

# Limits of Learning in Incomplete Networks

**Timothy LaRock**

[larock.t@husky.neu.edu](mailto:larock.t@husky.neu.edu)

In collaboration with



Timothy Sakharov



Sahely Bhadra



Tina Eliassi-Rad



**Northeastern University**  
*Network Science Institute*

Supported by NSF CNS-1314603 & NSF IIS-1741197

# Background: Incomplete Networks

- Network data is often incomplete
- Acquiring more data is often expensive and/or hard
- Research question: **Given a networked dataset and limited resources to collect more data, how can you get the most bang for your buck?**



# Two general approaches to network completion

Don't collect more data

Collect more data



# Two general approaches to network completion

## Don't collect more data

Assume a network model

Combine network model with incomplete data to get a model of the network structure

Infer missing data from this model

- [Kim et al. 2011]
- [Chen et al. 2018]

## Collect more data



# Two general approaches to network completion

## Don't collect more data

Assume a network model

Combine network model with incomplete data to get a model of the network structure

Infer missing data from this model

- [Kim et al. 2011]
- [Chen et al. 2018]

## Collect more data

Estimate Statistics from partially observed network

- [Soundarajan et al. 2015]
- [Soundarajan et al. 2016]

Utilize an explore-exploit approach

- [Pfeiffer III et al. 2014]
- [Soundarajan et al. 2017]
- [Murai et al. 2018]
- [Madhawa et al. 2018]
- This work!



## Our solution: Network Online Learning (NOL)

- Research Question: **Given a networked dataset and limited resources to collect more data, how can you get the most bang for your buck?**
- Learn to grow an incomplete network through sequential, optimal queries (to some API)
- Agnostic to both *data distributions* and *sampling method*
- Interpretable features that are computable online



# Assumptions

## We assume...

We know the API access model  
(complete vs incomplete queries).

The underlying network is static  
(probing the same node twice  
gives no new information).

## We do not assume...

A model of the underlying graph.

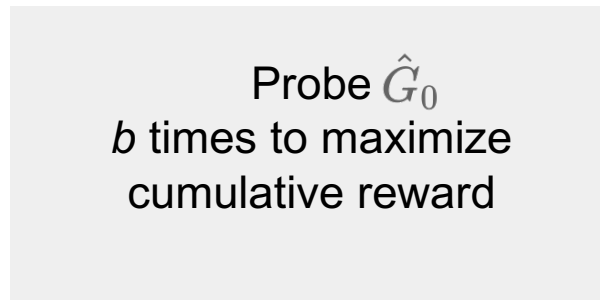
How the initial sample was  
collected.



# Network Online Learning (NOL)

## Inputs:

- $\hat{G}_0$ : Incomplete network
- $b$ : probing budget
- $r$ : reward function



## Output:

Network after  $b$  probes

$$\hat{G}_b \approx G$$

*Example reward function:*

- *number of new nodes observed*





# NOL algorithm

**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

- 1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$
- 2: **repeat**
- 3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}
- 4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^\top \phi_t(i)$  {Calculate estimated rewards}
- 5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}
- 6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^\top \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}
- 7: Probe node  $u_t$
- 8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$
- 9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$
- 10: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$
- 11: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$
- 12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$
- 13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$
- 14:  $t \leftarrow t + 1$
- 15: **until**  $t=b$
- 16: **return**  $\theta_b$  and  $\hat{G}_b$



# NOL algorithm

**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$

2: **repeat**

3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}

4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}

5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}

6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}

7: Probe node  $u_t$

8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$

9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$

10: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$

11: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$

12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$

13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$

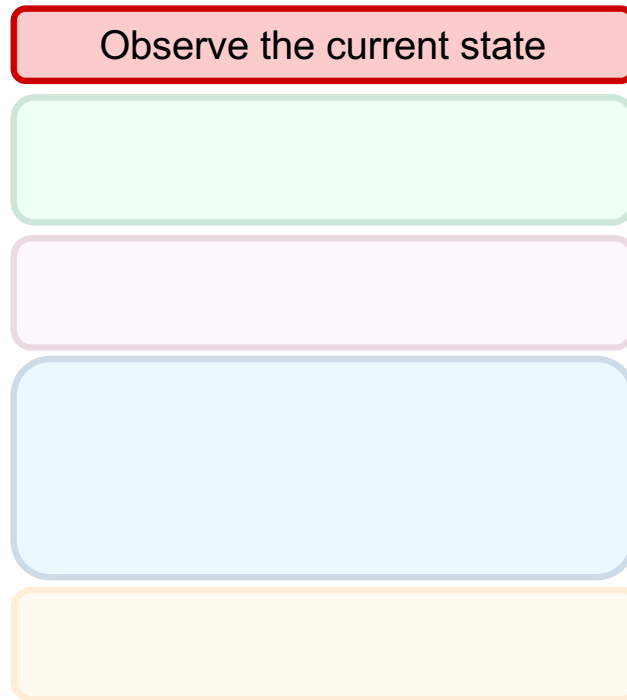
14:  $t \leftarrow t + 1$

15: **until**  $t=b$

16: **return**  $\theta_b$  and  $\hat{G}_b$



# NOL algorithm



**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$

2: **repeat**

3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}

4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}

5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}

6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}

7: Probe node  $u_t$

8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$

9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$

10: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$

11: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$

12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$

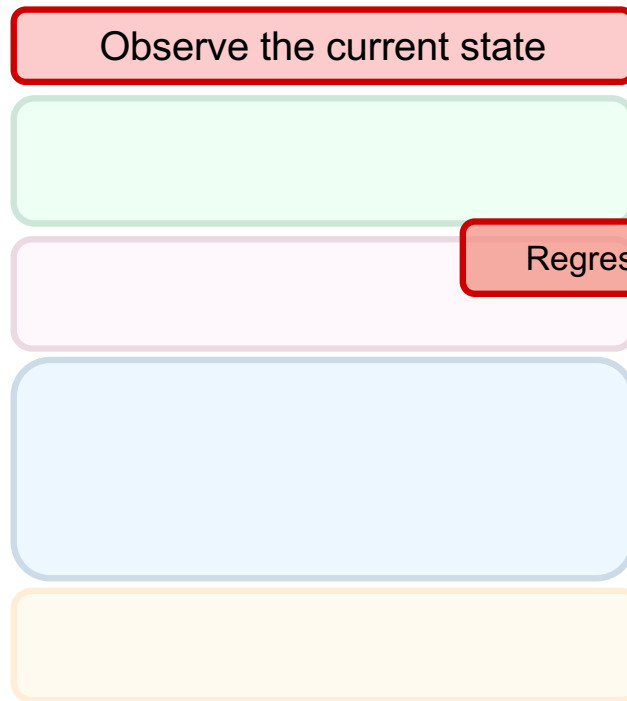
13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$

14:  $t \leftarrow t + 1$

15: **until**  $t=b$

16: **return**  $\theta_b$  and  $\hat{G}_b$

# NOL algorithm



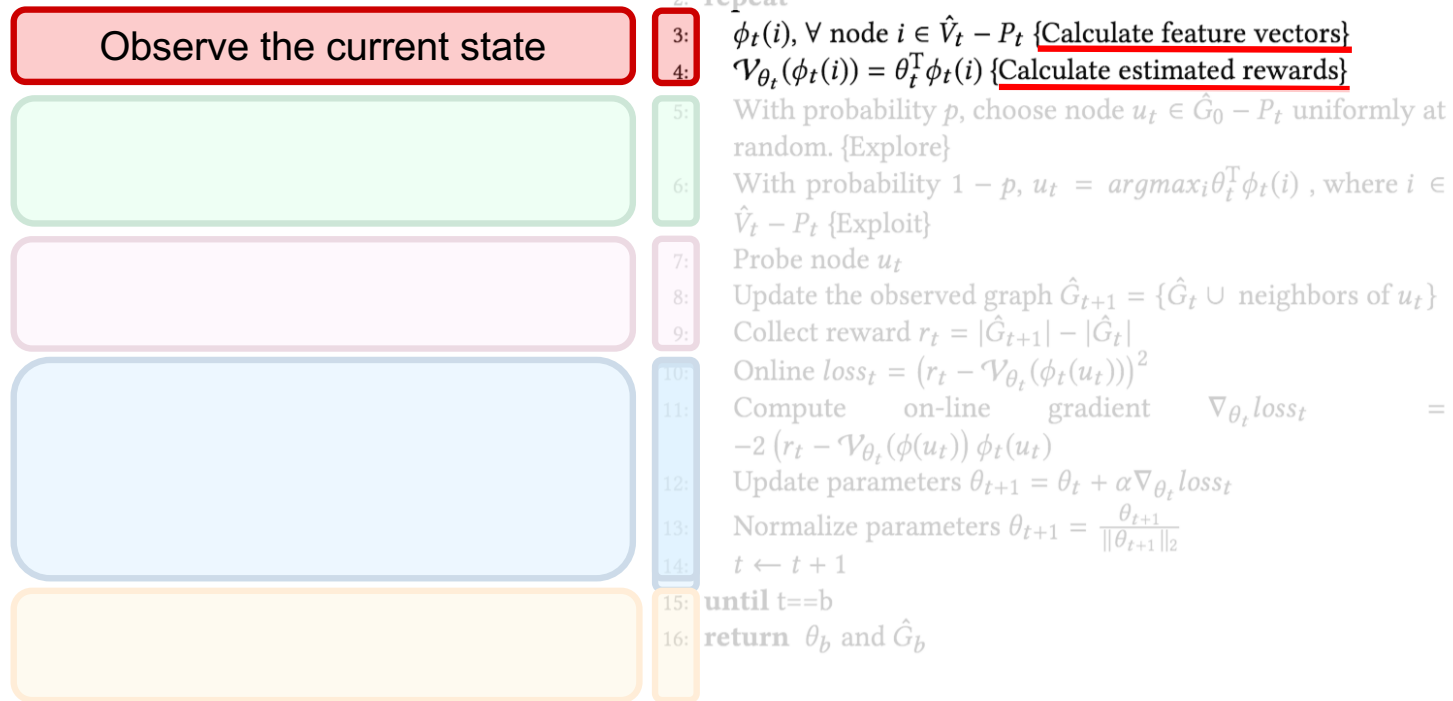
**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

