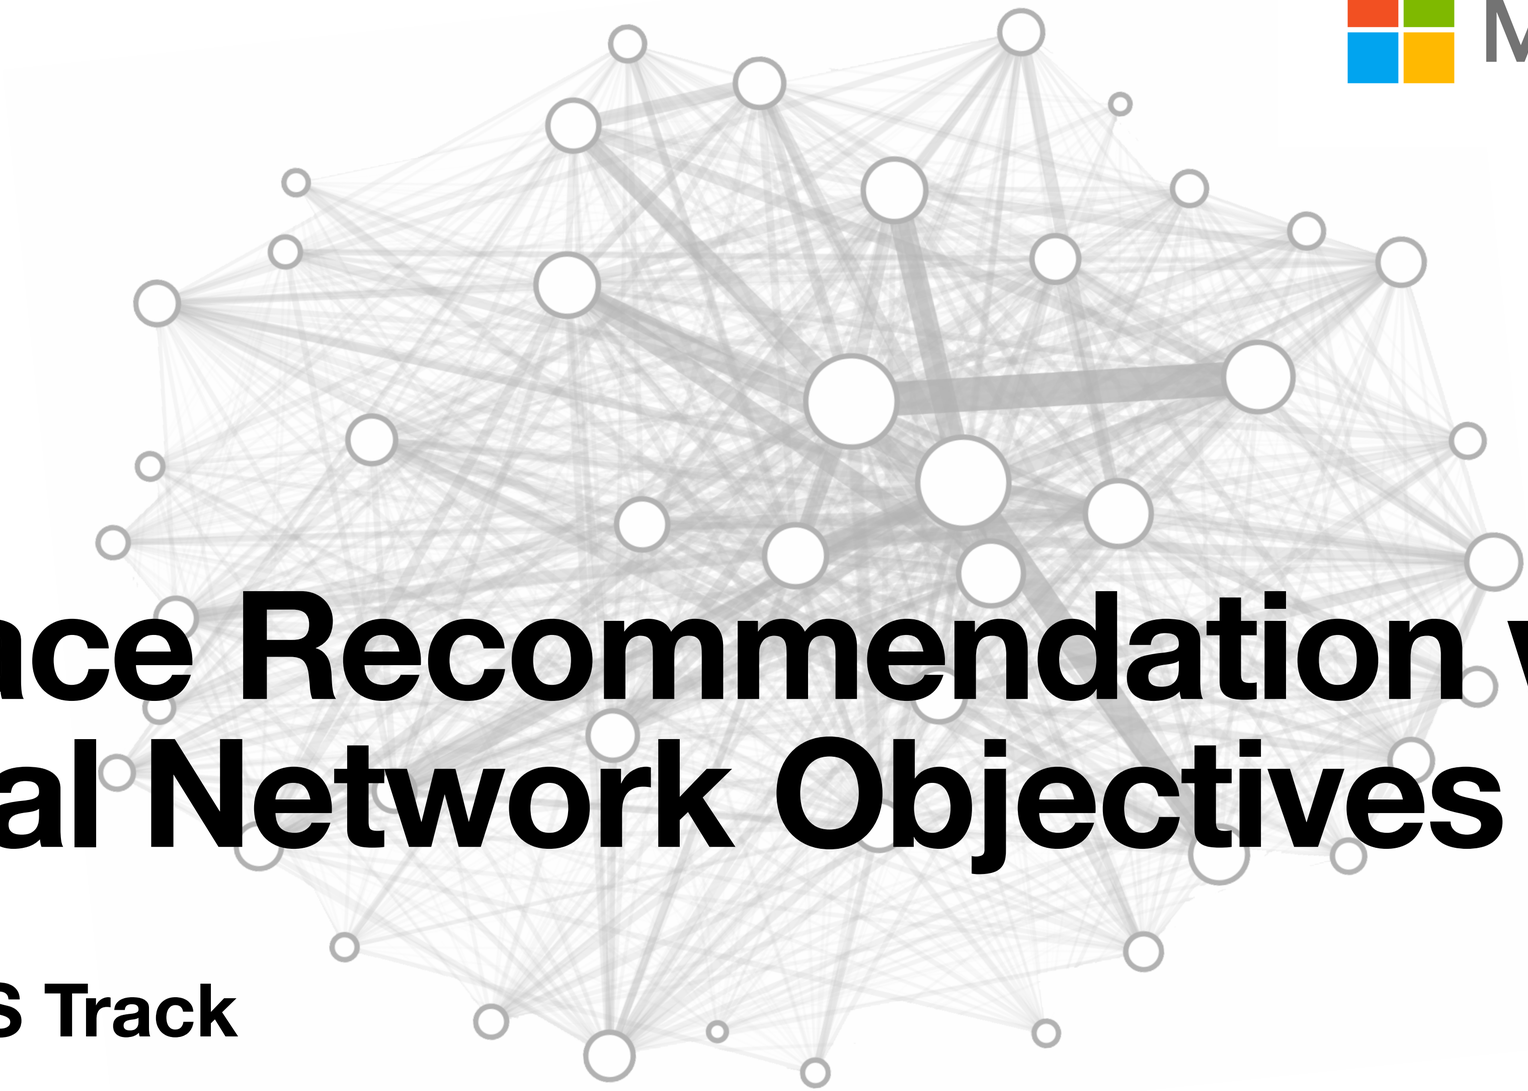

Workplace Recommendation with Temporal Network Objectives

KDD 2023, ADS Track

Kiran Tomlinson,[1] Jennifer Neville,[2] Longqi Yang,[2] Mengting Wan,[2] Cao Lu[2]
[1]Cornell University, [2]Microsoft

