19th International Conference on Autonomous Agents and Multi-Agent Systems 2020

AAMAS 2020

# Winning an Election:
# On Emergent Strategic Communication in Multi-Agent Networks

(Extended Abstract)

Shubham Gupta[1] and Ambedkar Dukkipati[1]

[1]Department of Computer Science and Automation
Indian Institute of Science, Bangalore – 560012, INDIA
[shubhamg, ambedkar]@iisc.ac.in

# Outline

- Proposed voting game
- Emergent communication
- Experiments and observations
- Future work

# Voting Game

# Voting Game

- Modelled as an augmented Markov Game with communication among networked agents
- Markov Game: $(S, \{A_n, r_n\}_{n=1}^{N}, T, \gamma)$
- Voting Game: $(S, \{A_n, r_n\}_{n=1}^{N}, T, G, V)$
- Communication via sequence of discrete symbols

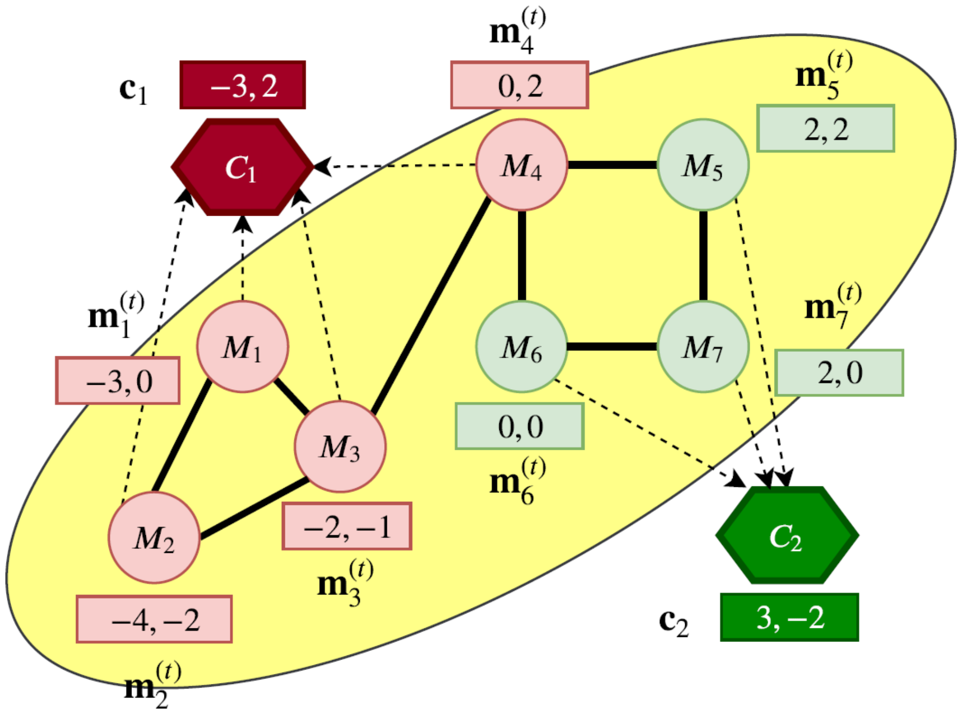

# Voting Game: Agents

- N agents connected via an underlying network (population members)
- Two special agents contesting an election (candidates)
- Let $G$ be the underlying (static) network that connects members

