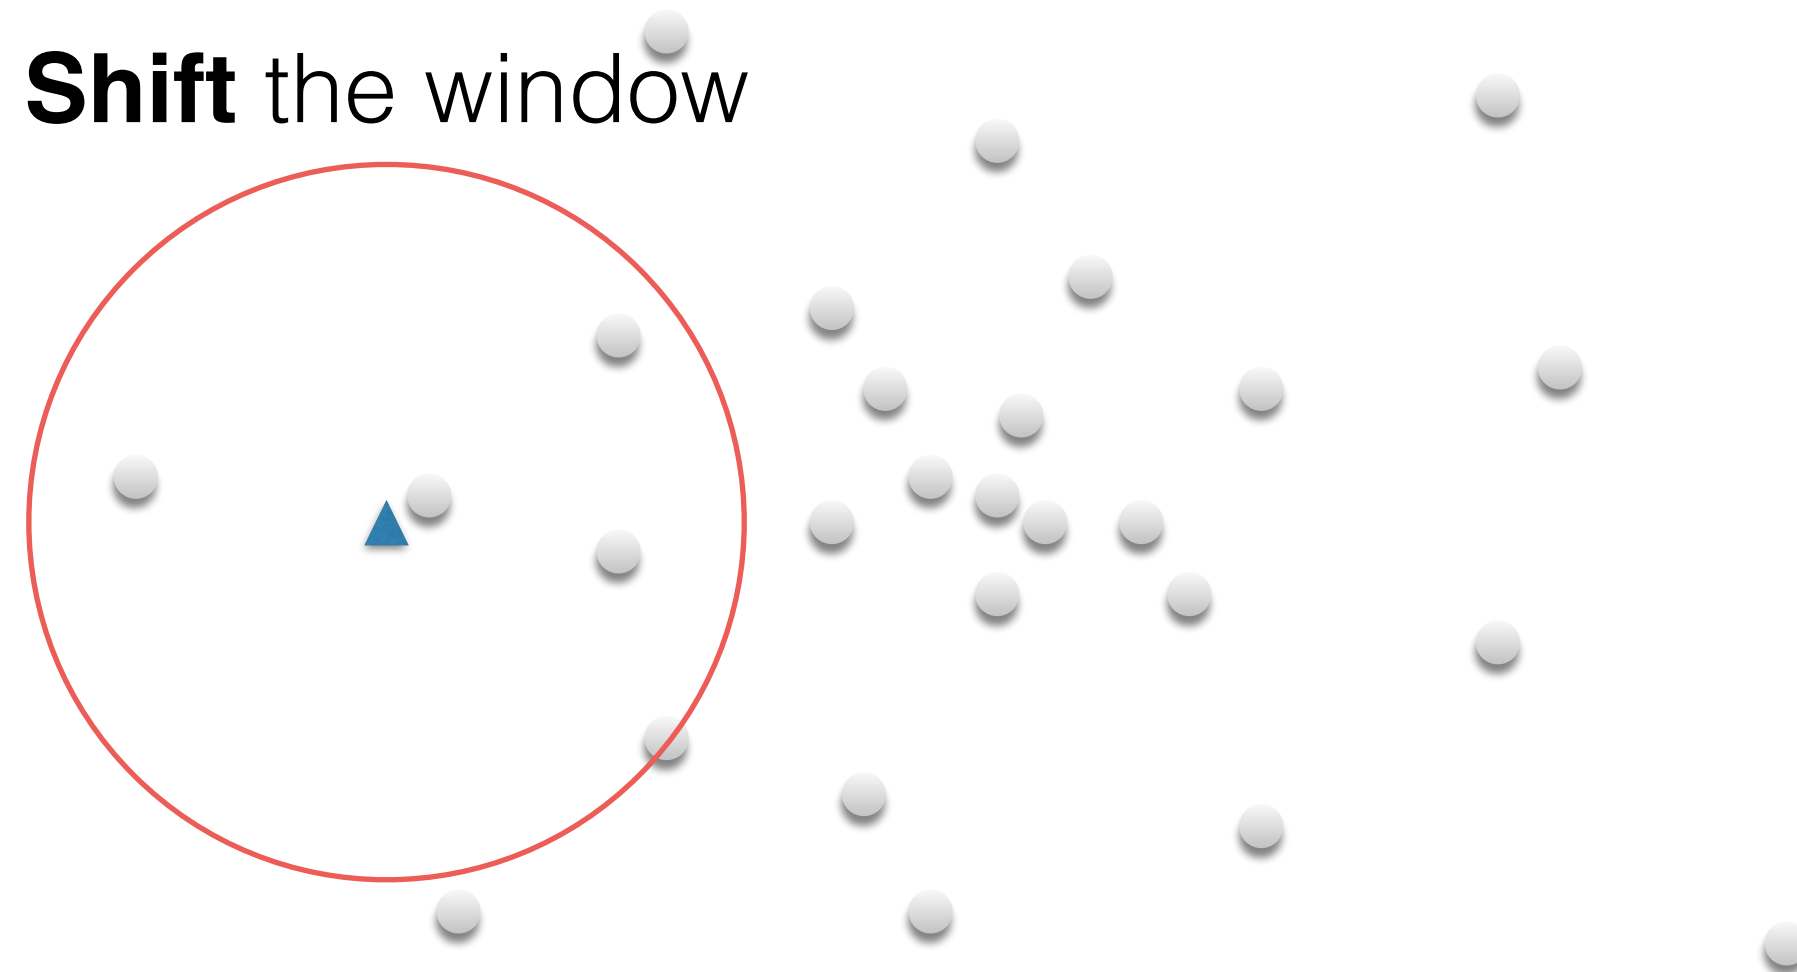
Compute the **mean**



Mean Shift Algorithm

A 'mode seeking' algorithm

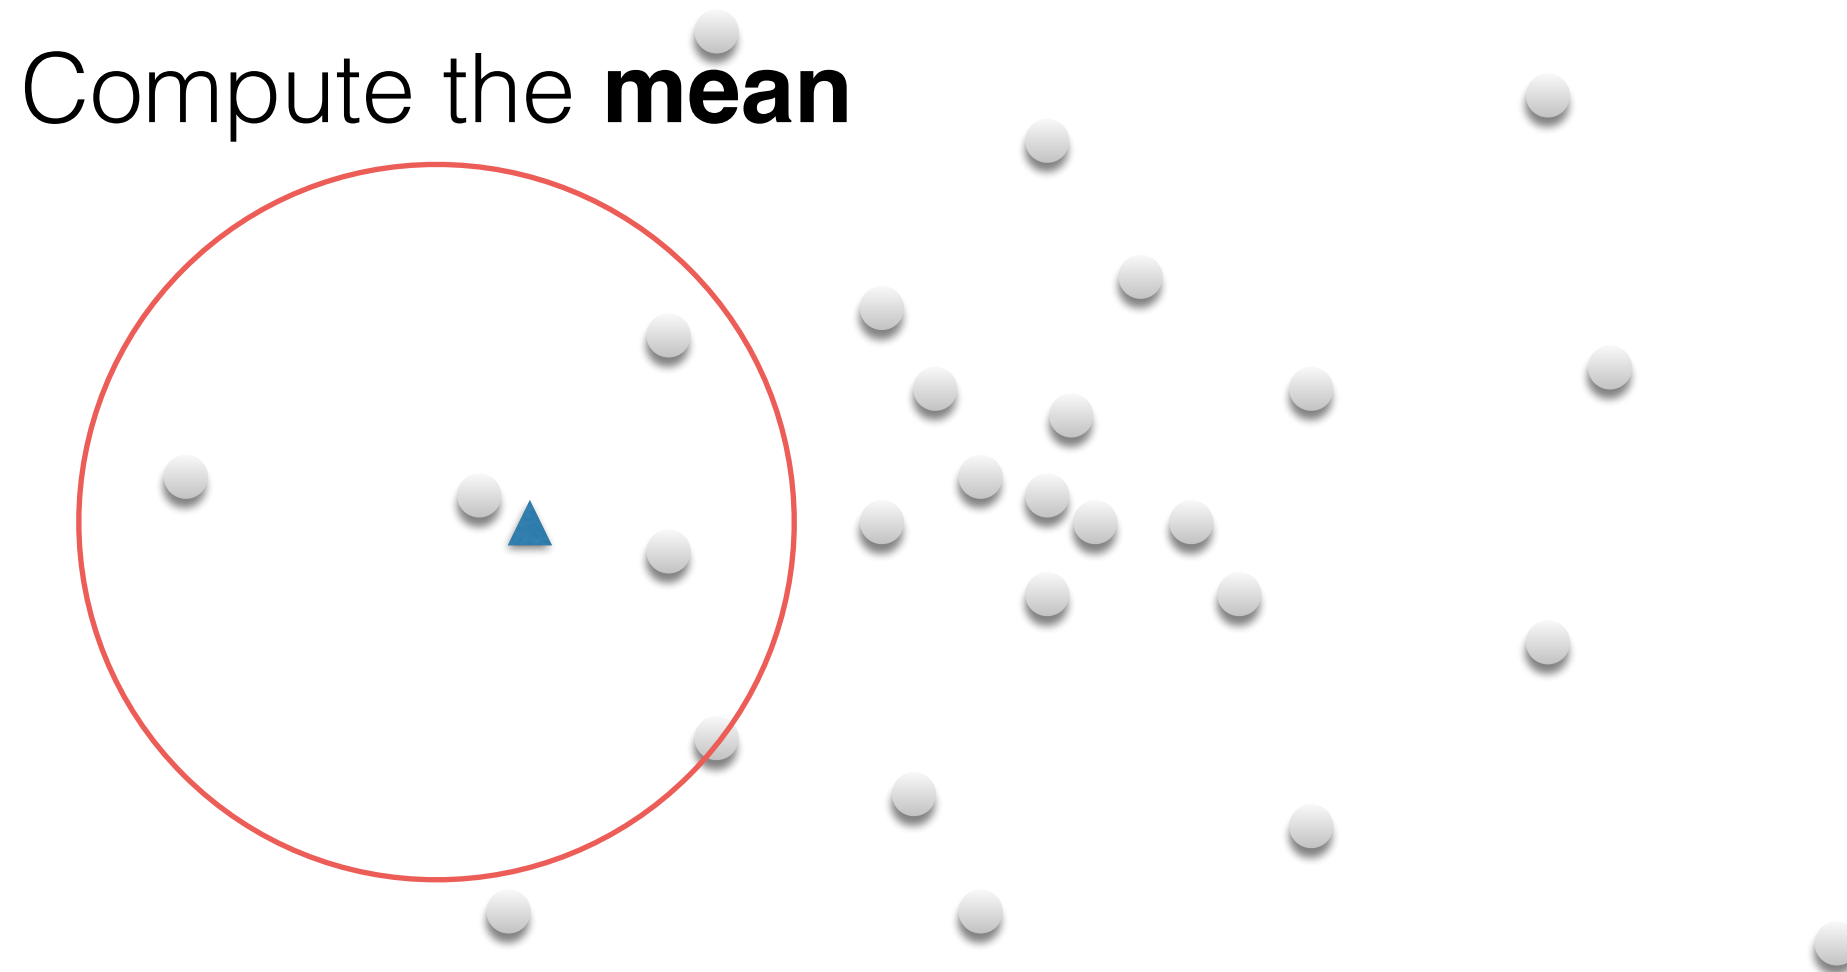
Fukunaga & Hostetler (1975)



