



SGDP: A Stream-Graph Neural Network Based Data Prefetcher

Yiyuan Yang*
University of Oxford
Huawei Noah's Ark Lab
Oxford, UK
yiyuan.yang@cs.ox.ac.uk

Rongshang Li*
University of Sydney
Huawei Noah's Ark Lab
Sydney, Australia
roli5128@uni.sydney.edu.au

Qiquan Shi
Huawei Noah's Ark Lab
Huawei
Shenzhen, China
shiqiquan@huawei.com


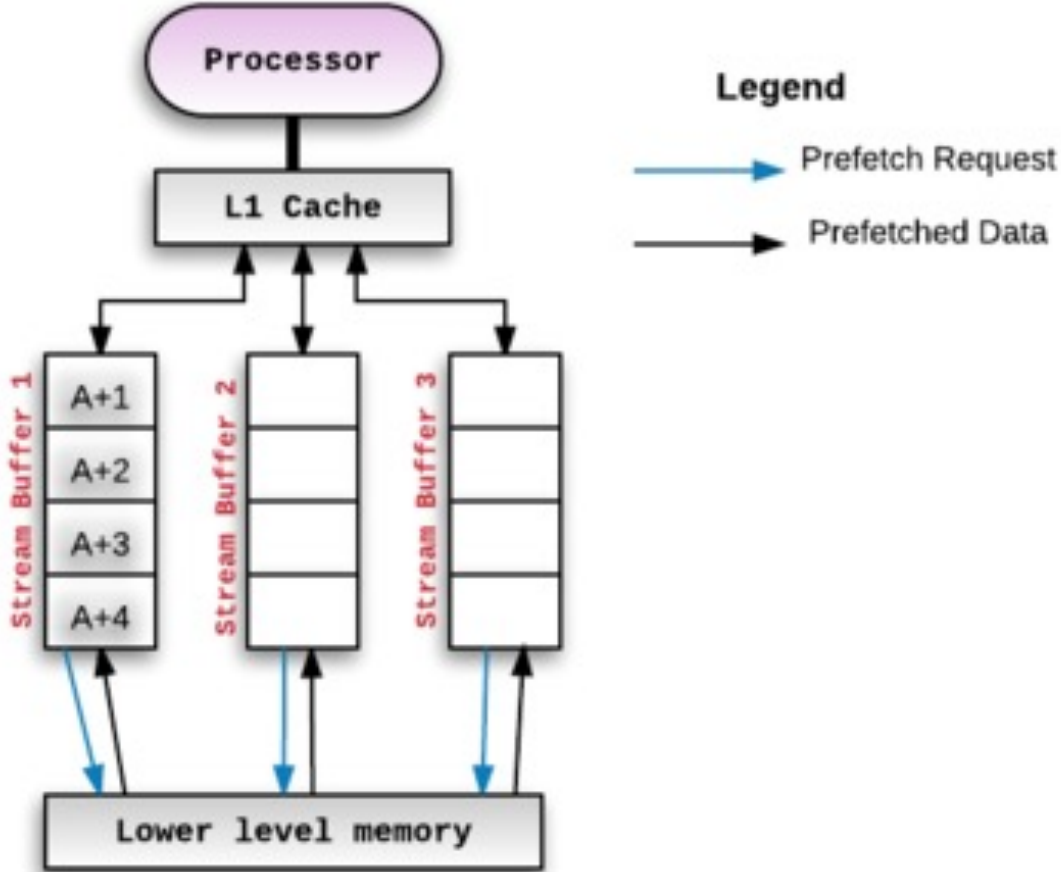
Xijun Li
MIRA Lab, USTC
Huawei Noah's Ark Lab
Shenzhen, China
xijun.li@huawei.com