- 1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$
- 2: **repeat**
- 3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}
- 4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}
- 5: With probability  $p$  choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}
- 6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}
- 7: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$
- 8: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$
- 9: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$
- 10: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$
- 11: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$
- 12: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$
- 13:  $t \leftarrow t + 1$
- 14: **until**  $t=b$
- 15: **return**  $\theta_b$  and  $\hat{G}_b$



# NOL algorithm



**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$

2: **repeat**

3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}

4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}

5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}

6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}

7: Probe node  $u_t$

8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$

9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$

10: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$

11: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$

12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$

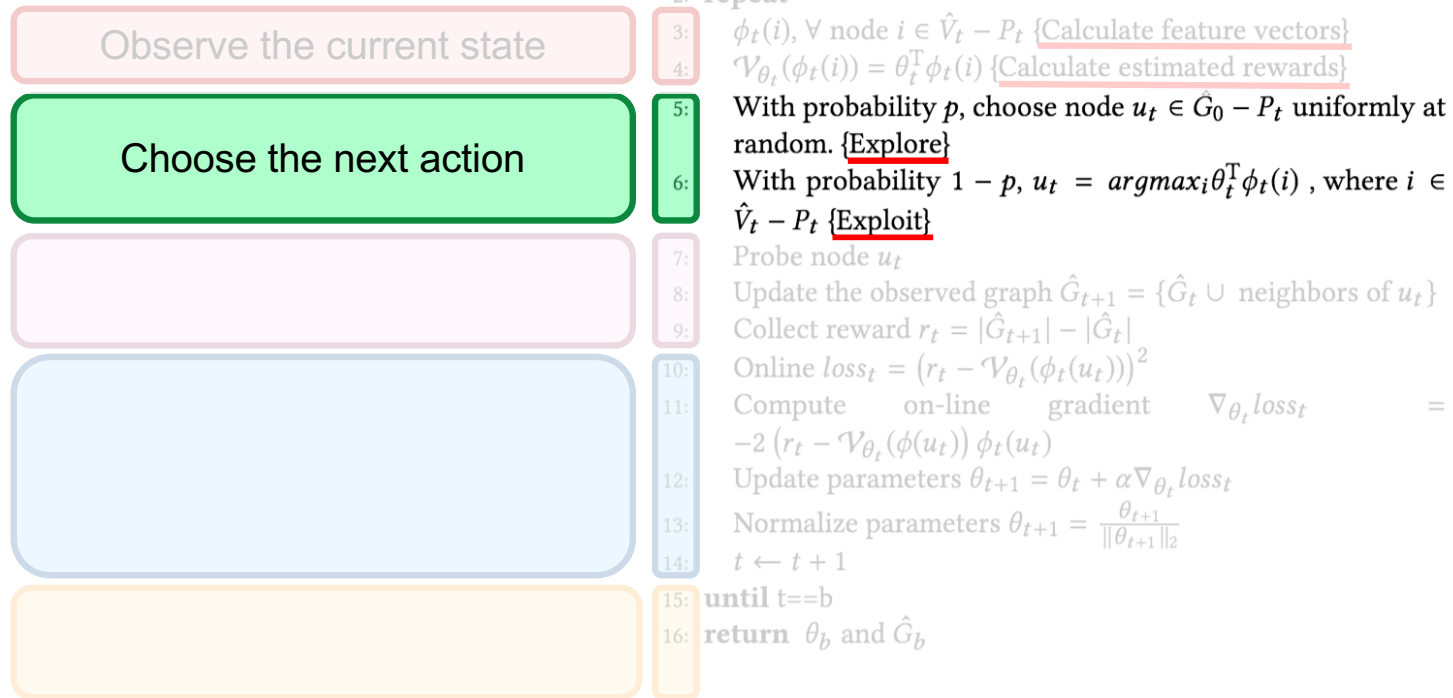
13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$

14:  $t \leftarrow t + 1$

15: **until**  $t=b$

16: **return**  $\theta_b$  and  $\hat{G}_b$

# NOL algorithm



**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$

2: repeat

3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}

4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}

5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}

6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}

7: Probe node  $u_t$

8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$

9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$

10: Online  $\text{loss}_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$

11: Compute on-line gradient  $\nabla_{\theta_t} \text{loss}_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$

12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} \text{loss}_t$

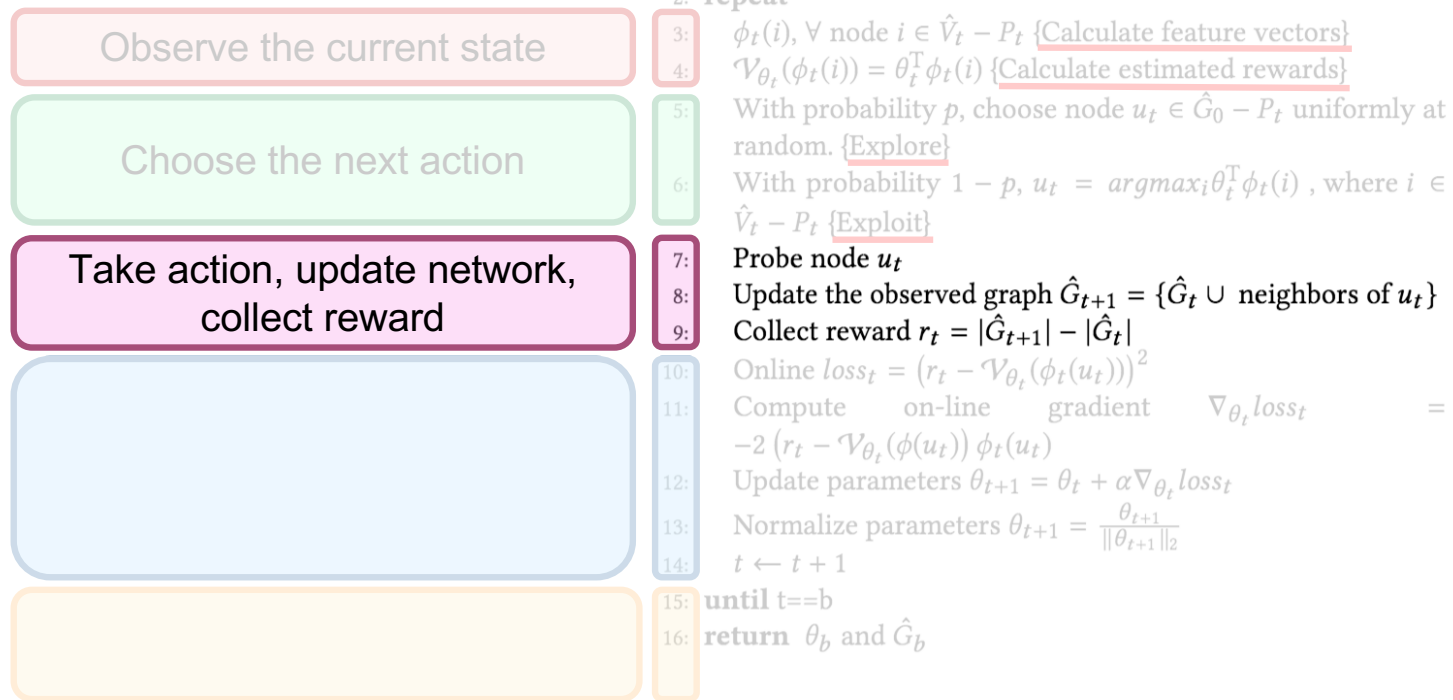
13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$

14:  $t \leftarrow t + 1$

15: until  $t=b$

16: return  $\theta_b$  and  $\hat{G}_b$

# NOL algorithm



**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$

2: **repeat**

3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}

4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}

5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}

6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}

7: Probe node  $u_t$

8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$

9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$

10: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$

11: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$

12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$

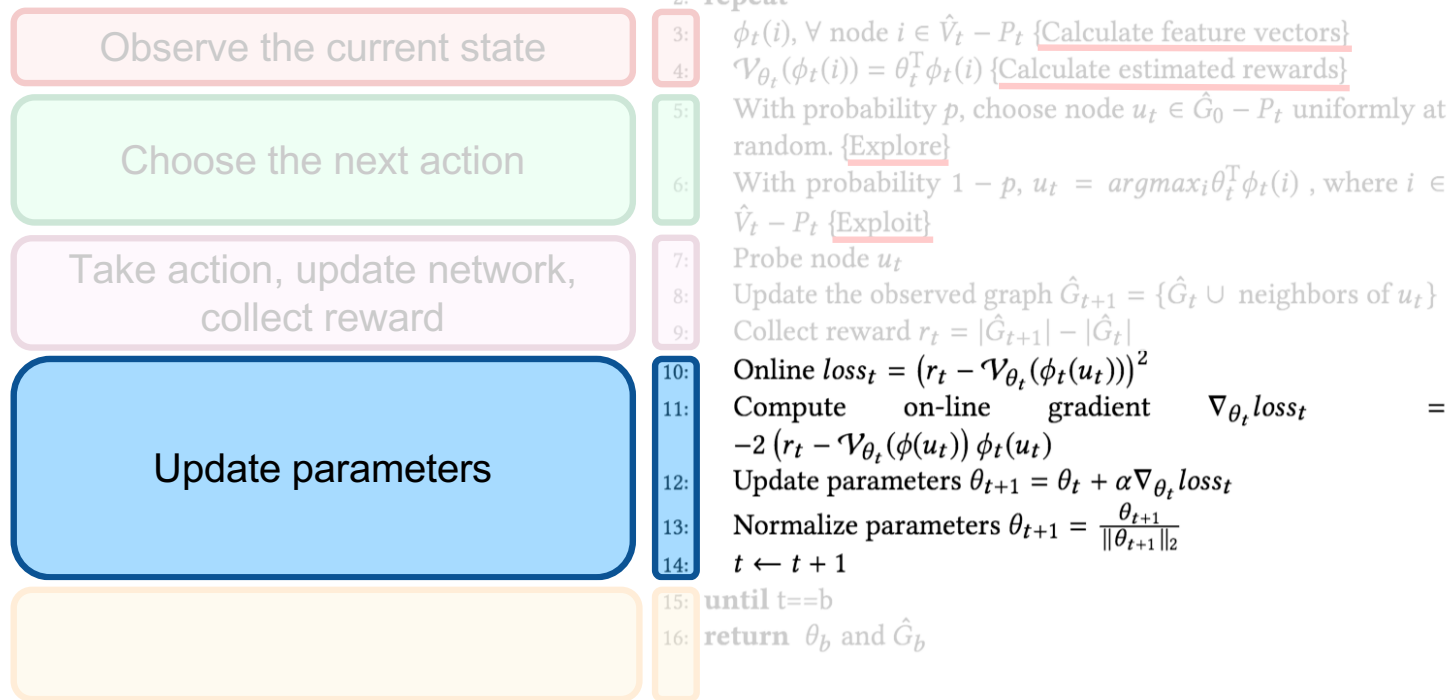
13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$

