



# Continuous CP Decomposition of Sparse Tensor Streams



Taehyung Kwon\*



Inkyu Park\*



Dongjin Lee




Kijung Shin

# Roadmap

---

- **Introduction <<**
- Problem Definition
- Proposed Method: **SliceNStitch**
  - Continuous Tensor Model
  - Optimization Algorithms
- Experiment Results
- Conclusions



# Tensors are Everywhere


---

- **Tensors**

- Multi-dimensional arrays
- Powerful tools to represent multi-aspect data

- **Tensor Streams**

- Tensor data are collected incrementally over time



(source, destination, 1)

**Traffic history data**



(user, product, quantity)

**Purchase history data**


# Tensors are Everywhere

- **Tensors**


- Multi-dimensional arrays
- Powerful tools to represent multi-aspect data

- **Tensor Streams**


- Tensor data are collected incrementally over time



**Traffic history data**



**Purchase history data**




# CP Decomposition

- **CANDECOMP/PARAFAC (CP) decomposition**

- Low-rank approximation of the tensor

- **Consider:**  $\mathcal{X} \in \mathbb{R}^{N_1 \times \dots \times N_M}$ , rank  $R$
- **To Find:** factor matrices  $A^{(1)}, \dots, A^{(M)}$
- **Minimize:**  $\|\mathcal{X} - \tilde{\mathcal{X}}\|_F$  where  $\tilde{\mathcal{X}} \equiv \sum_{r=1}^R A^{(1)}(:, r) \circ \dots \circ A^{(M)}(:, r)$



# CP Decomposition


- **Alternating Least Squares (ALS) [2]**

- Standard algorithm for computing CPD of static tensor


- **Input:** static tensor  $\mathcal{X} \in \mathbb{R}^{N_1 \times \dots \times N_M}$ , rank  $R$
- **Initialize** factor matrices  $\{A^{(m)}\}_{m=1}^M$
- While not converge:
  - For  $m = 1, \dots, M$ :
    - $A^{(m)} \leftarrow \operatorname{argmin}_{A^{(m)}} \|\mathcal{X} - \tilde{\mathcal{X}}\|_F$  (least-squares problem)
- **Output:** factor matrices  $\{A^{(m)}\}_{m=1}^M$

# Limitations of Common Tensor Modeling

- Tensor streams **grow once per period**
  - Outputs of CPD are also updated once per period
- To perform CPD continuously for real-time application,
  - **Granularity** of the time mode must be **extremely fine**