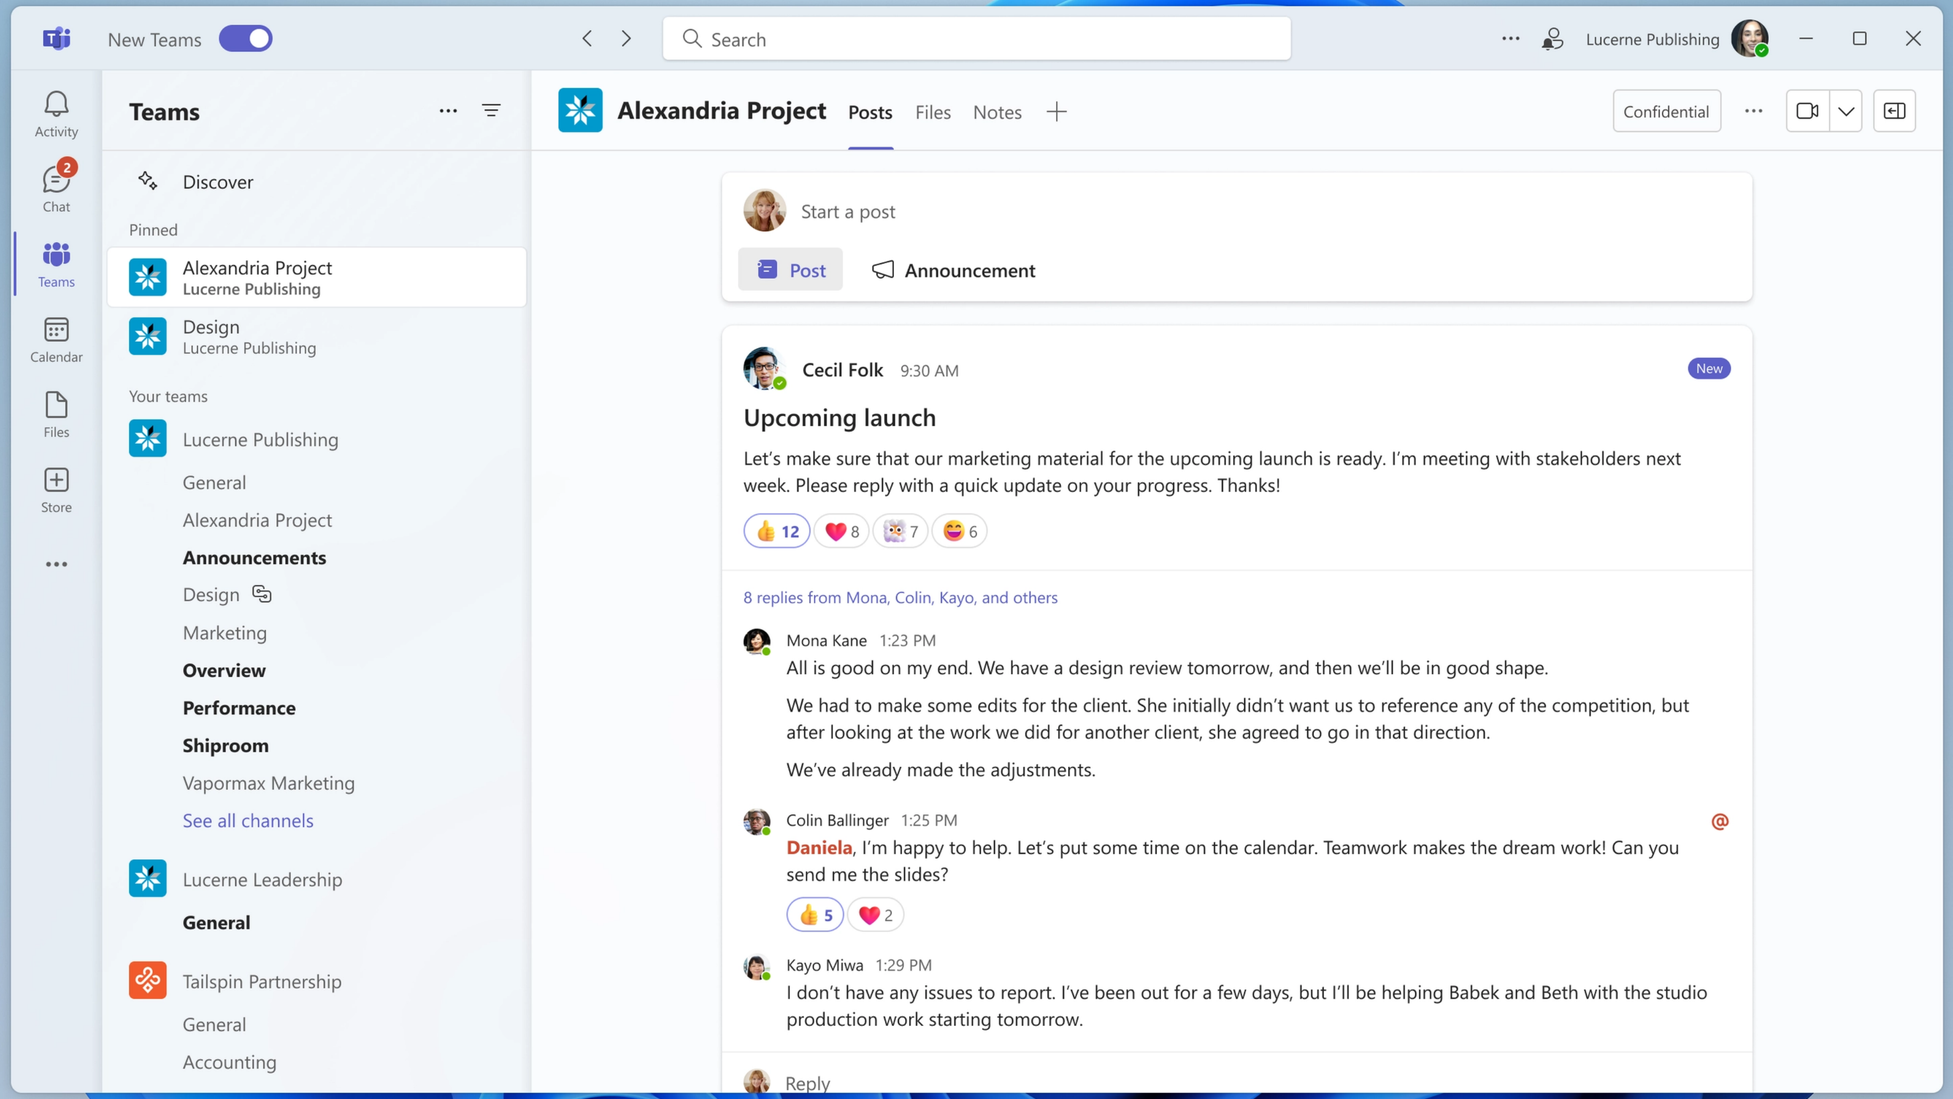
New Teams
Search
Lucerne Publishing
Activity
Chat
Teams
Calendar
Files
Store
Teams
Discover
Pinned
Alexandria Project
Lucerne Publishing
Design
Lucerne Publishing
Your teams
Lucerne Publishing
General
Alexandria Project
Announcements
Design
Marketing
Overview
Performance
Shiproom
Vapormax Marketing
See all channels
Lucerne Leadership
General
Tailspin Partnership
General
Accounting
Alexandria Project
Posts
Files
Notes
Confidential
Start a post
Post
Announcement
Cecil Folk 9:30 AM
New
Upcoming launch
Let's make sure that our marketing material for the upcoming launch is ready. I'm meeting with stakeholders next week. Please reply with a quick update on your progress. Thanks!
12
8
7
6
8 replies from Mona, Colin, Kayo, and others
Mona Kane 1:23 PM
All is good on my end. We have a design review tomorrow, and then we'll be in good shape.
We had to make some edits for the client. She initially didn't want us to reference any of the competition, but after looking at the work we did for another client, she agreed to go in that direction.
We've already made the adjustments.
Colin Ballinger 1:25 PM
Daniela, I'm happy to help. Let's put some time on the calendar. Teamwork makes the dream work! Can you send me the slides?
5
2
Kayo Miwa 1:29 PM
I don't have any issues to report. I've been out for a few days, but I'll be helping Babek and Beth with the studio production work starting tomorrow.
Reply
78°F
Sunny
9:28 AM
6/30/2023