14:  $t \leftarrow t + 1$

15: **until**  $t=b$

16: **return**  $\theta_b$  and  $\hat{G}_b$

# NOL algorithm



**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$

2: repeat

3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}

4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}

5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}

6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}

7: Probe node  $u_t$

8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$

9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$

10: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$

11: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$

12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$

13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$

14:  $t \leftarrow t + 1$

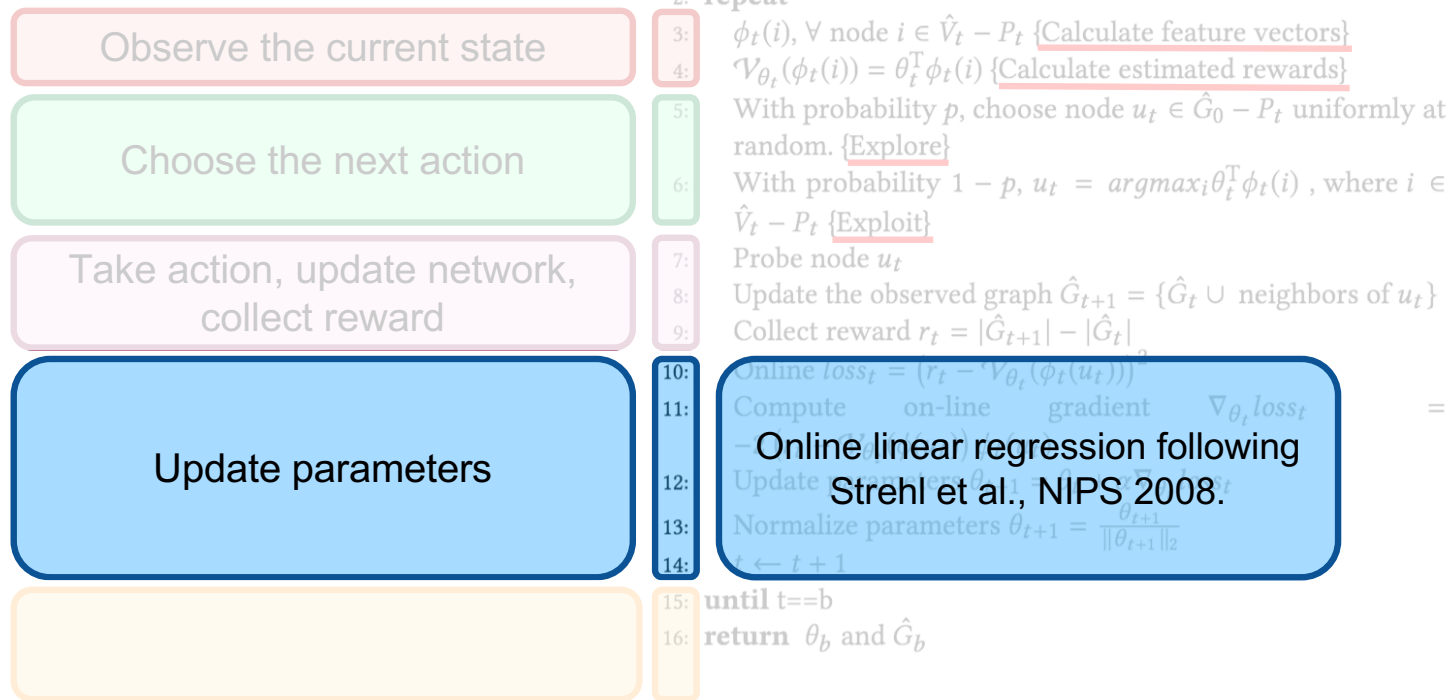
15: until  $t=b$

16: return  $\theta_b$  and  $\hat{G}_b$





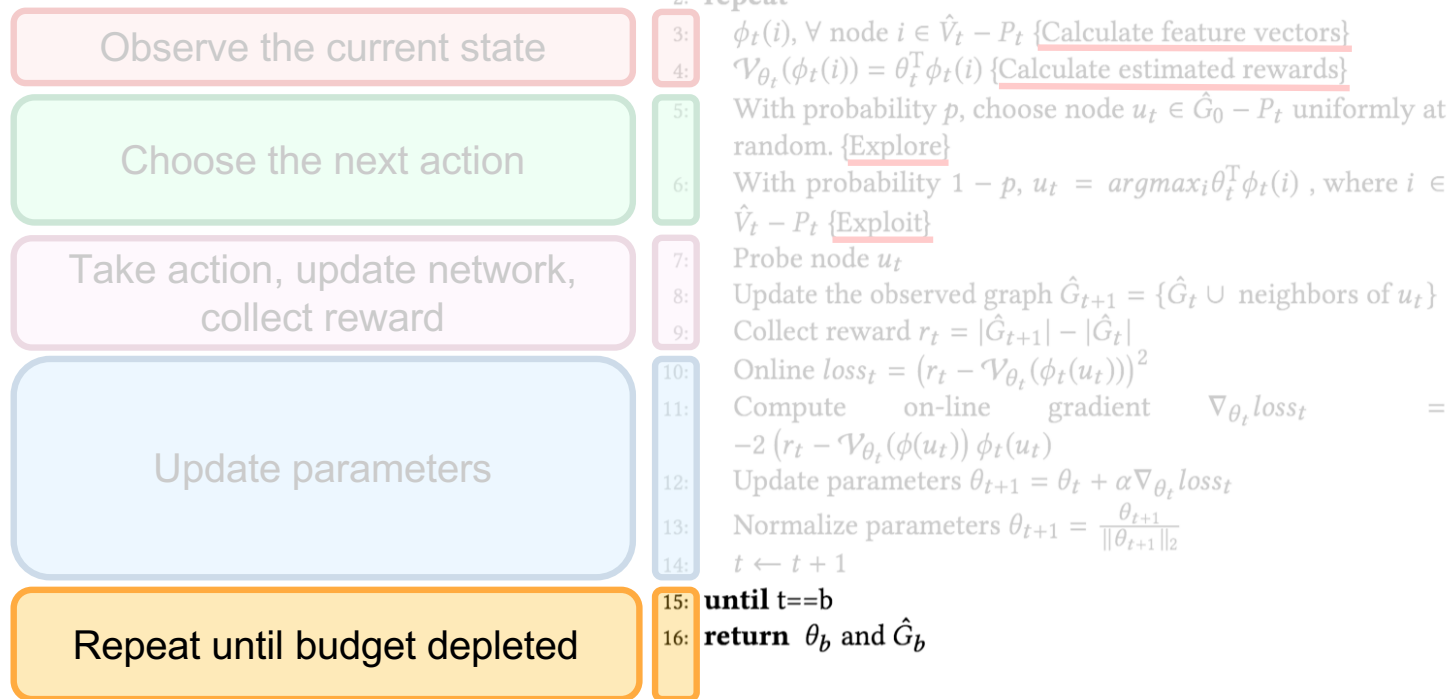
# NOL algorithm



Online linear regression following Strehl et al., NIPS 2008.



# NOL algorithm



**Input:**  $\hat{G}_0$  (initial incomplete network),  $b$  (probing budget),  $p$  (jump rate), and  $\alpha$  (step-size for gradient descent)

**Output:**  $\theta$  (parameters of the learning model),  $\hat{G}_b$  (network after  $b$  probes)

1: Initialize:  $\theta_0$  (randomly or heuristically);  $P_0 = \emptyset$

2: **repeat**

3:  $\phi_t(i), \forall \text{ node } i \in \hat{V}_t - P_t$  {Calculate feature vectors}

4:  $\mathcal{V}_{\theta_t}(\phi_t(i)) = \theta_t^T \phi_t(i)$  {Calculate estimated rewards}

5: With probability  $p$ , choose node  $u_t \in \hat{G}_0 - P_t$  uniformly at random. {Explore}

6: With probability  $1 - p$ ,  $u_t = \operatorname{argmax}_i \theta_t^T \phi_t(i)$ , where  $i \in \hat{V}_t - P_t$  {Exploit}

7: Probe node  $u_t$

8: Update the observed graph  $\hat{G}_{t+1} = \{\hat{G}_t \cup \text{neighbors of } u_t\}$

9: Collect reward  $r_t = |\hat{G}_{t+1}| - |\hat{G}_t|$

10: Online  $loss_t = (r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t)))^2$

11: Compute on-line gradient  $\nabla_{\theta_t} loss_t = -2(r_t - \mathcal{V}_{\theta_t}(\phi_t(u_t))) \phi_t(u_t)$

12: Update parameters  $\theta_{t+1} = \theta_t + \alpha \nabla_{\theta_t} loss_t$

13: Normalize parameters  $\theta_{t+1} = \frac{\theta_{t+1}}{\|\theta_{t+1}\|_2}$