Mean Shift Algorithm

A 'mode seeking' algorithm

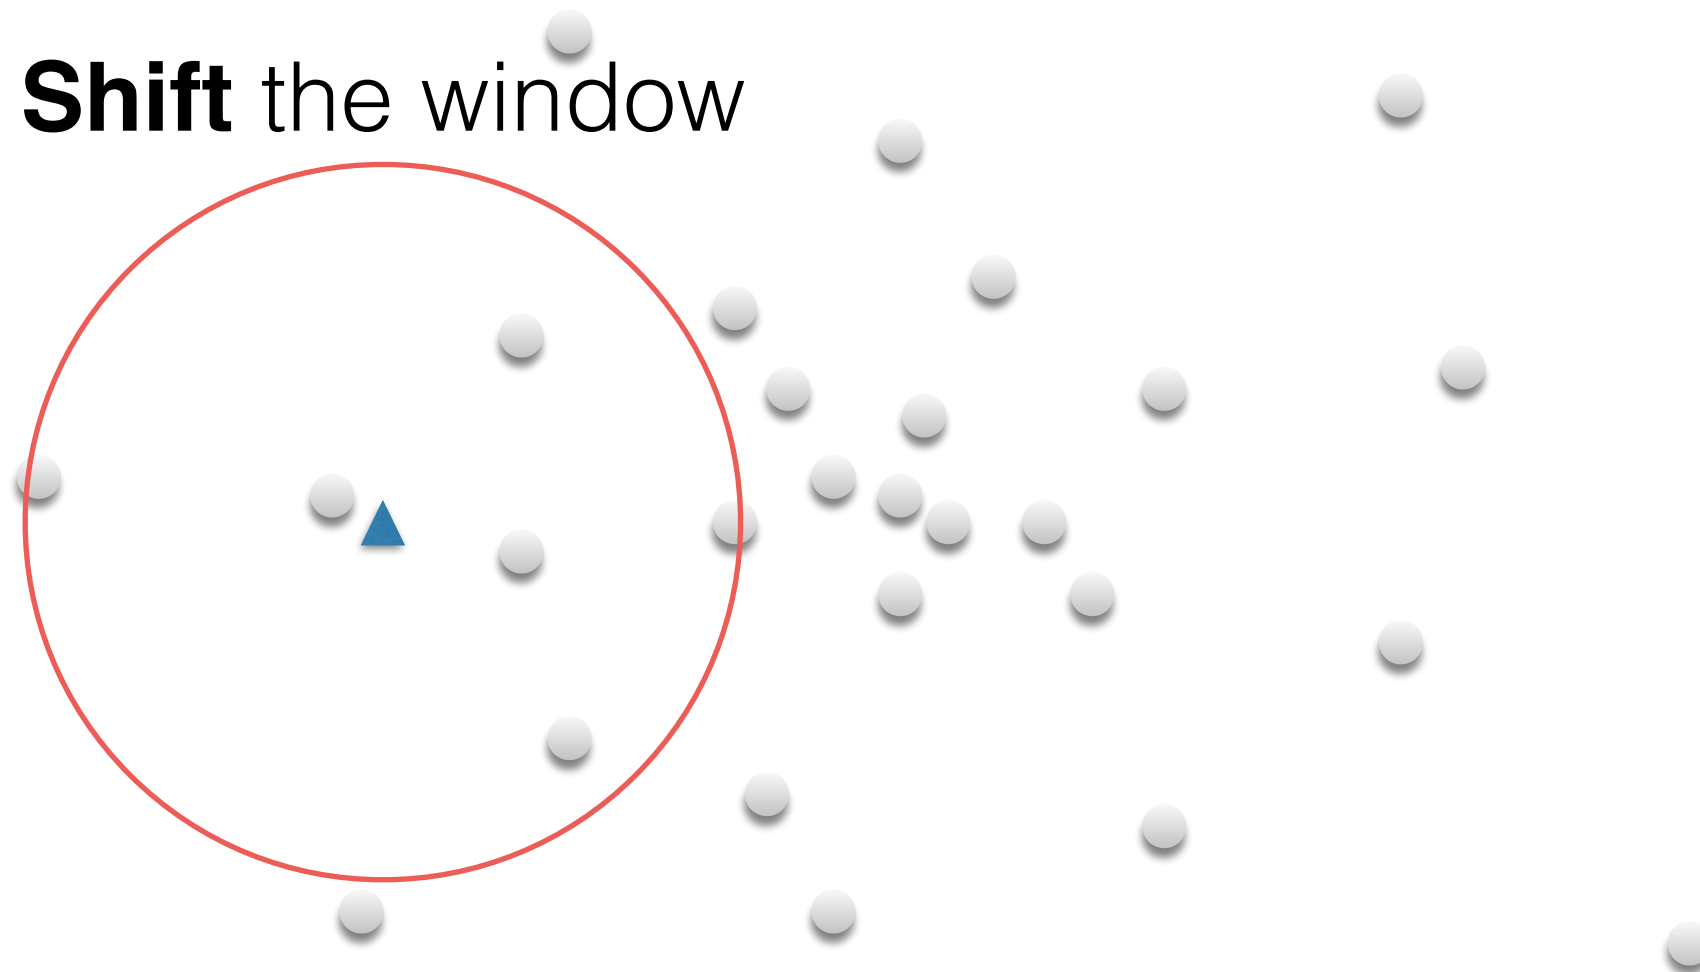
Fukunaga & Hostetler (1975)



