FoDA

L23 : Clustering

: Lloyd's Algorithm for k-means

K-Means Clustering Formulation

Input $X \subset \mathbb{R}^d$ $\quad X = \{x_1, x_2 \dots x_n\}$

$k \in \mathbb{Z}^+ \quad \in [1, 2, \dots]$

$d : X \times X \to \mathbb{R}_+ \quad d(x_1, x_2) = \|x_1 - x_2\|$

Goal Find s.t. $S = \{s_1, s_2, \dots s_k\} \subset \mathbb{R}^d$

minimize $\sum_{i=1}^{n} d\left(x_i, \phi_S(x_i)\right)^2$

↳ maps x_i to $\operatorname{argmin}_{s_j \in S} \|x_i - s_j\|$

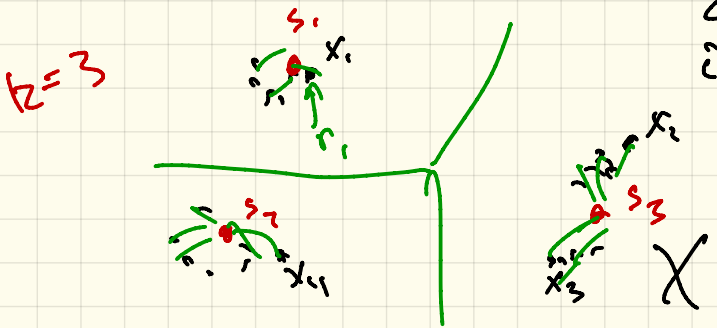

Lloyd's Algorithm → $X = \bigcup_{j=1}^{k} X_j$, $X_j \subset X$

0. Initialize k points $S \subset X$ (arbitrarily?)

1. repeat

a. assignment / Voronoi diagram: for all $x_i \in X$: assign x_i to X_j so $\phi_S(x_i) = s_j$

b. recenter / average: for all $s_j \in S$: update $s_j = \frac{1}{|X_j|} \sum_{x \in X_j} x$ = average (X_j)

until (the set S unchanged or change is small)