14:  $t \leftarrow t + 1$

15: **until**  $t=b$

16: **return**  $\theta_b$  and  $\hat{G}_b$

# NOL algorithm

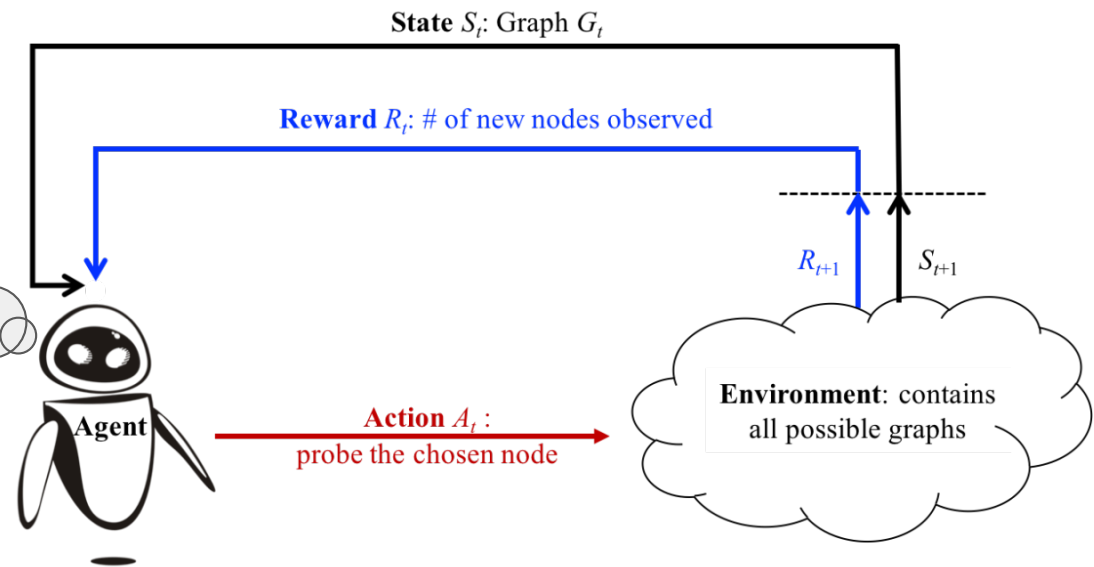
1. Observe the current state

2. Choose the next action


3. Take action, update network, collect reward

4. Update parameters


5. Repeat until budget depleted



# Features




# Features




## Feature: In-sample degree

$$\phi_{\mathbf{i},\mathbf{0}} = \hat{d}_i$$




# Feature: In-sample clustering coefficient

$$\phi_{i,1} = \hat{c}c_i$$



Feature: Normalized size of connected component


$$\phi_{i,2} = \text{CompSize}_i$$





Feature: Fraction of probed neighbors


$$\phi_{i,3} = \text{ProbedNeighbors}_i$$



## Feature: Lost Reward

$$\phi_{i,4} = \textit{LostReward}_i$$


Key idea: The order in which we probe nodes can impact the reward they yield.



# Feature: Lost Reward

$$\phi_{i,4} = \textit{LostReward}_i$$

Key idea: The order in which we probe nodes can impact the reward they yield.



Unobserved node  
(potential reward  
for u **or** v)

Partially observed nodes

Research Question: Given a networked dataset, and limited resources to collect more data, **how can you get the most bang for your buck?**

- Potential bang for your buck **depends on network structure!**



Heterogeneity in network properties creates a learning “spectrum”



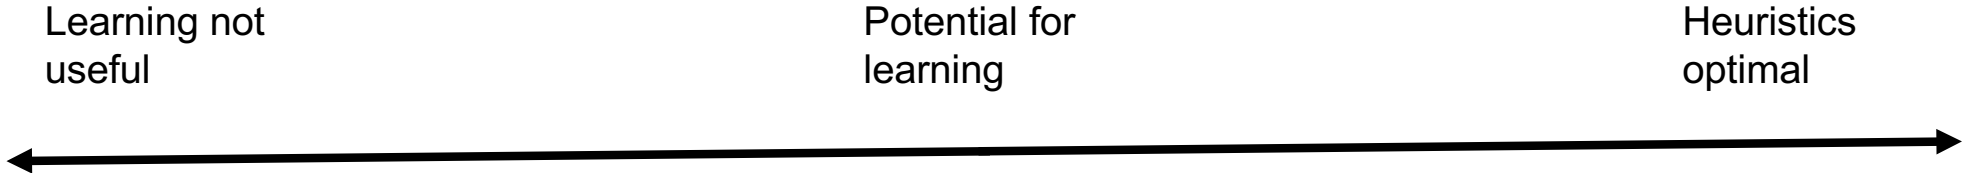
# Heterogeneity in network properties creates a learning “spectrum”

