Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks



Srijan Kumar
Stanford University
Georgia Institute of Technology

Xikun Zhang
UIUC

Jure Leskovec
Stanford University

Code and Data: https://snap.stanford.edu/jodie

Temporal Interaction Networks

Flexible way to represent time-evolving relations

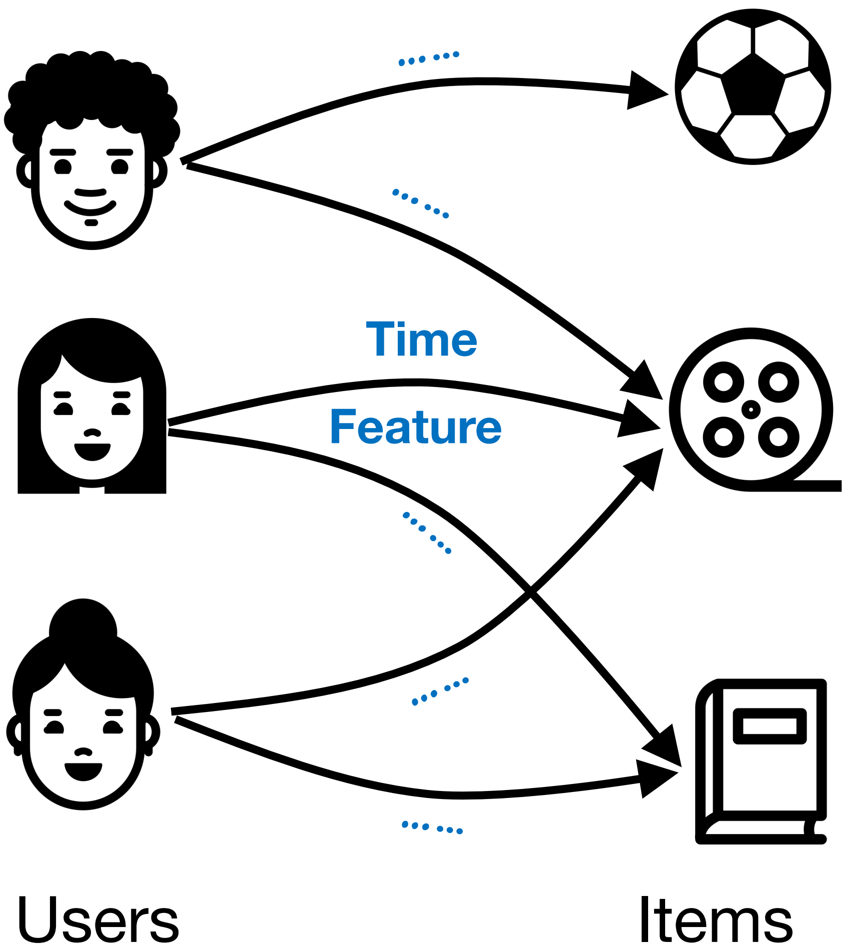

Represented as a sequence of interactions, sorted by time:

$$\underset{\text{interaction}}{S_r} = (\underset{\text{user}}{u_r}, \underset{\text{item}}{i_r}, \underset{\text{time}}{t_r}, \underset{\text{features}}{f_r})$$

Temporal Interaction Networks

[KDD'19]

Application domains

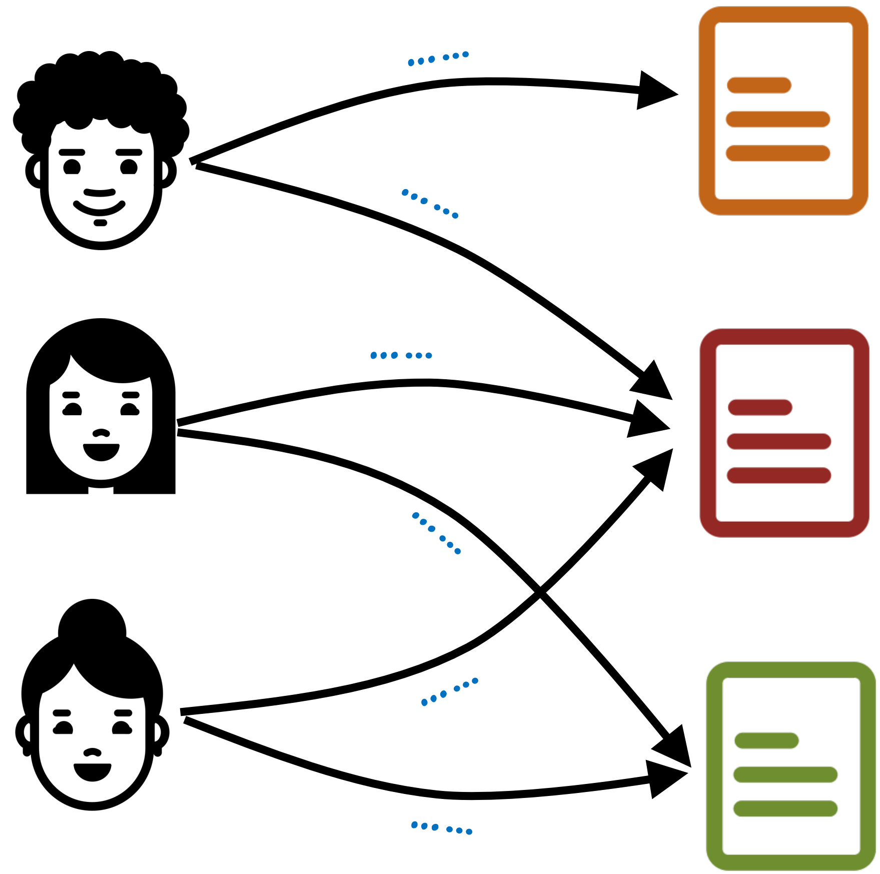

Accounts

Posts

[KDD'19]

