iTALS

Fast ALS-based tensor factorization for context-aware recommendation from implicit feedback

Balázs Hidasi – balazs.hidasi@gravityrd.com
Domonkos Tikk – domonkos.tikk@gravityrd.com

ECML/PKDD, 25th September 2012, Bristol

Overview

- Implicit feedback problem
- Context-awareness
  - Seasonality
  - Sequentaility
- iTALS
  - Model
  - Learning
  - Prediction
- Experiments

Feedback types

- Feedback: user-item interraction (events)
- Explicit:
  - Preferences explicitely coded
  - E.g.: Ratings
- Implicit:
  - Preferences not coded explicitely
  - E.g.: purchase history

Problems with implicit feedback

- Noisy positive preferences
  - E.g.: bought & disappointed
- No negative feedback available
  - E.g.: had no info on item

- Usually evaluated by ranking metrics
  - Can not be directly optimized

Why to use implicit feedback?

- Every user provides
- Some magnitudes larger amount of information than explicit feedback
- More important in practice
  - Explicit algorithms are for the biggest only

Context-awareness

- Context: any information associated with events
- Context state: a possible value of the context dimension
- Context-awareness
  - Usage of context information
  - Incorporating additional informations into the method
  - Different predictions for same user given different context states
  - Can greatly outperform context-unaware methods
    - Context segmentates items/users well

Seasonality as context

- Season: a time period
  - E.g.: a week
- Timeband: given interval in season
  - Context-states
  - E.g.: days
- Assumed:
  - aggregated behaviour in a given timeband is similar inbetween seasons
  - and different for different timebands
  - E.g.: daily/weekly routines

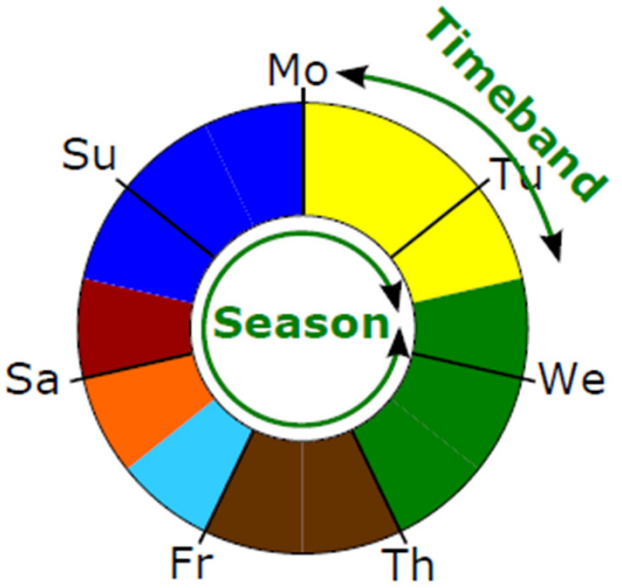

User	Item	Date	Context
1	A	12/07/2010	1
2	B	15/07/2010	3
1	B	15/07/2010	3
...	...	...	
1	A	19/07/2010	1

Sequentiality

- Bought A after B
  - B is the context-state of the user's event on A
- Usefullness
  - Some items are bought together
  - Some items bought repetetively
  - Some are not
- Association rule like information incorporated into model as context
  - Here: into factorization methods
- Can learn negated rules
  - If C then not A

iTALS - Model

- Binary tensor
  - D dimensional
  - User – item – context(s)
- Importance weights
  - Lower weights to zeroes (NIF)
  - Higher weights to cells with more events
- Cells approximated by sum of the values in the elementwise product of D vectors
  - Each for a dimension
  - Low dimensional vectors