Learning not  
useful

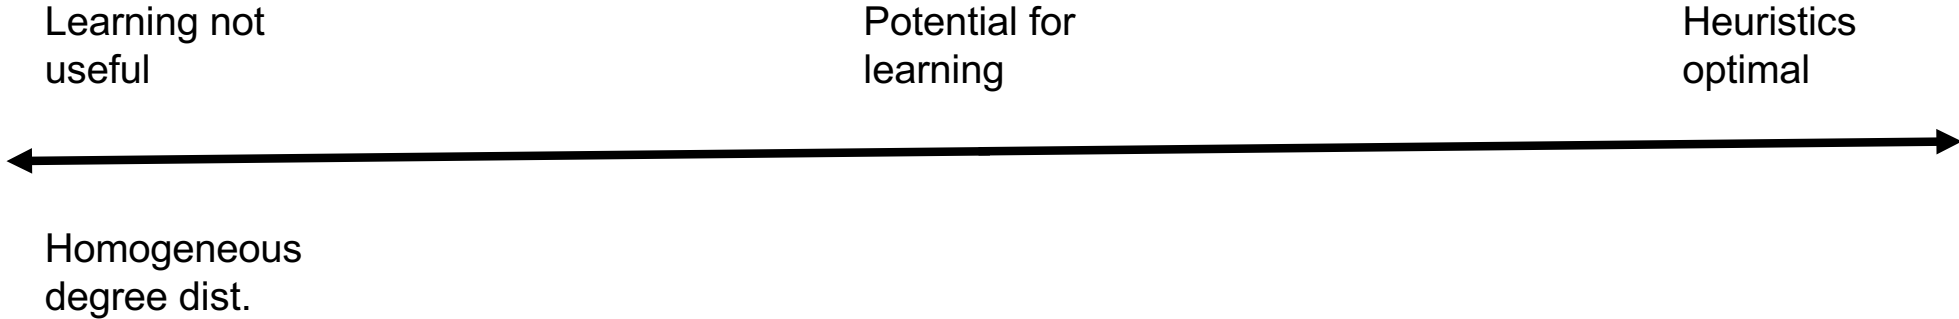
Heuristics  
optimal



# Heterogeneity in network properties creates a learning “spectrum”

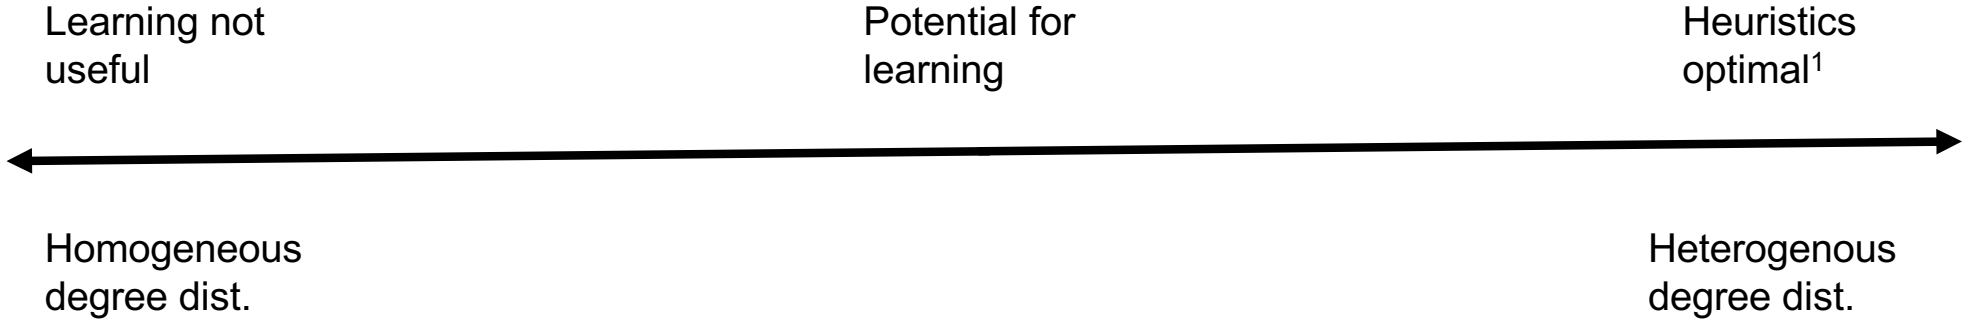


# Heterogeneity in network properties creates a learning “spectrum”



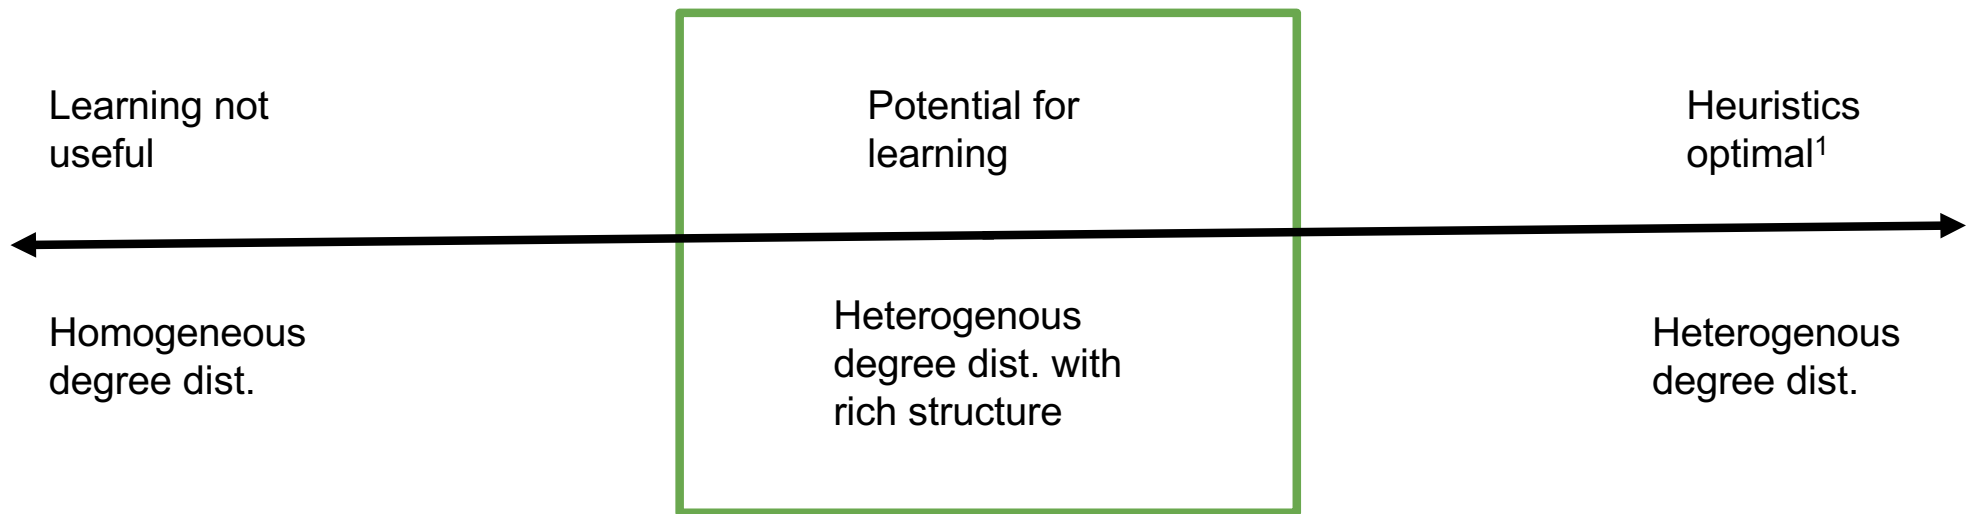


# Heterogeneity in network properties creates a learning “spectrum”



<sup>1</sup>Avrachenkov et al. *INFOCOM WKSHPs* (2014)

# Heterogeneity in network properties creates a learning “spectrum”



# Experiments



# Heuristic Baselines

- High degree
  - Probe the unprobed node with maximum degree
- High degree w/ jump
  - Probe the unprobed node with maximum degree, randomly jump with probability  $p$
- Low degree
  - Probe the unprobed node with minimum degree
- Random
  - Probe a node chosen uniformly at random from the unprobed nodes

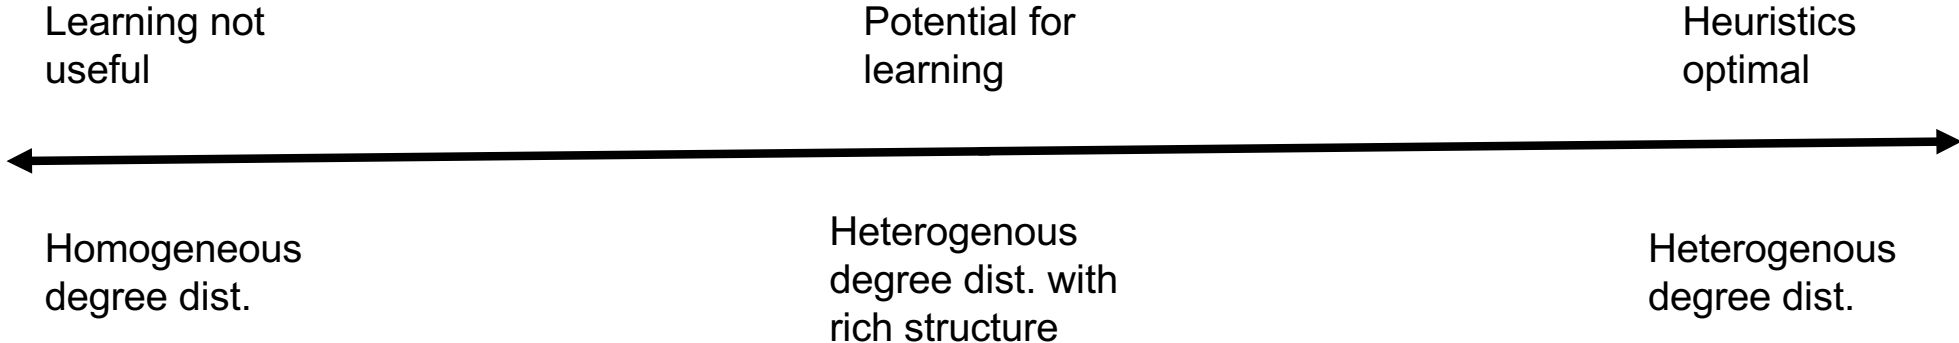


## iKNN-UCB [Madhawa+ ArXiv preprint, 2018]

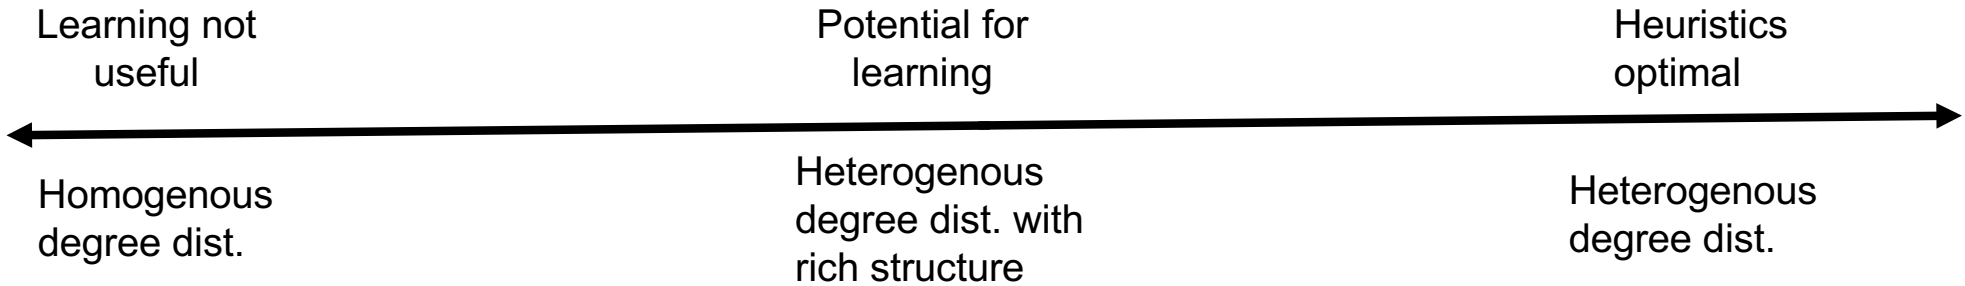
- K-Nearest Neighbors Upper Confidence Bound
  - Nonparametric multi-armed bandit approach
- Choose node to probe by combining nearest neighbor reward information (based on Euclidean distance between feature vectors) + extent of previous exploration of similar actions




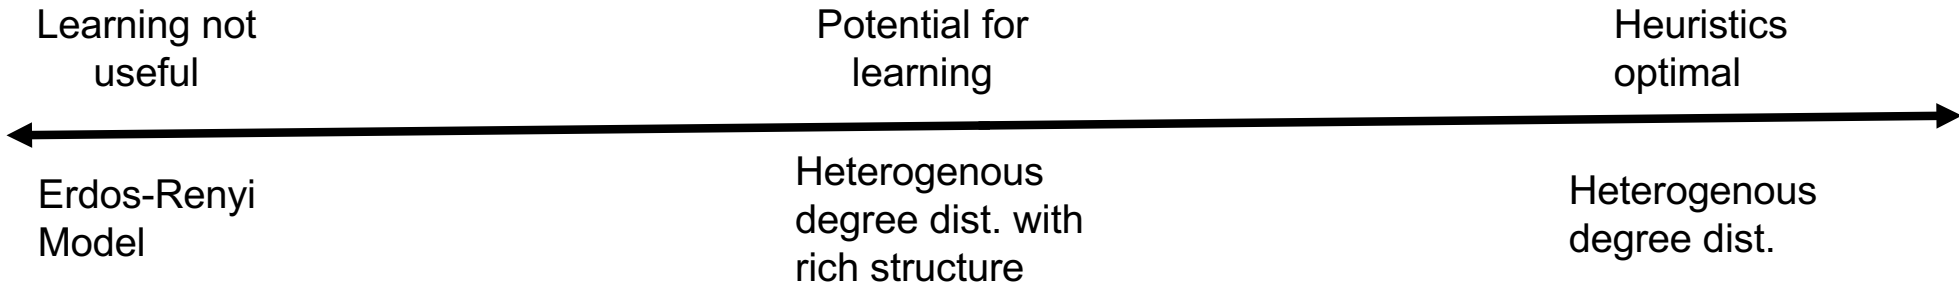
# Heterogeneity in network properties creates a learning “spectrum”