Mean Shift Algorithm

A 'mode seeking' algorithm

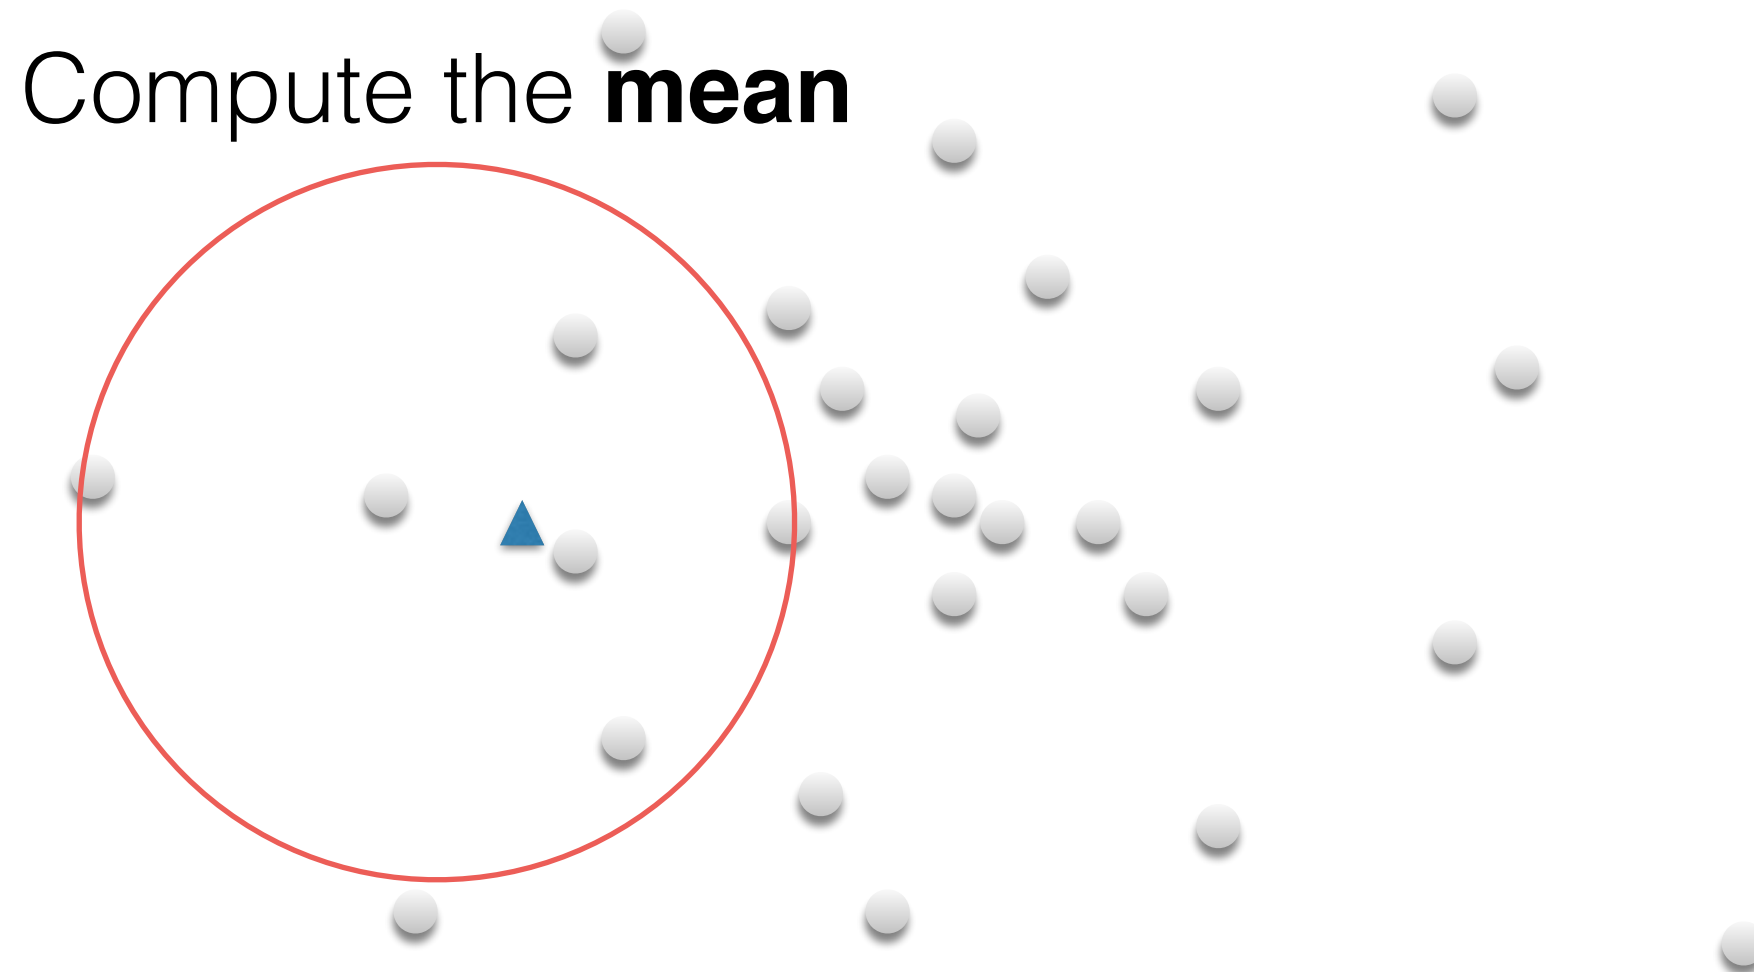
Fukunaga & Hostetler (1975)



Mean Shift Algorithm

A 'mode seeking' algorithm

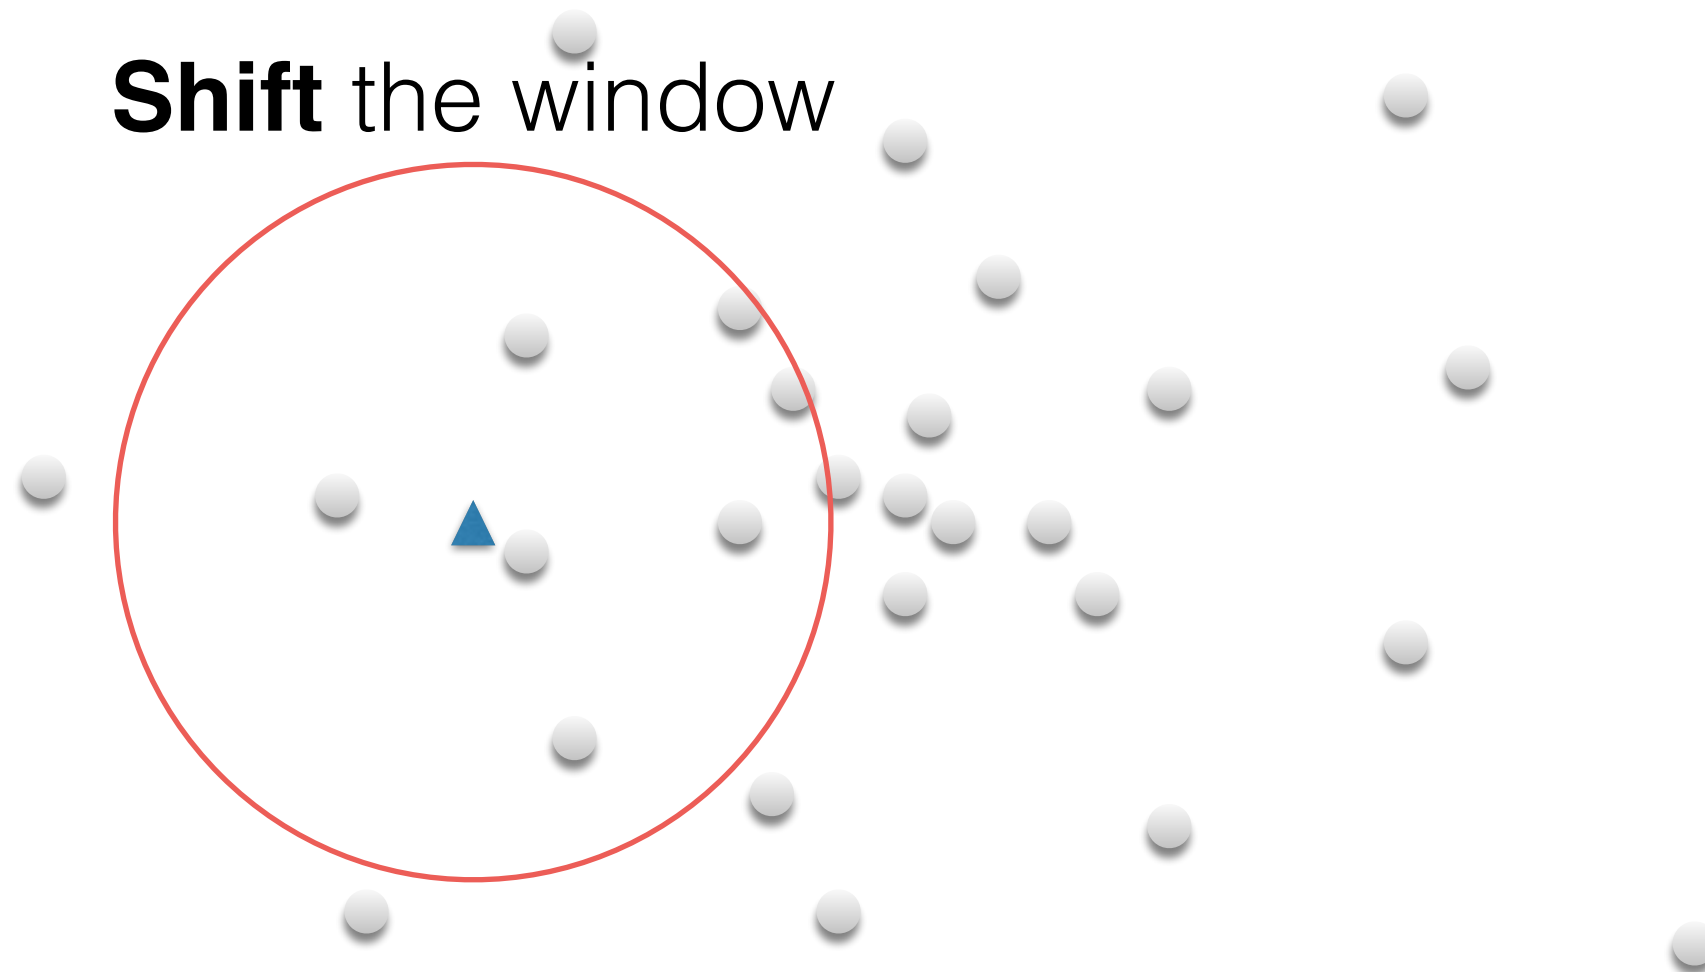
Fukunaga & Hostetler (1975)