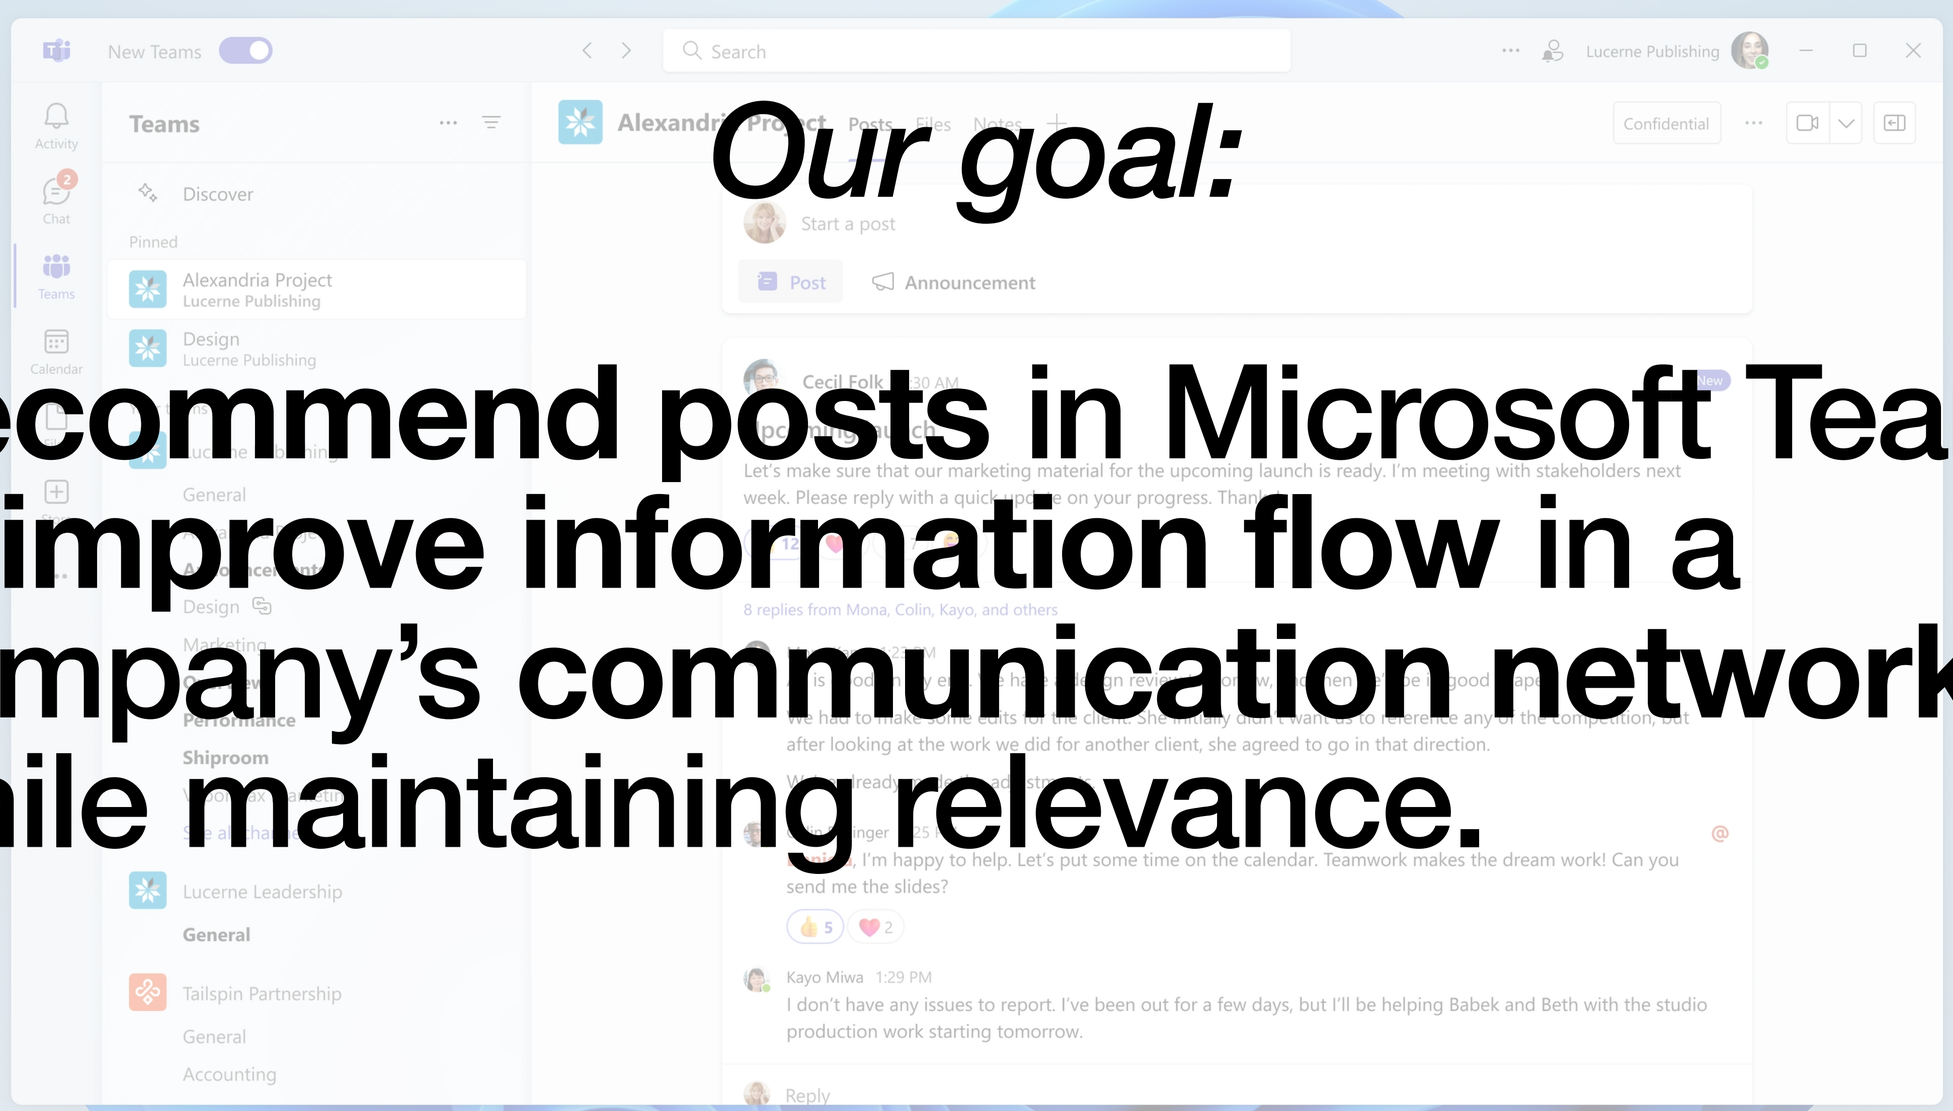

Our goal:

Recommend posts in Microsoft Teams to **improve information flow** in a company's **communication network**, while maintaining relevance.