Temporal Interaction Networks

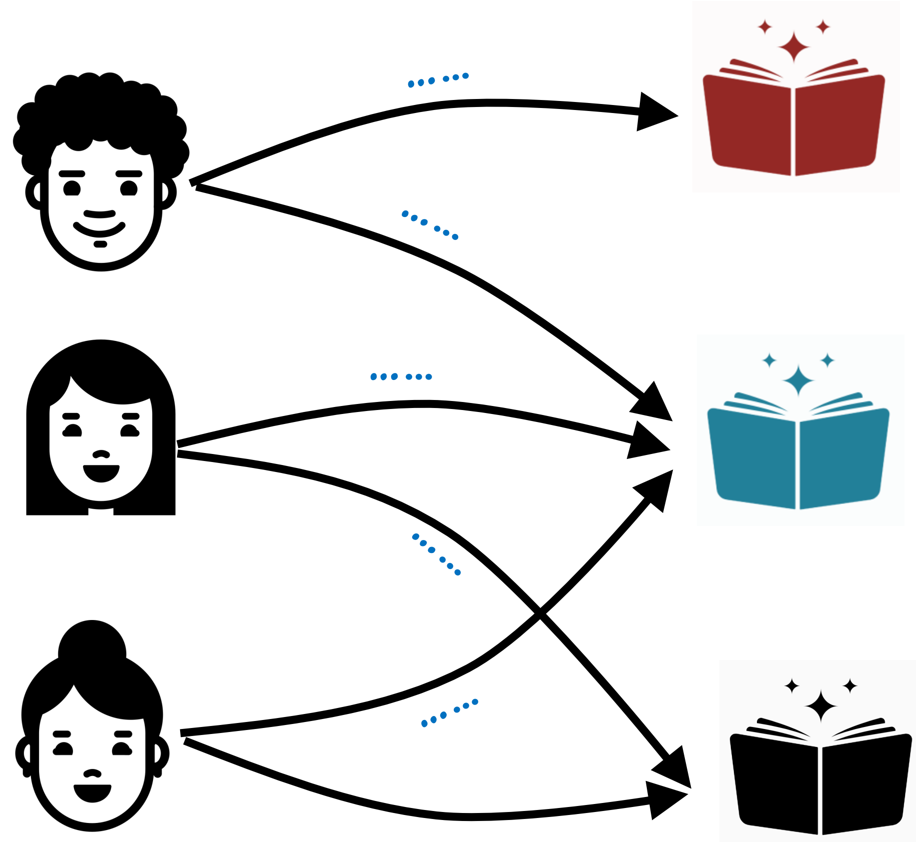

Application domains

Students

Courses

Problem Setup

Given a temporal interaction network

$$S_r = (u_r, i_r, t_r, f_r)$$

interaction user item time features

where $u_r \in \mathcal{U}, i_r \in \mathcal{I}, t_r \in \mathbb{R}^+, 0 < t_1 \leq \ldots \leq T, f_r \in \mathbb{R}^d$

generate an embedding trajectory of every user

$$\mathbf{u(t)} \in \mathbb{R}^n \ \forall u \in \mathcal{U}, \forall t \in [0, T]$$

and an embedding trajectory of every item

$$\mathbf{i(t)} \in \mathbb{R}^n \ \forall i \in \mathcal{I}, \forall t \in [0, T]$$

Goal: Generate Dynamic Trajectory

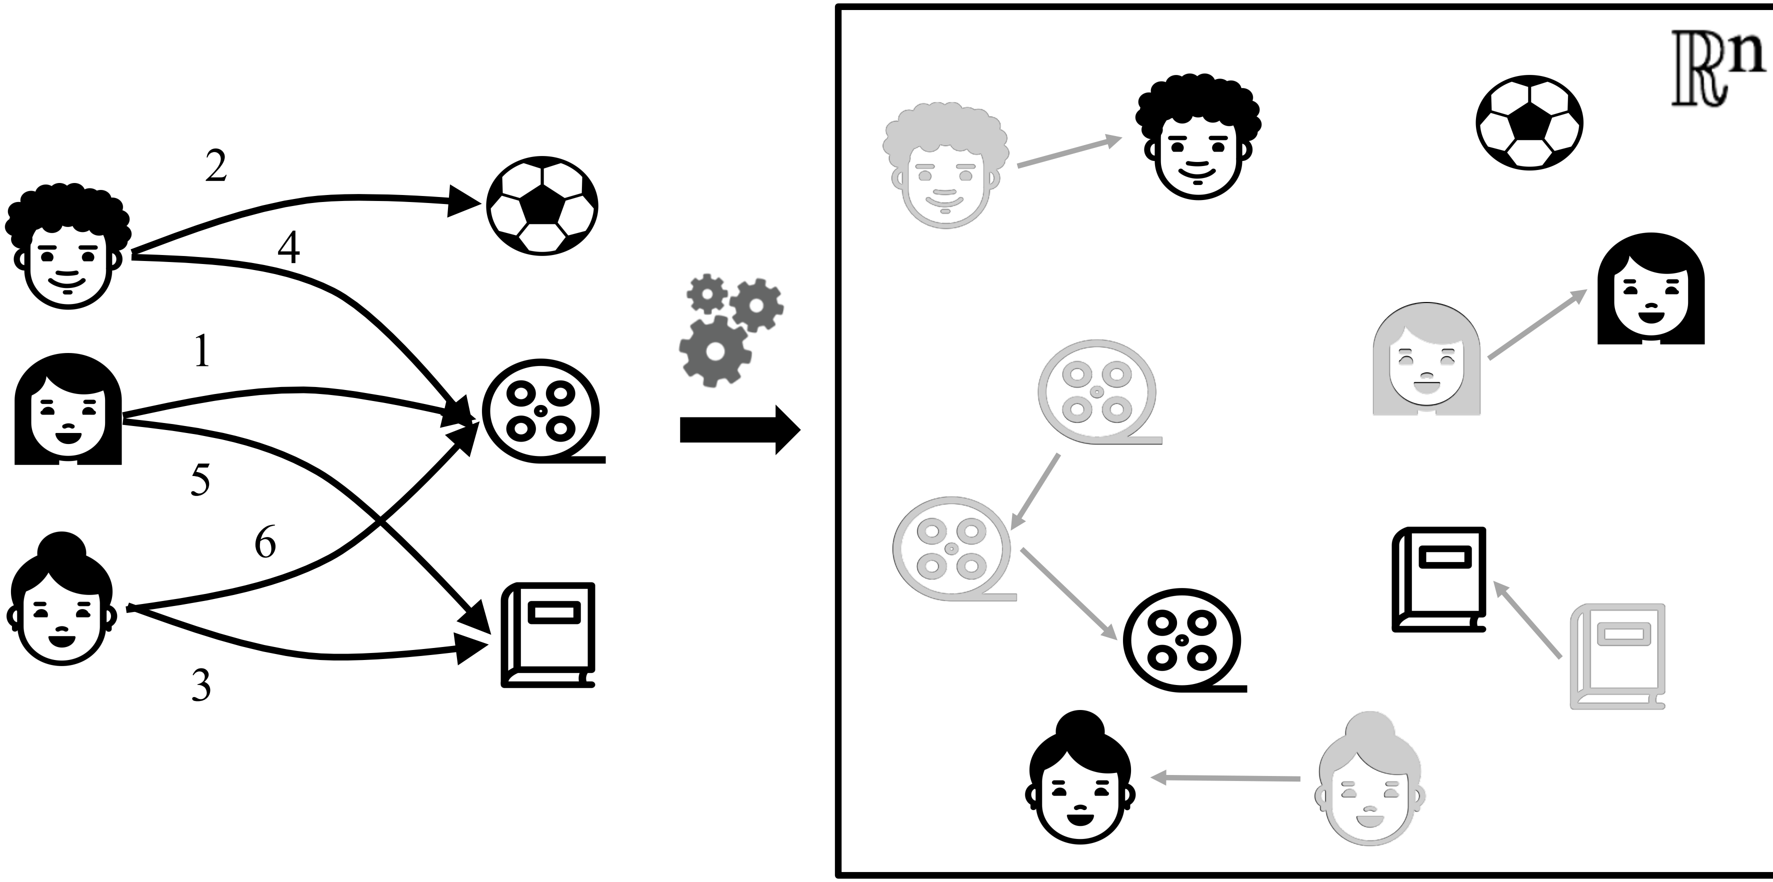

Input: Temporal interaction network

Output: Dynamic trajectory in embedding space

Challenges

Challenges in modeling:

- **C1:** How to learn inter-dependent user and item embeddings?
- **C2:** How to generate embedding for every point in time?