Mean Shift Algorithm

A 'mode seeking' algorithm

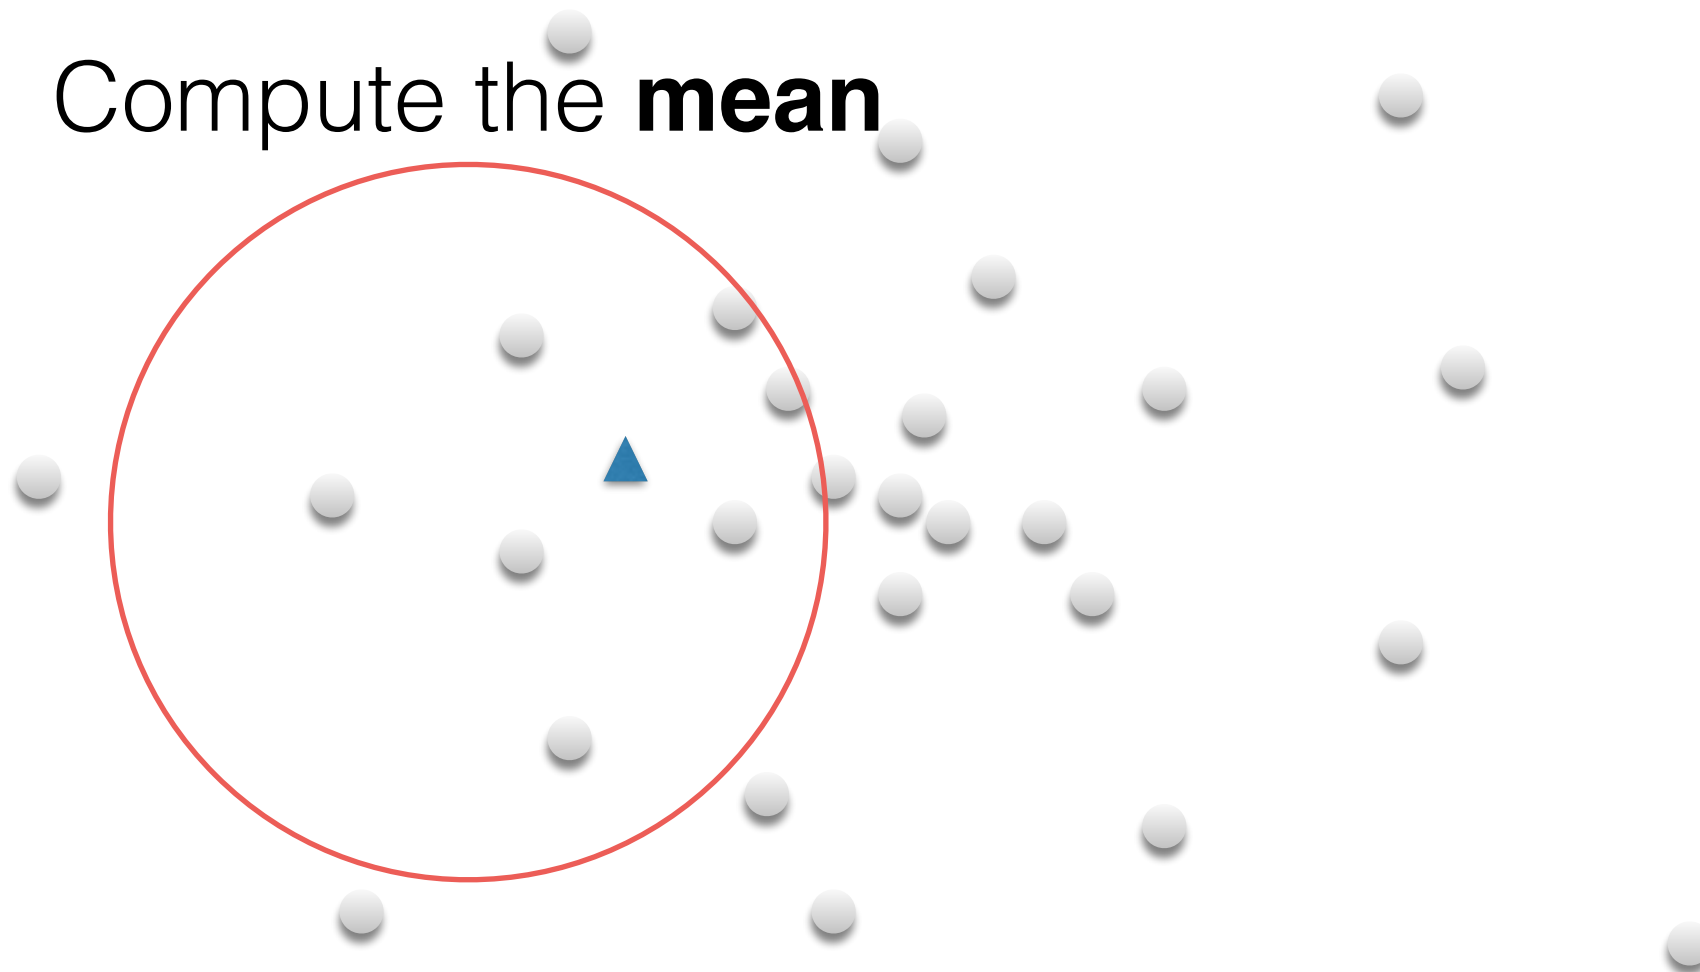
Fukunaga & Hostetler (1975)



Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

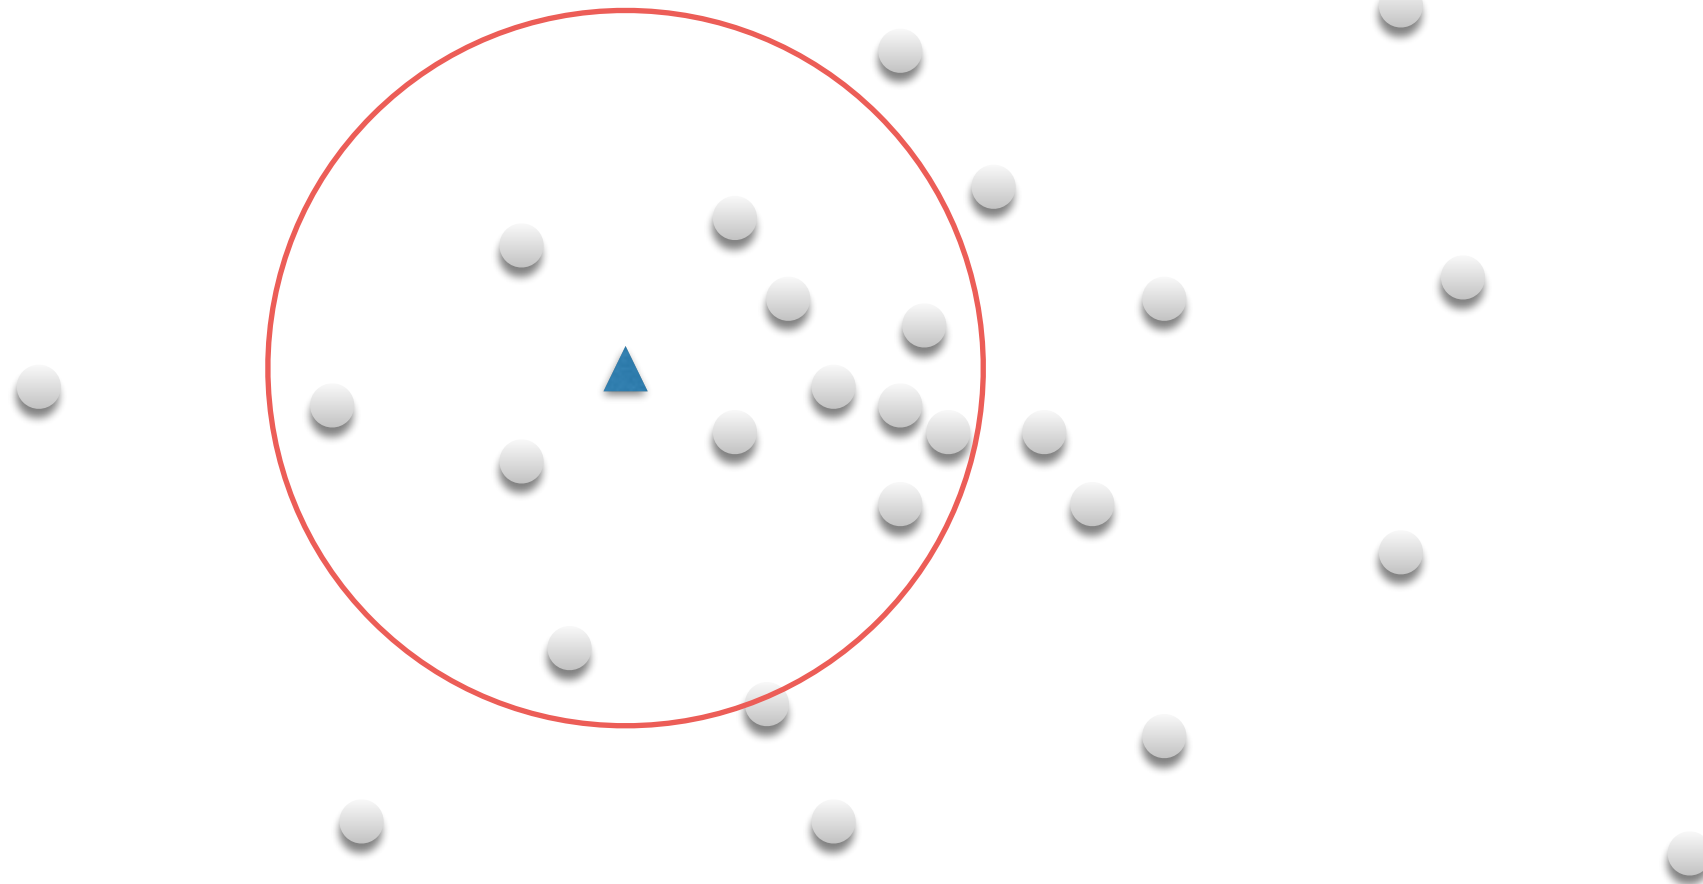


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

Shift the window

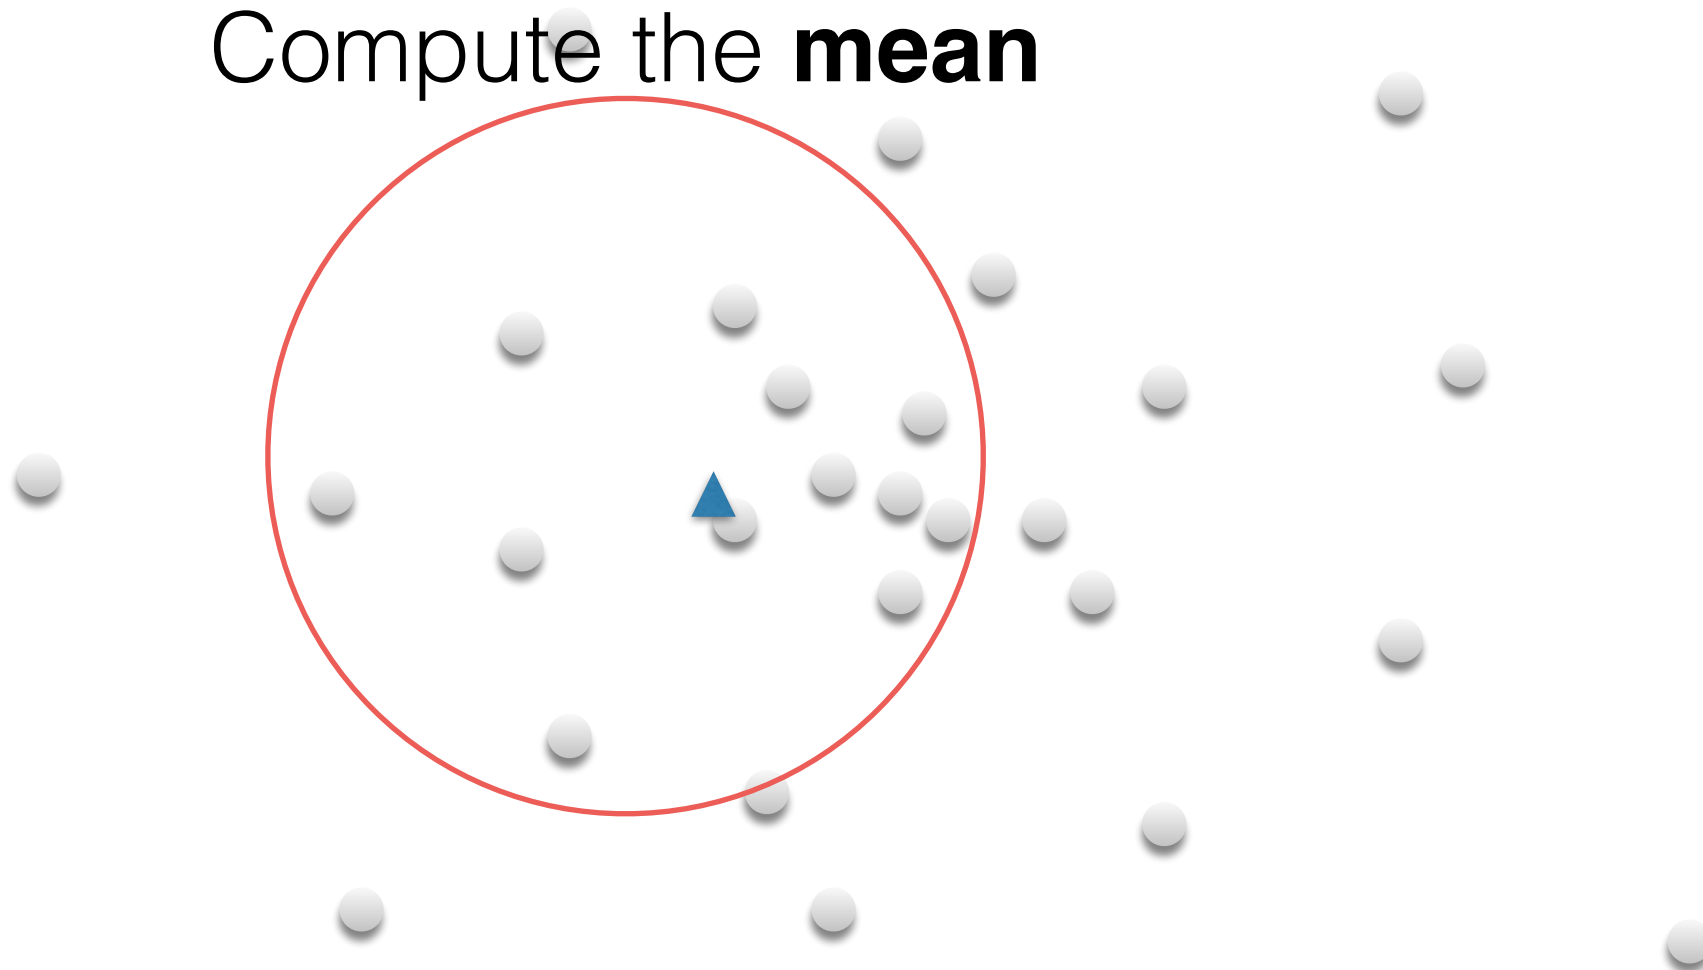


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