78°F
Sunny
9:28 AM
6/30/2023

Problem setting: Microsoft Teams

Organization

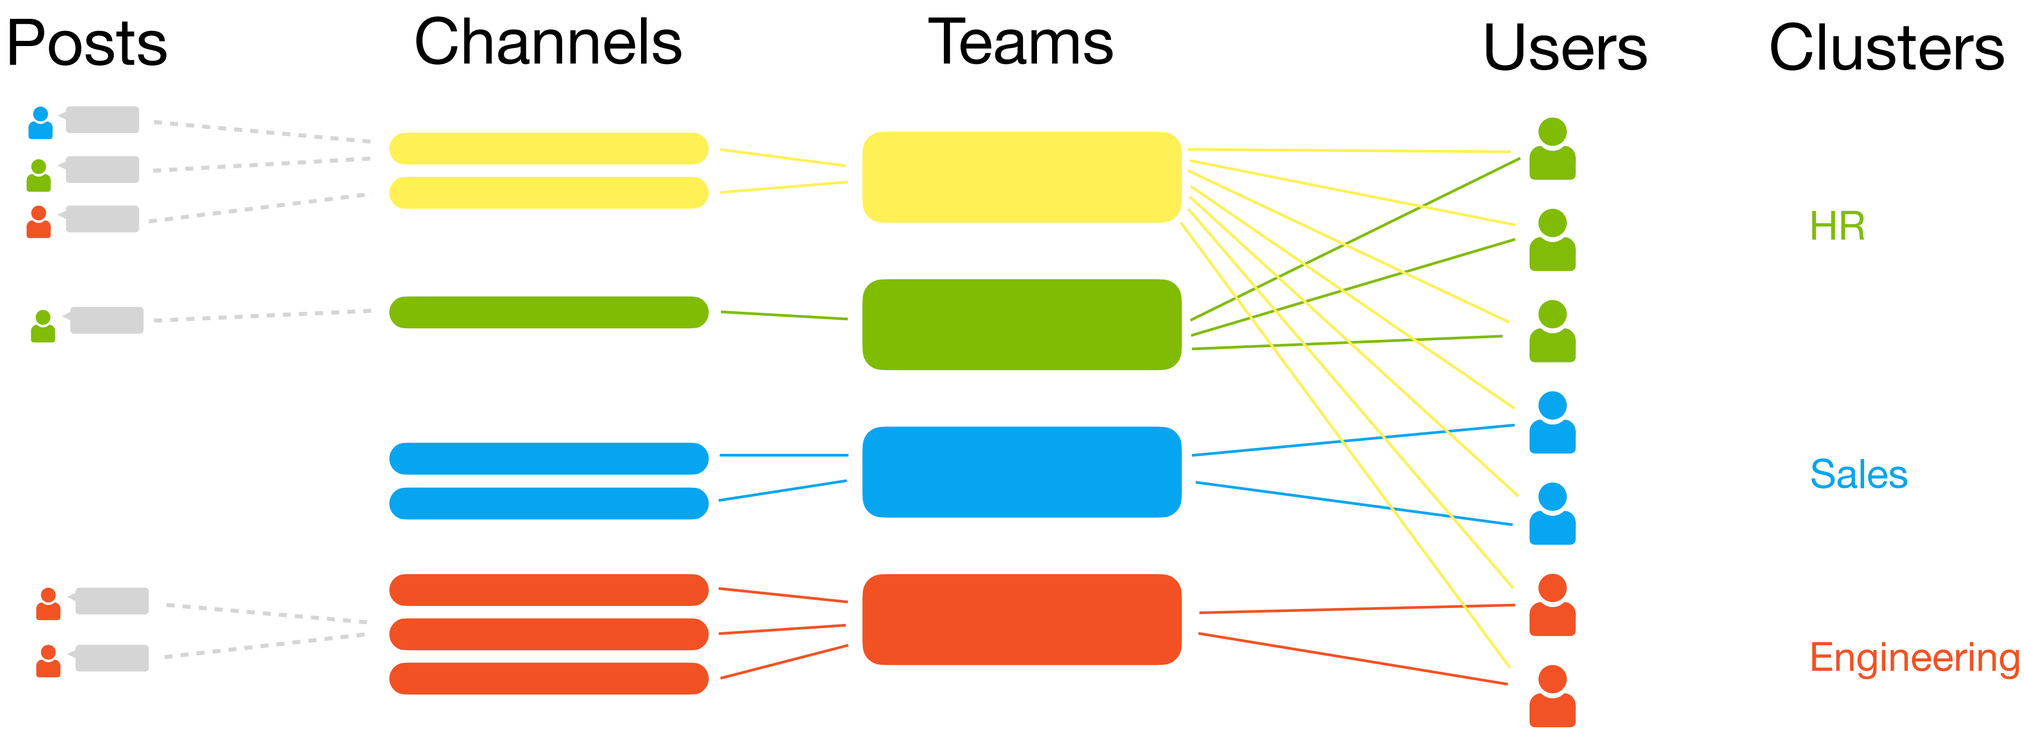

Modeling information spread in an organization

Temporal communication network

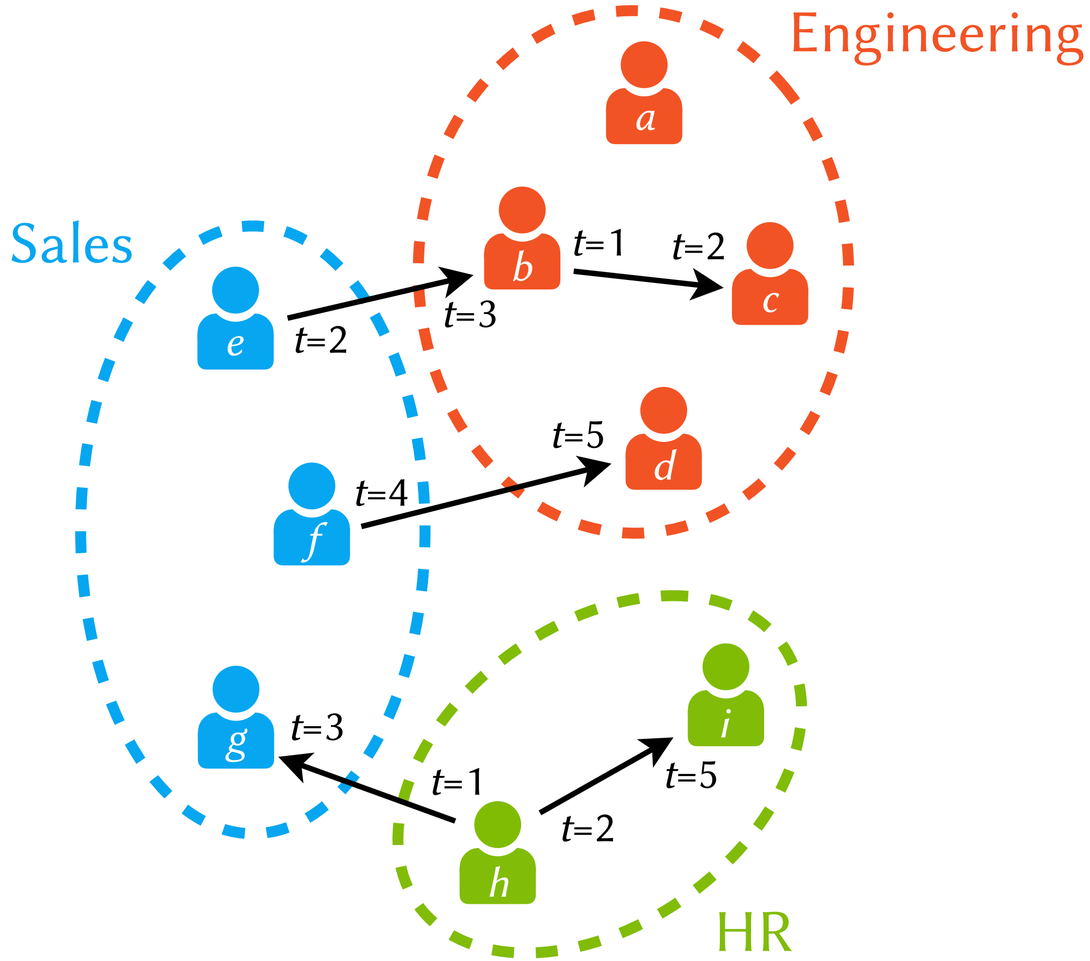

Interacting with posts influences the communication network

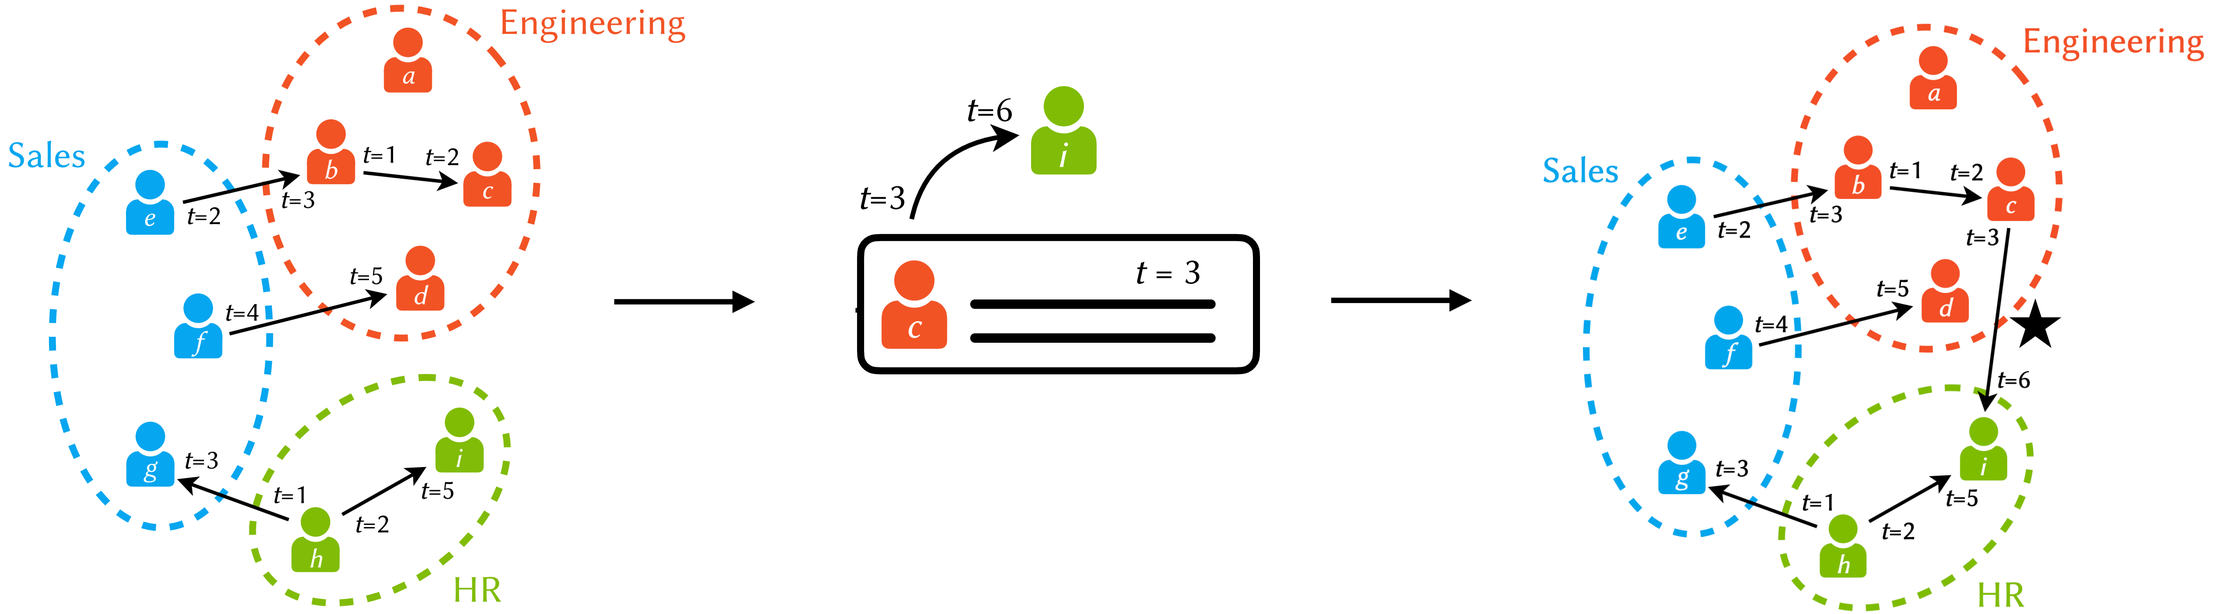

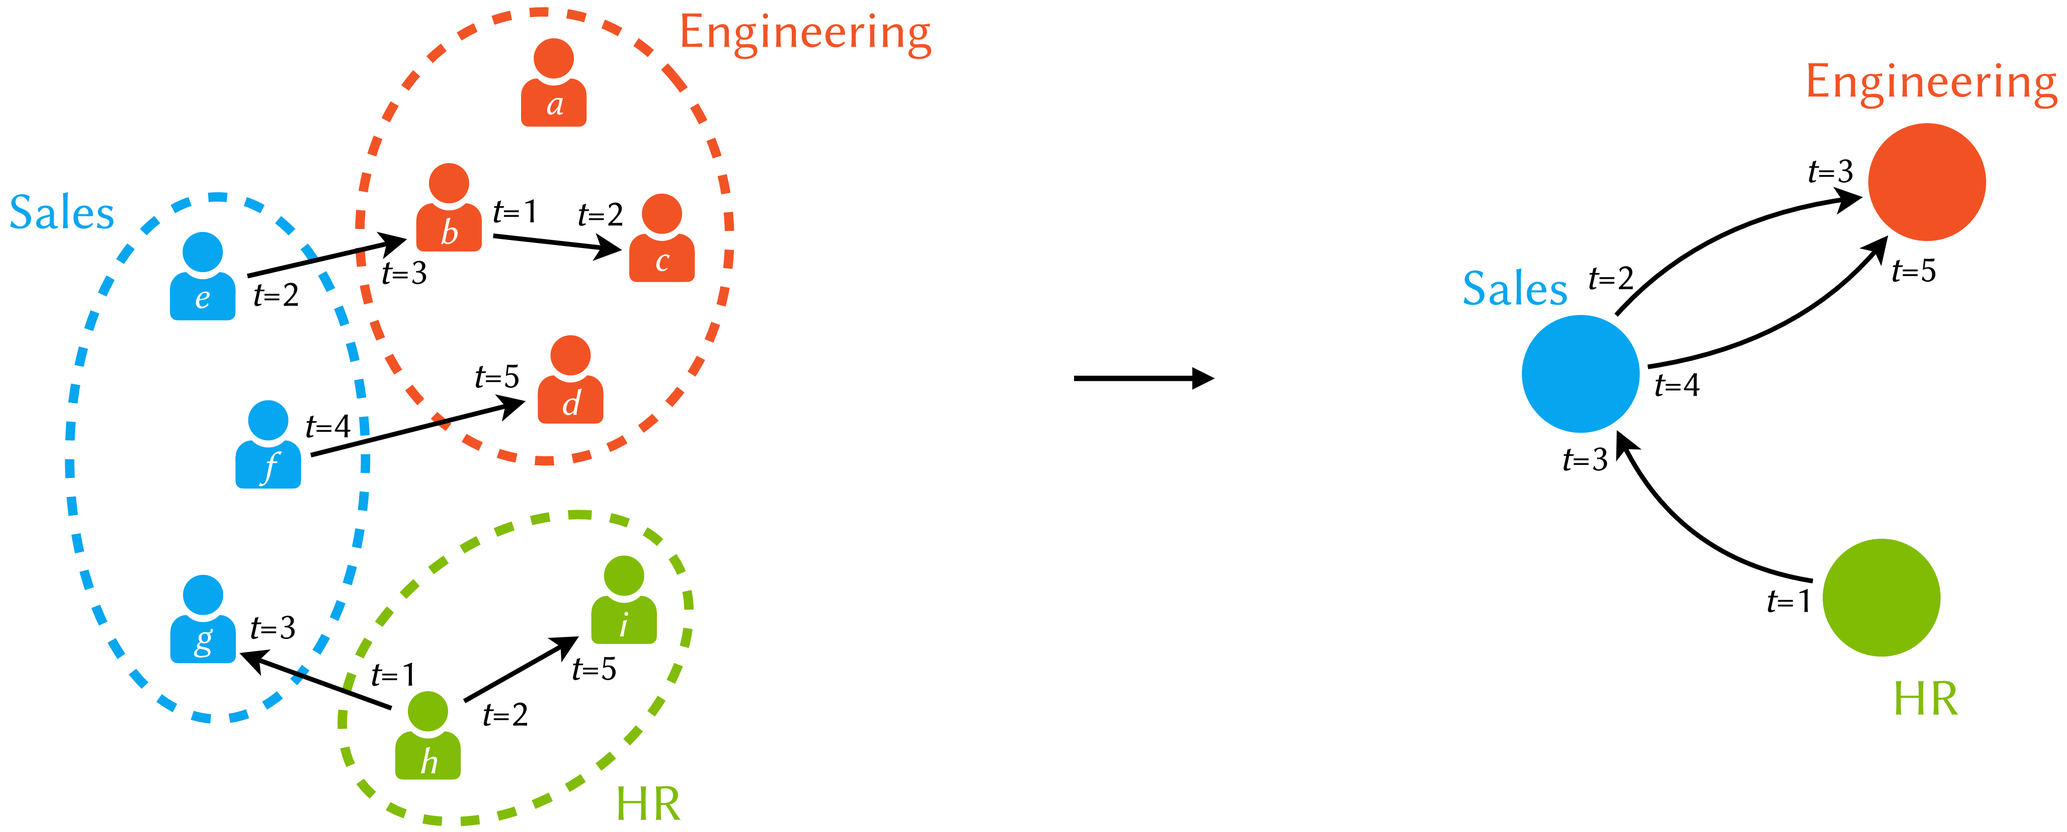

Motivation: remote work reduced cross-team communication

(Yang et al, *Nature Human Behavior* 2022)