Challenges in scalability:

- **C3:** How to scalably train models on temporal networks?

Existing Methods

	C1 Co-influence	C2 Embed any time	C3 Train in batches
Deep recommender systems • Time-LSTM (IJCAI 2017) • Recurrent Recommender Networks (WSDM 2017) • Latent Cross (WSDM 2018)			✓
Dynamic co-evolution • Deep Coevolve (DLRS, 2016)	✓		
Temporal network embedding • CTDNE (BigNet, 2018)	✓		✓
Our model: JODIE	✓	✓	✓

Our Model: JODIE

JODIE: Joint Dynamic Interaction Embedding

- Mutually-recursive recurrent neural network framework

Update Component

Project Component

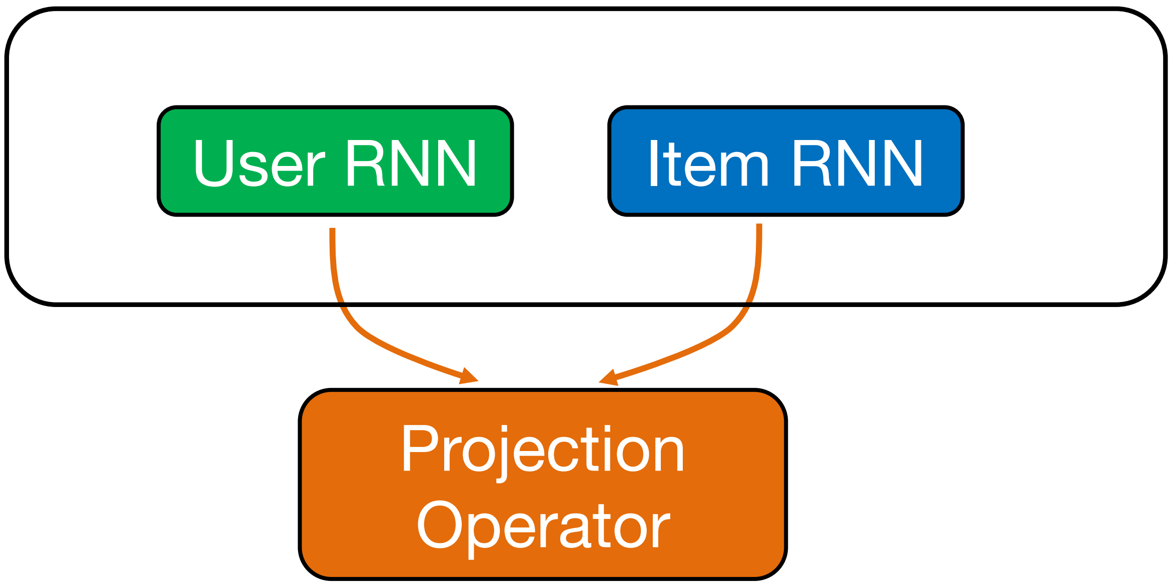

JODIE: Update Component

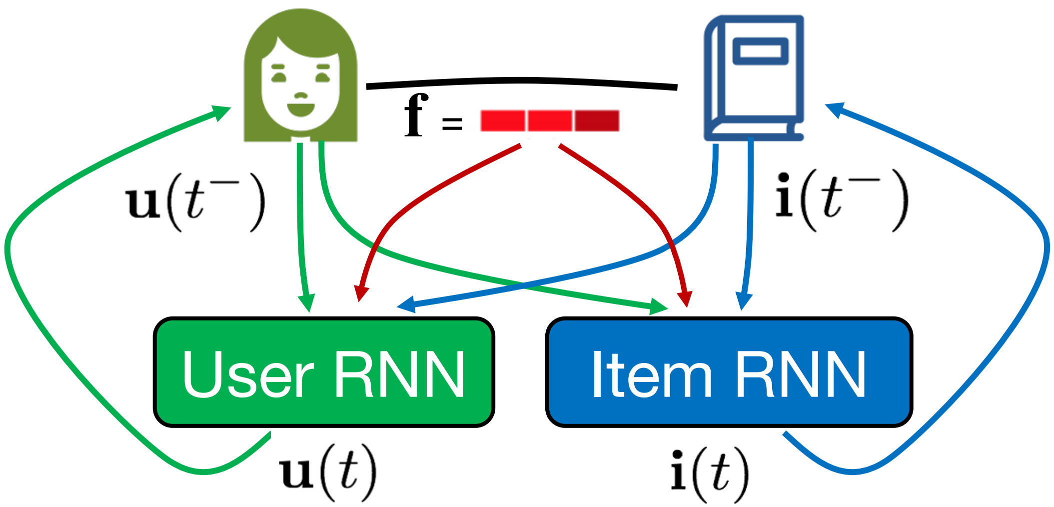

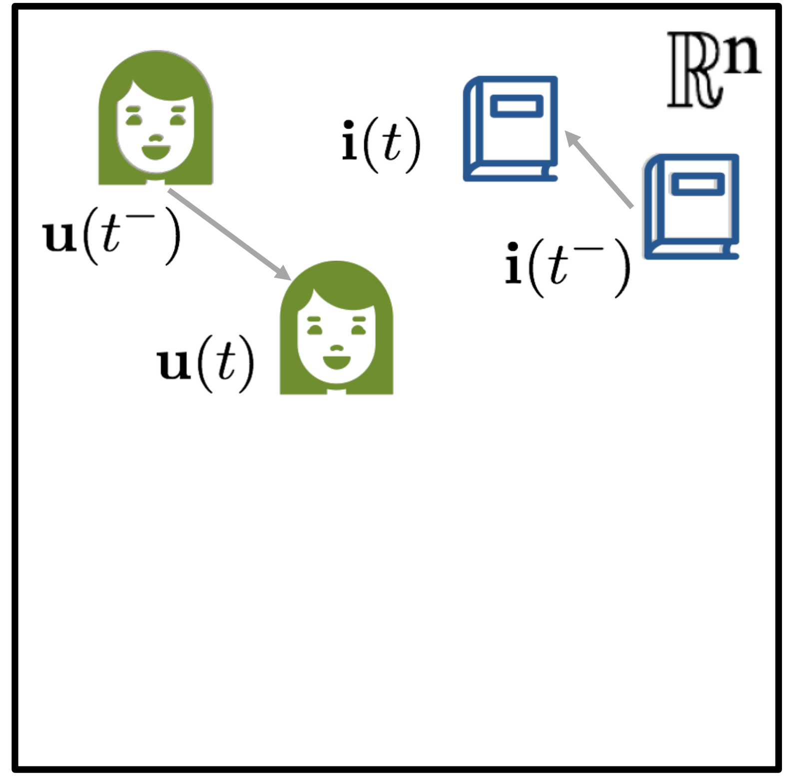

$$\left.\begin{aligned}\mathbf{u}(t) &= \sigma(W_1\mathbf{u}(t^-) + W_2\mathbf{i}(t^-) + W_3\mathbf{f})\\ \mathbf{i}(t) &= \sigma(W_4\mathbf{i}(t^-) + W_5\mathbf{u}(t^-) + W_6\mathbf{f})\end{aligned}\right\}$$

Weight matrices W are trainable

- All users share the User-RNN parameters. Similar for items.