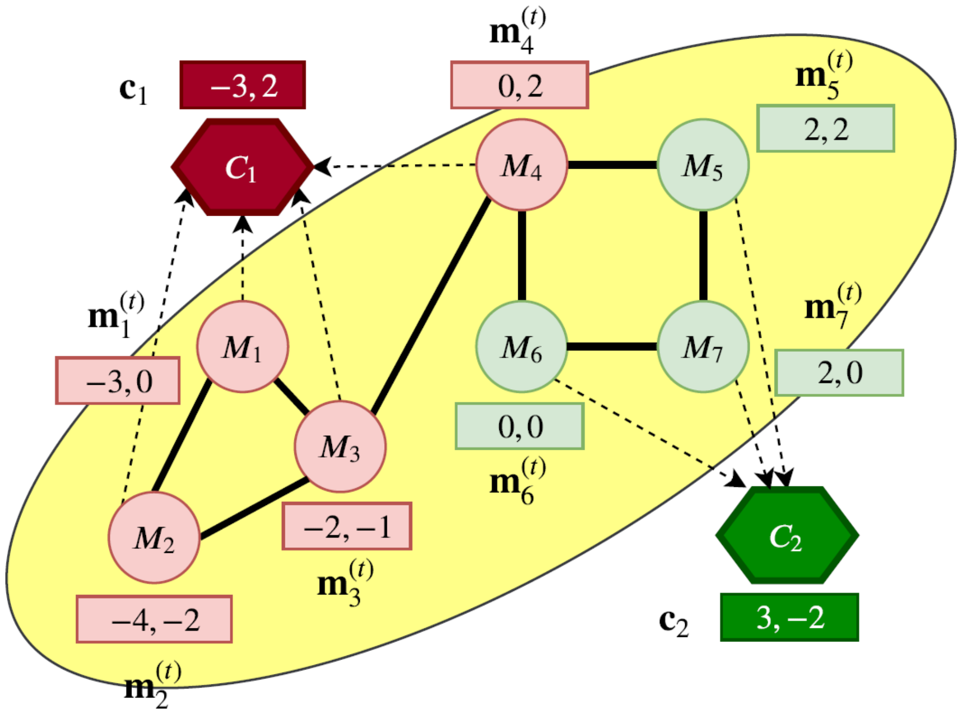

# Voting Game: State

- Each member has a preference vector and each candidate has a propaganda vector
- Members follow candidates, $F_n^{(t)} \in \{1,2\}$ denotes the choice of member $n$
- Voting is conducted after $T$ propaganda steps, $V_n \in \{1,2\}$ denotes the vote cast by member $n$
- Partial observability

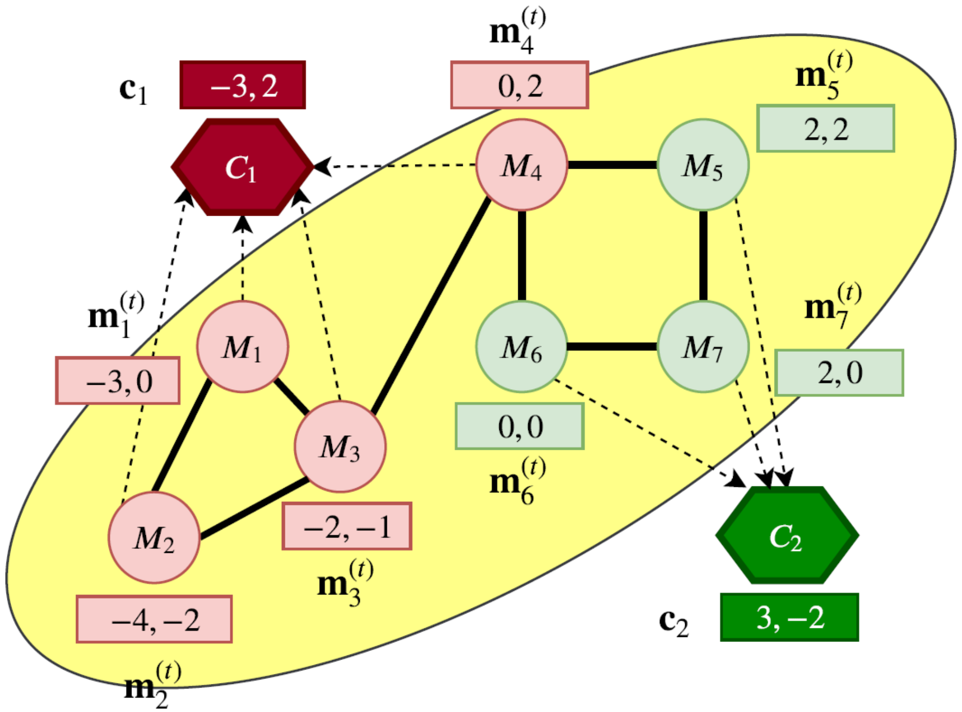

# Voting Game: Actions

- Communication action: Select and broadcast a sequence of discrete symbols from vocabulary $V$
- Modification action: Propose an update to preference vectors
- Candidates only take communication actions
- Actions are taken based on local observations and received messages.