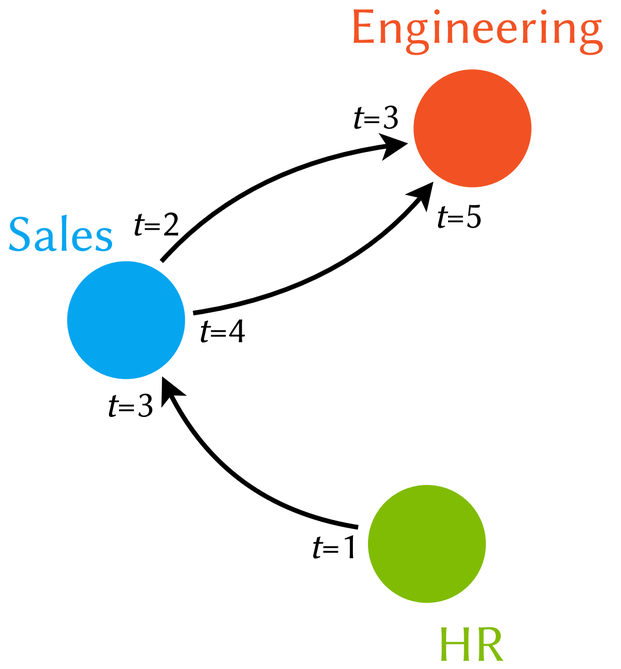

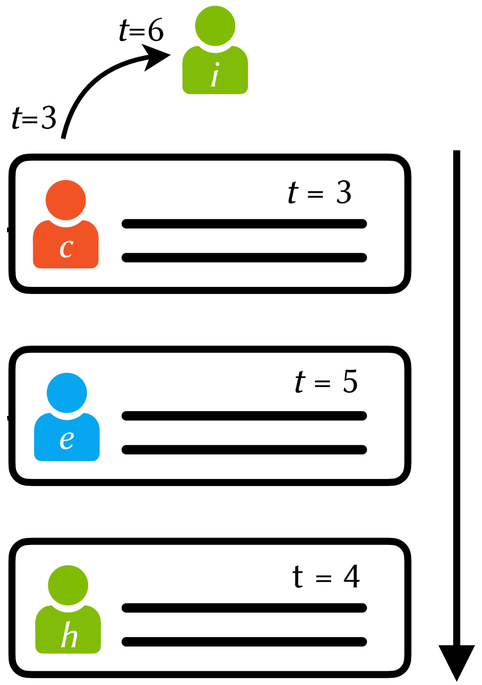

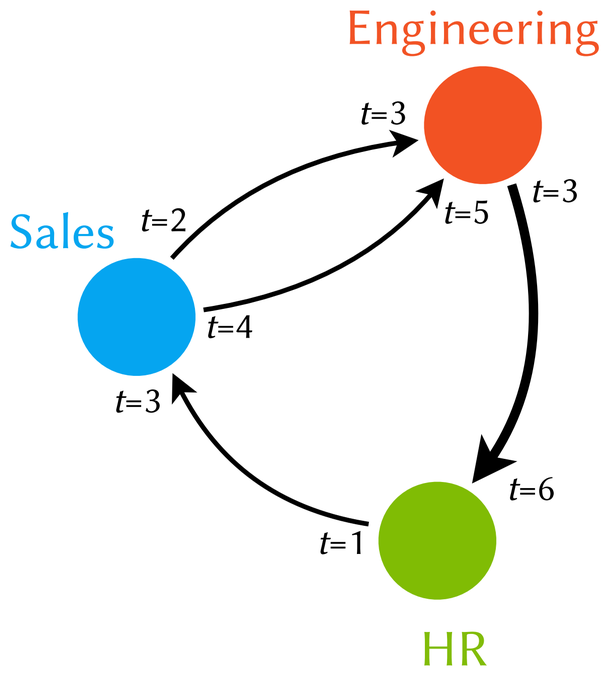

Can we improve cross-cluster information flow by recommending posts?

How do we measure information flow?

Information latency

(Kossinets, Kleinberg, Watts, *KDD '08*)

How recently could *v* have heard about *u*?

$$IL(G,t) = \sum_{u,v \in V} t - \max_{P \in \mathscr{P}_t(u,v)} d(P)$$

departure time of path *P*

temporal paths from *u* to *v* arriving no later than *t*

Total information

(Tomlinson et al., *KDD '23*)

What fraction of *u*'s state is known by *v*?

$$TI(G,t) = \sum_{u,v \in V} TI(u,v,t)$$

$$TI(u,u,t) = 1$$

$$TI(u,v,t) = \min\{1,\ \lambda TI(u,v,t-1) + \sum_{(z,v,d,a,w) \in E_t} w\lambda^{t-d} TI(u,z,t-1,d)\}$$

decayed prior knowledge

information transmitted along edge

edges into *v* at *t* that departed at *d* with weight *w*

Theory: optimizing temporal network objectives

Edge addition problems:
Add best *k* edges into/out of *u* to maximize IL/TI now/in the future

	In-edges	**Out-edges**
Myopic	NP-hard*	Easy! Greedy is optimal
Non-Myopic	NP-hard*	NP-hard*

* Greedy is $(1 - 1/e)$ approximation! (Objectives are submodular)

Temporal Information and Engagement Recommender (TIER)

Recommending to *u* at time *t*, candidate posts C

Network score (TI):

$$n_p = TI(G + (\text{author}(p), u, \text{time}(p), t, w), t) - TI(G, t)$$

how much would interacting with a post improve cross-cluster information? (Greedy, myopic)

→ Track TI/IL matrix over time (efficient algorithm)

Relevance score:

r_p from any traditional recommender

$$\text{Rank posts by: } \frac{n_p}{\max_{q \in C} n_q} + \alpha \frac{r_p}{\max_{q \in C} r_q}$$

Data: Microsoft internal communication

- ~ 180,000 users (full-time Microsoft employees)
- 1 month of communication (March 2022)
- Microsoft Teams, Outlook, and SharePoint
  - Posts, chats, emails, file sharing
- ~ 100M edges (u, v, d, a)
- clusters from org chart (or clustering alg)
- offline evaluation: simulate user actions on recommendations