JODIE: Project Component

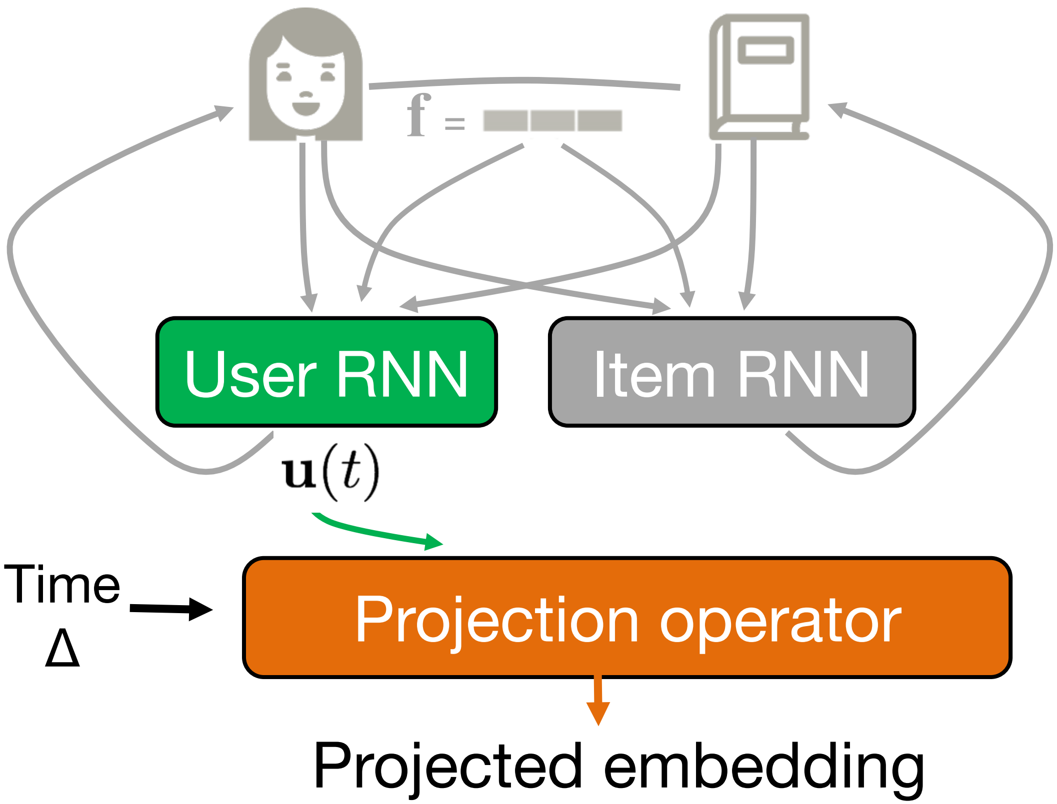

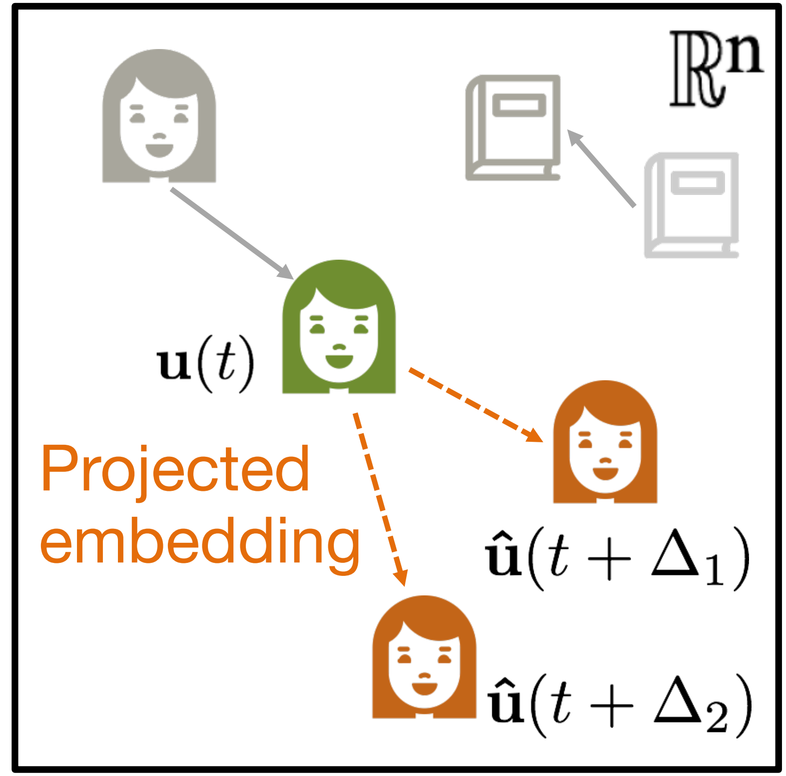

$$\hat{\mathbf{u}}(t+\Delta) = (1 + \mathrm{w}.\Delta) * \mathbf{u}(t)$$

How can we predict the next item?

- Rank items using distance in the embedding space

Summary: JODIE Formulation

Update embeddings:

$$\mathbf{u}(t) = \sigma(W_1 \mathbf{u}(t^-) + W_2 \mathbf{i}(t^-) + W_3 \mathbf{f})$$

$$\mathbf{i}(t) = \sigma(W_4 \mathbf{i}(t^-) + W_5 \mathbf{u}(t^-) + W_6 \mathbf{f})$$

Project user embedding:

$$\hat{\mathbf{u}}(t + \Delta) = (1 + \mathrm{w}.\Delta) * \mathbf{u}(t)$$

Predict next item:

$$\tilde{\mathbf{j}} = W_7 \hat{\mathbf{u}}(t + \Delta) + B$$

Loss:

$$\sum_{S_r:(u,i,t,f)} ||\tilde{\mathbf{j}}(t) - \mathbf{i}(t^-)||_2 + \lambda_U ||\mathbf{u}(t) - \mathbf{u}(t^-)||_2$$

$$+\lambda_I ||\mathbf{i}(t) - \mathbf{i}(t^-)||_2$$

Predicted next item is close to the real item embedding

Smoothness in evolving embeddings

[KDD'19]

Challenges in Dynamic Trajectories

Challenges in learning:

- **C1:** How to learn inter-dependent user and item embeddings? **Solution: Update component**
- **C2:** How to generate embedding for every point in time? **Solution: Project component**

Challenges in scalability:

- **C3:** How to scalably train models on temporal networks?