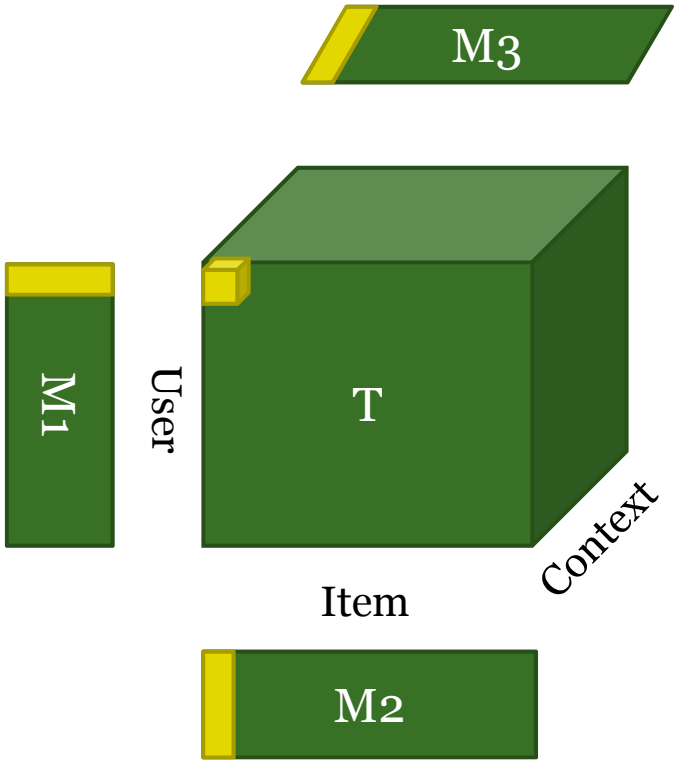

$$\hat{T}_{i_{1,\dots,i_D}} = 1^T \left({M^{(1)}}_{i_1} \circ \dots \circ {M^{(D)}}_{i_D} \right)$$

iTALS - Learning

- Optimizing weighted RMSE
  - Importance weights
- „Missing” values (zeroes) must be considered
  - Scalability issues with many teaching methods
- ALS used
  - Fixing all but one matrices and recompute that one
- Still requires speed-up steps
  - Computation in a non-trivial way (see paper)
- $O\left(K^3 \sum_{i=1}^{D} S_i + K^2 N^+\right)$

iTALS - Prediction

- Sum of values in elementwise product of vectors
  - User-item: scalar product of feature vectors
  - User-item-context: weighted scalar product of feature vectors
- Context-state dependent reweighting of features
- E.g.:
  - Third feature = horror movie
  - Context state1 = Friday night → third feature high
  - Context state2 = Sunday afternoon → third feature low

Experiments

- 5 databases
  - 3 implicit
    - Online grocery shopping
    - VoD consumption
    - Music listening habits
  - 2 implicitizied explicit
    - Netflix
    - MovieLens 10M
- Recall@20 as primary evaluation metric
- Baseline: context-unaware method in every context-state

