Interactive structure learning with structural query-by-committee

Christopher Tosh Sanjoy Dasgupta

Columbia University UC San Diego

Poster: **# 130**

The problem with traditional learning frameworks

Supervised learning issues:

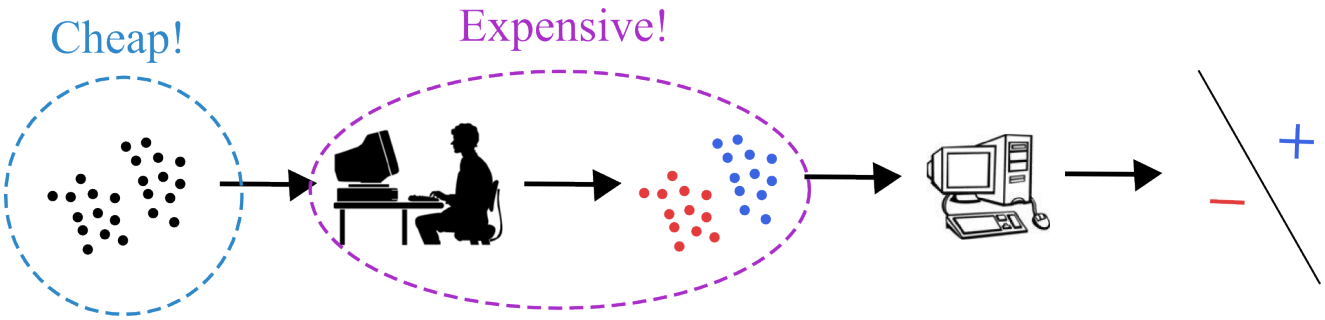

Unsupervised learning issues:

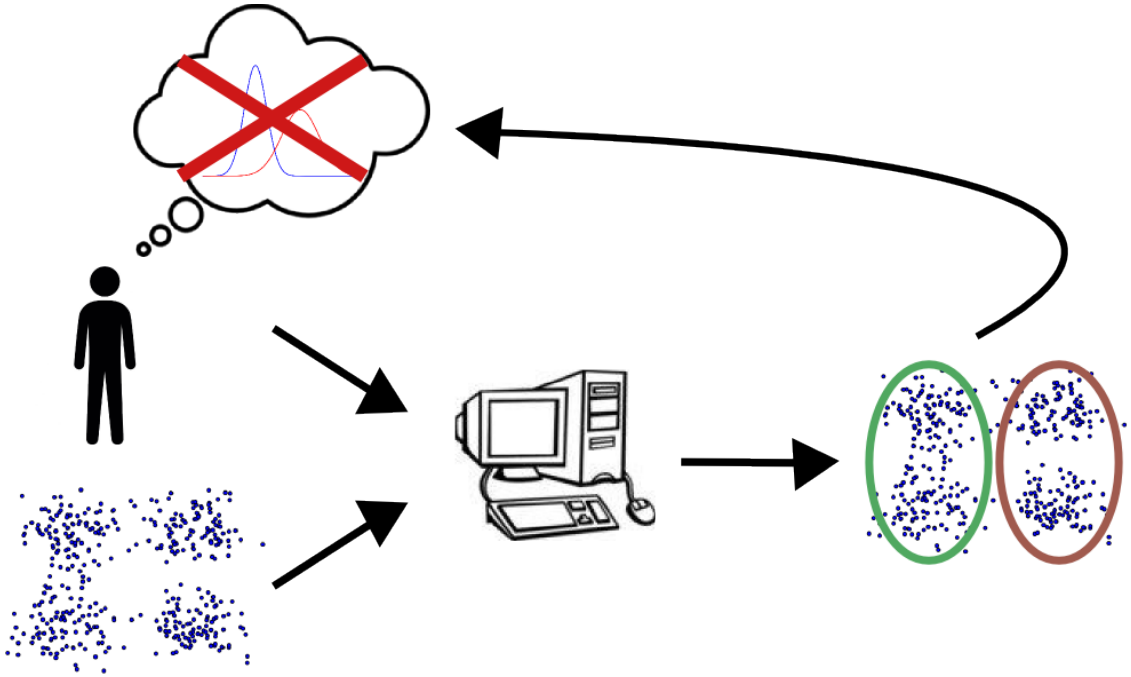

- Choosing model to fit
- Correcting obvious mistakes

Learning with interaction

Previous approaches:

Active learning

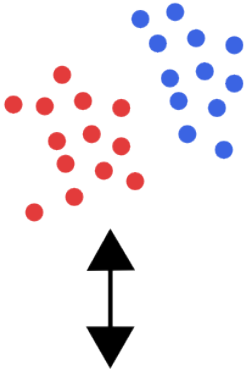

What is the label of this point?