Gang Hu
Huawei
Chengdu, China
hugang27@huawei.com

Xing Li
Huawei Noah's Ark Lab
Hongkong, China
li.xing2@huawei.com

Mingxuan Yuan
Huawei Noah's Ark Lab
Hongkong, China
Yuan.Mingxuan@huawei.com

Background



Contribution



1. SGDP can accurately learn complex access patterns by capturing the relations of LBA deltas in each stream. The relations are represented by sequential connect matrices and full-connect matrices using graph structures.
2. SGDP is the first work that utilizes the stream-graph structure of the LBA delta in the data prefetching problem. Using gated graph neural networks and attention mechanisms, we extract and aggregate sequential and global information for better prefetching.
3. As a novel solution in the hybrid storage system, SGDP can be generalized to multiple variants by different stream construction methods, which further enhances its robustness and expands its application to various real-world scenarios.
4. SGDP outperforms SOTA by 6.21% on hit ratio, 7.00% on effective prefetching ratio, and speeds up inference time by 3.13X. It has been verified in commercial hybrid storage systems in the experimental phase and will be deployed in the future product series.

Preliminaries

$$\langle lba_i \rangle_{i=1}^n = \langle lba_1, lba_2, \dots, lba_n \rangle, \quad (1)$$

$$ld_i = lba_{i+1} - lba_i, \quad (2)$$

$$\langle ld_i \rangle_{i=1}^{n-1} = \langle ld_1, ld_2, \dots, ld_{n-1} \rangle. \quad (3)$$