Scalability

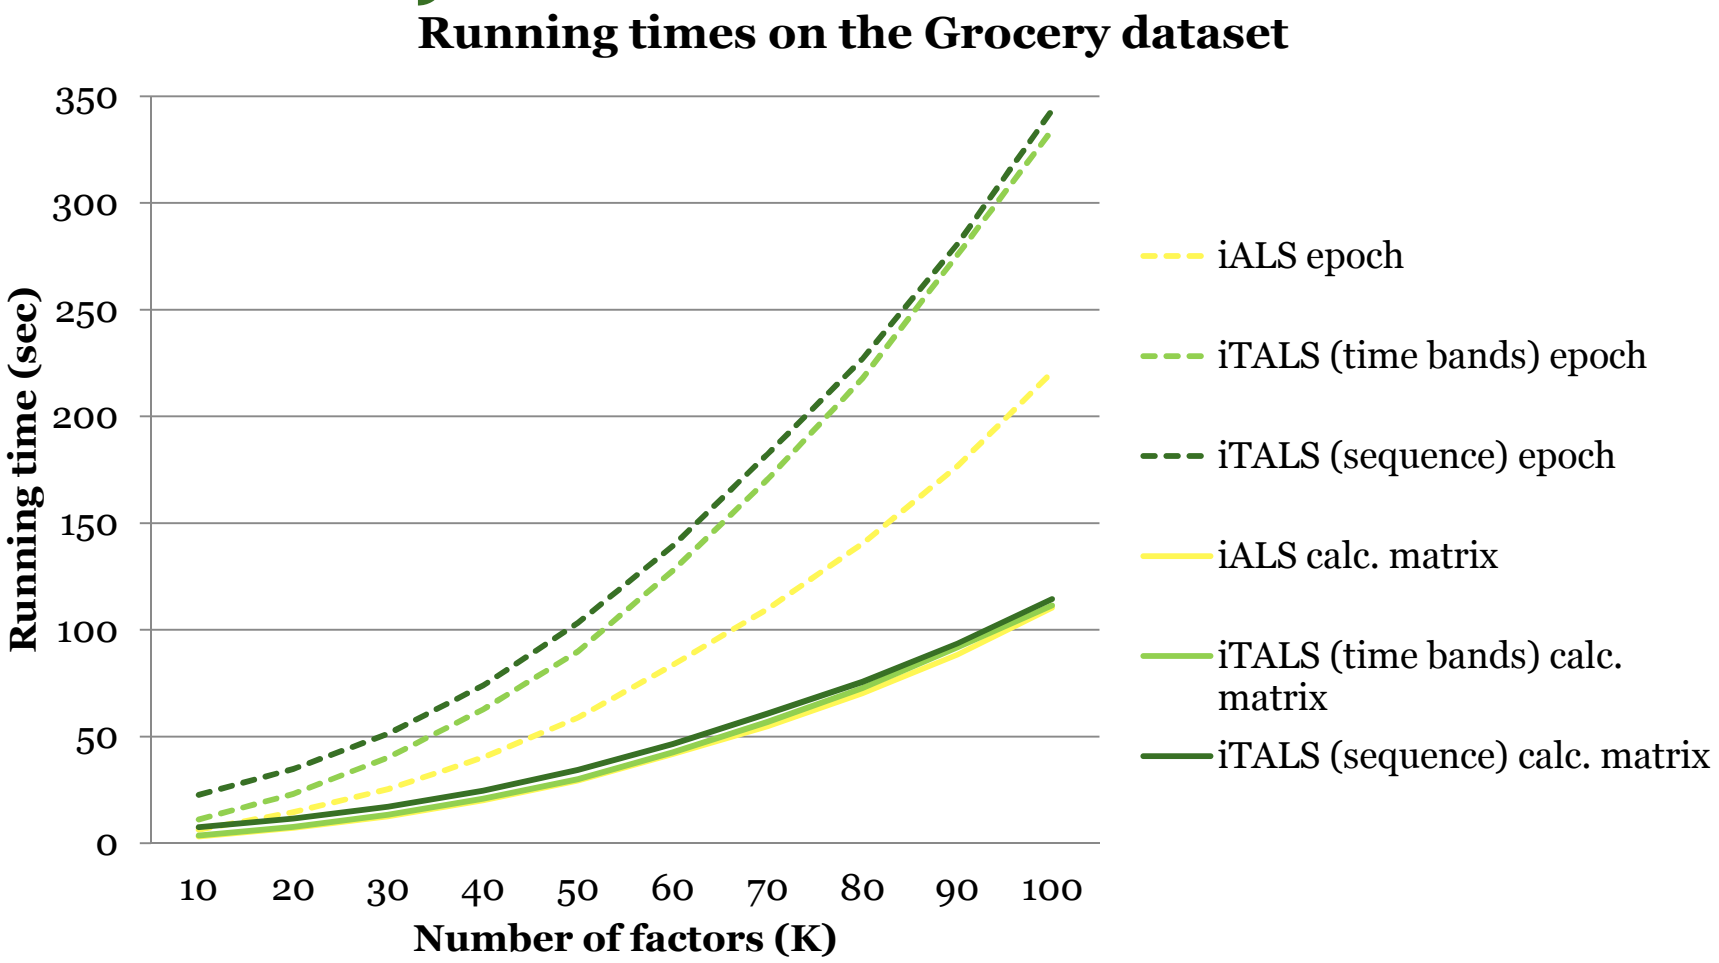

Results – Recall@20

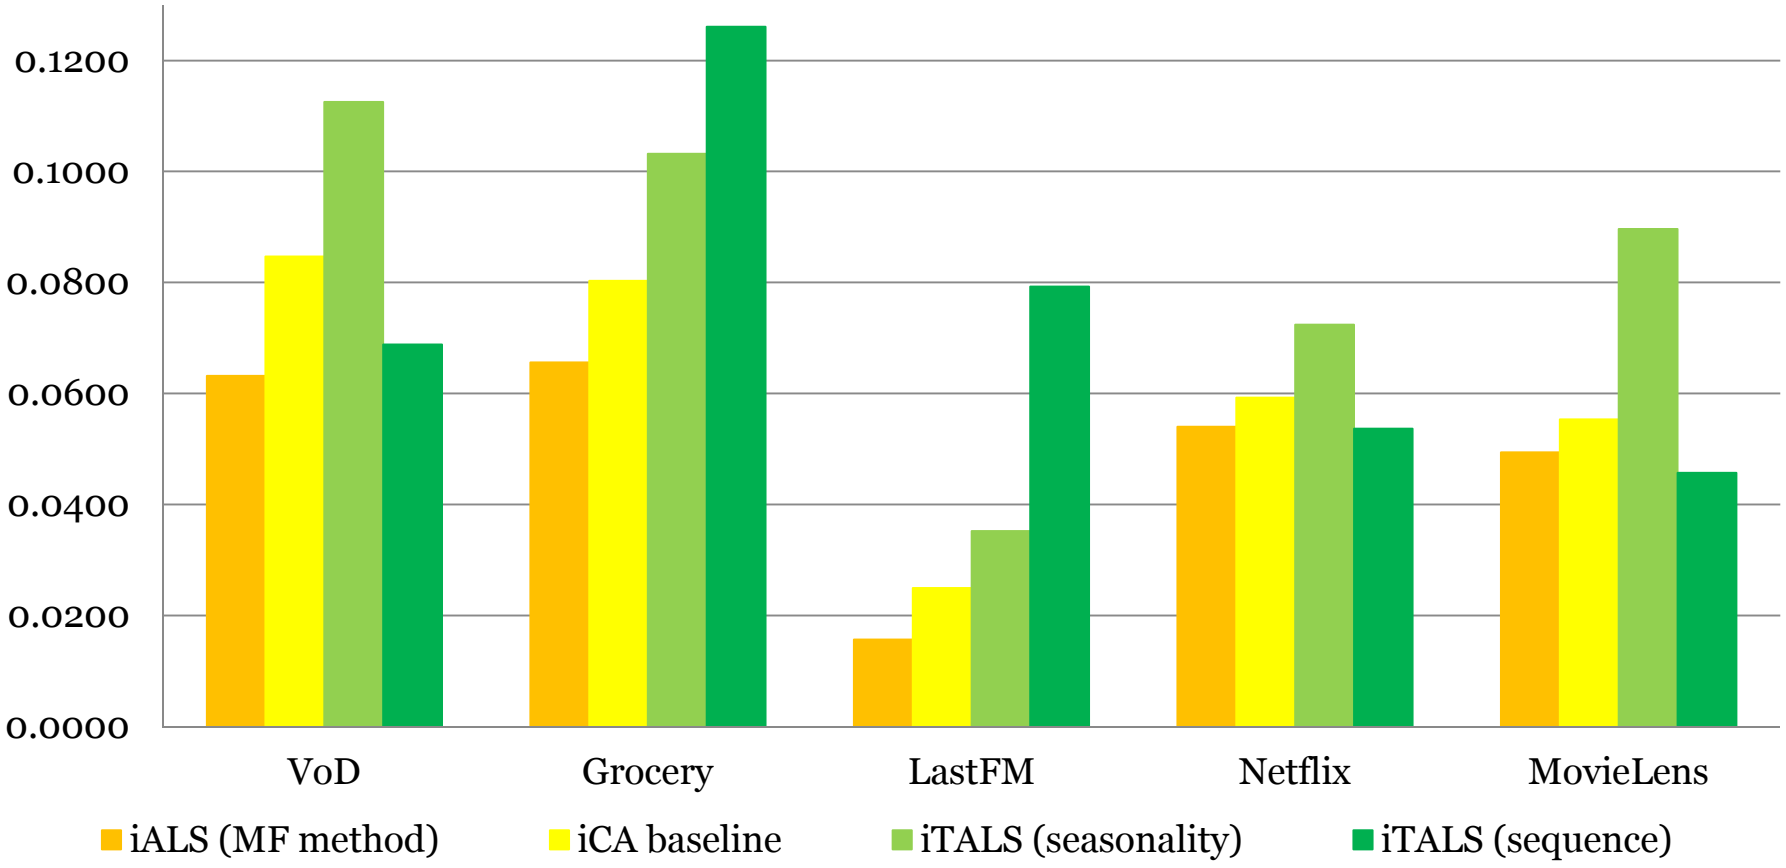

Results – Precision-recall curves

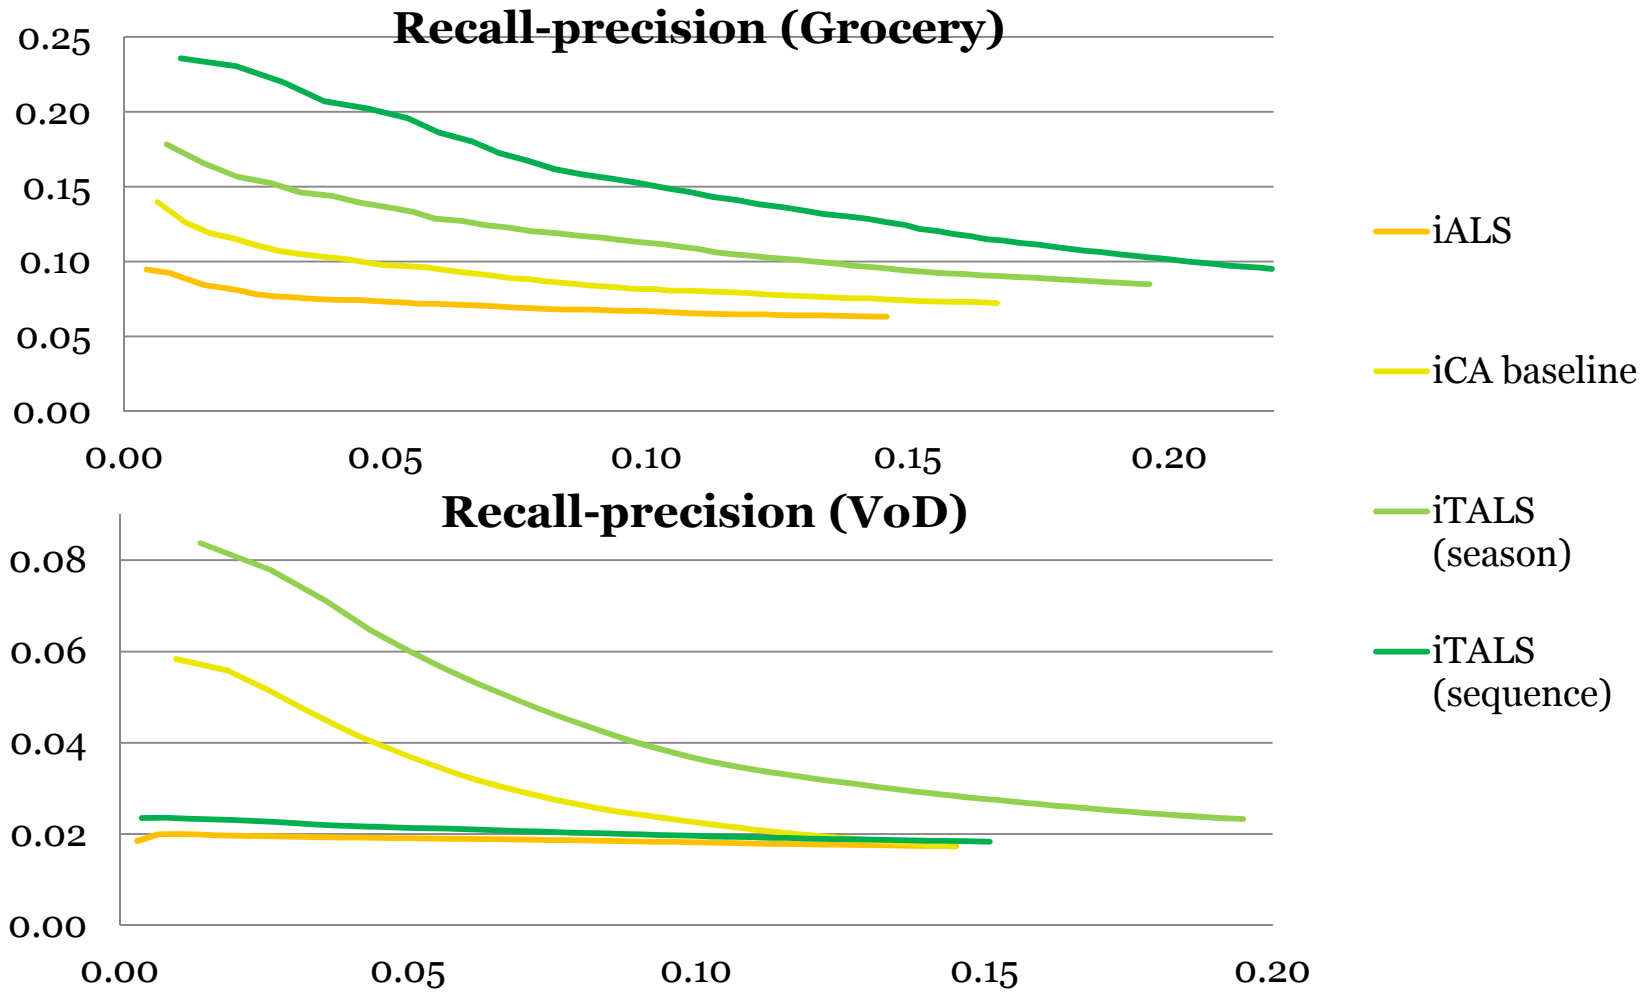

Summary

- iTALS is a
  - scalable
  - context-aware
  - factorization method
  - on implicit feedback data
- The problem is modelled
  - by approximating the values of a binary tensor
  - with elementwise product of short vectors
  - using importance weighting
- Learning can be efficiently done by using
  - ALS
  - and other computation time reducing tricks
- Recommendation accuracy is significantly better than iCA-baseline
- Introduced a new context type: sequentiality
  - association rule like information in factorization framework

Thank you for your attention!

For more of my recommender systems related research visit my website: http://www.hidasi.eu