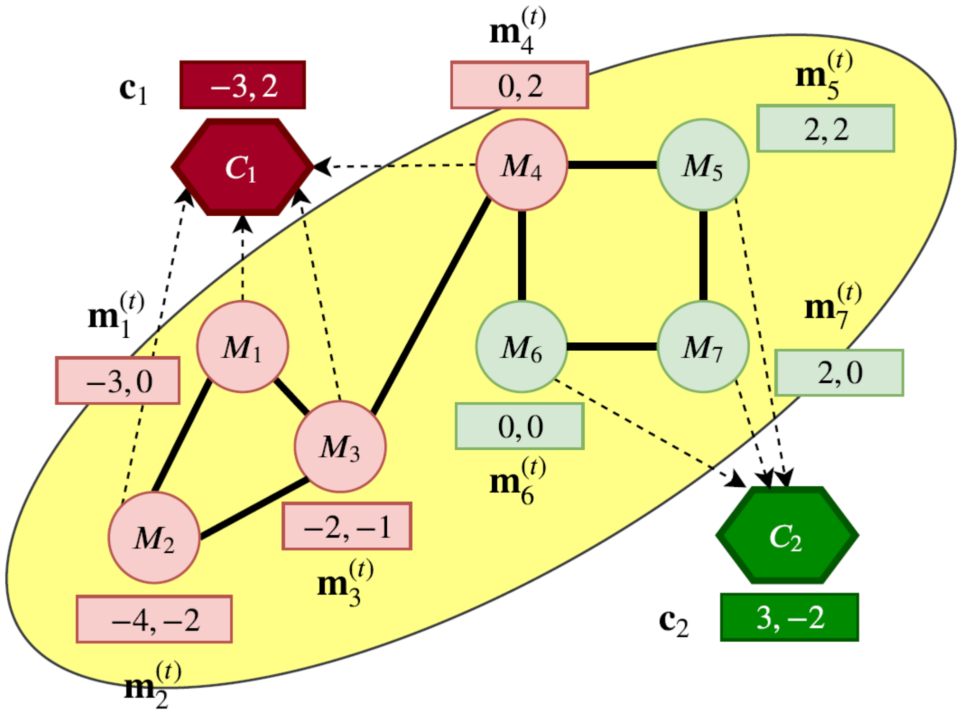

# Voting Game: Transitions

- Propaganda vectors remain fixed
- Preference vectors are updated based on proposed modification
- Underlying network $G$ remains fixed
- $F_n^{(t)}$ is sampled based on preference vector of member $n$ using Gumbel-Softmax trick