$$\widehat{lba}_{n+1} = lba_n + \widehat{ld}_n. \quad (4)$$

Methodology

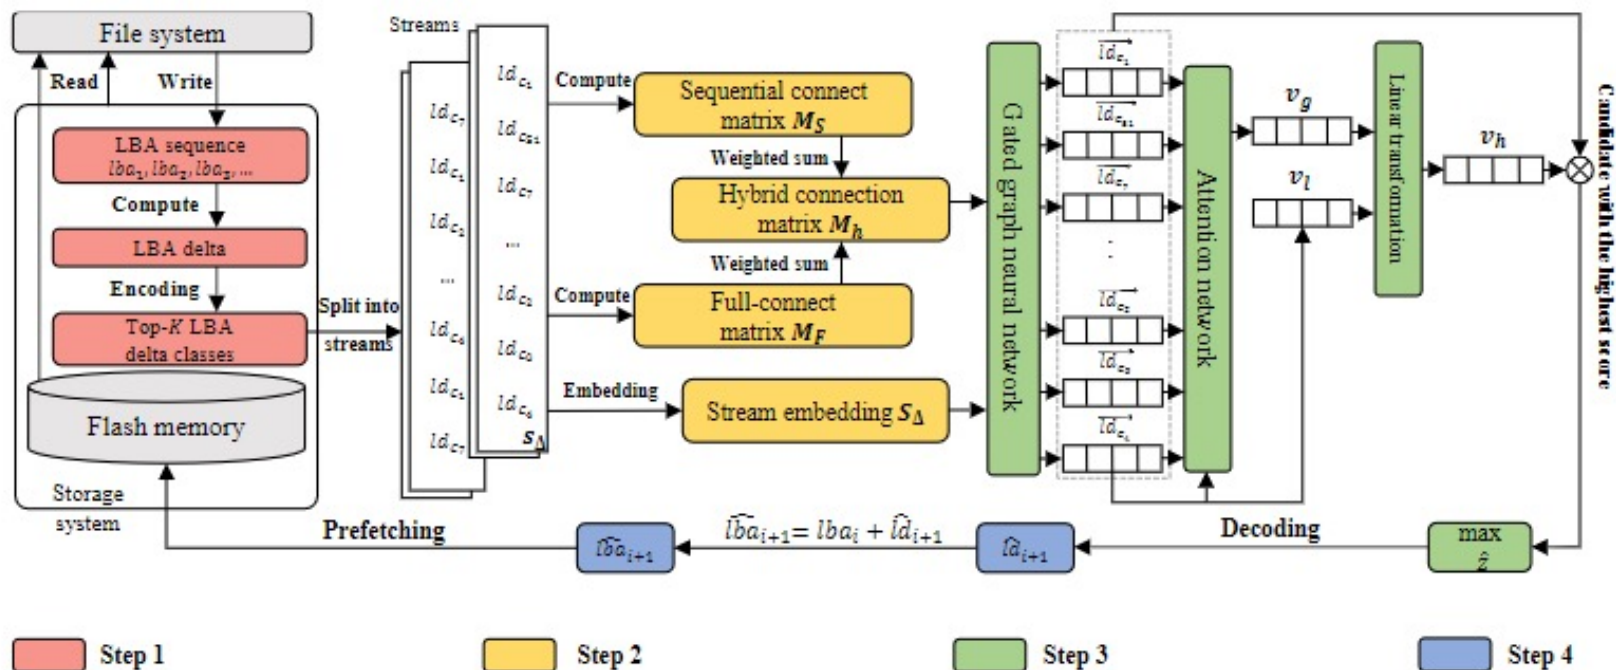


Fig. 1. The workflow of the SGDP framework. In **Step 1**, we compress the search space and reduce the learning complexity. In **Step 2**, we compute the hybrid connection matrix M_h with sequential and global information and embed the LBA delta stream into a matrix S_Δ . In **Step 3**, using gated graph neural networks to update the latent node vectors. Each stream is represented as the combination of the local preference v_l and global interaction v_g by an attention network. In **Step 4**, we predict the candidate with the highest score and decode it to get the next accessed LBA for prefetching. This framework corresponds to the four steps of **Algorithm 1**.

Methodology

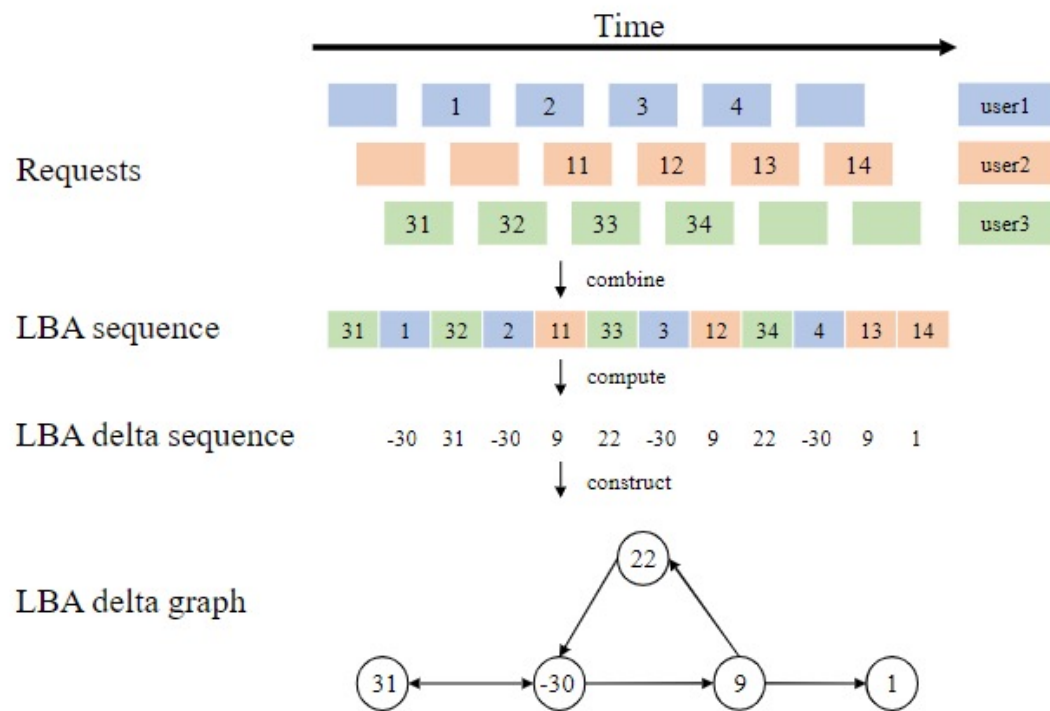


Fig. 2. Example of LBA delta and graph.