[KDD'19]

Standard Training Processes: N/A

Training must maintain temporal order

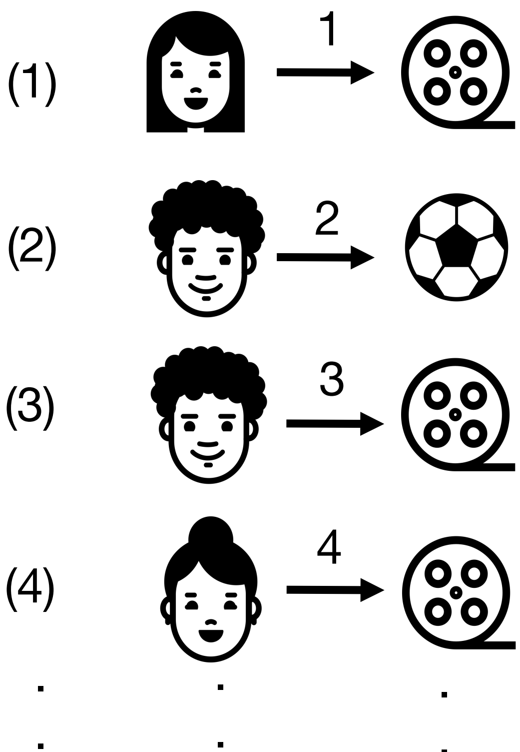

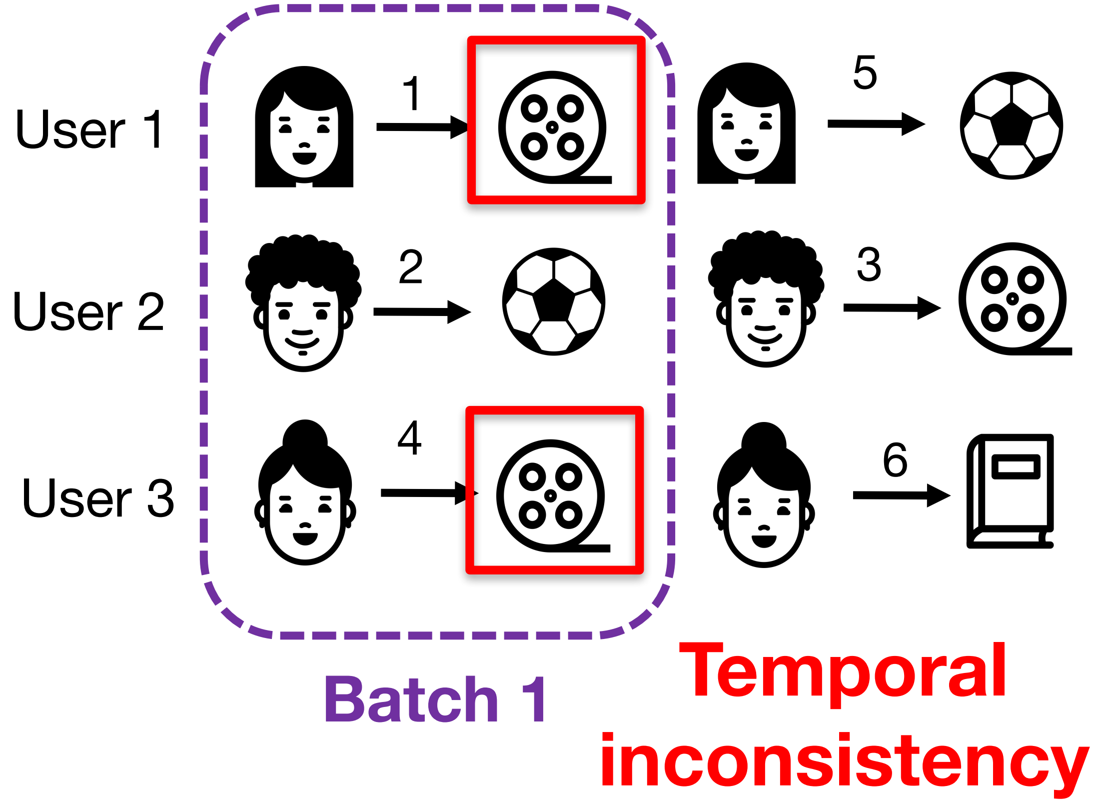

Sequential processing: not scalable

Split by user (or item): not allowed

T-batch: Batching for Scalability

[KDD'19]

T-batch: Temporal data batching algorithm

- **Main idea: create each batch as an independent edge set**

- **Create a sequence of batches**
  - Interactions in each batch are processed in parallel
  - Process the batches in sequence to maintain temporal ordering