# Heterogeneity in network properties creates a learning “spectrum”

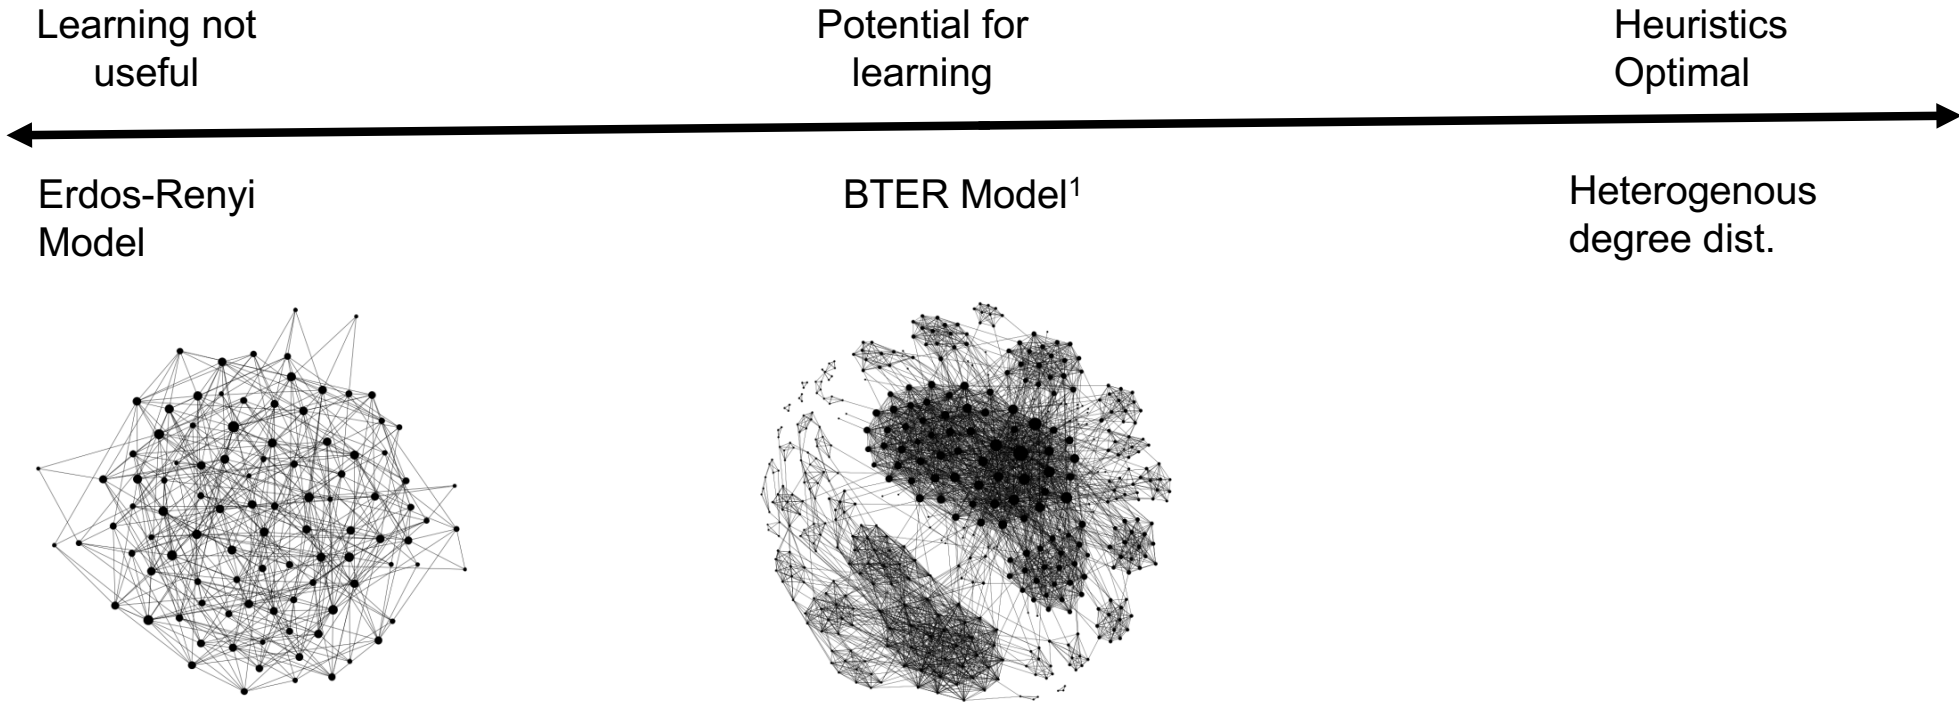


# Heterogeneity in network properties creates a learning “spectrum”



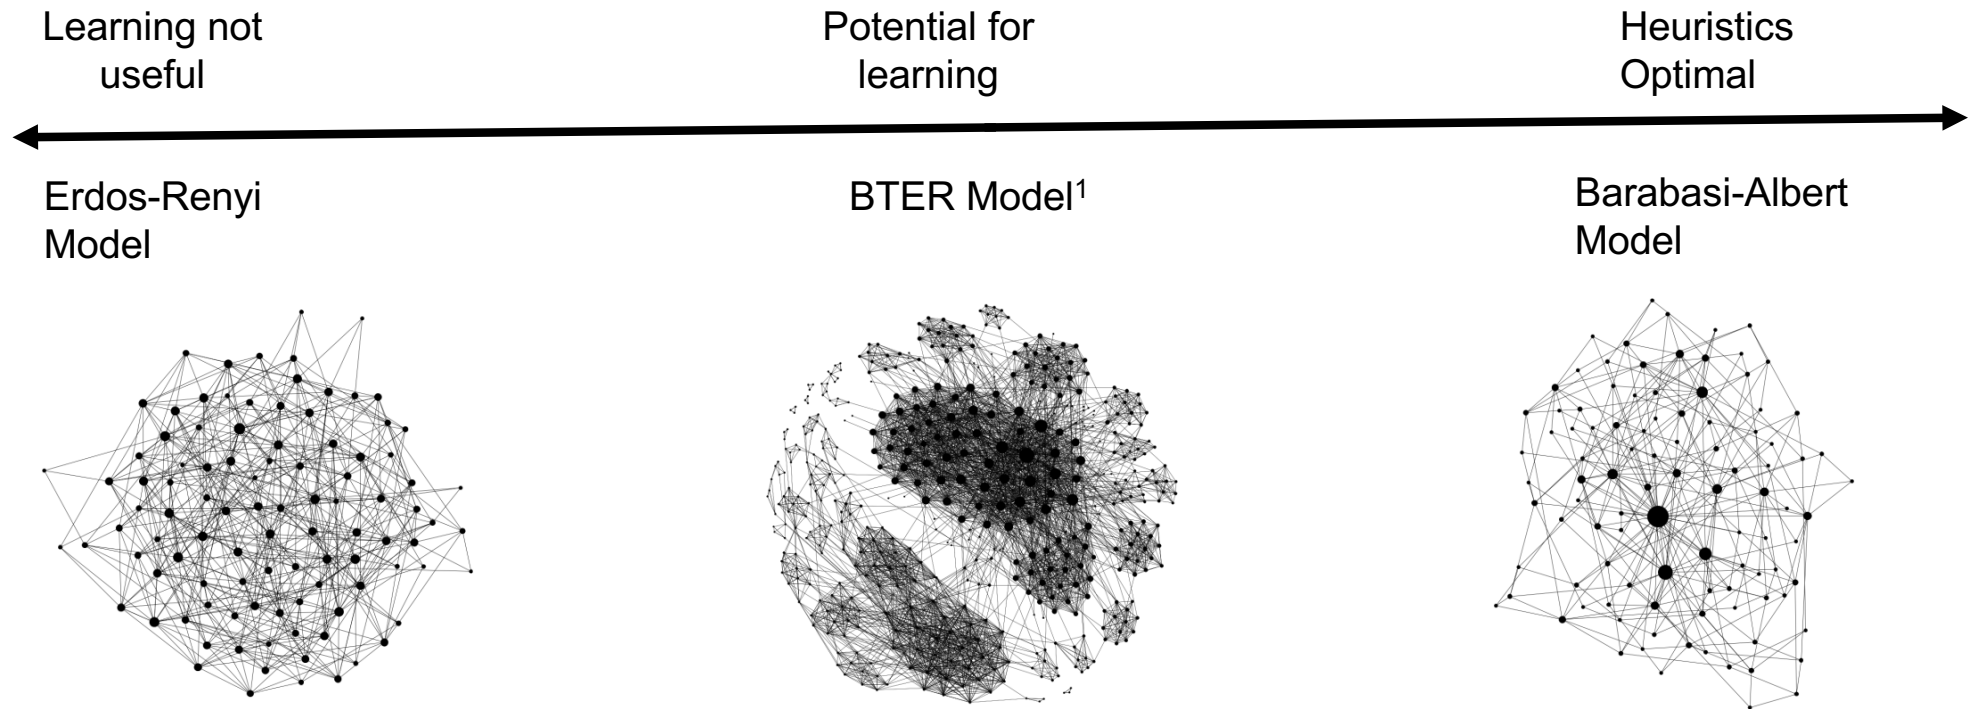


# Heterogeneity in network properties creates a learning “spectrum”




<sup>1</sup>C. Seshadhri et al. *Physical Review E* (2012)

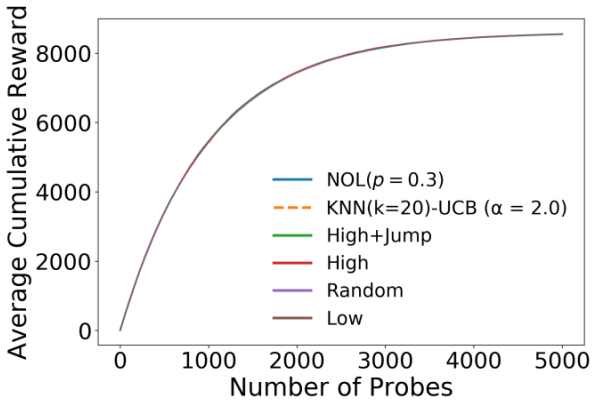
# Heterogeneity in network properties creates a learning “spectrum”