Interactive flat clustering

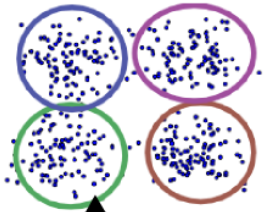

Do these two points belong to the same cluster?

Interactive hierarchical clustering

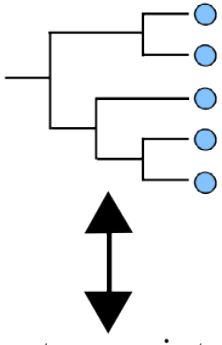

Should these two points be clustered before this other point?

Learning with interaction

Previous approaches:

Active learning

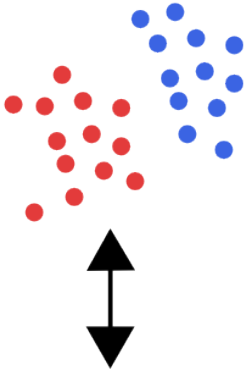

What is the label of this point?

Interactive flat clustering

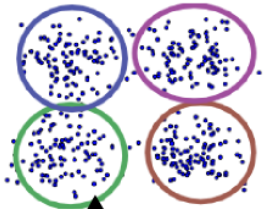

Do these two points belong to the same cluster?

Interactive hierarchical clustering

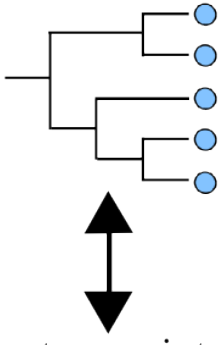

Should these two points be clustered before this other point?

Drawbacks:

- Specialized algorithms for different tasks
- Requires a user to answer every question