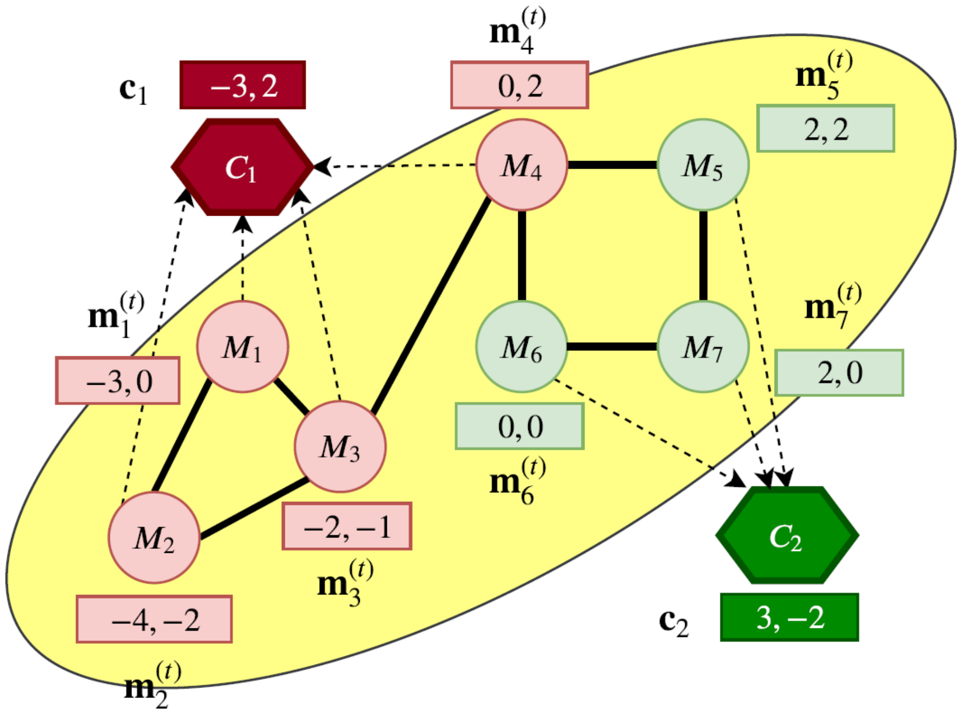

# Voting Game: Rewards

- For members: $r_n^{(t)} = -\left\|m_n^{(t)} - c_{V_n}\right\|^2$ if $t = T + 1$ and 0 otherwise
- For candidates: At $t = T + 1$
  - Competitive: $r_j^{(t)} = \sum_{n=1}^{N} 1\,\{V_n = j\}$
  - Cooperative: $r_1^{(t)} = \sum_{n=1}^{N} 1\,\{V_n = 1\}$ and $r_2^{(t)} = -\sum_{n=1}^{N} 1\,\{V_n = 2\}$
  - Reward is zero if $t \neq T + 1$

# Voting Game: Objective

- Find policies for agents to maximize the rewards
- We parameterize the policies using neural networks
- Candidates have their own policies
- Members share the policy
- Communication engine is shared by all agents

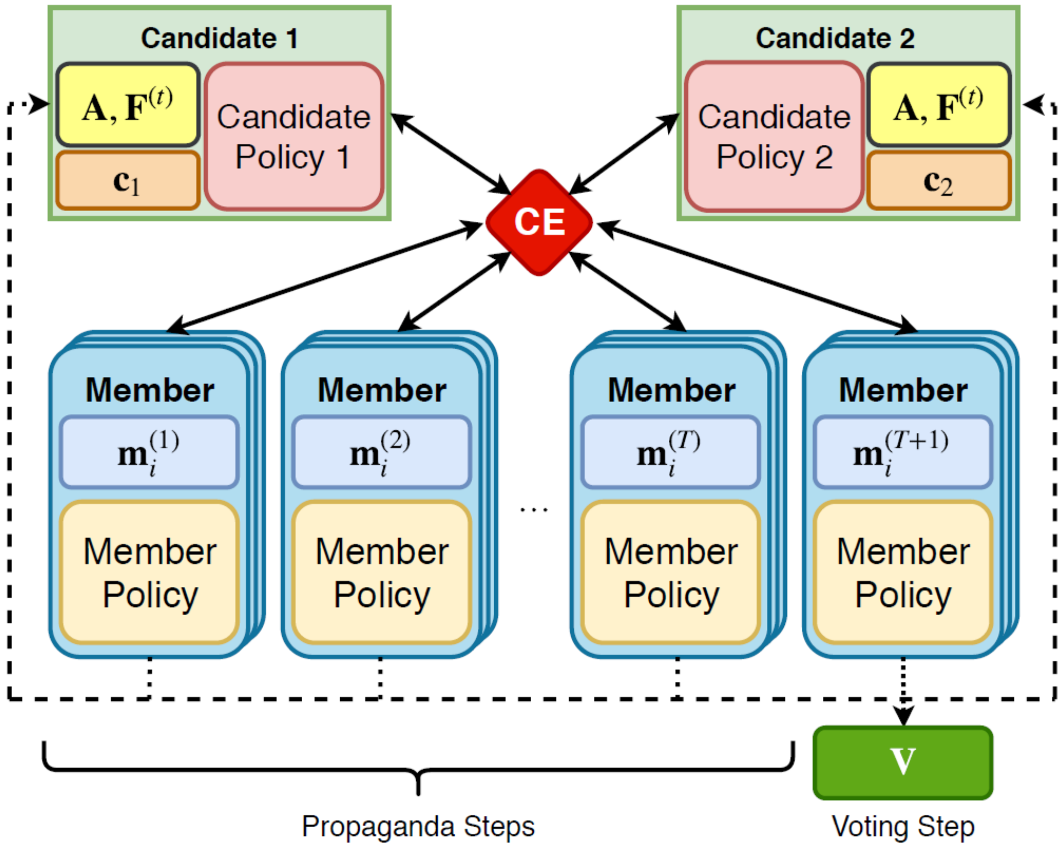

# Emergent Communication

# Emergent Communication

- Fixed vocabulary, no semantics attached to words apriori
- Agents have to evolve a language as they train (for spreading propaganda)
- Communication is needed to maximize rewards due to partial observability
- Used sequences of discrete symbols to ease analysis