T-batch: Batching for Scalability

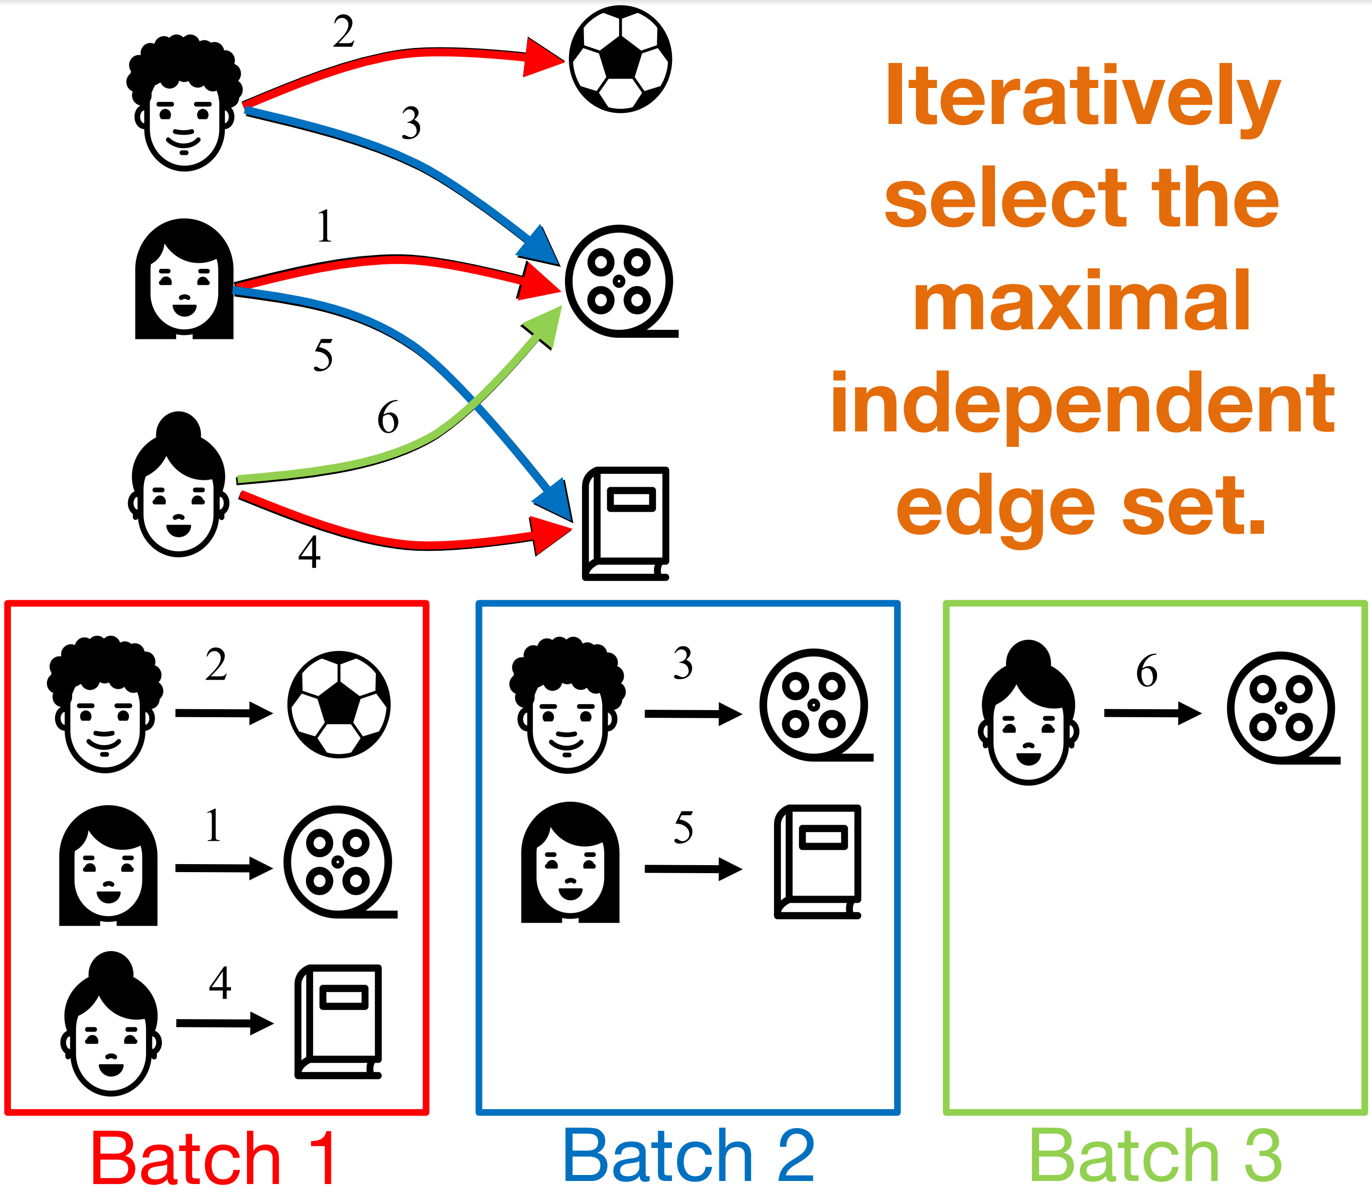

Iteratively select the maximal independent edge set.

Challenges in Dynamic Trajectories

[KDD'19]

Challenges in learning:

- **C1:** How to learn inter-dependent user and item embeddings? **Solution: Update component**
- **C2:** How to generate embedding for every point in time? **Solution: Project component**

Challenges in scalability:

- **C3:** How to scalably train models on temporal networks? **Solution: T-batch Algorithm**

Experiments: Prediction Tasks

- **Temporal Link Prediction:**
  - Which item $i \in I$ will user u interact with at time t?
- **Temporal Node Classification:**
  - Does a user u become anomalous after an interaction?
- **Settings:**
  - **Temporal Splits:** 80%, 10%, 10%
  - **Metrics:** Mean reciprocal rank, Recall@10, AUROC

Code and Data: https://snap.stanford.edu/jodie

Datasets

Dataset	Users	Items	Interactions	Temporal Anomalies	
Reddit	10,000	984	672,447	366	**NEW!**
Wikipedia	8,227	1,000	157,474	217	**NEW!**
LastFM	980	1,000	1,293,103	-	
MOOC	7,047	97	411,749	4,066	

Code and Data: https://snap.stanford.edu/jodie

Experiment 1: Link Prediction

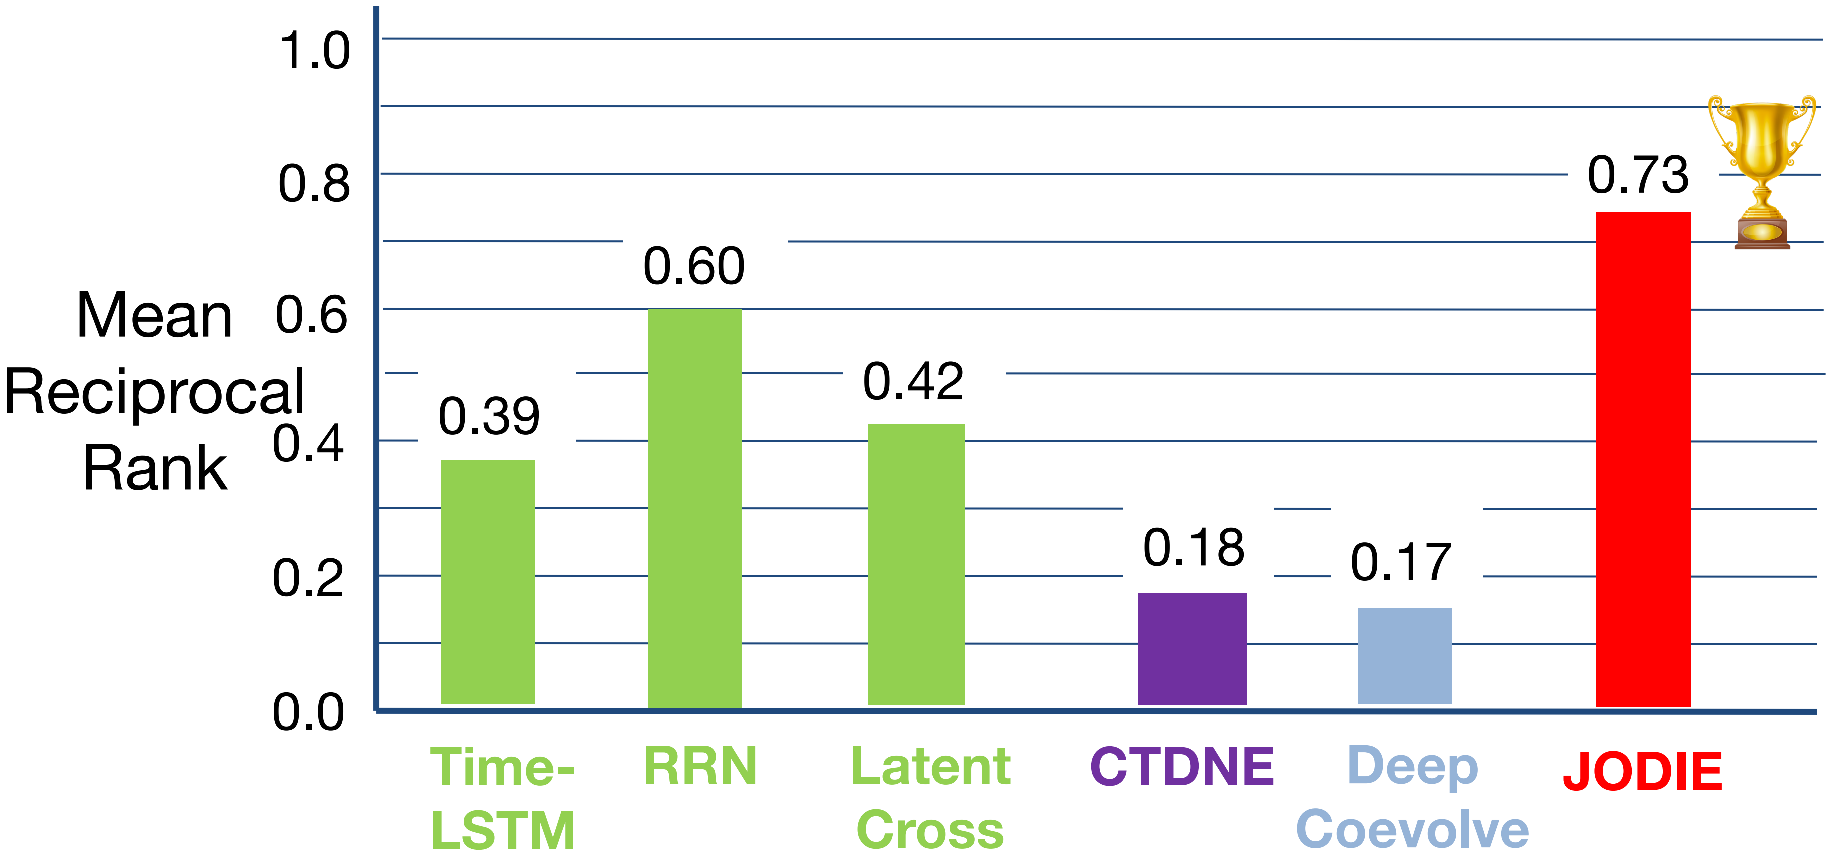

JODIE outperforms baselines by > 20%

[KDD'19]