Interactive structure learning

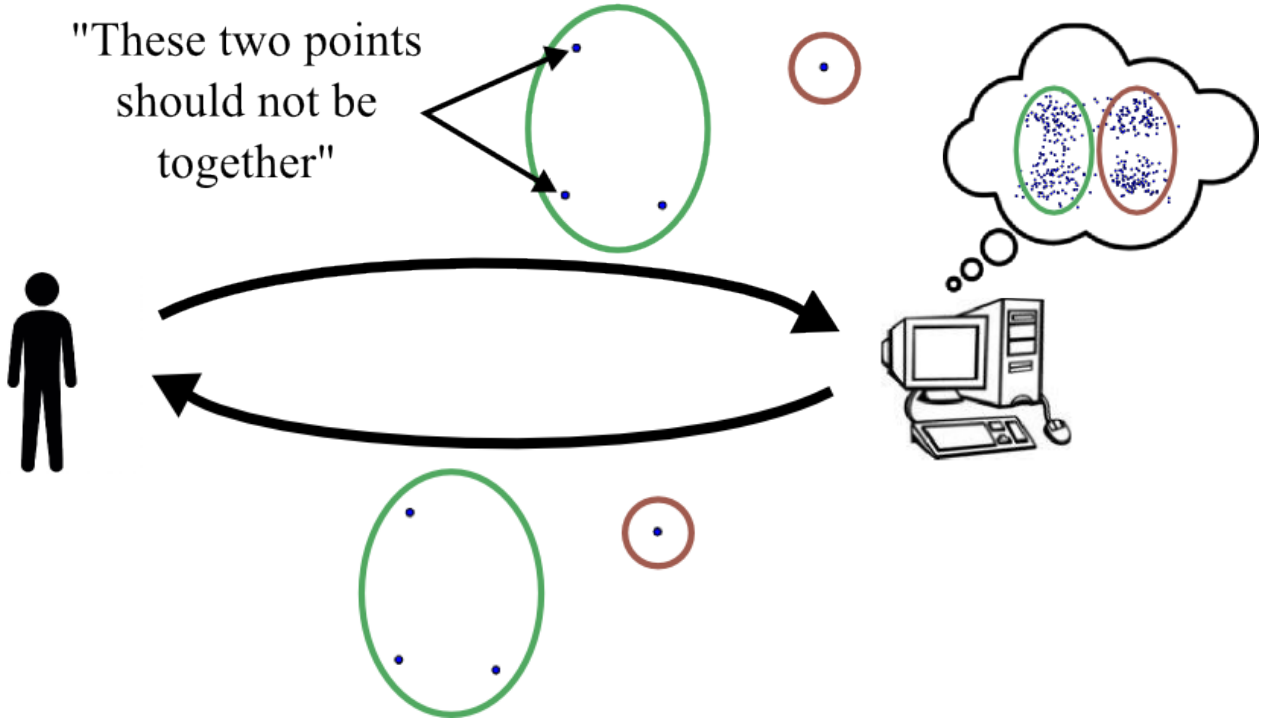

Interactive structure learning

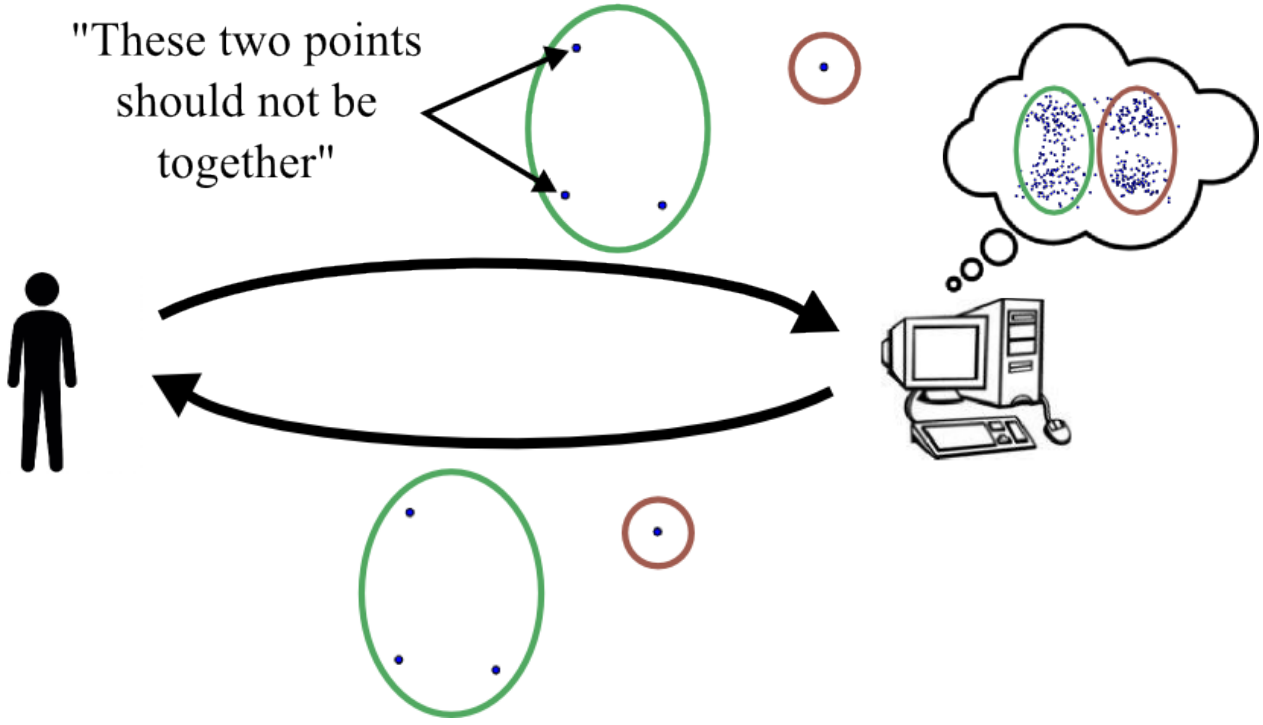

Features:

- Snapshots of structures (context)
- Partial correction

Interactive structure learning

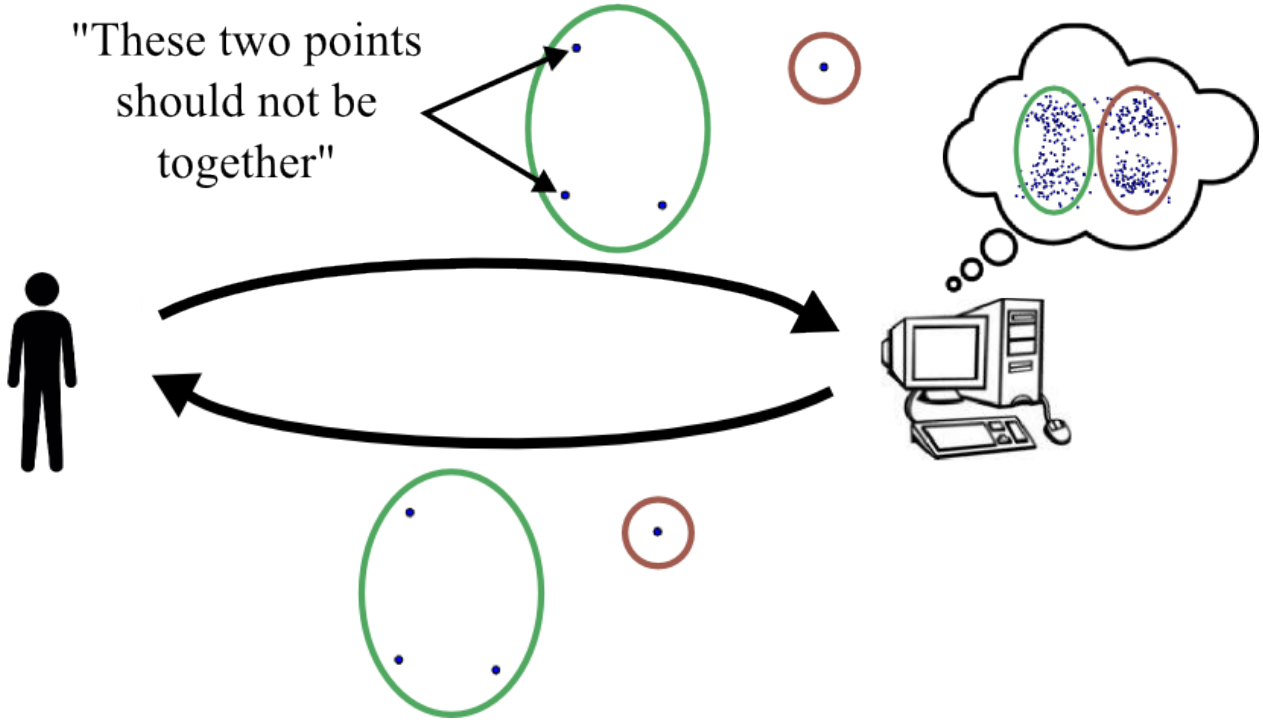

Generalizes:

- Active learning with label queries
- Interactive clustering with pairwise constraints
- Interactive hierarchical clustering with triplet constraints

Generic interactive learning scheme

Interactive structure learning protocol:

- Maintain a posterior π_t over structures
- For $t = 1, 2, \ldots$:
  - Select subset S and draw $g \sim \pi_{t-1}$.
  - Present user with $(S, g|_S)$, observe feedback.
  - Update posterior:
    $\pi_t(g) \propto \pi_{t-1}(g) \exp\left(-\beta \cdot \mathbb{1}[g \text{ is inconsistent with feedback}]\right)$

Generic interactive learning scheme

Interactive structure learning protocol:

- Maintain a posterior π_t over structures
- For $t = 1, 2, \ldots$:
  - Select subset S and draw $g \sim \pi_{t-1}$.
  - Present user with $(S, g|_S)$, observe feedback.
  - Update posterior:
    $\pi_t(g) \propto \pi_{t-1}(g) \exp\left(-\beta \cdot \mathbb{1}[g \text{ is inconsistent with feedback}]\right)$

Generalization of query-by-committee (SQBC).

- If the user's feedback occasionally fixes mistakes, then SQBC is consistent.
- Under stronger assumptions, SQBC converges on g^* after $\tilde{O}\left(\log \frac{1}{\pi(g^*)}\right)$ rounds.

Other results

- Generalized update rule

$$\pi_t(g) \;=\; \pi_{t-1}(g) \cdot \exp\left(-\beta \cdot \ell(g(x_t), y_t)\right)$$

  for efficiently samplable posteriors

- Efficiently kernelized querying rule

- Simulations

Other results

- Generalized update rule

$$\pi_t(g) \;=\; \pi_{t-1}(g) \cdot \exp\left(-\beta \cdot \ell(g(x_t), y_t)\right)$$

  for efficiently samplable posteriors

- Efficiently kernelized querying rule

- Simulations

Poster: **# 130**