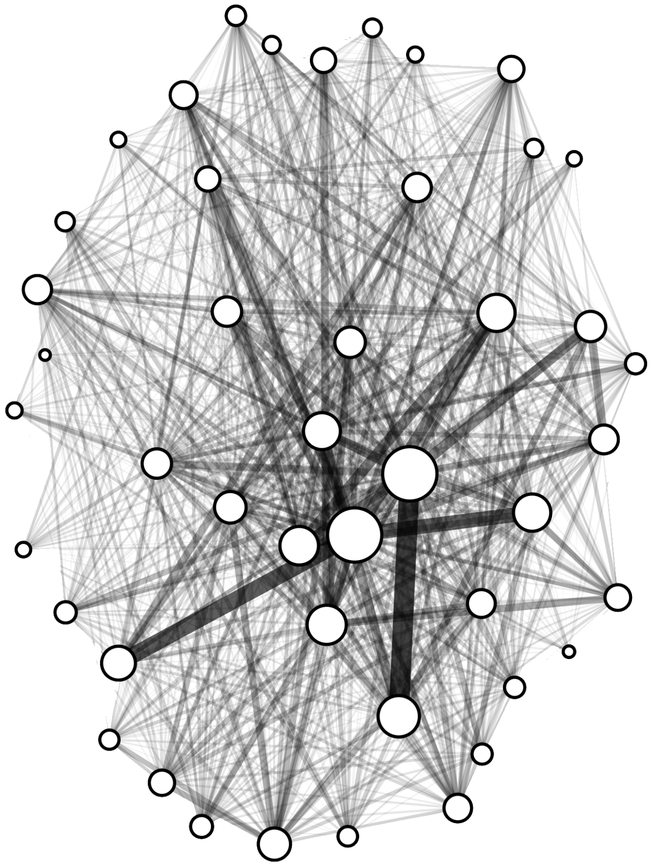

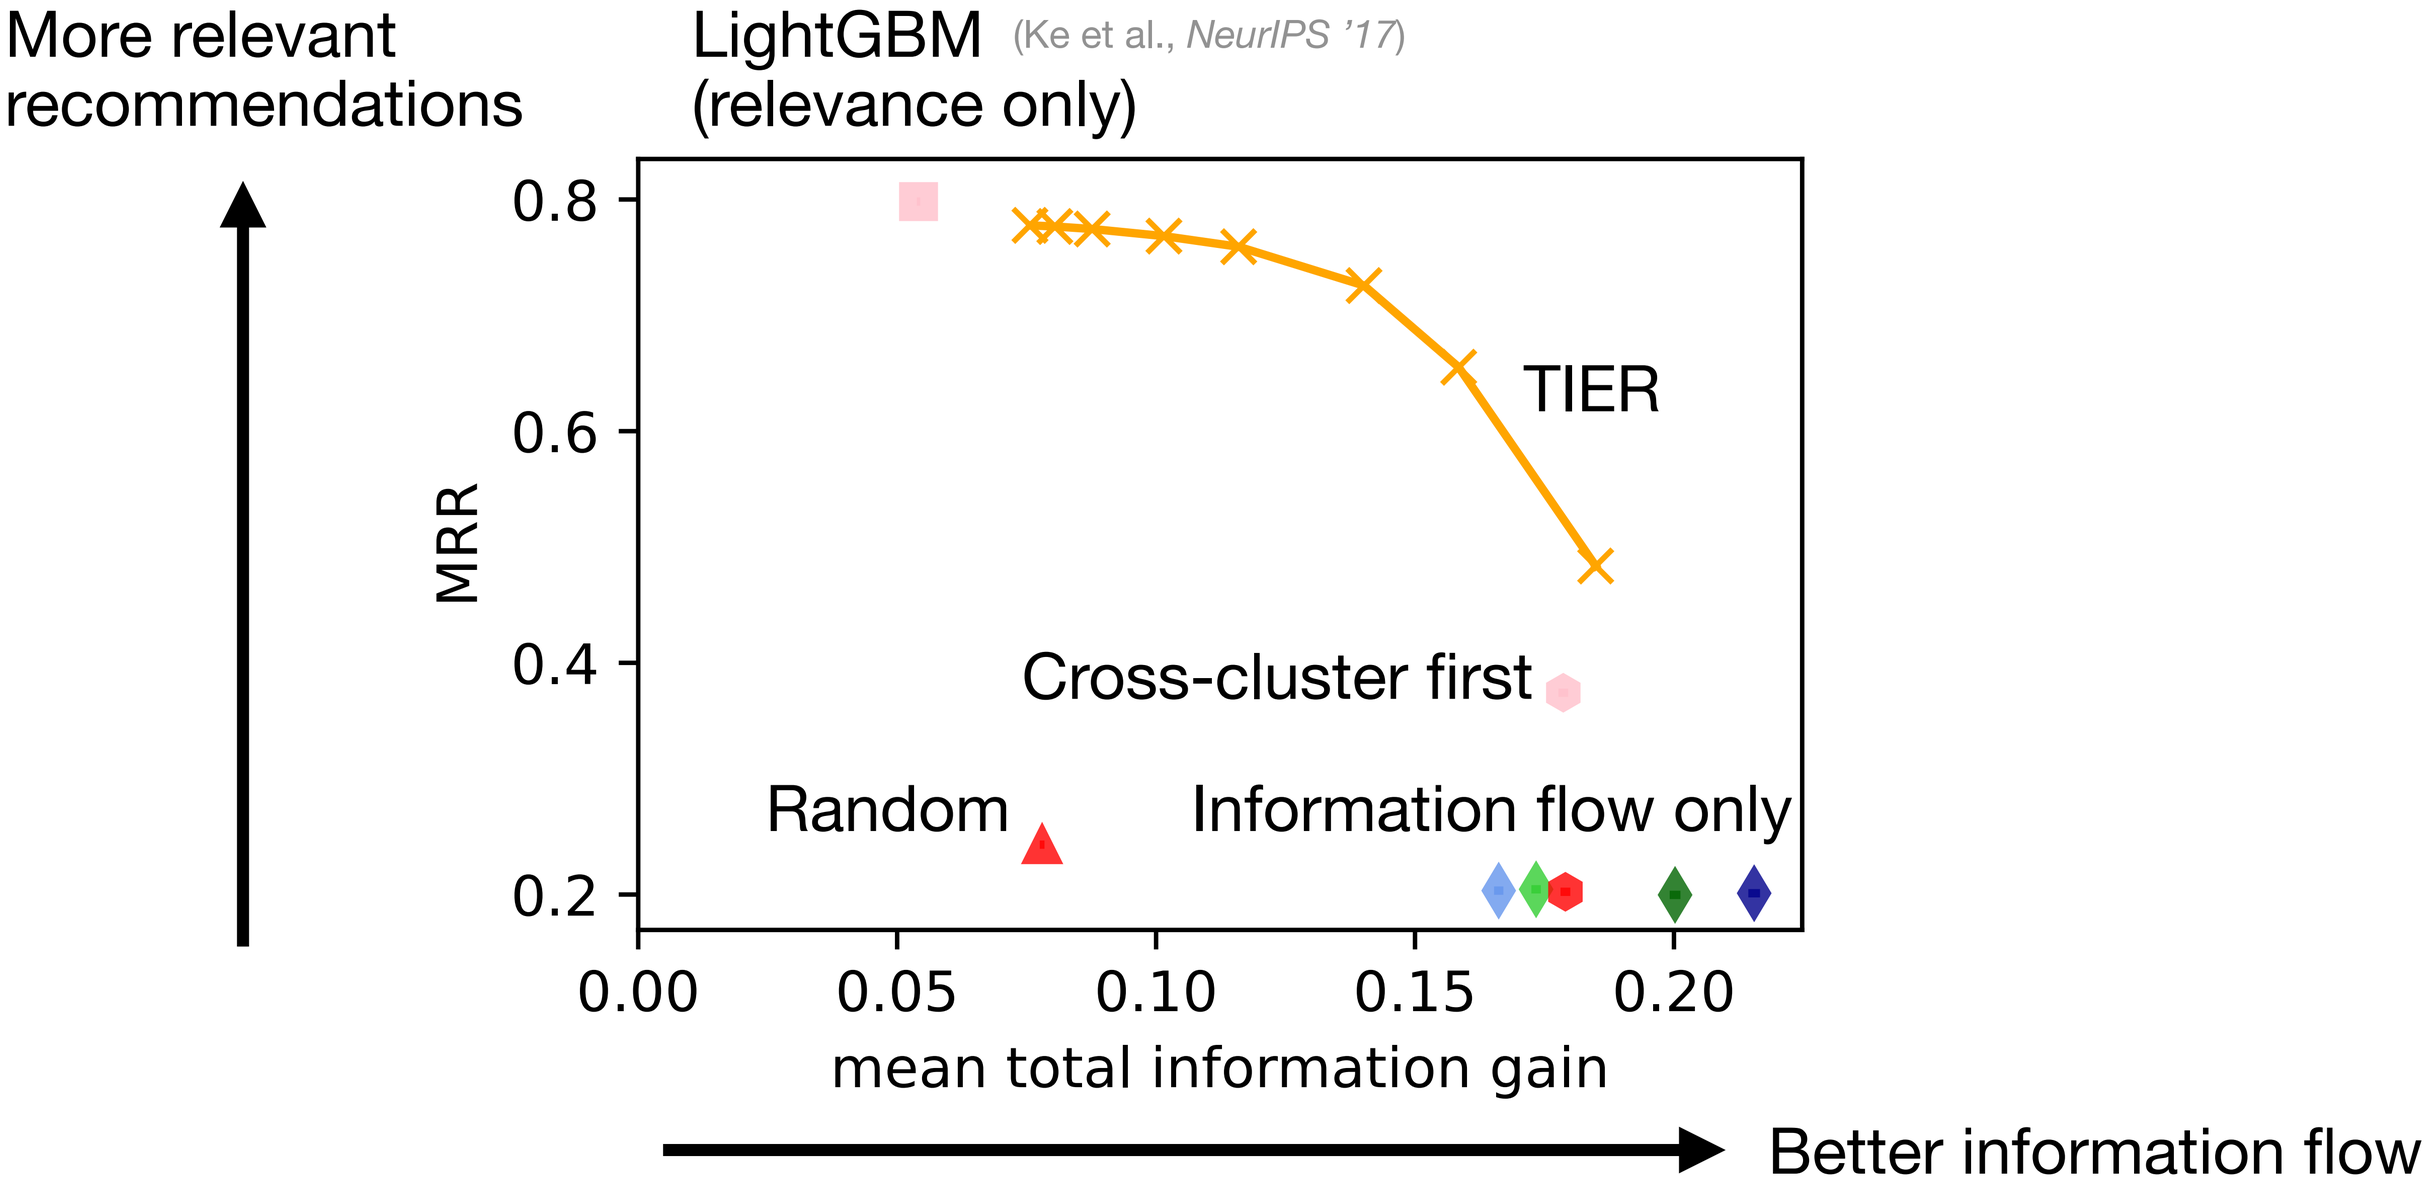
More relevant recommendations
LightGBM (Ke et al., NeurIPS '17) (relevance only)
0.8
0.6
0.4
0.2
MRR
TIER
Cross-cluster first
Random
Information flow only
0.00
0.05
0.10
0.15
0.20
mean total information gain
Better information flow