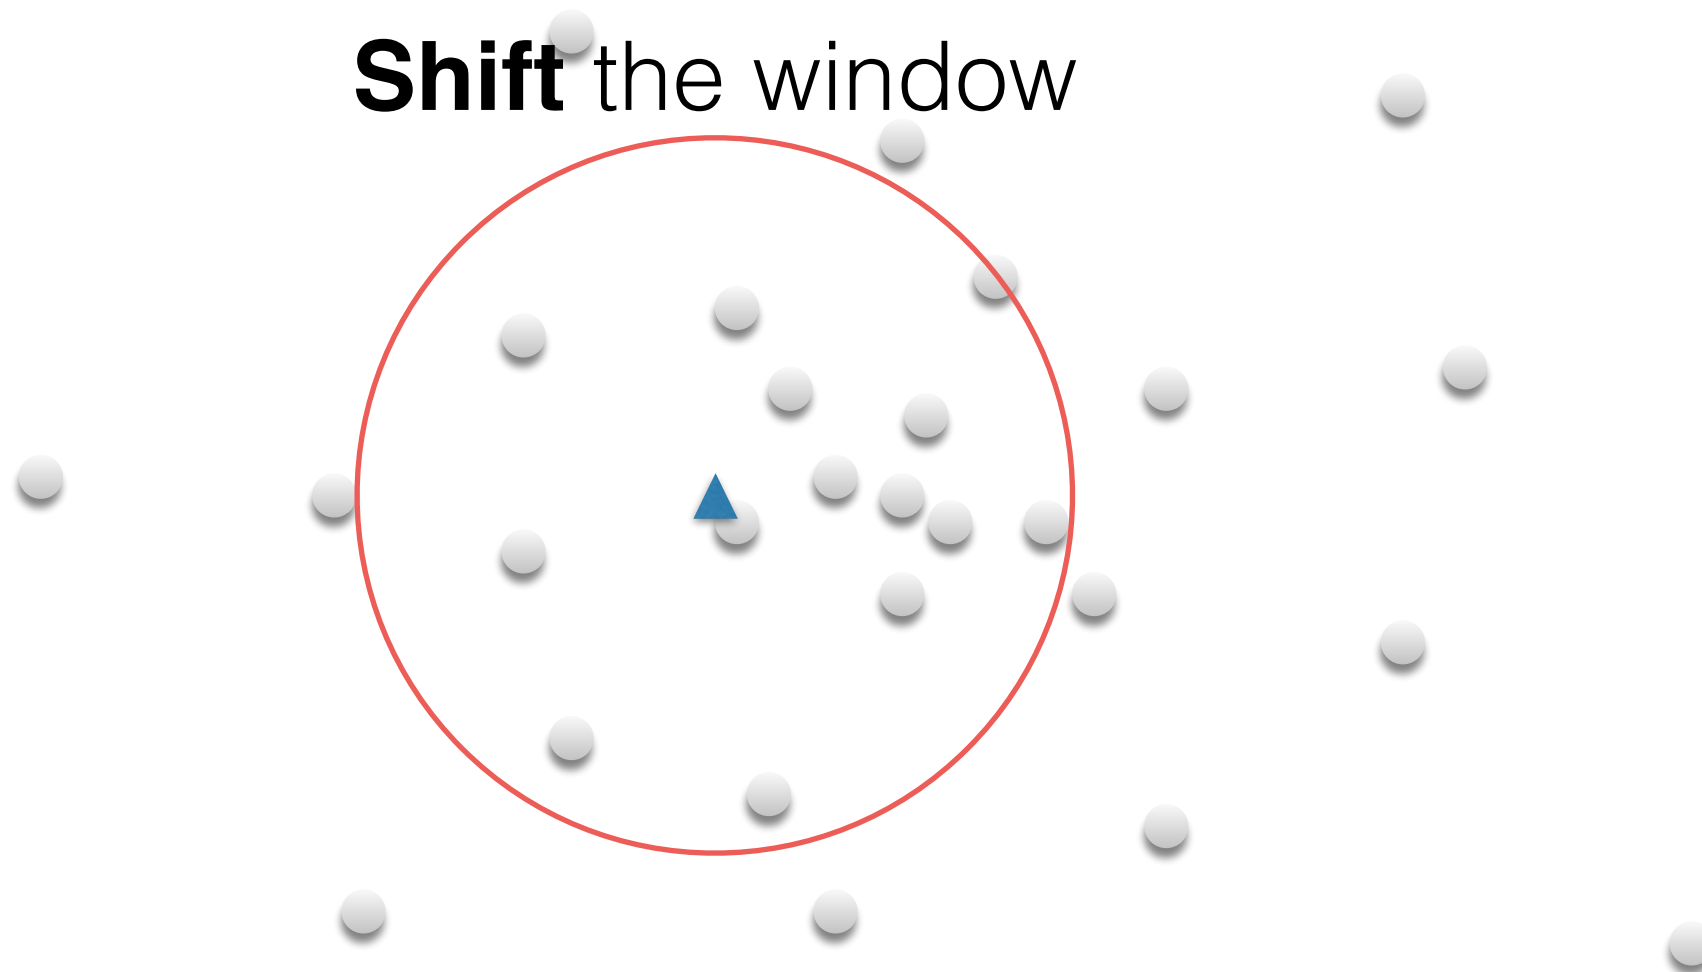
Compute the **mean**



Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

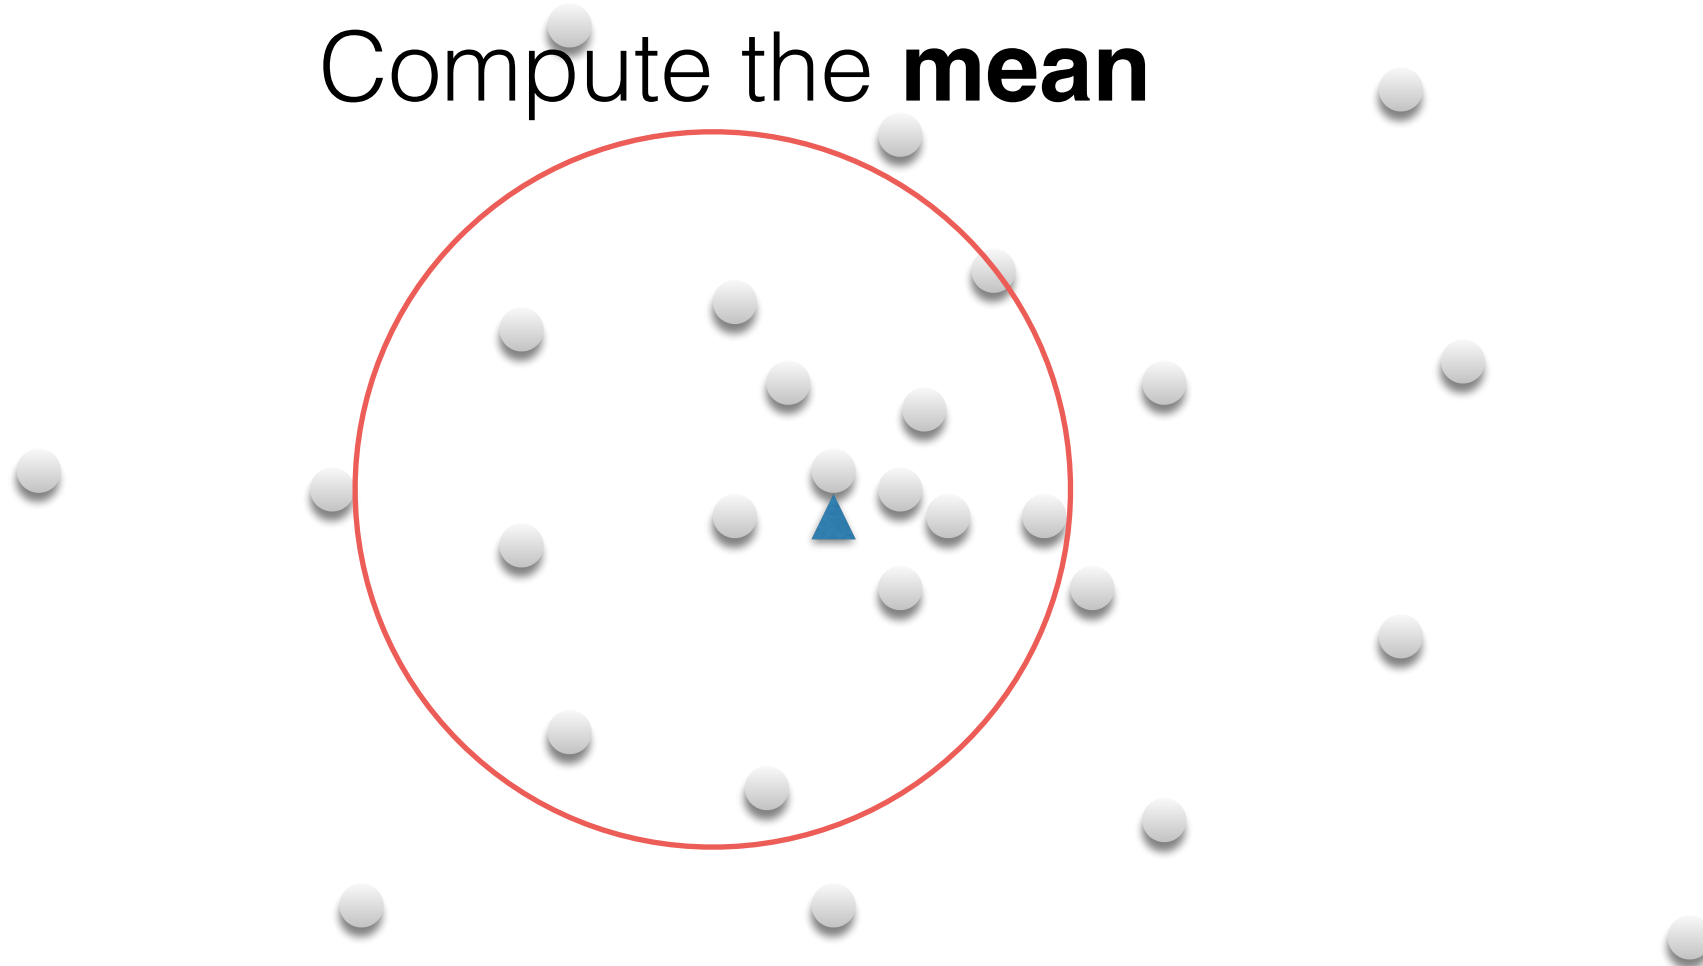


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)

