

# OPTIMIZING SOCIAL HIERARCHY



MILO TRUJILLO  
JUSTIN GRANA  
BRENDAN TRACEY  
DAVID WOLPERT

August, 2017

## Abstract

This project defines a social organization mathematically in terms of communication costs, constraints, and objectives. This allows one to solve for an “optimal” organization, and begins to answer questions of why hierarchies forms, and why some hierarchical structures are more advantageous than others.

## 1 INTRODUCTION

Consider an organization consisting of “agents”, and the “environment.” Agents may either listen to other agents, or the environment, to gain information. They may speak only to other agents in order to forward their information.

We assume that all communication has two costs: the speaker must put effort in to express themselves clearly, and the listener must put effort in to paying attention to the speaker. In the case where an agent is listening to the environment, the environment has no speaking cost, but the agent still has a cost to observe the environment clearly. Speaking is much cheaper than listening, in that a speaker can address an entire group, but listeners must listen to each speaker.

Agents can collectively distribute information through the organization through a series of speaking and listening interactions. This can be described as an optimization problem, where we seek to maximally distribute information, while minimizing communication costs. Many variants of this optimization problem exist, depending on the size of the organization and environment, the cost of interactions, and the objectives of individual agents.

## 2 MODEL

In this model, there are  $N$  agents, indexed by  $i = 1..N$ . There are also  $K$  random variables representing the state of the environment ( $\theta_1.. \theta_k$ ). Each agent can observe either environment variables or the output of other agents.

Each agent is described in terms of three vectors, called “listen weights”, “state weights”, and “speaking weights”. These vectors respectively represent how carefully the agent listens to each other agent or environment node, how it translates its observations in to an action or “state”, and how how it translates its observations in to a message it can send to other agents.

### 2.1 Receiving Observations

Let  $E_i$  be a  $1 \times (N + K)$  row vector, representing the information the agent receives. The observations of agent  $i$  are then given by:

$$O_i = E_i + W_i \odot \mathcal{N}_i^e(0, \Sigma_i^e)$$

Here,  $W_i$  is a  $1 \times (N + K)$  row vector of the form  $(\frac{1}{w_1}, \frac{1}{w_2}, \dots, \frac{1}{w_N})$  that represents the “listen weights”, or how agent  $i$  allocates its attention to each observation.

The term  $\mathcal{N}_i^e(0, \Sigma_i^e)$  is a  $1 \times (N + K)$  vector of normal random variables with standard deviation  $\Sigma_i^e$ . This term represents “random noise”, where a higher standard deviation indicates more attention is needed to overcome to noise and reproduce signal.

The  $\odot$  represents element-wise multiplication. Therefore, the observations agent  $i$  receives are equal to the information it receives (from its own observations or messages from other agents), corrupted by noise that is countered by the attention agent  $i$  allocates to the observation. In short, listening can be thought of as a Gaussian channel, where the listener can reduce noise by allocating more attention to the speaker.

### 2.2 Choosing an Action

Given what agent  $i$  observes,  $O_i$ , agent  $i$ 's action is given by:

$$A_i = O_i X_i$$

Here,  $X_i$  is a  $(N + K) \times D_i$  matrix. This means that  $A_i$  is a  $1 \times D_i$  row vector. More abstractly, this means that agent  $i$  can take  $D_i$  different actions, where the action is chosen by multiplying observations by a “state weight matrix”  $X_i$ .

### 2.3 Sending Messages

Given agent  $i$ 's observations  $O_i$ , what agent  $i$  says is given by:

$$M_i = O_i \Omega_i + \mathcal{N}_i^m(0, \Sigma_i^m)$$

Here,  $\Omega_i$  is a  $(N + K) \times F_i$  matrix and  $\mathcal{N}_i^m(0, \Sigma_i^m)$  is a  $1 \times F_i$  vector of normal random variables. That means  $M_i$  is a  $1 \times F_i$  vector.

Abstractly,  $F_i$  represents the number of distinct messages agent  $i$  can say,  $\Omega_i$  is a matrix translating from observations to outgoing messages, and  $\mathcal{N}_i^m(0, \Sigma_i^m)$  represents the noise when sending messages.

### 2.4 Optimization Problem

The exact details of the optimization will vary depending on the type of social organization being modeled. However, most optimization will follow a similar form, outlined here.

For convenience, let  $\Omega = [\Omega_1, \Omega_2, \dots, \Omega_N]$ ,  $X = [X_1, X_2, \dots, X_N]$ ,  $W = [W_1, W_2, \dots, W_N]$ , and  $A = [A_1, A_2, \dots, A_n]$ .