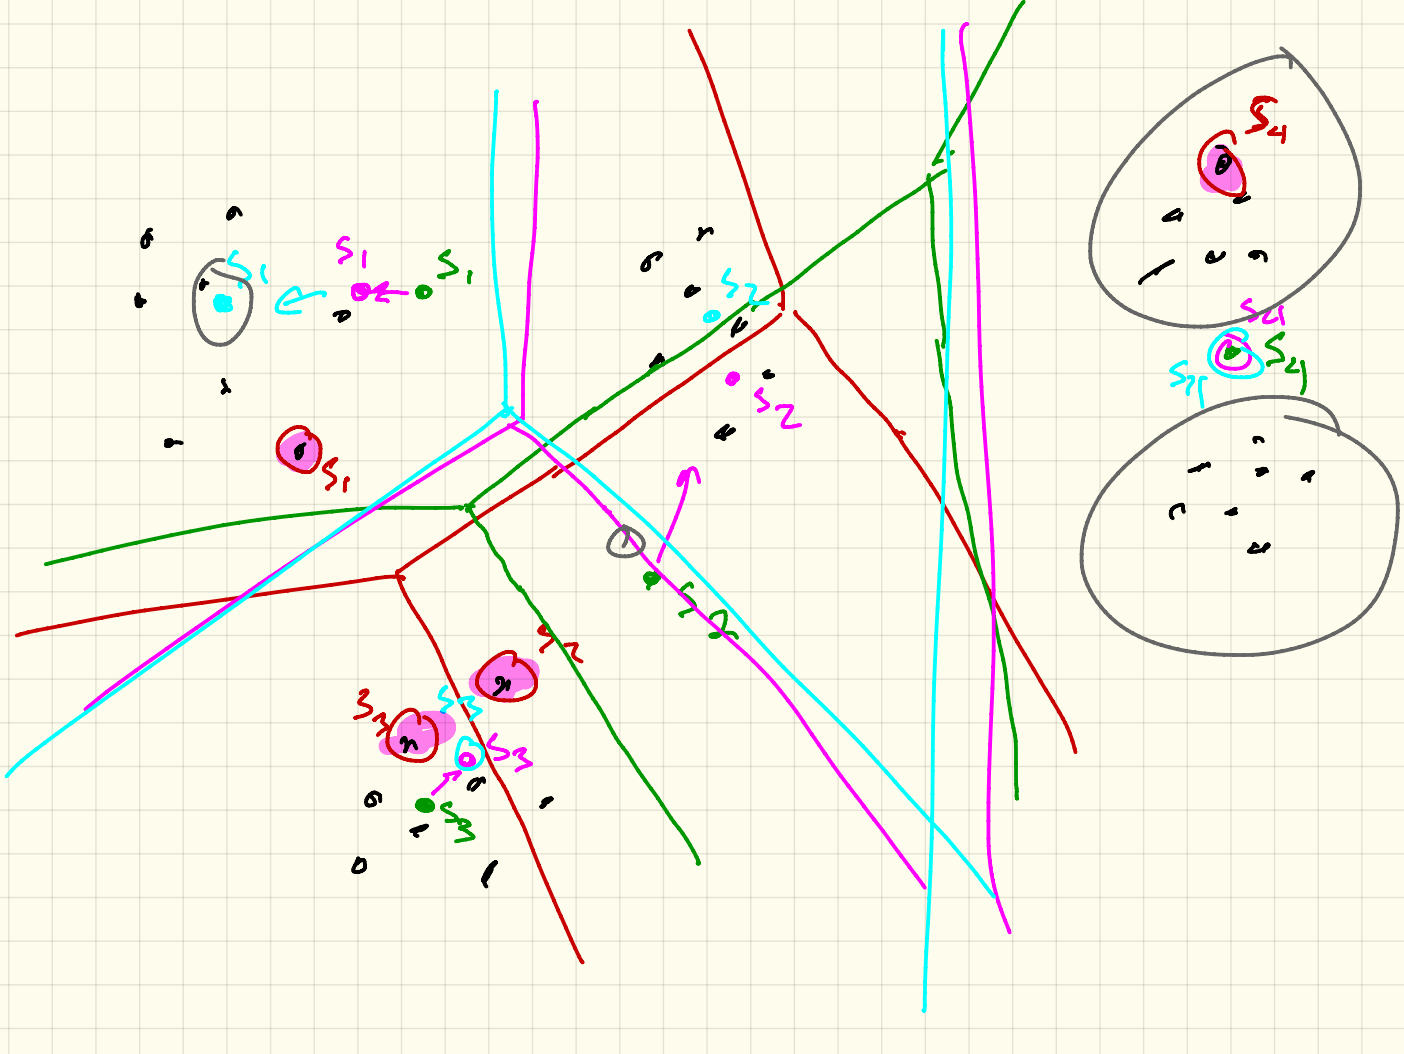
s_1
s_1
s_1
s_1
s_2
s_2
s_3
s_3
s_3
s_4
s_4
s_4