Compute the **mean**

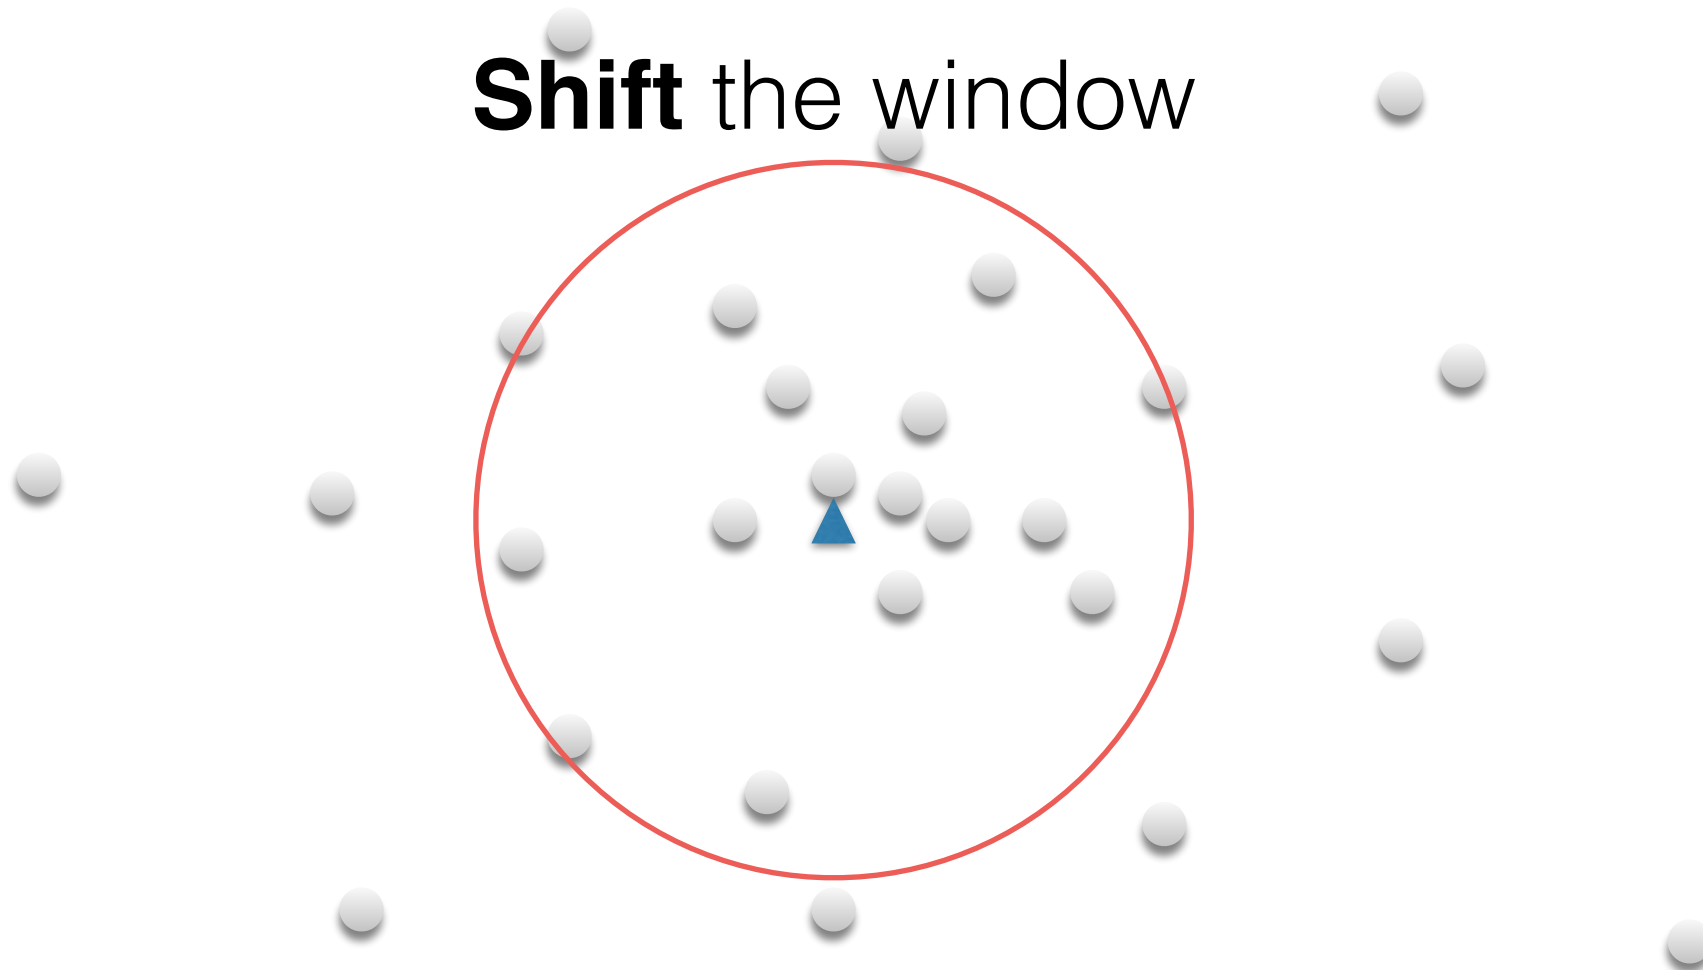


Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)


Shift the window



Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)




Mean Shift Algorithm

A 'mode seeking' algorithm

Fukunaga & Hostetler (1975)


Shift the window



To understand the mean shift algorithm ...

Kernel Density Estimation


Approximate the underlying PDF from samples



Put 'bump' on every sample to approximate the PDF

Kernel Density Estimation

Approximate the underlying PDF from samples from it



$$p(\mathbf{x}) = \sum_i c_i e^{-\frac{(\mathbf{x} - \mathbf{x}_i)^2}{2\sigma^2}}$$

Gaussian 'bump' aka 'kernel'


Put 'bump' on every sample to approximate the PDF

Kernel Function

$$K(x, x')$$


a 'distance' between two points

Epanechnikov kernel




$$K(\mathbf{x}, \mathbf{x}') = \begin{cases} c(1 - \|\mathbf{x} - \mathbf{x}'\|^2) & \|\mathbf{x} - \mathbf{x}'\|^2 \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

Uniform kernel



$$K(\mathbf{x}, \mathbf{x}') = \begin{cases} c & \|\mathbf{x} - \mathbf{x}'\|^2 \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

Normal kernel



$$K(\mathbf{x}, \mathbf{x}') = c \exp\left(-\frac{1}{2}\|\mathbf{x} - \mathbf{x}'\|^2\right)$$

Radially symmetric kernels

Radially symmetric kernels

...can be written in terms of its *profile*

$$K(\boldsymbol{x}, \boldsymbol{x}') = c \cdot k(\|\boldsymbol{x} - \boldsymbol{x}'\|^2)$$



profile

Connecting KDE and the Mean Shift Algorithm

Consider a set of points: $\{\mathbf{x}_s\}_{s=1}^S \quad \mathbf{x}_s \in \mathcal{R}^d$

Sample mean: $m(\mathbf{x}) = \frac{\sum_s K(\mathbf{x}, \mathbf{x}_s) \mathbf{x}_s}{\sum_s K(\mathbf{x}, \mathbf{x}_s)}$

Mean shift: $m(\mathbf{x}) - \mathbf{x}$

Mean shift algorithm

From each data point, move to its mean $\mathbf{x} \leftarrow m(\mathbf{x})$

Iterate until $\mathbf{x} = m(\mathbf{x})$

Where does this algorithm come from?

Consider a set of points: $\{\mathbf{x}_s\}_{s=1}^S \quad \mathbf{x}_s \in \mathcal{R}^d$

Sample mean: $m(\mathbf{x}) = \frac{\sum_s K(\mathbf{x}, \mathbf{x}_s) \mathbf{x}_s}{\sum_s K(\mathbf{x}, \mathbf{x}_s)}$

Mean shift: $m(\mathbf{x}) - \mathbf{x}$

*Where does this
come from?*



Mean shift algorithm

From each data point, move to its mean $\mathbf{x} \leftarrow m(\mathbf{x})$

Iterate until $\mathbf{x} = m(\mathbf{x})$

Where does this algorithm come from?

How is the KDE related to the mean shift algorithm?

Kernel density estimate
(radially symmetric kernels)

$$P(\mathbf{x}) = \frac{1}{N} c \sum_n k(\|\mathbf{x} - \mathbf{x}_n\|^2)$$

Gradient of the PDF is related to the mean shift vector

$$\nabla P(\mathbf{x}) \propto m(\mathbf{x})$$

The mean shift is a ‘step’ in the direction of the gradient of the KDE

Derivation

$$P(\boldsymbol{x}) = \frac{1}{N}c \sum_n k(\|\boldsymbol{x} - \boldsymbol{x}_n\|^2)$$

Gradient

$$\nabla P(\boldsymbol{x}) = \frac{1}{N}c \sum_n \nabla k(\|\boldsymbol{x} - \boldsymbol{x}_n\|^2)$$

expand derivative

$$\nabla P(\boldsymbol{x}) = \frac{1}{N}2c \sum_n (\boldsymbol{x} - \boldsymbol{x}_n)k'(\|\boldsymbol{x} - \boldsymbol{x}_n\|^2)$$

change of notation
(kernel-shadow pairs)

$$\nabla P(\boldsymbol{x}) = \frac{1}{N}2c \sum_n (\boldsymbol{x}_n - \boldsymbol{x})g(\|\boldsymbol{x} - \boldsymbol{x}_n\|^2)$$

$$k'(\cdot) = -g(\cdot)$$

$$\nabla P(\boldsymbol{x}) = \frac{1}{N} 2c \sum_n (\boldsymbol{x}_n - \boldsymbol{x}) g(\|\boldsymbol{x} - \boldsymbol{x}_n\|^2)$$

multiply it out

$$\nabla P(\boldsymbol{x}) = \frac{1}{N} 2c \sum_n \boldsymbol{x}_n g(\|\boldsymbol{x} - \boldsymbol{x}_n\|^2) - \frac{1}{N} 2c \sum_n \boldsymbol{x} g(\|\boldsymbol{x} - \boldsymbol{x}_n\|^2)$$

too long (enter short hand notation)

$$\nabla P(\boldsymbol{x}) = \frac{1}{N} 2c \sum_n \boldsymbol{x}_n g_n - \frac{1}{N} 2c \sum_n \boldsymbol{x} g_n$$

$$\nabla P(\mathbf{x}) = \frac{1}{N} 2c \sum_n \mathbf{x}_n g_n - \frac{1}{N} 2c \sum_n \mathbf{x} g_n$$

multiply by one!

$$\nabla P(\mathbf{x}) = \frac{1}{N} 2c \sum_n \mathbf{x}_n g_n \left(\frac{\sum_n g_n}{\sum_n g_n} \right) - \frac{1}{N} 2c \sum_n \mathbf{x} g_n$$

collecting like terms...

$$\nabla P(\mathbf{x}) = \frac{1}{N} 2c \sum_n g_n \left(\frac{\sum_n \mathbf{x}_n g_n}{\sum_n g_n} - \mathbf{x} \right)$$

Does this look familiar?

$$\nabla P(\mathbf{x}) = \frac{1}{N} 2c \sum_n g_n \underbrace{\left(\frac{\sum_n \mathbf{x}_n g_n}{\sum_n g_n} - \mathbf{x} \right)}_{\text{mean shift!}}$$

mean
shift


The mean shift is a ‘step’ in the direction of the gradient of the KDE

$$m(\mathbf{x}) = \left(\frac{\sum_n \mathbf{x}_n g_n}{\sum_n g_n} - \mathbf{x} \right) = \frac{\nabla P(\mathbf{x})}{\frac{1}{N} 2c \sum_n g_n}$$

Gradient ascent with adaptive step size

Dealing with images

Pixels for a lattice, spatial density is the same everywhere!



What can we do?