LBA delta stream

LBA delta class nodes

LBA delta stream graphs

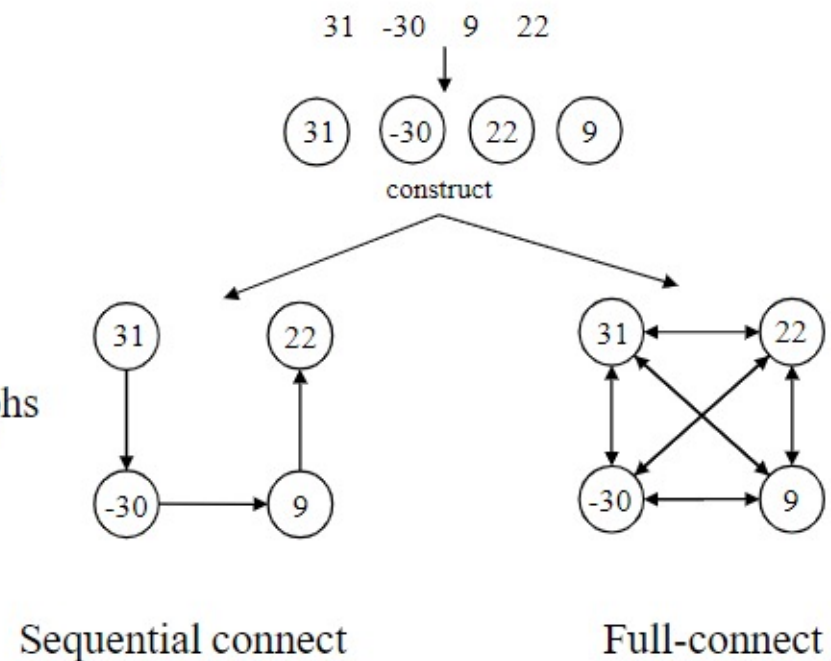


Fig. 3. Example of two kinds of graphs.

Experimental settings

TABLE I
DATASETS DESCRIPTION

Source	Dataset	Length	Memory (GB)	Function	Sequential (%)
MSRC	hm_1	1.08×10^6	6.36	Hardware monitoring	39.9
	mds_0	4.23×10^5	8.48	Media server	65.2
	proj_0	1.17×10^6	4.056	Project directories	57.3
	prxy_0	4.03×10^5	5.18	Firewall/web proxy	37.6
	src1_2	1.15×10^6	2.0	Source control	58.5
HW	hw_1	1.39×10^6	930.29	hybrid storage system	55.8
	hw_2	2.58×10^5	600.46	hybrid storage system	95.1
	hw_3	1.73×10^5	902.22	hybrid storage system	43.7

$$HR = \frac{\text{Cache Hits}}{\text{Cache Hits} + \text{Cache Misses}} \times 100\%.$$

$$EPR = \frac{\text{Correct Prefetchings}}{\text{All Prefetchings}} \times 100\%.$$

Performance & result

TABLE II

SINGLE-STEP RESULTS. THE RESULTS ARE IN PERCENTAGE, THE BEST RESULTS ARE IN **BOLD**, THE SECOND ONES ARE UNDERLINED, N IS CACHE SIZE.