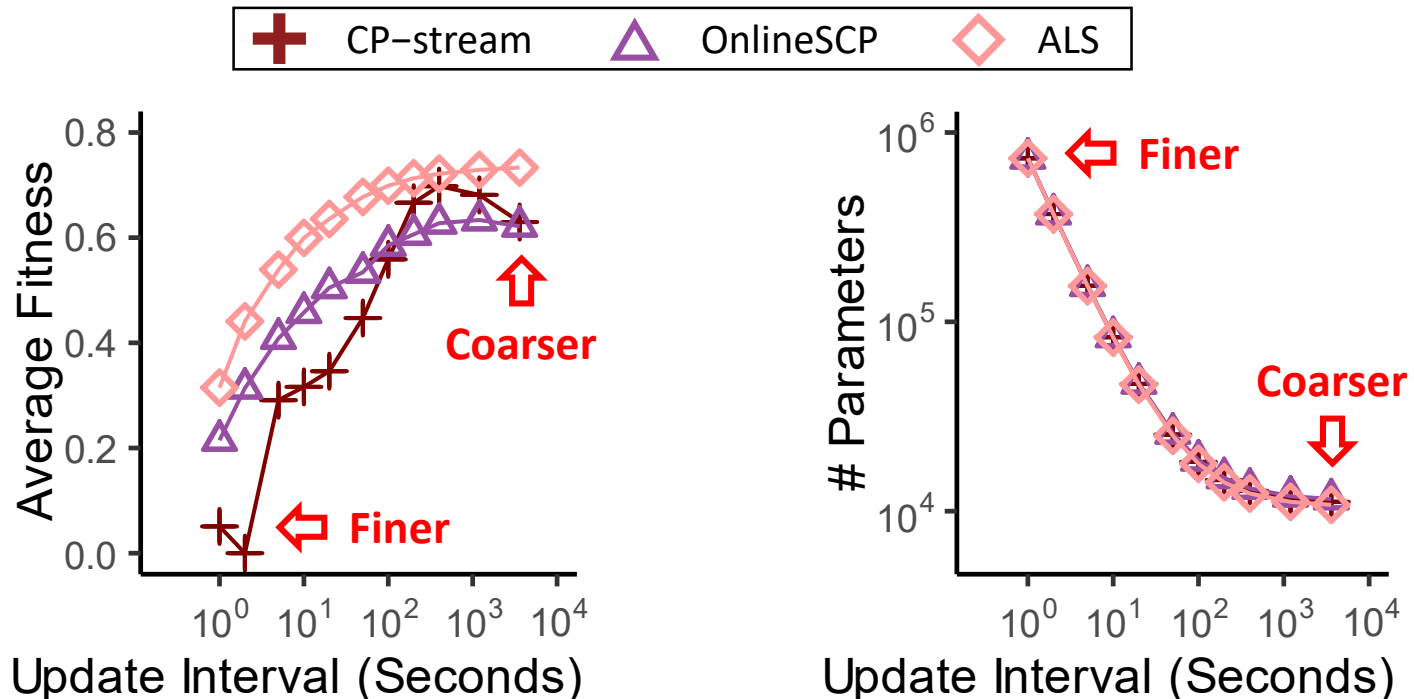
**Course-grained Tensor**



**Fine-grained Tensor**

# Limitations of Common Tensor Modeling

- Problems of fine-grained tensor modelings
  - Degradation of fitness
  - Increase the number of parameters





# In This Work

---

- We propose **SliceNStitch** for continuous CPD, which is **fast**, **space-efficient**, and **accurate**
  - New data model
  - Fast online algorithms


	Coarse-grained	Fine-grained	SliceNStitch (Proposed)
Update Interval	Long (👎)	Short (👍)	<b>Short (👍)</b>
Parameters	Few (👍)	Many (👎)	<b>Few (👍)</b>
Fitness	High (👍)	Low (👎)	<b>High (👍)</b>

**Remark:** This work has appeared at ICDE 2021 [1]

# Roadmap

---

- Introduction
- **Problem Definition <<**
- Proposed Method: **SliceNStitch**
  - Continuous Tensor Model
  - Optimization Algorithms
- Experiment Results
- Conclusions



# Problem Definition

---

- **Continuous CP decomposition**

- **Given:** a multi-aspect data stream
- **Update:** CP decomposition instantly in response to each new event in the stream
- **Without:** waiting for the current period to end

# Roadmap


---

- Introduction
- Problem Definition
- Proposed Method: **SliceNStitch**
  - **Continuous Tensor Model <<**
    - Optimization Algorithms
- Experiment Results
- Conclusions




# Continuous Tensor Model

- Tensor window and units evolve at each time




**Tensor window at time 3:00:00**



**Tensor window at time 3:00:01**

# Event-driven Implementation

- Each single data causes an event:
  - Move the quantity to the next tensor unit ( $\mathcal{X} \rightarrow \mathcal{X} + \Delta\mathcal{X}$ )



# Roadmap

---

- Introduction
- Problem Definition
- Proposed Method: **SliceNStitch**
  - Continuous Tensor Model
  - **Optimization Algorithms <<**
- Experiment Results
- Conclusions



# SliceNStitch-Matrix ( $\text{SNS}_{\text{MAT}}$ )

- Factor matrices of previous window are good initial points
- Single iteration of ALS is enough to achieve high accuracy

