Learning Policies for Effective Incentive Allocation in Unknown Social Networks

Shiqing Wu[1], Quan Bai[1], and Weihua Li[2]
[1]School of Technology, Environments, and Design, University of Tasmania, Hobart, Australia
[2]Department of Computer Science, Auckland University of Technology, Auckland, New Zealand

INTRODUCTION

Incentivizing users to take behaviors that are profitable or beneficial to the incentive providers is a crucial problem to many fields. Such processes are computationally modeled as *incentive allocation problem*, where the goal is to incentivize users with effective incentives under a budget limitation, such that the number of users who take the behavior that the incentive provider expects is maximized.

CHALLENGE

- To design a reasonable incentive structure and the pricing policies.
  - Overpricing the incentive would waste the budget, whereas underpricing could fail to incentivize the user.
  - Users' attributes, such as users' preferences and skill abilities can be unavailable for generating incentives.
- To engage influential users and exploit their social influence to incentivize more users.
  - Recognizing influential users in a social network is difficult in practice.
  - Influence probabilities on edges are unknown.

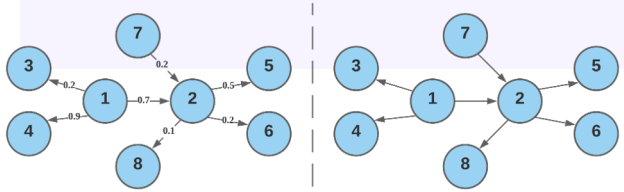

METHODOLOGY

We propose a Reinforcement Learning-based framework, named **Geometric Actor Critic (GAC),** to tackle the aforementioned challenges.

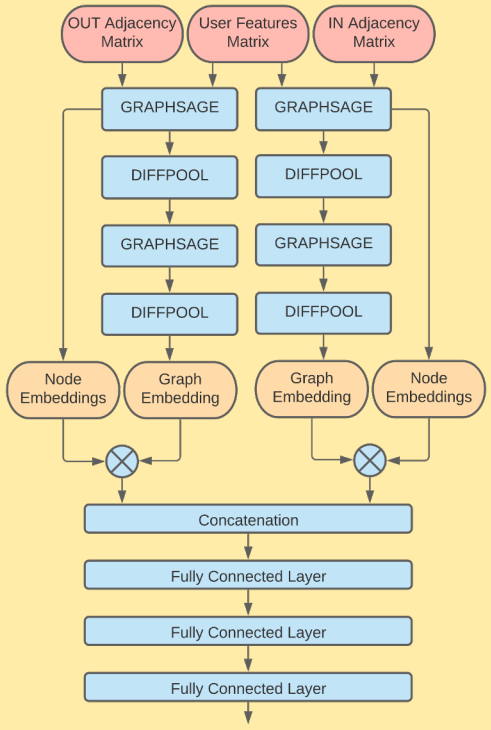

EXPERIMENTAL SETUP

We conduct the simulation-based experiments to evaluate the performance of the propsoed GAC comparing with four baseline approches. Three real-world social network datasets are deployed.

Dataset	#Nodes	#Edges	Avg. Degree
Dolphins	62	159	5.1
Twitter	236	2478	21.0
Wiki-Vote	889	2914	6.6

RESULTS

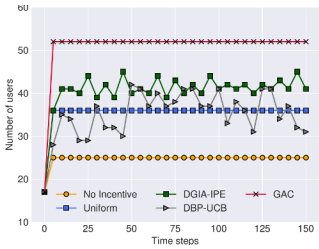

Dolphins, Budget=3

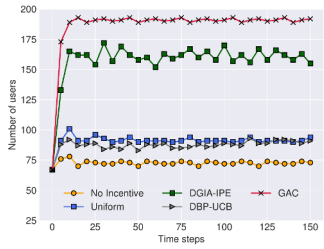

Twitter, Budget=20

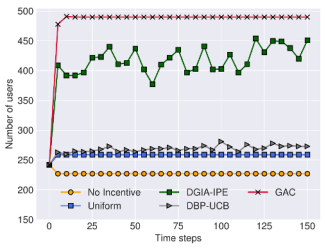

Wiki-Vote, Budget=40

CONCLUSION

Pros:
- GAC is able to solve the incentive allocation problem in unknown social networks, where only limited prior knowledge is required.
- GAC learns to represent the network from both global and local perspectives.
- The trained GAC outperforms compared baseline appraoches in all three deployed datasets under a fixed budget constraint.

Cons:
- GAC requires well-trained beforehand to obtain superior performance.
- The cost for training GAC could be exploded with the increasing size of the network, as it requires two adjacency matrices as input.