Experiment 2: Node Classification

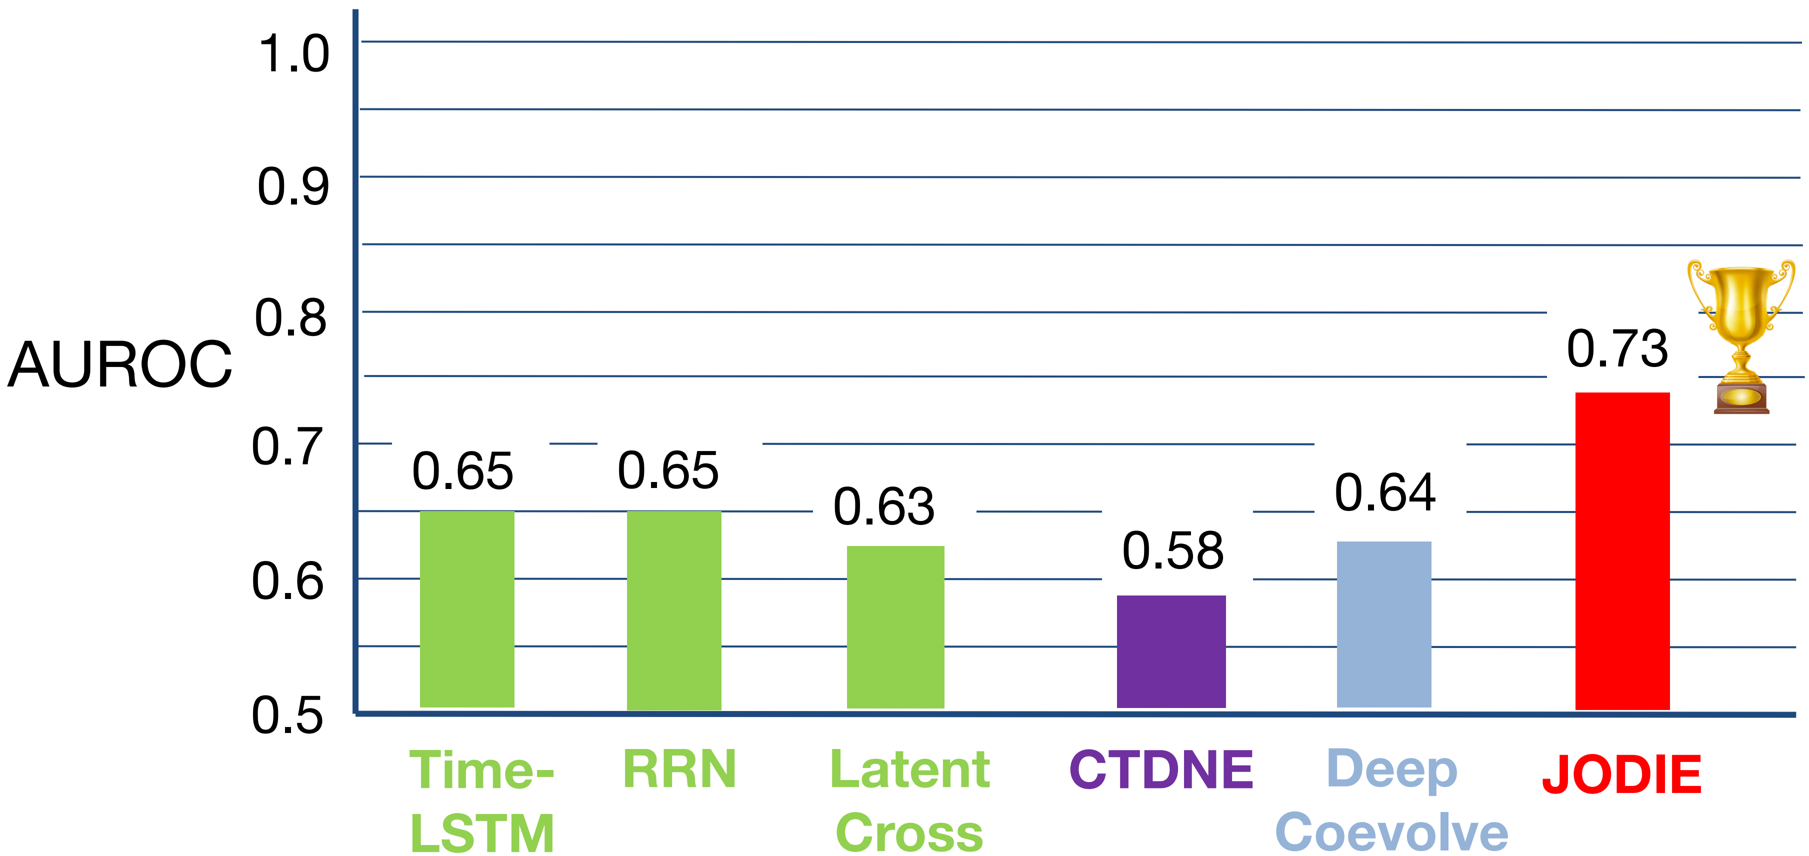

JODIE outperforms all baselines by >12%

Experiment 3: T-batch Speed-up

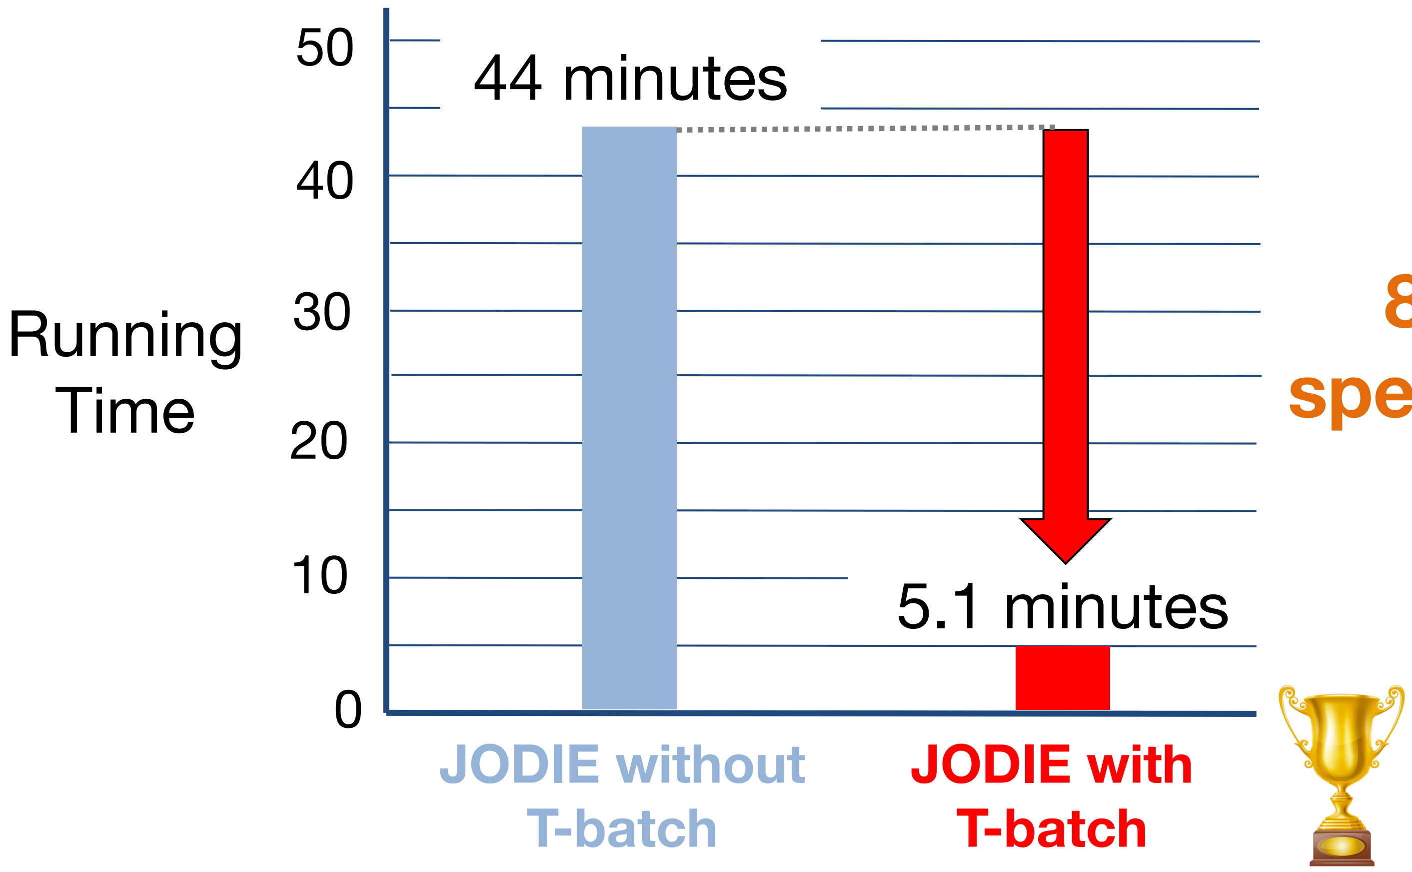

8.5x speed-up

T-batch leads to 8.5x speed-up in training

Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks

Srijan Kumar, Xikun Zhang, Jure Leskovec

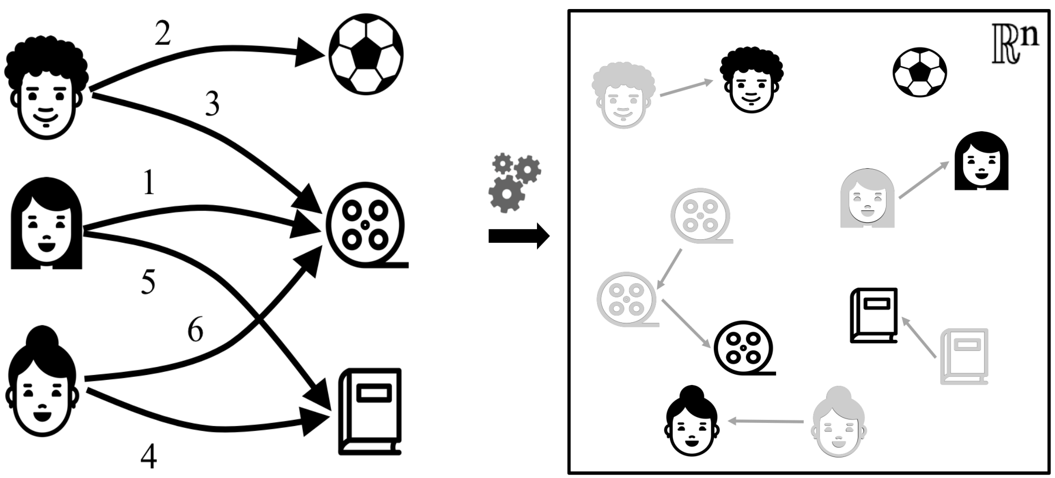

JODIE generates and projects embedding trajectories

- **JODIE: a mutually-recursive RNN** framework
- **T-batch:** 8.5x training speed-up
- **Efficient** in temporal link prediction and node classification
- **Extendible** to > 2 entity types

Code and Data: https://snap.stanford.edu/jodie