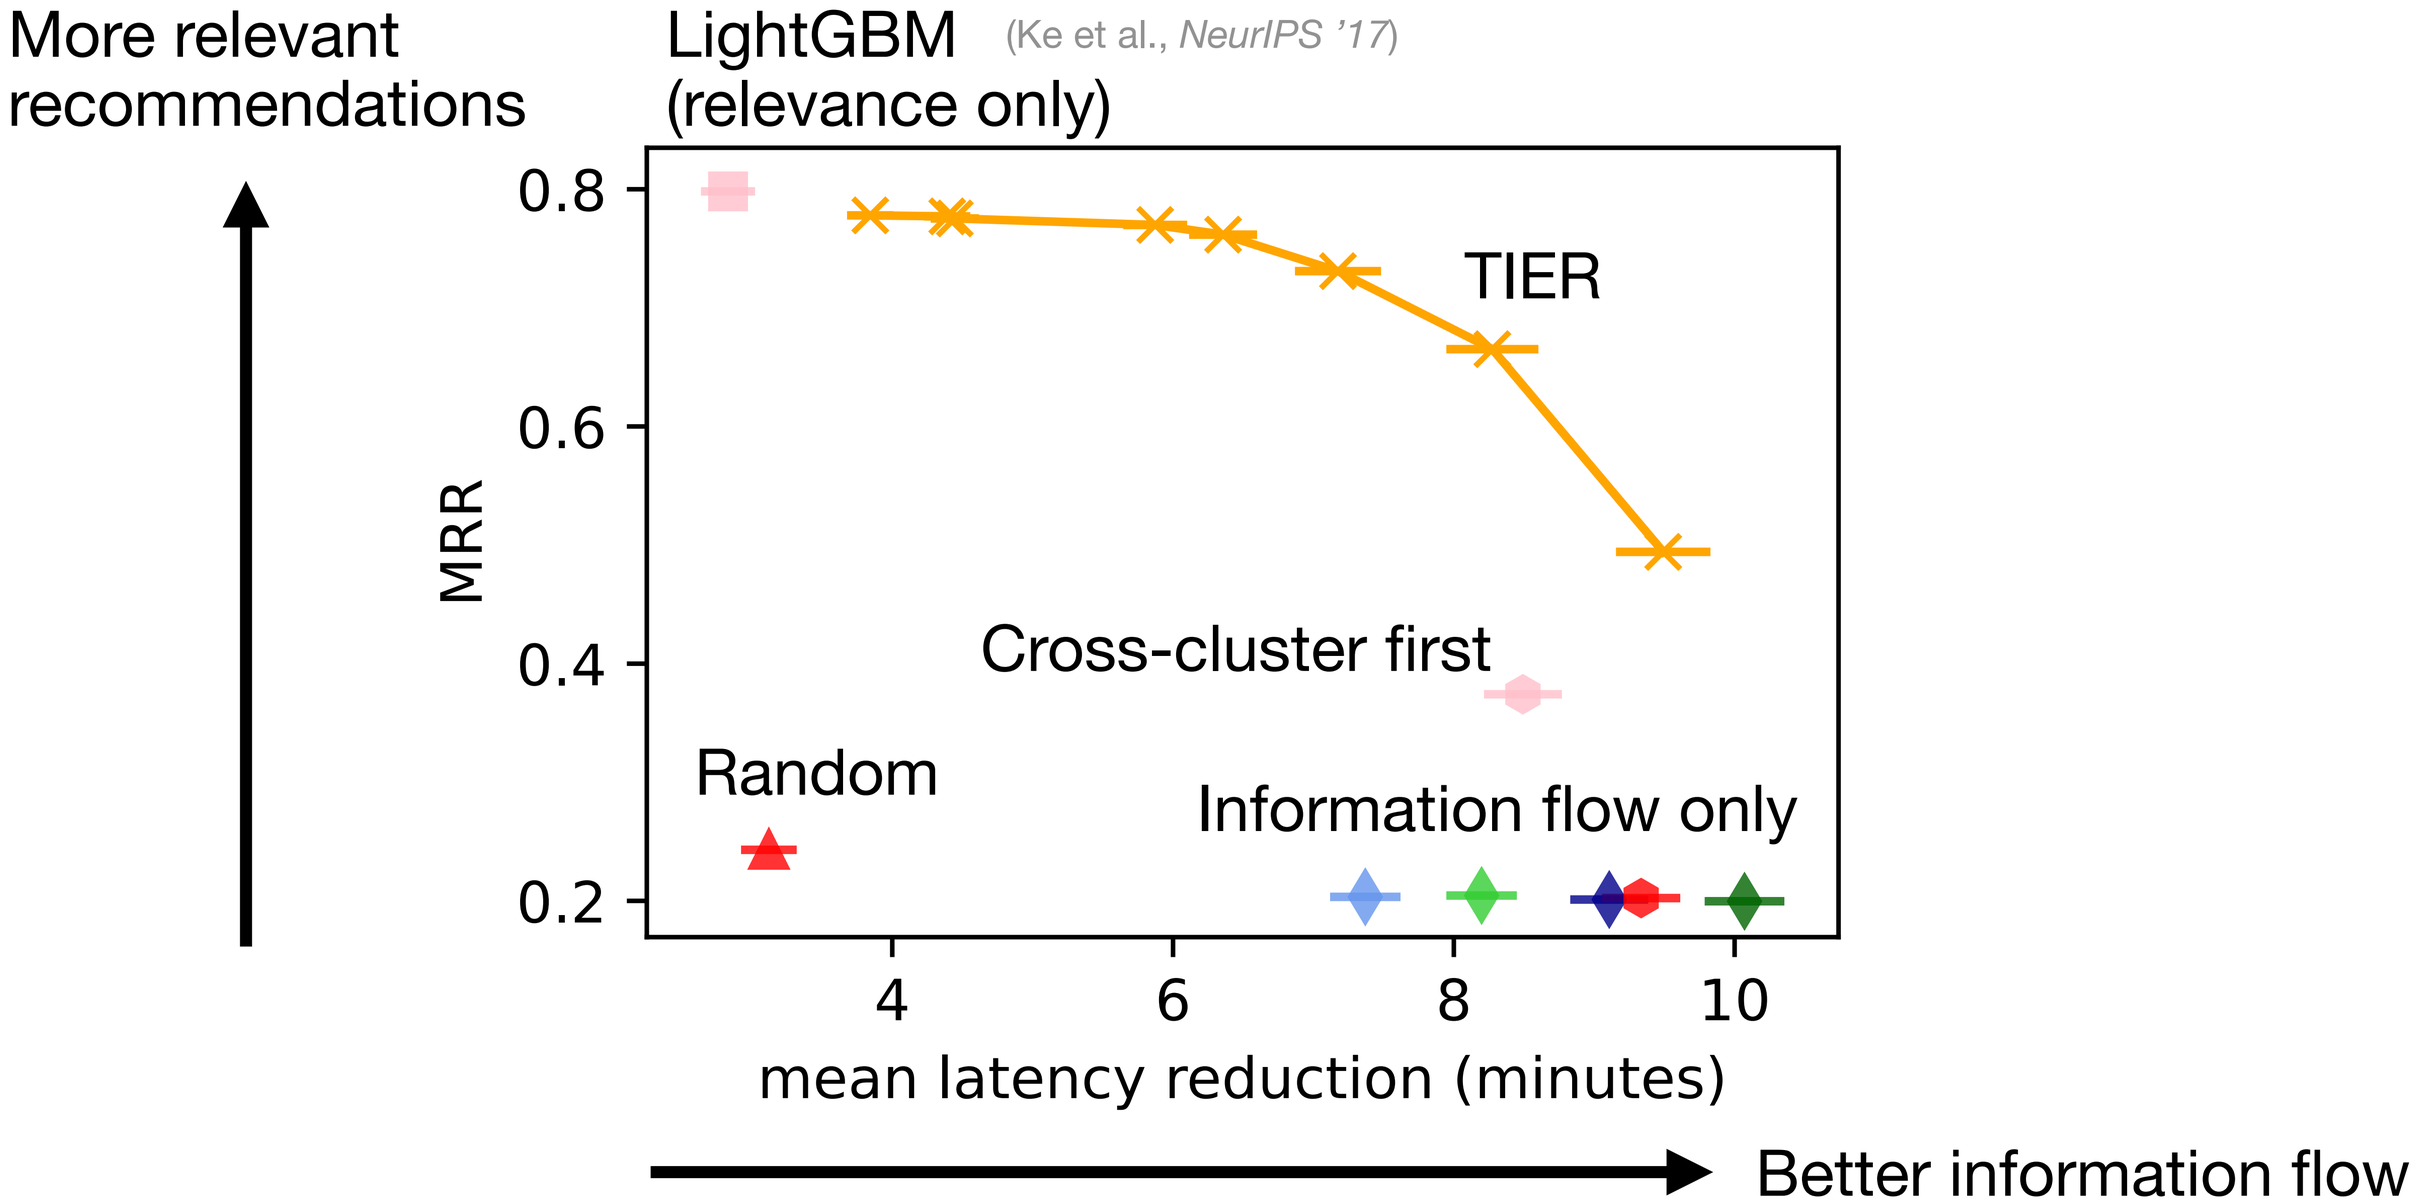
More relevant recommendations
LightGBM (relevance only)
(Ke et al., NeurIPS '17)
0.8
0.6
0.4
0.2
MRR
TIER
Cross-cluster first
Random
Information flow only
4
6
8
10
mean latency reduction (minutes)
Better information flow