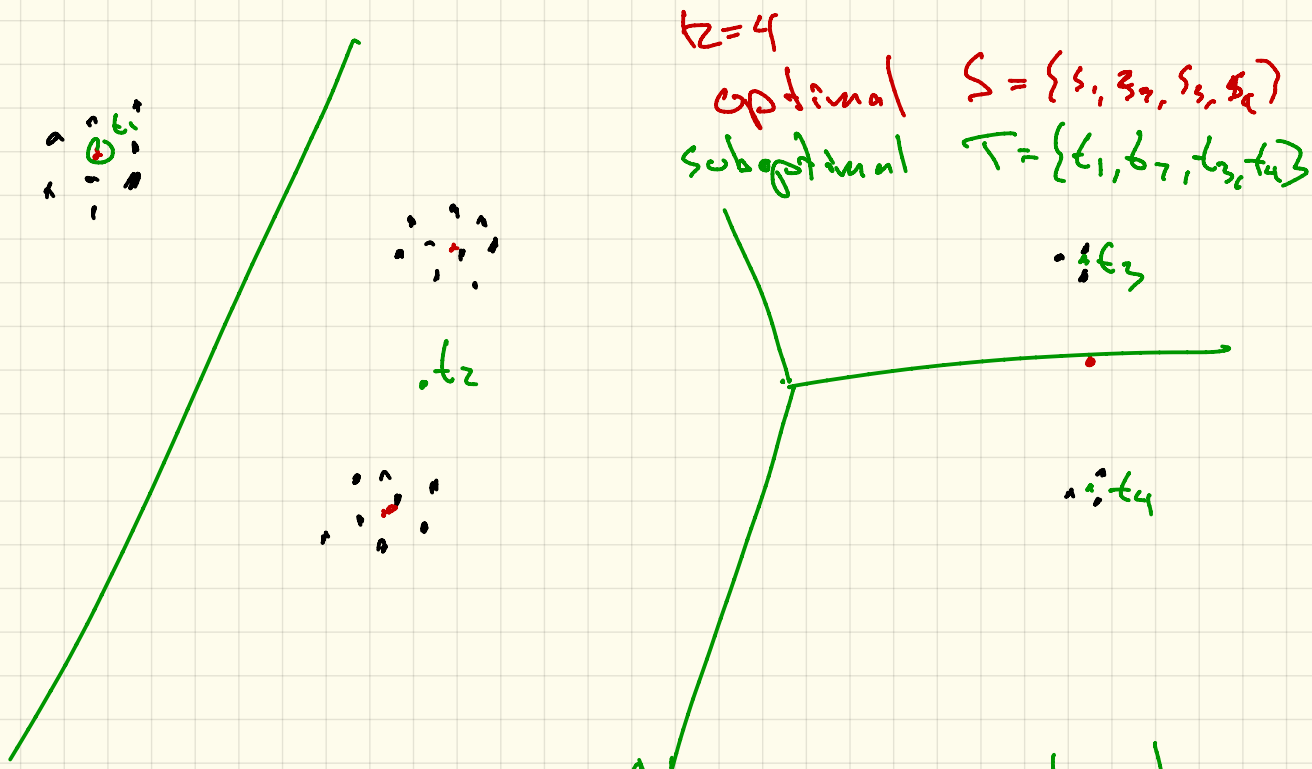

Lloyd's Algo is stuck

Most of the time Lloyd's works well
(with a little help)

Tricks to help Lloyd's

- Random Restarts
  - a. Randomly initialize S.
  - b. Run Lloyd's → compute cost (SSE)
  - c. Repeat (a, b) say 10 times
  - d. Return final S w/ lowest cost.

  (×10)

- Better Initialize
  - Gonzalez Alg.
  - k-means++

Corner Cases

- Might be site w/ no points closest to it

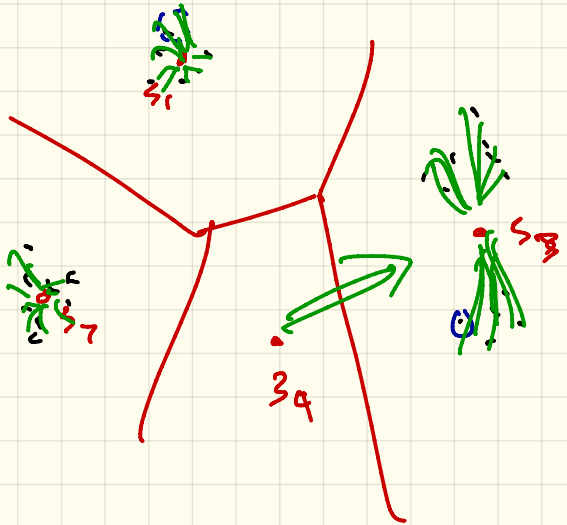

→ Randomly assign s_4 (w/ no points in X_4) to some $x_i \in X$.

from X_j w/ most variance.