# Results - Learning Spectrum




Erdos-Renyi Model



BTER Model



Barabasi-Albert Model



# Results - Learning Spectrum

Learning not useful

Potential for learning



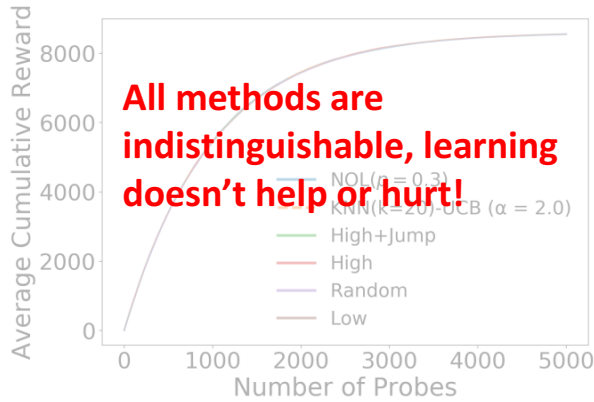
Heuristics Optimal



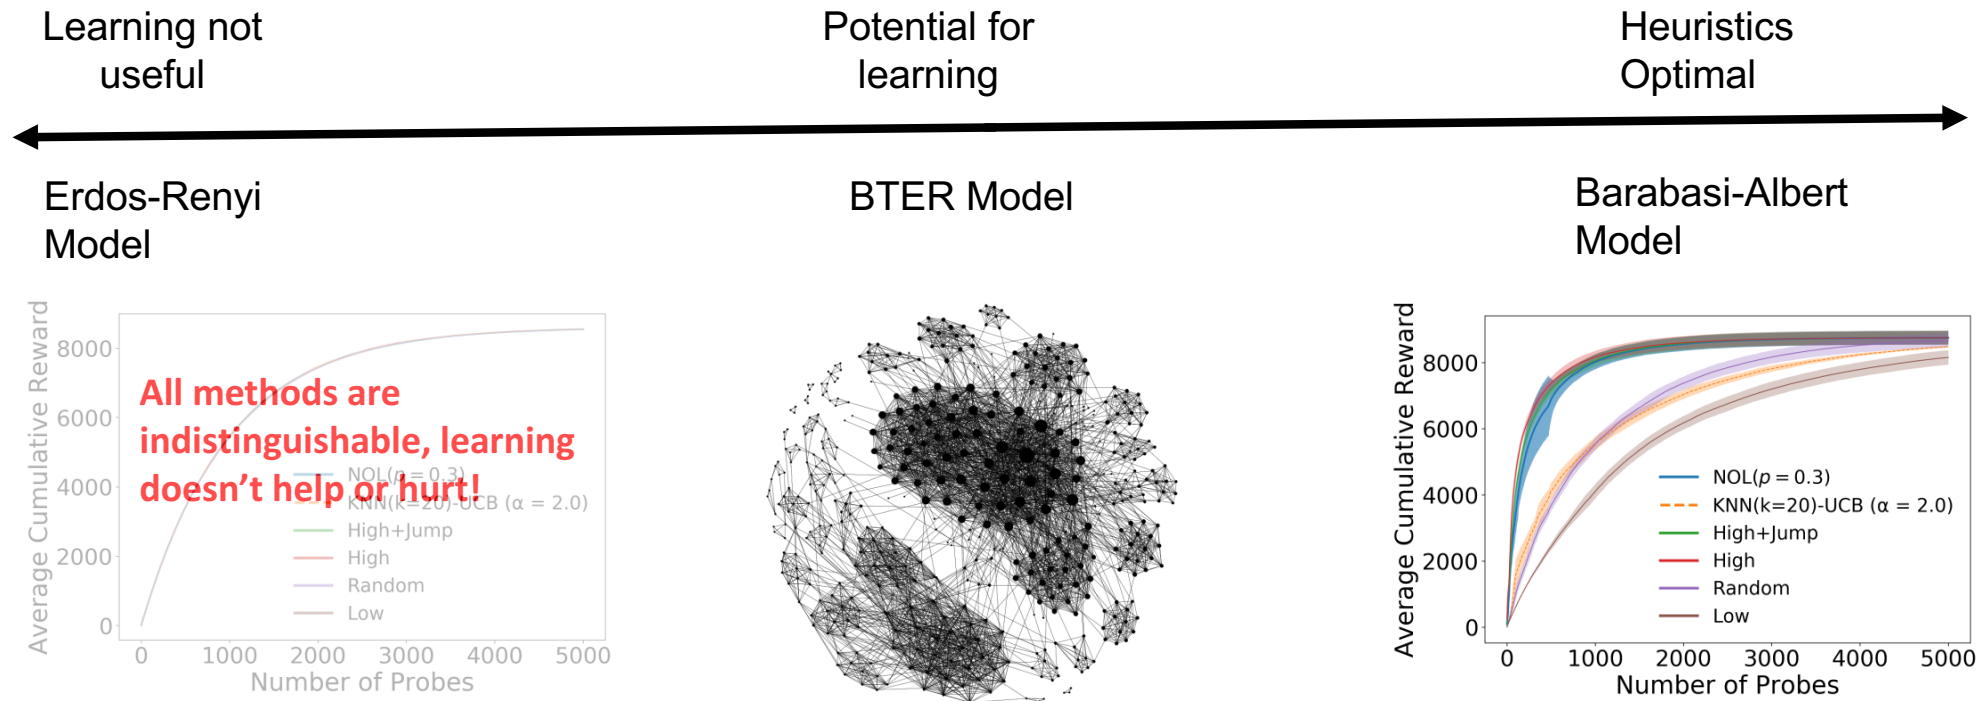
Erdos-Renyi Model

BTER Model

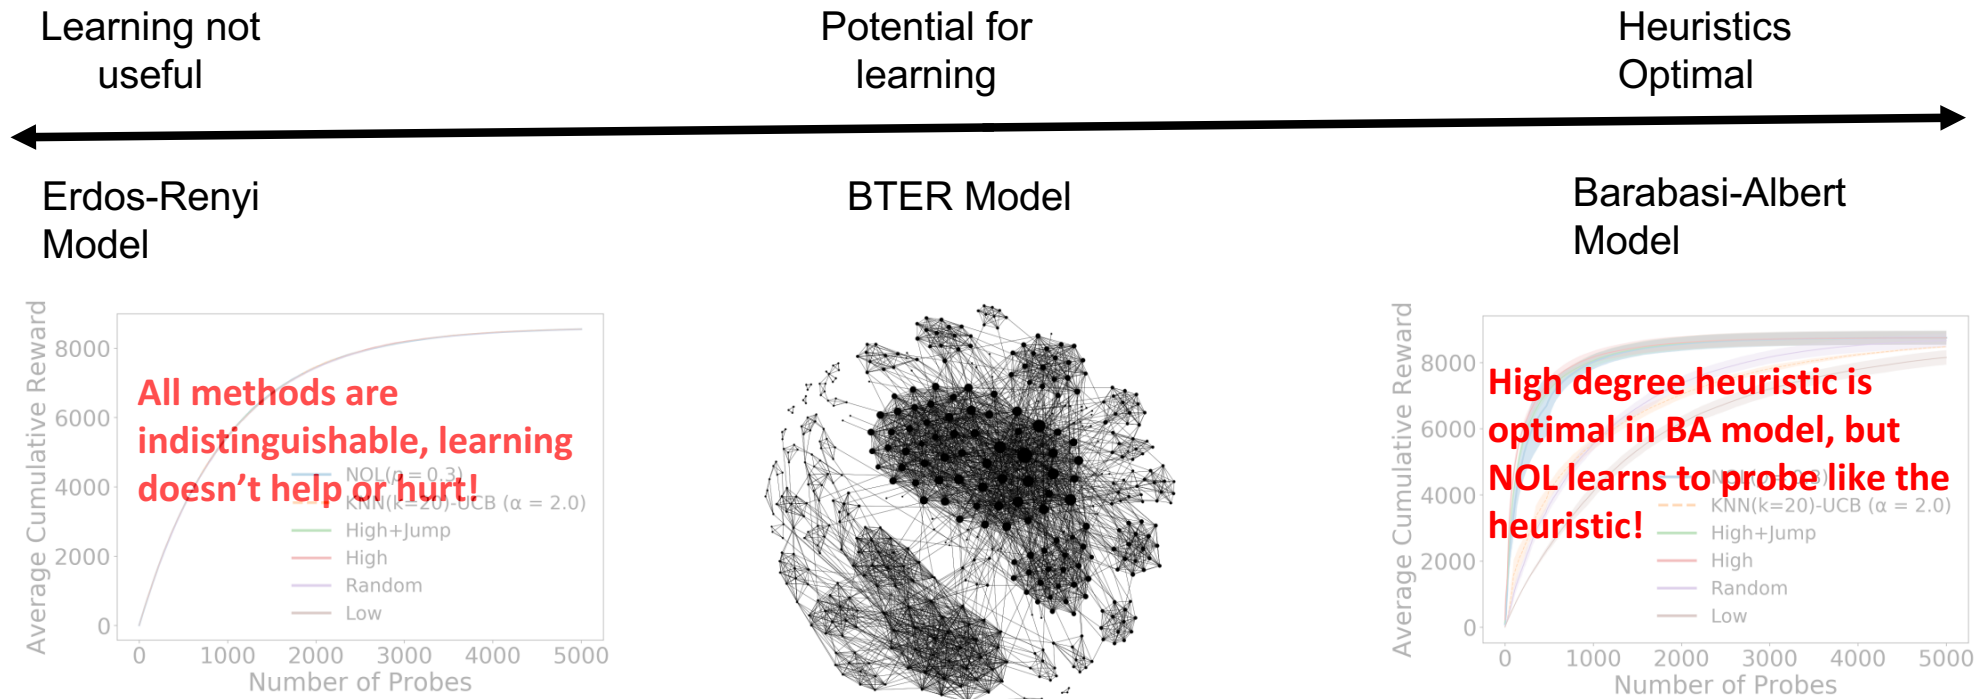
Barabasi-Albert Model



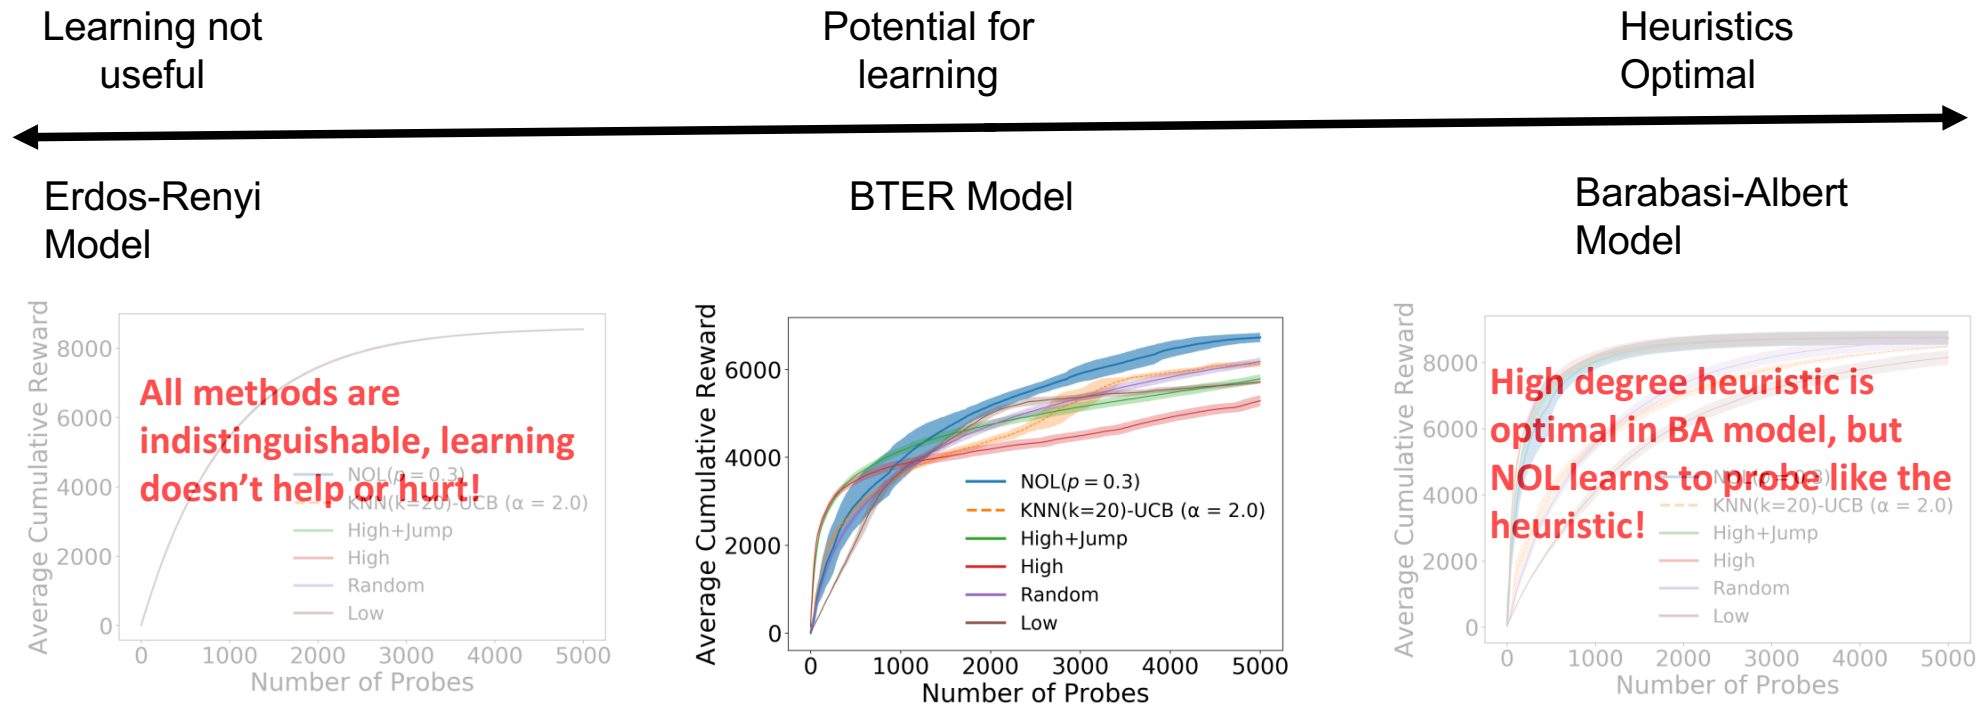
# Results - Learning Spectrum



# Results - Learning Spectrum

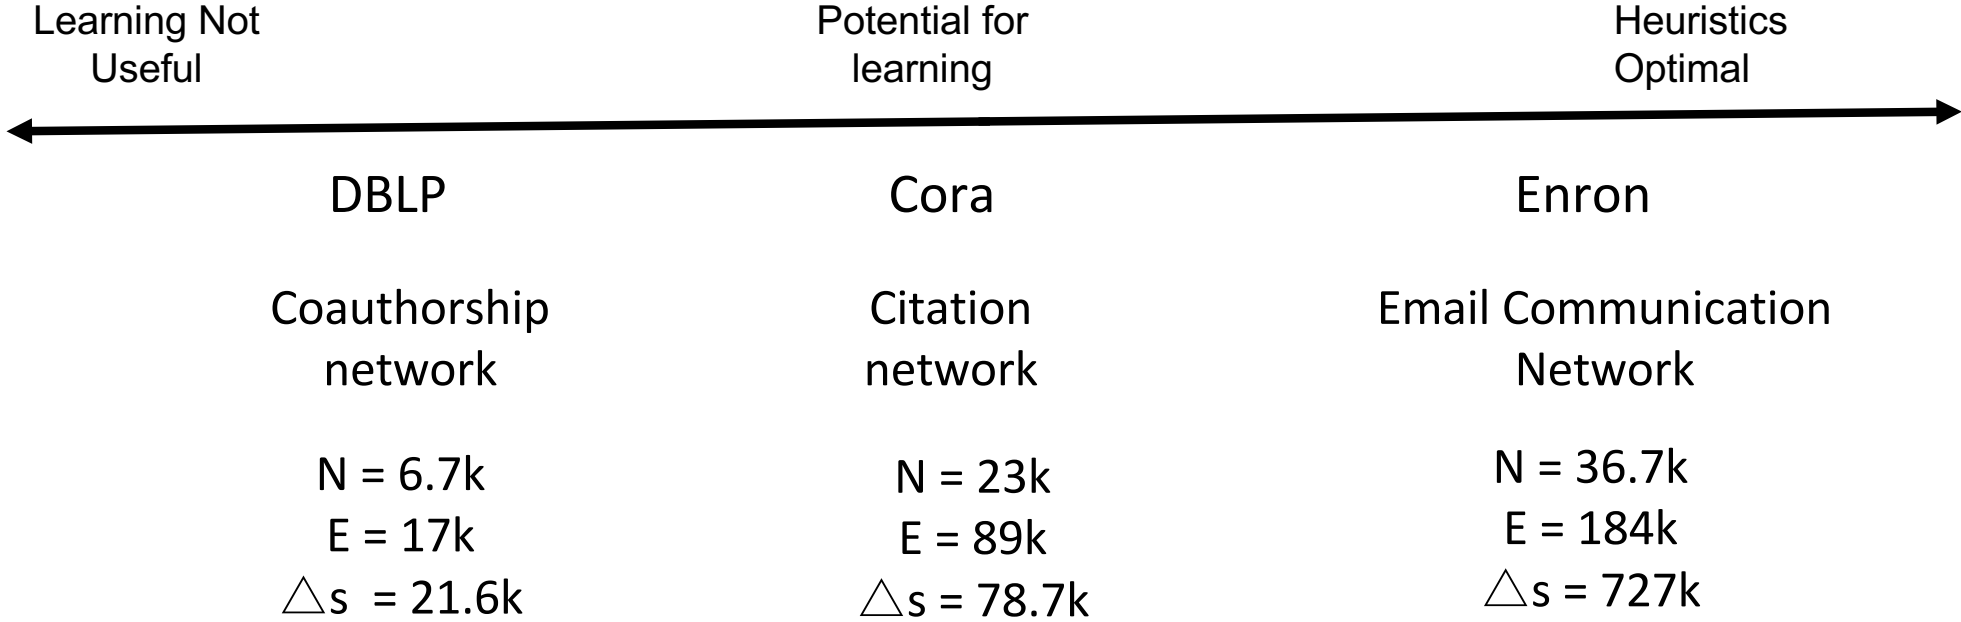


# Results - Learning Spectrum





# Heterogeneity in network properties creates a learning “spectrum”





# Heterogeneity in network properties creates a learning “spectrum”

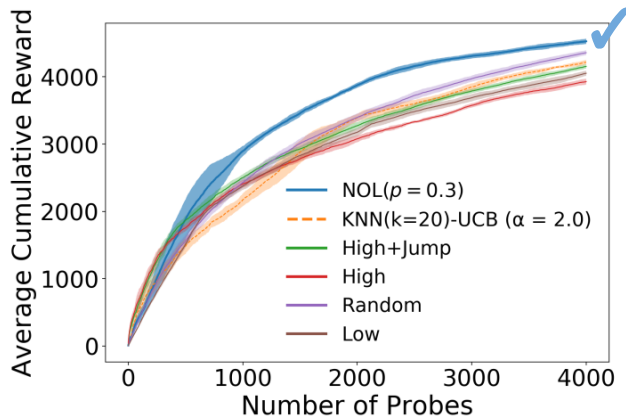
Learning Not Useful

Potential for learning


Heuristics Optimal




### DBLP



### Cora



### Enron



# Summary

- **Network Online Learning** can learn to probe online with minimal assumptions
- Success is tied to properties of the underlying network:
  - Spectrum based on objective function being maximized (degree distribution in these experiments)
  - NOL can learn to behave like the optimal heuristic
- Preliminary experiments suggest some real world complex networks fall in the “learnable” category



# Thanks!

Tim LaRock

[larock.t@husky.neu.edu](mailto:larock.t@husky.neu.edu)

---