Dataset		hw_1						hw_2						hw_3						hm_1					
Metric		HR@N			EPR@N			HR@N			EPR@N			HR@N			EPR@N			HR@N			EPR@N		
Method		10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000
No_pre		0.0	0.3	54.2	0.0	0.0	0.0	1.0	1.1	1.1	0.0	0.0	0.0	0.0	0.1	1.3	0.0	0.0	0.0	2.7	25.3	98.3	0.0	0.0	0.0
Naïve		57.5	58.0	63.2	63.3	64.5	64.5	92.5	92.6	92.7	93.3	93.7	94.0	47.7	47.9	48.8	48.0	48.3	48.7	31.7	43.8	97.4	30.5	31.2	5.6
Stride		43.7	44.0	65.8	80.5	81.1	80.6	91.0	91.1	91.1	99.1	99.2	99.2	38.4	38.6	39.6	81.6	82.0	82.3	27.1	47.0	99.1	<u>82.3</u>	<u>84.4</u>	88.4
ARIMA		1.9	4.0	8.8	1.9	4.3	6.2	82.8	82.9	83.0	85.9	86.2	86.4	0.3	0.3	1.3	0.2	0.3	0.3	3.5	19.0	95.2	<u>2.7</u>	<u>5.2</u>	2.5
Informer		0.2	0.9	5.8	0.3	0.9	2.9	1.0	1.1	1.1	0.0	0.0	0.0	0.0	0.0	0.9	0.0	0.0	0.0	1.1	14.0	90.4	0.1	0.7	0.7
DeepPrefetcher		74.3	74.6	79.2	75.4	75.9	76.5	92.2	92.5	92.8	93.4	94.0	94.5	50.4	50.7	51.7	50.4	50.7	51.2	38.5	59.1	99.3	38.5	56.0	46.1
Delta-LSTM		74.4	74.8	79.3	75.5	76.0	76.6	92.5	92.8	93.1	93.7	94.2	94.7	56.4	56.8	57.9	66.2	66.8	67.2	30.0	50.6	99.3	57.7	72.8	87.6
SGDP		79.2	79.5	85.8	82.9	83.5	81.6	93.0	93.0	93.1	97.5	97.7	97.8	76.0	76.6	77.5	88.9	89.5	90.1	38.1	55.7	99.4	87.8	90.1	86.2
SGDP _l		78.5	78.8	84.9	82.1	82.7	80.6	92.9	93.1	<u>93.2</u>	97.0	97.2	97.4	78.5	79.0	79.8	83.6	84.2	84.7	43.1	61.4	99.1	46.3	60.8	24.4
SGDP _p		75.7	78.2	83.6	<u>77.6</u>	80.4	79.6	93.7	94.0	94.2	94.4	95.0	95.4	48.1	48.3	49.6	72.1	73.1	75.1	43.9	62.9	<u>99.4</u>	46.8	63.8	34.8
Dataset		mds_0						proj_0						prxy_0						srcl_2					
Metric		HR@N			EPR@N			HR@N			EPR@N			HR@N			EPR@N			HR@N			EPR@N		
Method		10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000	10	100	1000
No_pre		13.2	35.0	61.0	0.0	0.0	0.0	6.1	28.7	35.2	0.0	0.0	0.0	20.1	40.7	48.8	0.0	0.0	0.0	3.9	34.8	48.2	0.0	0.0	0.0
Naïve		54.3	68.2	85.2	47.8	51.1	52.2	61.1	70.1	74.3	58.7	59.7	60.8	46.4	64.3	72.7	35.1	38.5	40.9	60.5	73.0	80.8	59.9	63.1	66.3
Stride		47.3	62.2	79.8	82.3	90.6	89.8	51.0	61.1	65.4	82.5	88.1	88.3	40.3	56.5	63.8	69.6	<u>81.1</u>	<u>81.4</u>	48.3	63.8	73.4	81.0	89.6	92.0
ARIMA		16.6	37.4	58.3	8.6	9.2	12.0	12.9	33.5	39.3	12.0	10.1	10.4	19.9	42.2	52.3	6.5	7.2	8.2	14.6	42.0	54.8	19.5	17.7	19.2
Informer		9.6	28.3	54.5	0.3	1.2	5.2	3.9	19.8	34.7	0.1	0.5	2.3	13.7	32.1	46.9	0.0	0.0	0.2	1.7	22.5	45.3	0.0	0.1	0.6
DeepPrefetcher		60.7	73.7	88.5	66.9	77.5	83.3	72.6	79.1	82.8	75.0	78.6	81.5	57.0	70.2	77.4	63.5	70.4	73.9	74.5	82.9	89.0	76.2	80.9	87.0
Delta-LSTM		57.3	69.6	86.2	80.2	<u>87.8</u>	<u>89.8</u>	62.3	69.1	73.3	84.3	86.2	87.4	52.2	64.2	71.3	<u>75.7</u>	79.3	80.9	70.0	79.6	86.2	77.5	81.3	87.2
SGDP		66.0	76.3	91.6	80.2	87.0	88.4	73.4	78.5	82.1	84.0	87.6	88.2	62.2	73.2	79.9	76.3	83.3	84.1	<u>75.4</u>	83.1	88.8	82.5	88.5	90.8
SGDP _l		<u>66.1</u>	<u>77.5</u>	<u>92.1</u>	65.4	73.9	79.6	75.5	81.1	84.6	79.8	83.6	85.6	64.1	76.5	83.0	65.7	74.3	78.4	76.3	83.9	89.4	81.5	87.4	89.3
SGDP _p		67.4	79.8	92.6	68.7	81.9	87.9	<u>73.7</u>	81.3	85.2	74.5	80.2	83.9	63.9	<u>76.2</u>	83.0	64.9	72.6	76.1	74.9	84.8	89.8	76.1	84.0	87.3