Most Data Looks Like

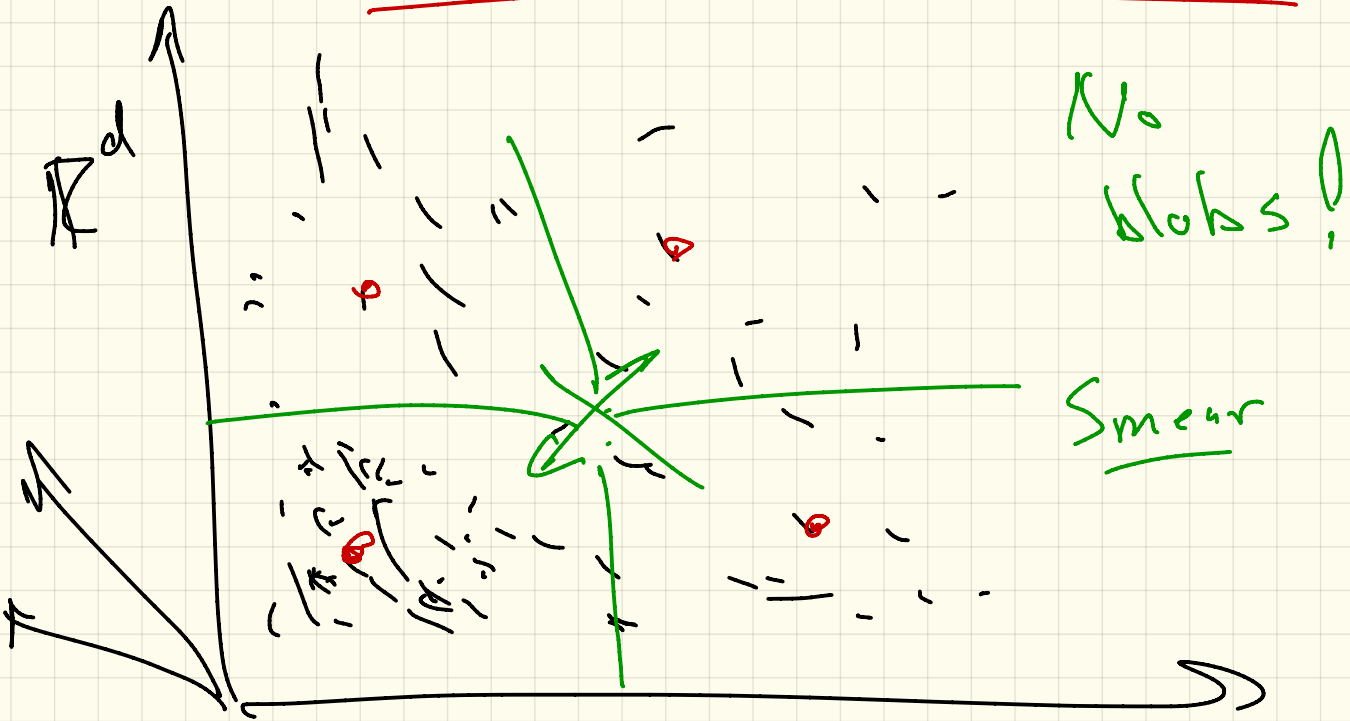

Number of Clusters. k?