# Emergent Communication: Other Approaches

## Functional view of communication: beyond referential games

- Approaches like [Lazaridou *et al.*, 2017; Das *et al.*, 2017; Havrylov and Titov, 2017; Lazaridou *et al.*, 2018] use variants of Lewis's Signalling Game [Lewis, 1969]
- An agent has learned a language if it can use it to accomplish certain goals [Gauthier and Mordatch, 2016; Mordatch and Abbeel, 2018; Cao *et al.*, 2018]
- Can agents develop a language for devising and executing abstract strategies?

# Emergent Communication: Other Approaches

## Network restricted communication:

- Emergent communication with a subset of agents [Das *et al.*, 2019]
- Fixed communication protocol with networked agents [Zhang *et al.*, 2018]
- A fixed topology restricts the communication among agents
- Topology influences the emergent language

# Experiments and Observations

# Experiments: Evaluation Procedure

- Either use a fixed underlying network or sample a random network
- Train the agents to maximize their rewards
- A active candidate is the one who is allowed to broadcast messages
- Observe the various effects of having only one/both candidates active

# Experiments: Learned Strategies

| Network Used | Active | Biased Training | Unbiased Training |
|---|---|---|---|
| **RGG** | $C_1$ | (**0.94**, 0.04) | (**0.69**, 0.25) |
| | $C_2$ | (**0.96**, 0.02) | (0.22, **0.73**) |
| | Both | (**0.99**, 0.01) | (0.31, **0.60**) |
| **NS** | $C_1$ | (**0.92**, 0.08) | (**0.62**, 0.38) |
| | $C_2$ | (**0.94**, 0.06) | (0.37, **0.63**) |
| | Both | (**0.96**, 0.04) | (0.46, **0.54**) |

- Active candidate wins in competitive (unbiased) setting
- Candidate 1 wins in cooperative (biased) setting
- Candidate 1 wins with a smaller margin in competitive (unbiased) setting when only it is active

# Experiments: Properties of Emergent Communication

- Candidates use same high-frequency words
- Candidates differ in their usage of low-frequency words
- Language usage correlates with structural communities