The welfare of the model can now be defined as:

$$F = U(\theta, A) - \|\Omega\|^d - \|W\|^d$$

Here,  $\|\Omega\|^d$  and  $\|W\|^d$  represent the  $L^d$  norm of each matrix, and  $U$  is a function comparing the action of each agent to the environment state, determining how far from optimal the agents' actions are.

Intuitively,  $F$  is a trade-off between accuracy of information distributed, and the cost of distributing information.

The optimization problem can now be written simply as:

$$\max_{\Omega, X, W} F$$

### 3 METHODS

We implemented the above model as a neural network, where the network can train the listen, state, and speaking weights of each agent. In our first simulation, we tasked each agent with getting an estimate of the average of all environment nodes. That is, information about each environment node must be distributed through the entire group, and the “action” of an agent is its estimate of the average environment.

The interpretation of this network is somewhat unusual, in that with most neural networks one cares about the output, or how the network approximates some function. The links between nodes in the network are traditionally ignored, or may even be unintelligible in a deep neural net with hidden layers. By contrast, we do not care about the output of the network (which determines only how well it is performing), but instead focus on the edges between nodes, which describe the communication structure of the modeled organization.

The resulting network can be rendered for human-viewing as a directed graph, with vertices representing agents or environment variables, edges describing information flow from one agent to another, and edge weights describing the listening weights for those connections. For any edge weight below a threshold, we do not include an edge and say the agent is “not listening”. This cutoff threshold is largely arbitrary, and is based on the level of noise, and therefore, average edge weight magnitude.

There are many local optima for social organizations, and several strategies are necessary to find the global optima. For each training iteration we evaluated with 1000 sets of random data for the environment values, we repeated training with several different algorithms (the most successful were Adaptive Delta Gradient-Descent and Root Mean Square Propagation), and restarted each trial multiple times with different seeds.

*Note: While the mathematical model supports agents with multiple actions ( $D_i$ ) and messages ( $F_i$ ), our implementation only has a single action and message per agent. This makes reasoning about results simpler, but introduces limitations that will be addressed later.*

#### 3.1 Implementation Caveat

In most neural network software, communication must be acyclic. That is, agent 1 cannot listen to agent 2 if agent 2 also listens to agent 1, since this creates a circular dependency. To remedy this problem, we define each agent as having multiple *layers*. The first layer of each agent is permitted only to listen to the environment. The second layer is permitted only to listen to the output of agents in layer 1. Each agent includes the observations of its counterpart in the layer below it in its own observations *without* noise or listen weights.

An arbitrary number of layers may be used, following this pattern. In the above mathematical model, we coalesce the layers of the agent in to a single agent  $i$  for convenience. As a result, all agents may listen to all other agents, and the environment, effectively making  $O_i$  a  $1 \times (N + K)$  matrix.

This is not an over-simplification, as the welfare of the model relies only on the actions of the outermost layer, and the listen and speaking costs of all layers, which will be correctly described by the mathematical model. The actions of intermediate layers may safely be ignored.

This caveat means there are two ways to view the resulting network: as an acyclic graph that shows the steps in the conversation between agents, or as

a “collapsed” and potentially-cyclic graph, which describes the interactions of an entire agent rather than a single layer of that agent.

### 3.2 Objective Variants

In addition to the trivial “average all environment nodes” objective, we created a number of alternative welfare functions to emphasize different aspects of human organizations.

#### 1. Selfish

In a human organization, one person doing 100 units of work is not equivalent to two people doing 50 units of work each. To disincentivize unnatural organizations where one person does all the work and distributes their answers to the rest of the group, we made the cost of listening to the environment exponential per agent. That is, there is now an incentive to distribute the labor of listening to the environment.

#### 2. Robustness

For long-lasting organizations, single points of failure should be avoided. In a human community, that single point of failure may get sick or be fired. In a computer network, that central system may crash. In this model we create  $N + 1$  versions of the network - one normal, and  $N$  in which one agent has been disabled and sends no messages.

Listen weights are consistent across these parallel network versions, but state and speaking weights may vary. For a real-world scenario, consider a company where “Fred” has just been fired. The rest of the organization cannot instantly schedule alternate meetings, restructure the hierarchy, and replace Fred, but they *can* stop waiting for any information they were supposed to receive from him.

The total welfare of the network is the sum of the welfare for each version. This means an optimal network will have no single points of failure, so each version of the network will perform satisfactorily.


#### 3. Binary Goals

In this model, half the agents are responsible for calculating the average of even environment variables, and half are responsible for odd variables. These varied goals mean a single information hierarchy is no longer optimal, since it would involve many unnecessary connections.