- **Input:** (1) tensor  $\mathcal{X} + \Delta\mathcal{X} \in \mathbb{R}^{N_1 \times \dots \times N_M}$   
(2) factor matrices  $\{A^{(m)}\}_{m=1}^M$  of previous window  $\mathcal{X}$
- For  $m = 1, \dots, M$ :
  - $A^{(m)} \leftarrow \underset{A^{(m)}}{\operatorname{argmin}} \|(\mathcal{X} + \Delta\mathcal{X}) - \tilde{\mathcal{X}}\|_F$  (least-squares problem)
- **Output:** updated factor matrices  $\{A^{(m)}\}_{m=1}^M$



# SliceNStitch-Matrix ( $\text{SNS}_{\text{MAT}}$ )

---

- **Pros:**

- High-quality solution, since it uses all non-zero entries


- **Cons:**

- High computational cost, since it uses all non-zero entries <<

- **Solution:** Update only the rows of factor matrices which are used for approximating changed entries

# Outline of Improved Algorithms

- Update **time mode** factor matrix:




3<sup>rd</sup> Mode Factor Matrix  $A^{(3)}$



Rows to be updated in  $A^{(3)}$

# Outline of Improved Algorithms

- Update **non-time mode** factor matrices:




# SliceNStitch-Vector ( $\text{SNS}_{\text{VEC}}$ )

- **Time mode:**

- Approximate  $\mathcal{X}$  as  $\tilde{\mathcal{X}}$  in the least squares solution
- Computational complexity  $\propto \text{NNZ}(\Delta\mathcal{X})$

- **Non-time mode:**

- Approximate the least squares problem
- $\min_{A^{(m)}} \|(\mathcal{X} + \Delta\mathcal{X}) - \tilde{\mathcal{X}}\|_F \rightarrow \min_{A^{(m)}(:, i_m)} \|(\mathcal{X} + \Delta\mathcal{X}) - \tilde{\mathcal{X}}\|_F$
- Computational complexity  $\propto \text{NNZ}(\text{Slice which contains } \Delta\mathcal{X})$



# SliceNStitch-Vector ( $\text{SNS}_{\text{VEC}}$ )

---

- **Pros:**

- Significantly faster than  $\text{SNS}_{\text{MAT}}$

- **Cons:**

- Numerically unstable ( $\because$  no normalization)
- Slow if many non-zeros are in the slice <<

- **Solution:** Use random sampling with smaller sample size if too many non-zeros are in the slice

# SliceNStitch-Random ( $\text{SNS}_{\text{RND}}$ )

---

- If  $\text{NNZ}(\text{Slice contains } \Delta\mathcal{X}) \leq \theta$ :

- Use  $\text{SNS}_{\text{VEC}}$

- Otherwise:

- Randomly select  $\theta$  indices ( $= S$ ) from the slice contains  $\Delta\mathcal{X}$
- Define  $\bar{\mathcal{X}}_S$  s.t.  $\tilde{\mathcal{X}}(J) + \bar{\mathcal{X}}_S(J) = \begin{cases} \mathcal{X}(J), & \text{if } J \in S \\ \tilde{\mathcal{X}}(J), & \text{otherwise} \end{cases}$  : Original tensor : Approximation
- Approximate  $\mathcal{X}$  as  $\tilde{\mathcal{X}} + \bar{\mathcal{X}}_S$  in the least squares solution of  $\text{SNS}_{\text{VEC}}$
- Computational complexity  $\propto \theta$

# SliceNStitch-Random ( $\text{SNS}_{\text{RND}}$ )

---

- **Pros:**

- Computational complexity is constant

- **Cons:**

- Numerically unstable ( $\because$  no normalization) <<



- **Solution:** Use clipping to prevent the extreme value which yields numerical instability

# SliceNStitch-Stable ( $\text{SNS}_{\text{VEC}}^+$ , $\text{SNS}_{\text{RND}}^+$ )


---

- **Update entries one by one**
  - **Clip** each value if the absolute value is larger than  $\eta$

$\text{SNS}_{\text{VEC}}$  and  $\text{SNS}_{\text{RND}}$ :  
Update the row **at once**