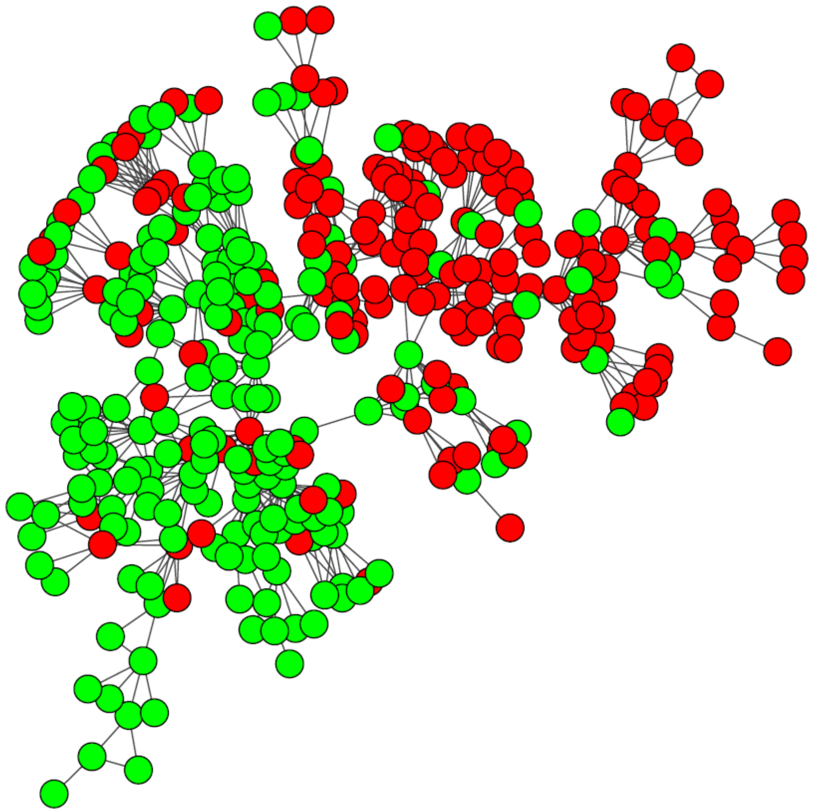

# Future Work

# Future Work

- What if the underlying network is not static?
- What is the effect of incentivizing candidates to learn about network structure?
- What if one-to-one communication is allowed?
- Can we explain emergence of communities using our framework?
- Can we discover more connections between fields like network science, game theory and multi-agent reinforcement learning?

# Summary

- Proposed a novel voting game
- Studied emergent communication over networked agents
- Answered the following questions:
  - Do candidates learn meaningful strategies?
  - Is the emergent communication useful?
  - What is the relationship between the community structure in the network and emergent communication?
- Our framework can serve as stepping stone for future research in this direction

Thank you!

# References

[Lazaridou et al., 2017] Angeliki Lazaridou, Alexander Peysakhovich, and Marco Baroni. Multi-agent cooperation and the emergence of (natural) language. 5th International Conference on Learning Representations (ICLR), 2017.

[Das et al., 2017] A. Das, S. Kottur, J. M. F. Moura, S. Lee, and D. Batra. Learning cooperative visual dialog agents with deep reinforcement learning. In 2017 IEEE International Conference on Computer Vision (ICCV), pages 2970–2979, 2017.

[Havrylov and Titov, 2017] Serhii Havrylov and Ivan Titov. Emergence of language with multi-agent games: Learning to communicate with sequences of symbols. In Advances in Neural Information Processing Systems 30, pages 2149–2159. Curran Associates, Inc., 2017.

[Lazaridou et al., 2018] Angeliki Lazaridou, Karl Moritz Hermann, Karl Tuyls, and Stephen Clark. Emergence of linguistic communication from referential games with symbolic and pixel input. In 6th International Conference on Learning Representations (ICLR), 2018.

# References (Contd.)


[Lewis, 1969] David Lewis. Convention. Harvard University Press, Cambridge, MA, 1969.

[Gauthier and Mordatch, 2016] Jon Gauthier and Igor Mordatch. A paradigm for situated and goal driven language learning. NIPS 2016 Machine Intelligence Workshop, 2016.

[Mordatch and Abbeel, 2018] Igor Mordatch and Pieter Abbeel. Emergence of grounded compositional language in multi-agent populations. In Thirty-Second AAAI Conference on Artificial Intelligence, 2018.

[Cao et al., 2018] Kris Cao, Angeliki Lazaridou, Marc Lanctot, Joel Z Leibo, Karl Tuyls, and Stephen Clark. Emergent communication through negotiation. In 6th International Conference on Learning Representations (ICLR), 2018.

# References (Contd.)

[Das et al., 2019] Abhishek Das, Theophile Gervet, Joshua ´ Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. Tarmac: Targeted multi-agent communication. Proceedings of the 36th International Conference on Machine Learning, PMLR, 97:1538–1546, 2019.

[Zhang et al., 2018] Kaiqing Zhang, Zhuoran Yang, Han Liu, Tong Zhang, and Tamer Basar. Fully decentralized multi-agent reinforcement learning with networked agents. Proceedings of the 35th International Conference on Machine Learning, PMLR, 80:5872–5881, 2018.

[Sukhbaatar et al., 2016] Sainbayar Sukhbaatar, Arthur Szlam, and Rob Fergus. Learning multiagent communication with backpropagation. In Advances in Neural Information Processing Systems 29, pages 2244–2252. Curran Associates, Inc., 2016.

# References (Contd.)

[Foerster et al., 2016] Jakob Foerster, Ioannis Alexandros Assael, Nando de Freitas, and Shimon Whiteson. Learning to communicate with deep multi-agent reinforcement learning. In Advances in Neural Information Processing Systems 29, pages 2137–2145. Curran Associates, Inc., 2016.