## 4 RESULTS

### 4.1 Trivial Objective

In the average-all-environment trials we see many “funnel” structures, where a single agent monitors all the environment variables, summarizes them, and distributes that information to all other agents. This minimizes the total number of edges in the graph, and means there is only a single speaker, lowering costs drastically.

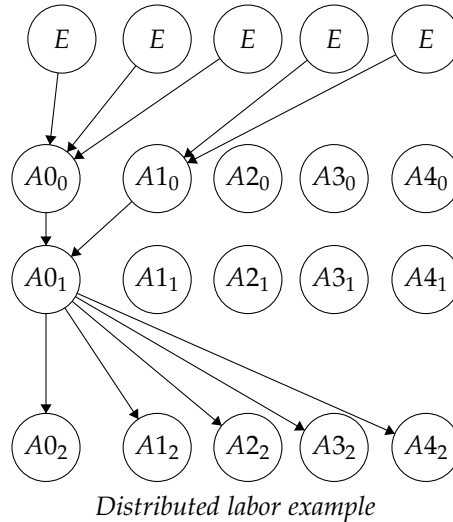


*The funnel structure*

As we increase the cost of listening or speaking to other agents we see a spontaneous shift where it becomes cheaper for all agents to monitor the environment directly, and they cease speaking to one another. Adjusting other variables (number of environment or agent nodes, or randomness of the environment state) has no effect, although increasing the cost of listening to the environment eventually leads to a critical scenario where it is cheaper to listen to nothing, and suffer the welfare consequences, then to observe the environment and get the right answer.

#### 4.2 *Selfish*

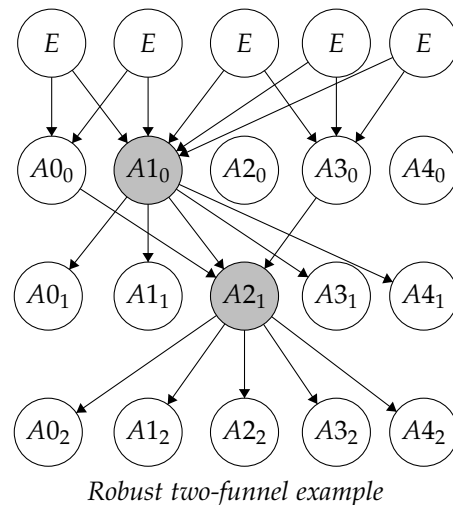
In the selfish model, there is an incentive *not* to put all of the environment load on a single agent. Instead, we see distribution of labor leading up to a funnel in the middle layer of communication.



Otherwise, the model is nearly-identical to the trivial objective, except that the cost threshold for switching all agents to reading the environment directly is much higher.

#### 4.3 *Robustness*

In the robustness trials there is an incentive not to have any single critical nodes, making the funnel structures from the trivial and selfish objectives unlikely. Instead, unsurprisingly, we see *two* funnels (highlighted below), so either can fail with no ill-effects.



In the above example, all agents receive their estimates in layer 1, from agent 1. However, if agent 1 is removed, agents will still receive estimates in layer 2, via the redundant agent 2 funnel.

#### 4.4 Binary Goals

Giving agents different goals yielded disastrous consequences. The welfare of the network fluctuated heavily, and the neural network produced *no* viable organizations. It appears that the Neural Network lowered all listening weights to near-zero to minimize costs, and agents proceeded to randomly guess at the environmental average by sampling the noise on the channels with minimal listen weights.

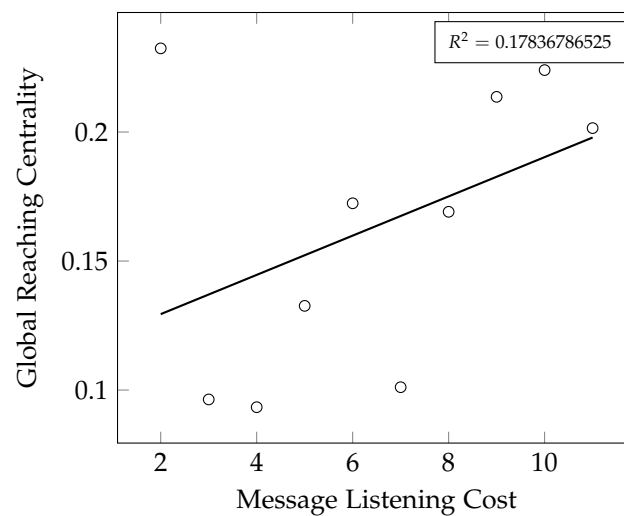
There are a number of potential explanations for this behavior. Since the organization begins with all agents listening to all environment variables, the “default” listen and speaking weights will create outgoing messages that combine information about all environment variables. This makes messages worthless for agents that only care about half of the environment.