$$\text{Cost}(X, S) = \sum_{i=1}^{n} \| x_i - \phi_S(x_i) \|^2$$

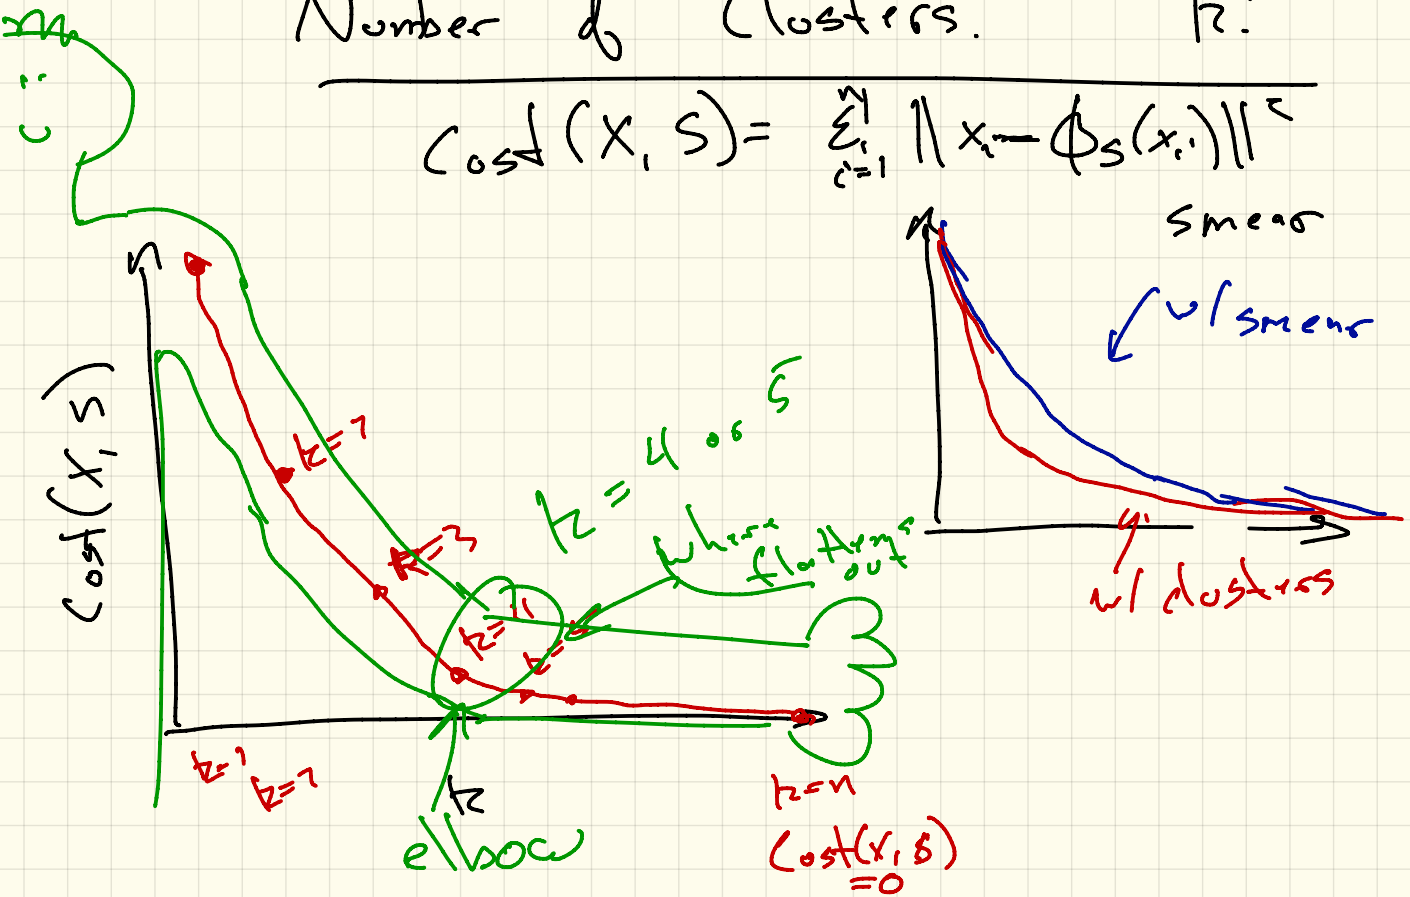

Why Lloyd's Algo converges?

$$\text{cost}(X, S) = \sum_{x \in X} \|x - \phi_S(x)\|^2 \quad \leftarrow (a)$$

$$= \sum_{s_j \in S} \left[\sum_{x \in X_j} \|x - s_j\|^2 \right] \quad \leftarrow (b)$$

(a) assignment

(b) centering