TABLE III

AVERAGE RESULTS OF MULTI-STEP PREFETCHING. THE RESULTS ARE IN PERCENTAGE, THE BEST RESULTS ARE IN **BOLD**, AND THE CACHE SIZE IS 100.

Metric		HR@100										EPR@100									
Step		1	2	3	4	5	6	7	8	9	10	1	2	3	4	5	6	7	8	9	10
Method																					
DeepPrefetcher		74.0	76.1	76.7	77.1	77.2	77.1	77.1	77.0	77.0	76.9	73.1	60.7	52.5	46.5	42.0	38.3	35.3	32.8	30.7	28.8
Delta-LSTM		70.2	74.9	77.1	78.0	78.5	78.9	79.3	79.6	79.8	80.0	81.5	72.3	65.9	60.6	56.3	52.6	49.5	46.8	44.5	42.3
SGDP		77.0	78.3	78.9	79.3	79.6	79.8	80.0	80.1	80.3	80.3	88.4	80.5	74.0	68.9	64.6	60.9	57.7	54.9	52.7	50.5
SGDP _l		78.9	80.7	81.4	81.8	82.2	82.4	82.5	82.6	82.7	82.7	80.5	70.2	62.9	57.5	53.1	49.6	46.6	44.0	41.8	39.8
SGDP _p		75.7	77.4	78.1	78.6	78.9	79.1	79.2	79.3	79.5	79.5	78.9	68.3	60.8	55.1	50.6	46.8	43.7	41.1	38.8	36.7