Unfortunately, this theory does not explain why agents would not listen to the appropriate environment variables themselves. It is conceivable that since all the environment variables are drawn from a normal distribution centered on the same mean, there is no clear positive feedback to increasing listening weights for useful variables or decreasing weights on irrelevant variables. However, this only satisfies why agents would listen to the *wrong* variables, not why they would cease listening to the environment altogether. We continue to investigate this scenario.

#### 4.5 Parameter Sweeps


In addition to developing the above models, we sought to find a quantifiable correlation between different communications costs and optimal structure types. To crudely describe the hierarchy developed, we rely on a measure of *Global Reaching Centrality* [4]. In layman’s terms the GRC measures how far most agents are from the most connected agent, such that a low GRC is extremely unbalanced (one agent is much more connected than the others), and a high GRC is extremely balanced (all agents have a similar centrality score).

We conducted a parameter sweep, increasing the inter-agent message listening costs, environmental listening costs, and speaking costs, and measured the GRC of each generated network. These costs were adjusted by increasing the standard-distribution of noise for difference types of communication, thereby requiring a higher listen or speaking weight to convey information.




The most efficient model for a low listening cost network is the “funnel structure”, where one agent listens to the environment and synthesizes information, reducing costs for all other agents. However, as we increase the cost of listening to other agents the funnel network becomes cost-ineffective,

and agents switch to reading the environment variables directly. When this occurs all agents have the same centrality score, since none are speaking to one another.



Increasing the environmental listening cost further centralizes the network, by providing an added incentive to avoid duplication of labor. This guarantees a funnel structure, unless the environmental cost is so high that all agents cease listening to the environment (not pictured above).



Increasing the speaking cost has an almost identical effect to increasing environmental listening costs. In both cases there is an incentive to centralize work to a single agent, so only that agent needs to speak. The speaking cost graph differs slightly from environmental costs because speakers can address all other agents at once (a one-to-many connection), while listeners must listen to each speaker individually (one-to-one connection). This means that with the same level of noise on listening or speaking, speaking will always be cheaper.

Increasing the speaking cost to an extreme has the same effect as a high message listening cost: The agents cease speaking to each other and read from the environment directly.

## 5 CONCLUSION

Early experiments have been successful in making information flow graphs with minimal waste. When given inhibiting constraints, like a high cost to centralizing work, or a penalty for single points of failure, our neural network succeeded in designing optimal flow graphs with distributed labor and no single points of failure.

There is considerable work to be done with more complex scenarios, such as only requiring that the leaders of the organization receive specific information, or having leaders make decisions which must then be passed down through the organization. Some of these scenarios have been investigated by other researchers [2], but most research has focused on the effects of hierarchy rather than the communications constraints leading *to* hierarchy.

There are also many ways to expand our parameter sweeps. The Global Reaching Centrality is a non-ideal measurement, since a fully connected and fully disconnected network will have the same GRC, and the centrality of an agent does not capture some interesting phenomenon, such as a hub-agent with few connections that is critical for passing information. Future work may use measurements like Betweenness Centrality or Information Centrality [1] to better describe network structure, particularly for non-intuitive objective functions like robustness. These measurements need to be adapted, since many assume a single layer network, or handle multi-stage networks by either flattening the stages or treating each stage independently. Fortunately, there is some research in this area [3].

In addition, there are several limitations to reducing social organizations to information-flow diagrams. Our model leaves no room for differences in the personalities of agents, or constraints on the design of the organization imposed by outside rules, such as corporate tax law, or anti-union regulations. However, we believe this research will show there are fundamental benefits to specific hierarchies, and that these hierarchies will emerge “organically” without planning, as a consequence of information and labor constraints alone.

#### REFERENCES

- [1] Paolo Crucitti, Vito Latora, and Sergio Porta. Centrality measures in spatial networks of urban streets. *Phys. Rev. E*, 73:036125, Mar 2006.
- [2] Tamás Dávid-Barrett and RIM Dunbar. Cooperation, behavioural synchrony and status in social networks. *Journal of theoretical biology*, 308:88–95, 2012.
- [3] Xiaowen Dong, Pascal Frossard, Pierre Vandergheynst, and Nikolai Nefedov. Clustering with multi-layer graphs: A spectral perspective. *IEEE Transactions on Signal Processing*, 60(11):5820–5831, 2012.
- [4] Enys Mones, Lilla Vicsek, and Tamás Vicsek. Hierarchy measure for complex networks. *PloS one*, 7(3):e33799, 2012.