Consider a set of points: $\{\mathbf{x}_s\}_{s=1}^S \quad \mathbf{x}_s \in \mathcal{R}^d$

Associated weights: $w(\mathbf{x}_s)$

Sample mean:
$$m(\mathbf{x}) = \frac{\sum_s K(\mathbf{x}, \mathbf{x}_s) w(\mathbf{x}_s) \mathbf{x}_s}{\sum_s K(\mathbf{x}, \mathbf{x}_s) w(\mathbf{x}_s)}$$


Mean shift: $m(\mathbf{x}) - \mathbf{x}$

Mean shift algorithm


From each data point, move to its mean $\mathbf{x} \leftarrow m(\mathbf{x})$

Iterate until $\mathbf{x} = m(\mathbf{x})$


For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight




For images, each pixel is point with a weight



For images, each pixel is point with a weight

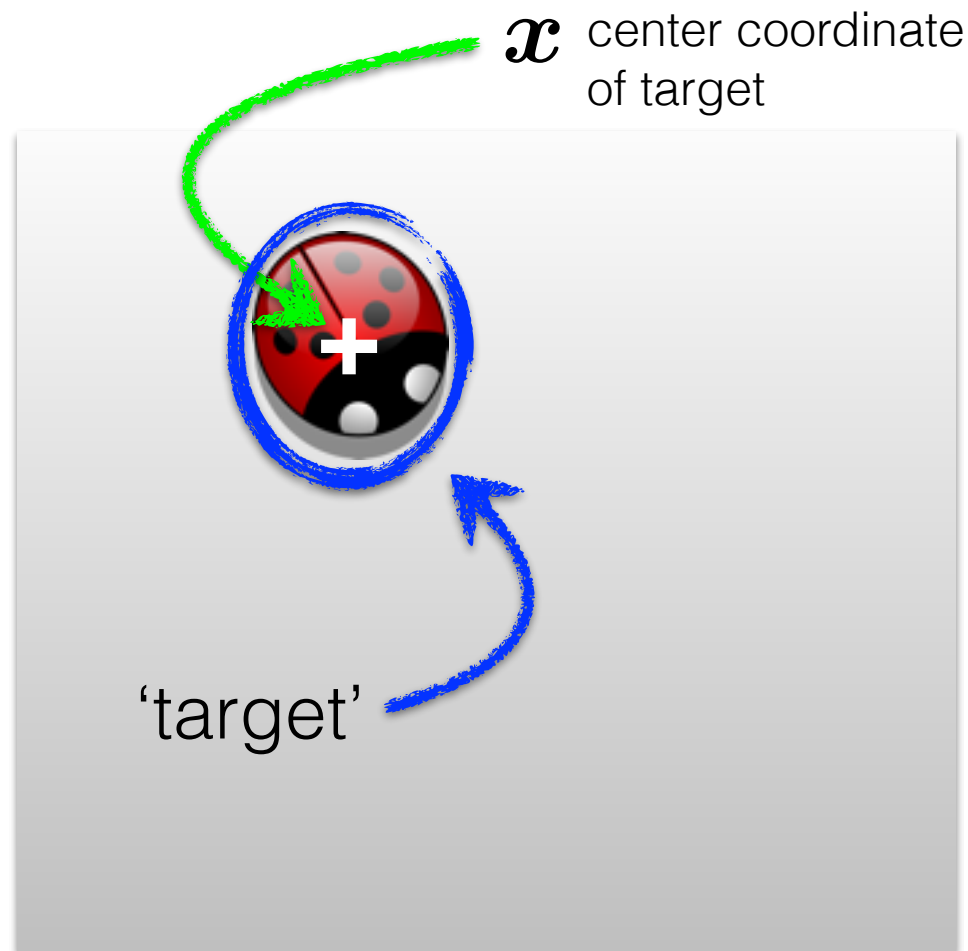


For images, each pixel is point with a weight




Finally... mean shift tracking in video

Goal: find the best candidate location in frame 2




Frame 1




Frame 2

Use the mean shift algorithm
to find the best candidate location

Non-rigid object tracking




Compute a descriptor for the target



Target


Search for similar descriptor in neighborhood in next frame



Target


Candidate

Compute a descriptor for the new target



Target

Search for similar descriptor in neighborhood in next frame



Target

Candidate

How do we model the target and candidate regions?

Modeling the target



M-dimensional **target** descriptor

$$\mathbf{q} = \{q_1, \dots, q_M\}$$


(centered at target center)

$$q_m = C \sum_n k(\|\mathbf{x}_n\|^2) \delta[b(\mathbf{x}_n) - m]$$

Normalization factor

Kronecker delta function

n function of inverse distance (weight)



Modeling the candidate


M-dimensional **candidate** descriptor

$$\mathbf{p}(\mathbf{y}) = \{p_1(\mathbf{y}), \dots, p_M(\mathbf{y})\}$$

(centered at location \mathbf{y})

$$p_m = C_h \sum_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right) \delta[b(\mathbf{x}_n) - m]$$

bandwidth



Similarity between the target and candidate

Distance function


$$d(\mathbf{y}) = \sqrt{1 - \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]}$$

Bhattacharyya Coefficient


$$\rho(\mathbf{y}) \equiv \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] = \sum_m \sqrt{p_m(\mathbf{y}) q_m}$$

Just the Cosine distance between two unit vectors


$$\rho(\mathbf{y}) = \cos \theta_{\mathbf{y}} = \frac{\mathbf{p}(\mathbf{y})^\top \mathbf{q}}{\|\mathbf{p}\| \|\mathbf{q}\|} = \sum_m \sqrt{p_m(\mathbf{y}) q_m}$$




Now we can compute the similarity between a target and multiple candidate regions




image



similarity over image



image



similarity over image

Objective function

$$\min_{\mathbf{y}} d(\mathbf{y}) \quad \text{same as} \quad \max_{\mathbf{y}} \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]$$

Assuming a good initial guess

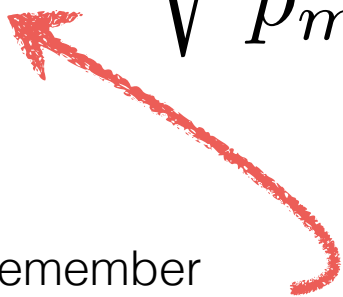
$$\rho[\mathbf{p}(\mathbf{y}_0 + \mathbf{y}), \mathbf{q}]$$

Linearize around the initial guess (Taylor series expansion)

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m} + \frac{1}{2} \sum_m p_m(\mathbf{y}) \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}}$$

function at specified value derivative

Linearized objective

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m} + \frac{1}{2} \sum_m p_m(\mathbf{y}) \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}}$$


$$p_m = C_h \sum_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right) \delta[b(\mathbf{x}_n) - m]$$

Remember definition of this?

Fully expanded

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m} + \frac{1}{2} \sum_m \left\{ C_h \sum_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right) \delta[b(\mathbf{x}_n) - m] \right\} \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}}$$

Fully expanded linearized objective

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m} + \frac{1}{2} \sum_m \left\{ C_h \sum_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right) \delta[b(\mathbf{x}_n) - m] \right\} \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}}$$

Moving terms around...

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \underbrace{\frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m}}_{\text{Does not depend on unknown } \mathbf{y}} + \underbrace{\frac{C_h}{2} \sum_n w_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right)}_{\text{Weighted kernel density estimate}}$$

where $w_n = \sum_m \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}} \delta[b(\mathbf{x}_n) - m]$

Weight is bigger when $q_m > p_m(\mathbf{y}_0)$

OK, why are we doing all this math?

We want to maximize this

$$\max_{\boldsymbol{y}} \rho[\boldsymbol{p}(\boldsymbol{y}), \boldsymbol{q}]$$

We want to maximize this

$$\max_{\mathbf{y}} \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]$$

Fully expanded linearized objective

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m} + \frac{C_h}{2} \sum_n w_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right)$$

$$\text{where } w_n = \sum_m \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}} \delta[b(\mathbf{x}_n) - m]$$

We want to maximize this

$$\max_{\mathbf{y}} \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]$$

Fully expanded linearized objective

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m} + \frac{C_h}{2} \sum_n w_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right)$$

doesn't depend on unknown \mathbf{y}

$$\text{where } w_n = \sum_m \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}} \delta[b(\mathbf{x}_n) - m]$$

We want to maximize this

$$\max_{\mathbf{y}} \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]$$

only need to
maximize this!

Fully expanded linearized objective

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \underbrace{\frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m}}_{\text{doesn't depend on unknown } \mathbf{y}} + \frac{C_h}{2} \sum_n w_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right)$$

where $w_n = \sum_m \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}} \delta[b(\mathbf{x}_n) - m]$

We want to maximize this

$$\max_{\mathbf{y}} \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]$$

Fully expanded linearized objective

$$\rho[\mathbf{p}(\mathbf{y}), \mathbf{q}] \approx \underbrace{\frac{1}{2} \sum_m \sqrt{p_m(\mathbf{y}_0) q_m}}_{\text{doesn't depend on unknown } \mathbf{y}} + \frac{C_h}{2} \sum_n w_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right)$$

where $w_n = \sum_m \sqrt{\frac{q_m}{p_m(\mathbf{y}_0)}} \delta[b(\mathbf{x}_n) - m]$

what can we use to solve this weighted KDE?

Mean Shift Algorithm!

$$\frac{C_h}{2} \sum_n w_n k \left(\left\| \frac{\mathbf{y} - \mathbf{x}_n}{h} \right\|^2 \right)$$

the sample of mean of this KDE is


$$\mathbf{y}_1 = \frac{\sum_n \mathbf{x}_n w_n g \left(\left\| \frac{\mathbf{y}_0 - \mathbf{x}_n}{h} \right\|^2 \right)}{\sum_n w_n g \left(\left\| \frac{\mathbf{y}_0 - \mathbf{x}_n}{h} \right\|^2 \right)} \quad \text{(this was derived earlier)}$$

(new candidate location)

Mean Shift Tracking procedure


1. Initialize location \mathbf{y}_0
Compute \mathbf{q}
Compute $\mathbf{p}(\mathbf{y}_0)$
2. Derive weights w_n
3. Shift to new candidate location (mean shift) \mathbf{y}_1
4. Compute $\mathbf{p}(\mathbf{y}_1)$
5. If $\|\mathbf{y}_0 - \mathbf{y}_1\| < \epsilon$ return
Otherwise $\mathbf{y}_0 \leftarrow \mathbf{y}_1$ and go back to 2

Compute a descriptor for the target



Target
 q

Search for similar descriptor in neighborhood in next frame




Target

Candidate


$$\max_{\mathbf{y}} \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]$$

Compute a descriptor for the new target



Target
 q

Search for similar descriptor in neighborhood in next frame



Target

Candidate

$$\max_{\mathbf{y}} \rho[\mathbf{p}(\mathbf{y}), \mathbf{q}]$$