## References

Kim et al. (2011). The Network Completion Problem: Inferring Missing Nodes and Edges in Networks. *SIAM International Conference on Data Mining*.

Seshadhri et al. (2012). Community structure and scale-free collections of Erdos-Renyi graphs, *Physical Review E*.

Avrachenkov et al. (2014). Pay few, influence most: Online myopic network covering. *Proceedings - IEEE INFOCOM*.

Pfeiffer III et al. (2014). Active Exploration in Networks: Using Probabilistic Relationships for Learning and Inference. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*.

Soundarajan et al. (2015). MaxOutProbe: An Algorithm for Increasing the Size of Partially Observed Networks. *The 2015 NIPS Workshop on Networks in the Social and Information Sciences*.

Soundarajan et al. (2016). MaxReach: Reducing network incompleteness through node probes. *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*.

Soundarajan et al. (2017).  $\epsilon$ -WGX: Adaptive Edge Probing for Enhancing Incomplete Networks. In *Proceedings of the 2017 ACM on Web Science Conference*.

Murai et al. (2018). Selective Harvesting Over Networks. *Data Mining and Knowledge Discovery*.

Madhawa et al. (2018). Exploring Partially Observed Networks with Nonparametric Bandits. arXiv:1804.07059.

Chen et al. (2018). Flexible Model Selection for Mechanistic Network Models via Super Learner. arXiv:1804.00237.

Slides at [http://eliassi.org/larock\\_netsci18.pdf](http://eliassi.org/larock_netsci18.pdf)