$\text{SNS}_{\text{VEC}}^+$  and  $\text{SNS}_{\text{RND}}^+$ :  
Update row entries **one by one**





# Roadmap

---


- Introduction
- Problem Definition
- Proposed Method: **SliceNStitch**
  - Continuous Tensor Model
  - Optimization Algorithms
- **Experiment Results <<**
- Conclusions



# Experimental Settings

---

- 4 real-world sparse time series datasets



Divvy Bikes



Chicago Crimes




New York Taxi



Ride Austin


- 4 baselines that update CPD periodically
  - **ALS** [2], **onlineSCP** [3], **CP-stream** [4], **NeCPD** [5]

# Exp 1. SliceNStitch is Accurate




$$\text{Relative Fitness} = \frac{\text{Fitness}_{\text{target}}}{\text{Fitness}_{\text{ALS}}}$$


where  $\text{Fitness} = 1 - \frac{\|\tilde{\mathcal{X}} - \mathcal{X}\|_F}{\|\mathcal{X}\|_F}$




New York Taxi



Chicago Crimes




Divvy Bikes



Ride Austin


# Exp 1. SliceNStitch is Accurate

- SliceNStitch achieve 72 - 100% relative fitness compared to the most accurate baseline



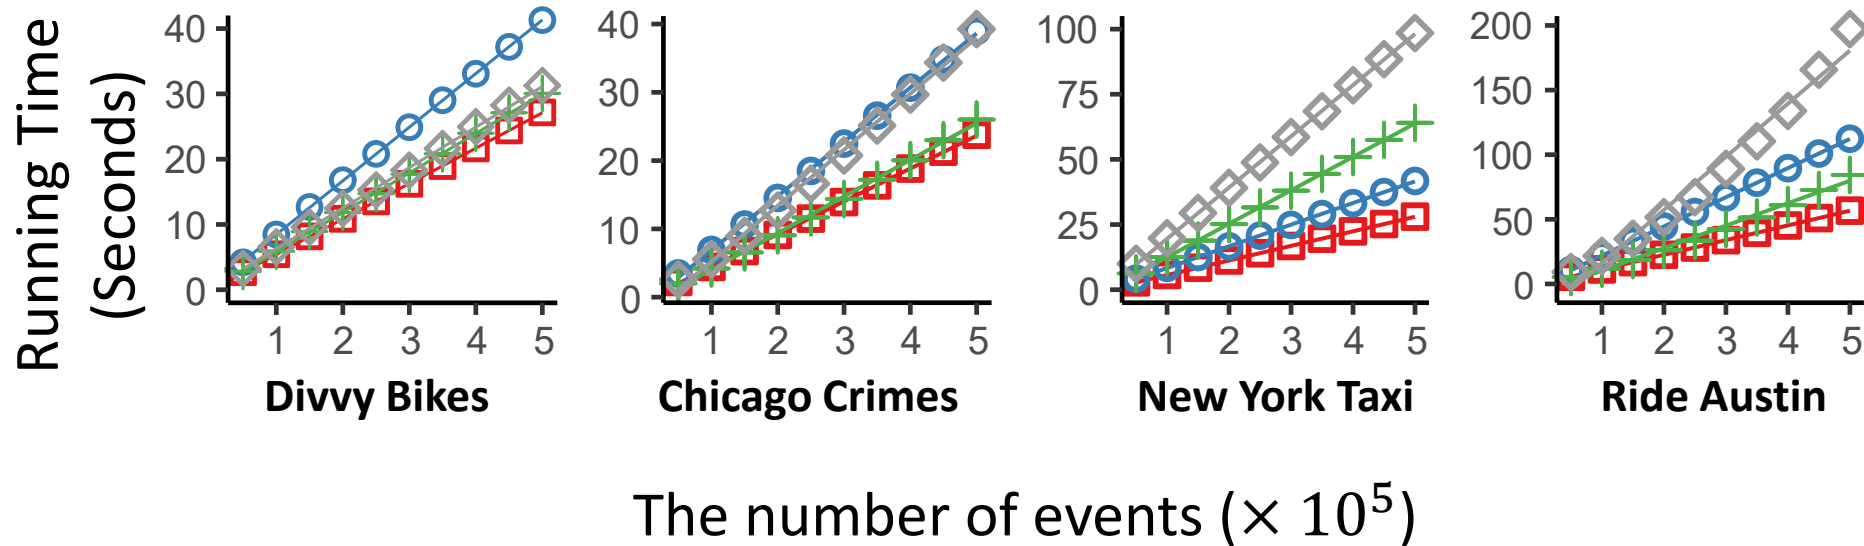

# Exp 2. SliceNStitch is Fast

- $\text{SNS}_{\text{RND}}^+$  is up to  $464 \times$  faster than CP-stream




# Exp 3. SliceNStitch is Scalable

- Total runtime of SliceNStitch is **linear** in the number of events



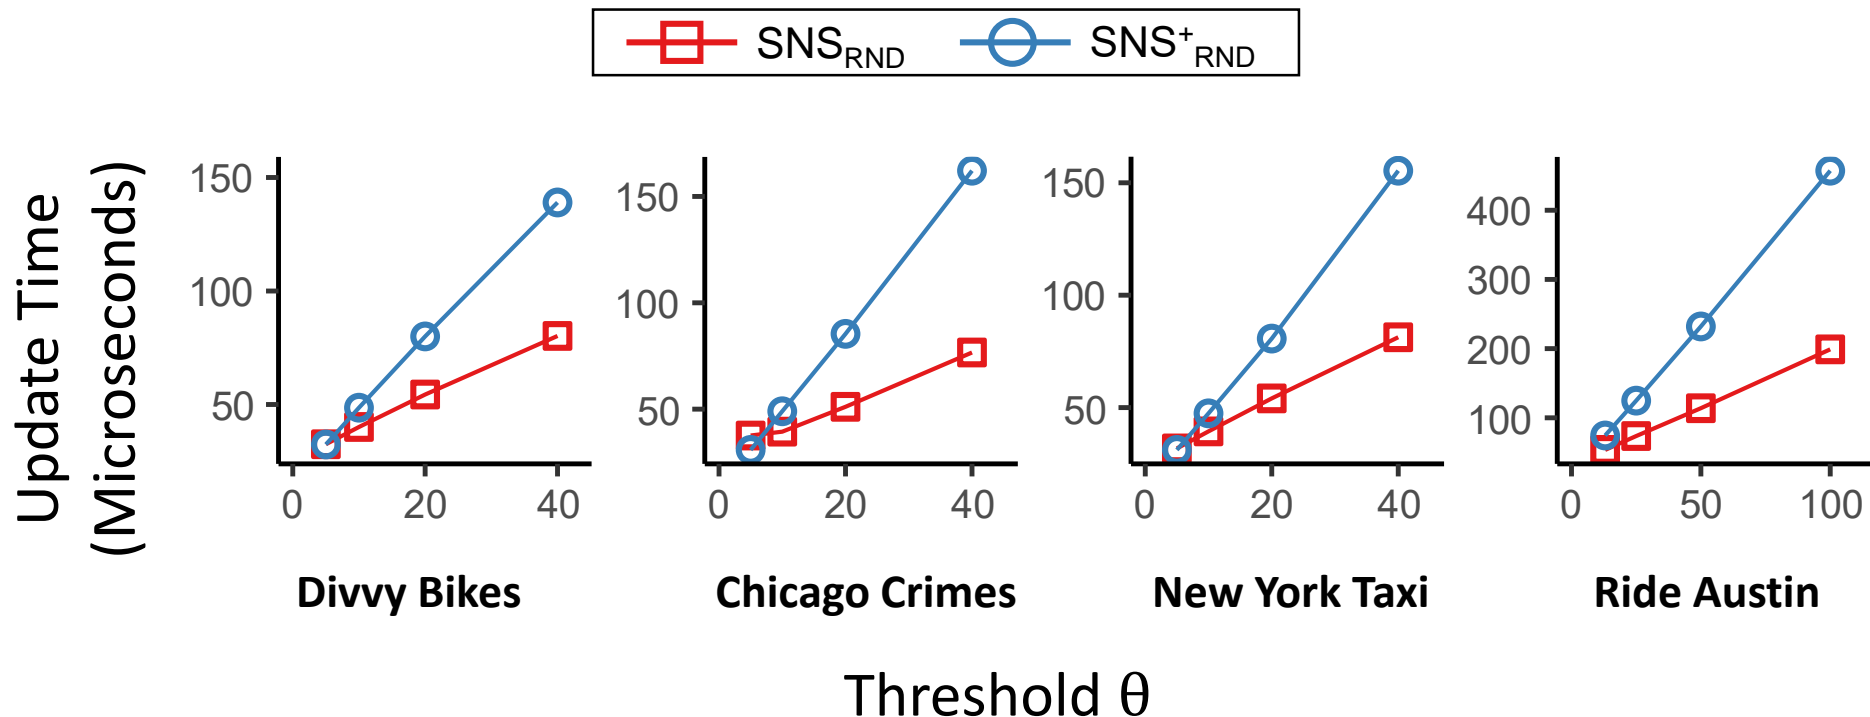
# Exp 4. Effect of Sampling Threshold $\theta$

