Trading off Mistakes and I-Don't-Know Responses in Online Learning

Amin Sayedi — Carnegie Mellon
Morteza Zadimoghaddam — MIT
Avrim Blum — Carnegie Mellon

NIPS 2010

A natural theoretical question

What happens if allow "don't know" in online learning:

- World presents example (is this spam?)
- Algorithm makes prediction or says "I don't know"
- Algorithm is given correct answer.

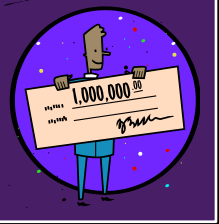

Q: If we view saying "I don't know" as (much) better than making a mistake, how can you trade them off? [Li-Littman-Walsh: *only* don't knows]

Two interesting ends of spectrum:

- Just a few mistakes allowed (KWIK + a few mistakes)
- Convert as many mistakes as possible to O(1) don't knows.

Examine information-theoretically and algorithmically

Intuition

Information-theoretically, close connection to classic "egg dropping puzzle":

Want to figure out highest floor you can drop egg without breaking. (N floors total).

- If unlimited supply, just do binary search.
  - O(log N) drops.
- If only 1, need to do linear search.
  - O(N) drops.
- What about 2 eggs? k eggs?
  - 2*sqrt(N), $kN^{1/k}$ drops.

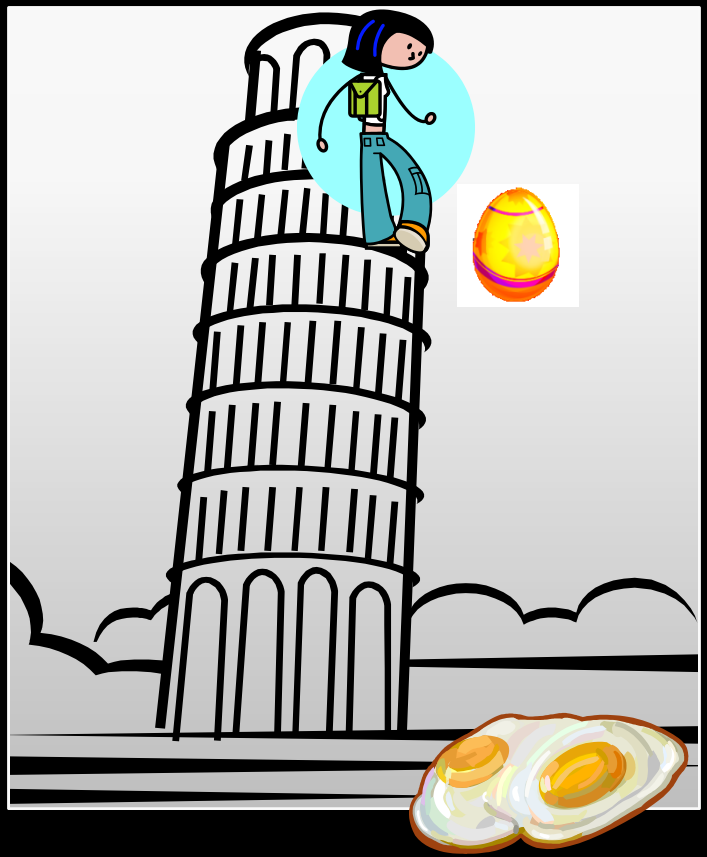

Use this to show: if allow k-1 mistakes, can learn any $f \in C$ with $k|C|^{1/k}$ don't knows.

Helpful example: initial intervals of 1,…,N.

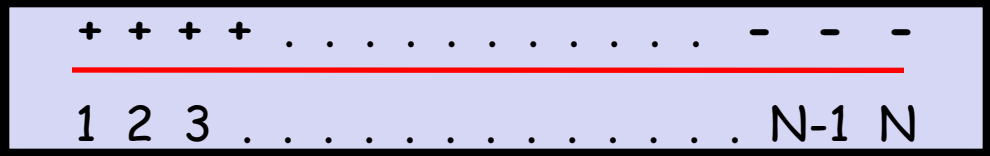

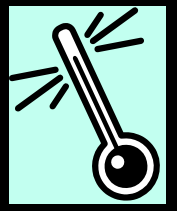

Algorithmic Results

Disjunctions: classic bound of n mistakes (necessary and sufficient in worst case)

- Reduce to n/2 mistakes + O(n) don't knows.
- Reduce to n/3 mistakes + O(n) don't knows.

Linear separators

- Use random sampling to estimate volume of version space voting + or -.
- Bounds a bit more complex: see poster.