Fast Cross-Validation for Incremental Learning

Pooria Joulani, András György, Csaba Szepesvári

Department of Computing Science
University of Alberta
Edmonton, Alberta

July 11, 2015

Appearing in the
International Joint Conference on Artificial Intelligence,
Buenos Aires, Argentina, July 2015.

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!
  - Leave-One-Out in $O(\log n)$!

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!
  - Leave-One-Out in $O(\log n)$!
- Does not rely on a specific

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!
  - Leave-One-Out in $O(\log n)$!
- Does not rely on a specific
  - type of the learning problem (classification, regression, density estimation, etc.);

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!
  - Leave-One-Out in $O(\log n)$!
- Does not rely on a specific
  - type of the learning problem (classification, regression, density estimation, etc.);
  - inner structure of the algorithm (e.g., QP, influence matrix, etc.);

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!
  - Leave-One-Out in $O(\log n)$!
- Does not rely on a specific
  - type of the learning problem (classification, regression, density estimation, etc.);
  - inner structure of the algorithm (e.g., QP, influence matrix, etc.);
  - loss function used for CV (accuracy, F-measure, etc.).

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!
  - Leave-One-Out in $O(\log n)$!
- Does not rely on a specific
  - type of the learning problem (classification, regression, density estimation, etc.);
  - inner structure of the algorithm (e.g., QP, influence matrix, etc.);
  - loss function used for CV (accuracy, F-measure, etc.).
- Easy parallelization / distributed computing.

TreeCV: Fast CV for Incremental Learning

A new cross-validation algorithm: TreeCV!

- Speed up CV for incremental, single-pass algorithms.
  - k-fold CV: running time penalty $O(\log k)$ instead of $O(k)$!
  - Leave-One-Out in $O(\log n)$!
- Does not rely on a specific
  - type of the learning problem (classification, regression, density estimation, etc.);
  - inner structure of the algorithm (e.g., QP, influence matrix, etc.);
  - loss function used for CV (accuracy, F-measure, etc.).
- Easy parallelization / distributed computing.
- Theoretical bounds and experimental results on the speed and accuracy.

TreeCV in action: Leave-One-Out CV estimation

- SVM Classification with PEGASOS (Shalev-Shwartz et al., 2011).
  - CV over the 0-1 loss.
- Least-square regression with SGD (Nemirovski et al., 2009).
  - CV over the squared loss.