Open Positions @ Georgia Tech

- **Hiring multiple Ph.D. students**
- **Research areas:**
  - Machine Learning for Networks
  - Safety, Integrity, and Anti-Abuse
  - Computational Social Science
- **Collaborations**

Georgia Tech

Contact: srijan@cs.stanford.edu

Predicting Dynamic Embedding Trajectory in Temporal Interaction Networks

Srijan Kumar, Xikun Zhang, Jure Leskovec

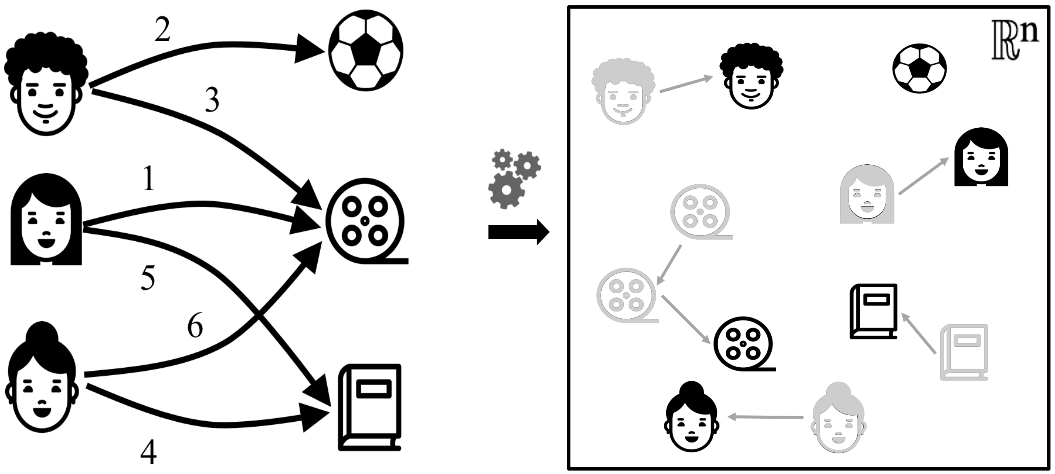

JODIE generates and projects embedding trajectories

- **JODIE: a mutually-recursive RNN** framework
- **T-batch:** 8.5x training speed-up
- **Efficient** in temporal link prediction and node classification
- **Extendible** to > 2 entity types

Code and Data: https://snap.stanford.edu/jodie