Performance & result

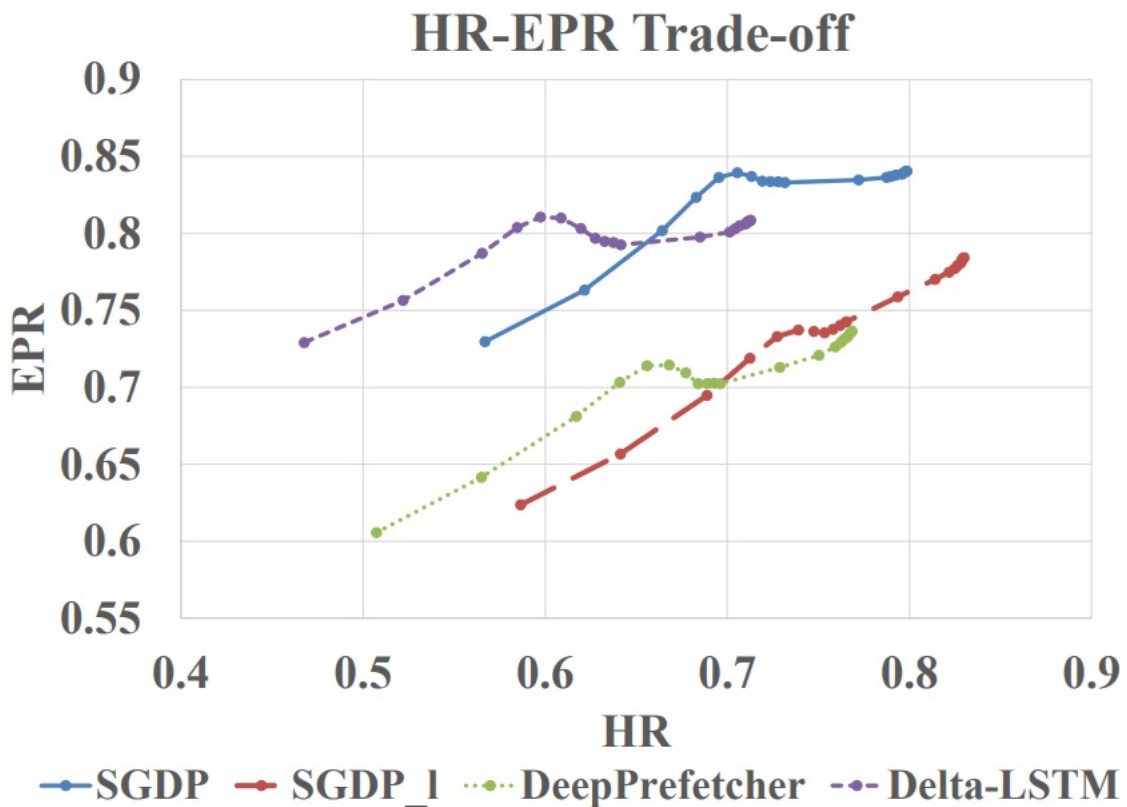



Fig. 4. HR-EPR trade-off.

TABLE IV
 THE NUMBER OF PREDICTIONS INFERRED PER SECOND BY
 LEARNING-BASED METHODS.

Method \ Dataset	hw_1	hw_2	hw_3	hm_1	mds_0	proj_0	prxy_0	src1_2	avg
Delta-LSTM	89.4	87.4	94.5	92.4	90.7	91.5	88.4	95.1	91.2
DeepPrefetcher	208.2	154.5	194.2	160.1	248.4	178.4	187.9	249.6	197.7
SGDP	644.5	692.4	666.1	515.2	543.5	553.3	470.0	550.7	579.5
SGDP _l	634.7	686.9	614.7	500.1	651.4	663.9	526.3	670.7	618.6
SGDP _p	599.5	645.6	593.9	567.0	491.7	529.3	574.8	558.7	570.1



Many Thanks

Yiyuan Yang*

University of Oxford
Huawei Noah's Ark Lab
Oxford, UK
yiyuan.yang@cs.ox.ac.uk

Rongshang Li*

University of Sydney
Huawei Noah's Ark Lab
Sydney, Australia
roli5128@uni.sydney.edu.au

Qiquan Shi

Huawei Noah's Ark Lab
Huawei
Shenzhen, China
shiqiquan@huawei.com

Xijun Li

MIRA Lab, USTC
Huawei Noah's Ark Lab
Shenzhen, China
xijun.li@huawei.com

Gang Hu

Huawei
Chengdu, China
hugang27@huawei.com

Xing Li

Huawei Noah's Ark Lab
Hongkong, China
li.xing2@huawei.com

Mingxuan Yuan

Huawei Noah's Ark Lab
Hongkong, China
Yuan.Mingxuan@huawei.com