- As  $\theta$  increases,
  - Fitness increases with diminishing returns
  - Runtime grows linearly



# Exp 4. Effect of Sampling Threshold $\theta$


- As  $\theta$  increases,
  - Fitness increases with diminishing returns
  - Runtime grows linearly





# Exp 5. Effect of Clipping Value $\eta$


- Fitness is insensitive to  $\eta$  as long as  $\eta$  is small enough



# Practitioner's Guide

---


- We do not recommend  $\text{SNS}_{\text{VEC}}$  and  $\text{SNS}_{\text{RND}}$  due to instability issues
- We recommend using the most accurate version within your runtime budget



- If  $\text{SNS}_{\text{RND}}^+$  is chosen, increase the sampling threshold  $\theta$  enough within your runtime budget

# Application: Anomaly Detection

- SliceNStitch detects anomalies much faster with comparable accuracy



	Precision @ Top 20	Time Gap between Occurrence and Detection
SNS <sup>+</sup> <sub>RND</sub>	<b>0.80</b>	<b>0.0015 sec</b>
OnlineSCP	<b>0.80</b>	1601.00 sec
CP-stream	0.70	1424.57 sec

# Roadmap





---

- Introduction
- Problem Definition
- Proposed Method: **SliceNStitch**
  - Continuous Tensor Model
  - Optimization Algorithms
- Experiment Results
- **Conclusions <<**



# Conclusion

- We propose **SliceNStitch** for continuous CPD with
  - **Near-instant updates**
  - **High fitness**
  - **Small number of parameters**



# References

---

- [1] Taehyung Kwon\*, **Inkyu Park\***, Dongjin Lee, and Kijung Shin, “SliceNStitch: Continuous CP Decomposition of Sparse Tensor Streams,” in ICDE 2021.
- [2] R. A. Harshman, “Parafac2: Mathematical and Technical Notes,” UCLA Working Papers in Phonetics, vol. 22, 30-44, 1972.
- [3] Shuo Zhou, Sarah M. Erfani, and James Bailey, “Online CP Decomposition for Sparse Tensors,” in ICDM, 2018.
- [4] Shaden Smith, Kejun Huang, Nicholas D. Sidiropoulos, and George Karypis, “Streaming Tensor Factorization for Infinite Data Sources,” in SDM, 2018.
- [5] Ali Anaissi, Basem Suleiman, and Seid Miad Zandavi, “NeCPD: An Online Tensor Decomposition with Optimal Stochastic Gradient Descent,” arXiv preprint arXiv:2003.08844, 2020.