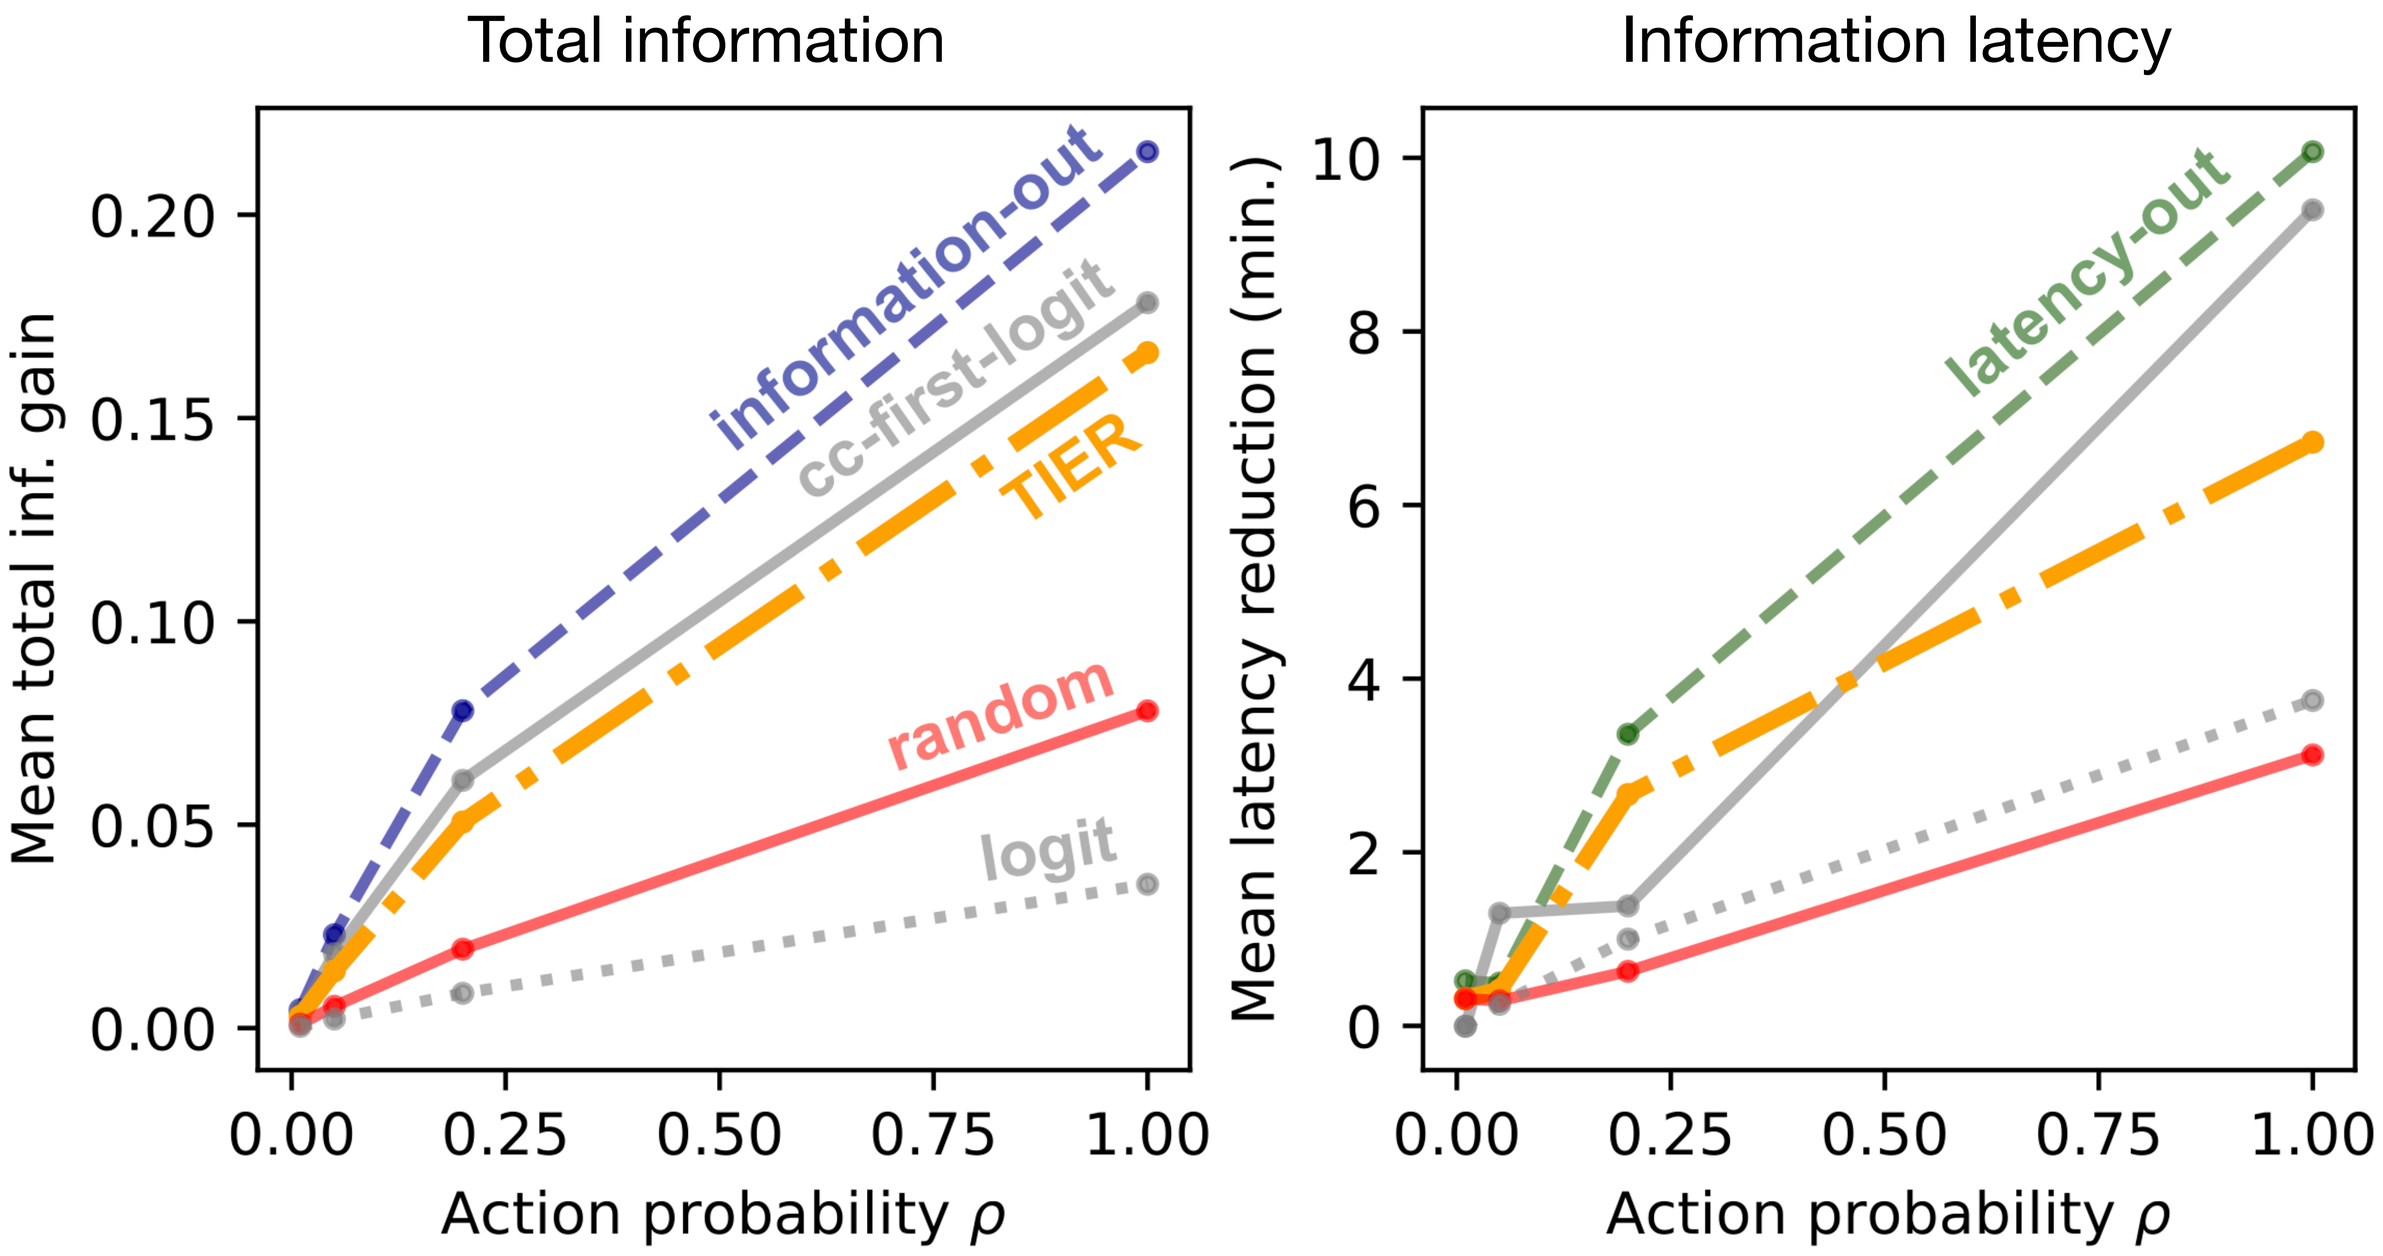
Total information
Mean total inf. gain
0.00
0.05
0.10
0.15
0.20
0.00
0.25
0.50
0.75
1.00
Action probability ρ
information-out
cc-first-logit
TIER
random
logit
Information latency
Mean latency reduction (min.)
0
2
4
6
8
10
0.00
0.25
0.50
0.75
1.00
Action probability ρ
latency-out

Workplace Recommendation with Temporal Network Objectives

Kiran Tomlinson
Cornell University
Ithaca, NY, USA
kt@cs.cornell.edu

Jennifer Neville
Microsoft
Redmond, WA, USA
jenneville@microsoft.com

Longqi Yang
Microsoft
Redmond, WA, USA
longqi.yang@microsoft.com

Mengting Wan
Microsoft
Redmond, WA, USA
mengting.wan@microsoft.com

Cao Lu
Microsoft
Redmond, WA, USA
cao.lu@microsoft.com

Thank you!

Microsoft

CORNELL UNIVERSITY FOUNDED A.D. 1865

cs.cornell.edu/~kt/

@kiran_tomlinson

kt@cs.cornell.edu