A standard Schelling model

Wan Dang, Adithya Prabha, Joanne Romo, and Ben Sarr, mentored by Dr. Michael Jarret and Anthony E. Pizzimenti

What's a Schelling model?

A ***Schelling model*** studies how local preferences can lead to global segregation in a social system [1]. Many variants of this ***agent-based model*** exist, each with different vocabulary and underlying assumptions, making them hard to compare directly [2].

In this project, ***we develop a standard framework for Schelling-type models using MCMC (Monte Carlo Markov chain)***.

Re-imagining agent moves as updates to a graph coloring, we argue that ***most "Schelling" models are just different versions of the same random process.***

Definitions and setup

We work on a finite graph $G = (V, E)$ where each vertex is assigned one of q labels from $[q] = \{0, \ldots, q-1\}$. A ***configuration*** is a map $f : V \to [q]$. The ***demography*** is the vector $\varepsilon = (\varepsilon_0, \ldots, \varepsilon_{q-1})$, where ε_i is the proportion of vertices with label i. The global ***unhappiness*** is

$$\mathbf{H}(f) = -\sum_{v \in V} c(v),$$

where $c(v)$ counts how many of v's neighbors have a different label than v.

Our goal is to ***simulate a Markov chain on all configurations with the same demography***.

A basic update step in the Markov chain changes a configuration f by swapping labels.

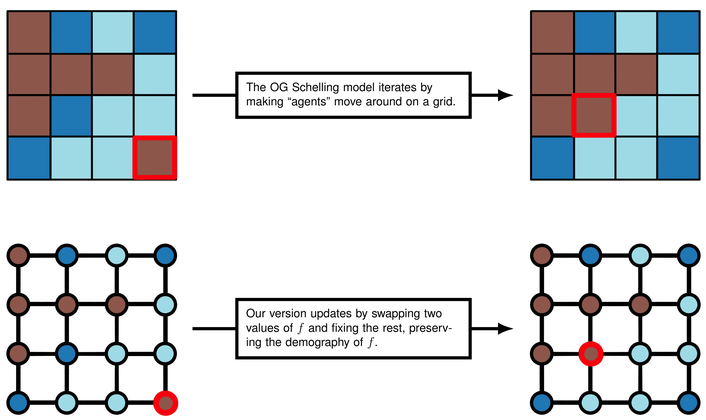

Our setup does the same thing as the classical "agent-based" model by representing local changes ("agent moves") as global ones (changes to configurations).

Algorithmic implementation

We use a ***Metropolis-Hastings strategy*** to simulate a Markov chain on the space of configurations [3]. Given a proposal distribution φ, an acceptance distribution α, and an initial configuration f_0, set $t := 0$. Then,

Algorithm 1 *(Schelling dynamics).*

1. Sample a candidate configuration f^* from $\varphi(- \mid f_t)$.
2. Sample $\beta \in [0, 1]$ uniformly. Then,

$$f_{t+1} := \begin{cases} f^* & \beta < \min\left\{1, \frac{\varphi(f^* \mid f_t)}{\varphi(f_t \mid f^*)} \frac{\alpha(f^* \mid f_t)}{\alpha(f_t \mid f^*)}\right\} \\ f_t & \text{otherwise,} \end{cases}$$

   which chooses f_{t+1} according to the ***Metropolis-Hastings criterion***.
3. If a stopping condition is met, go to Step **4**. Otherwise, set $t := t + 1$ and go to Step **1**.
4. Report f_t.

Algorithm 1 is flexible enough that ***most variants of the classical Schelling model can be captured just by modifying*** $\mathbf{H}$***,*** φ***, and*** α. Here, φ is an exponential distribution proportional to $e^{-\zeta \mathbf{H}(f)}$.

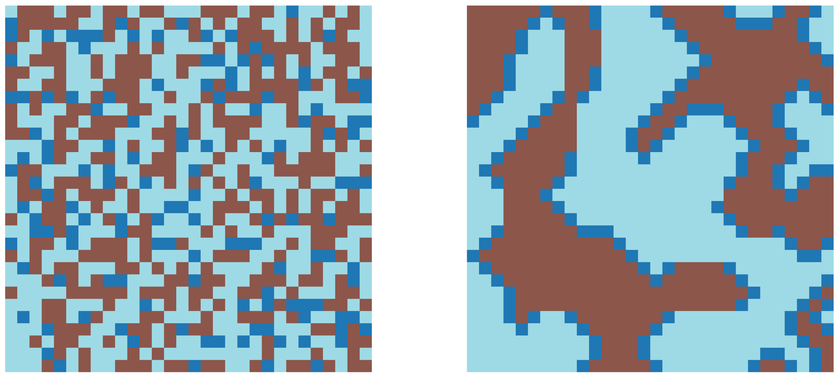

Figure 2. Initial (left) and final (right) configurations after ten thousand iterations of Algorithm 1.

Figure 2 shows how ***configurations rearrange to decrease*** $\mathbf{H}$ ***over time, resulting in segregated clusters of distinct groups***.

MCMC Dynamics

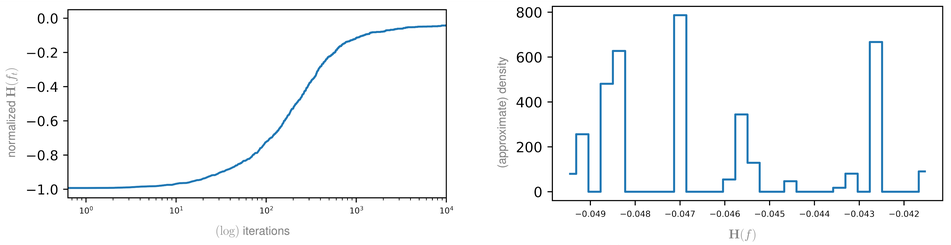

Figure 3. Using data from the run in Figure 2: normalized $\mathbf{H}(f_t)$ against log-scaled time (left); the distribution of $\mathbf{H}(f_t)$, truncated to the last five thousand iterations (right).

Since φ is proportional to $e^{-\zeta \mathbf{H}(f_t)}$, we expect to see exponential decay in $\mathbf{H}$ as t grows. Figure 3 shows just that: $\mathbf{H}$ exponentially decreases and then stabilizes once the Markov chain finds a local minimum.

Error Analysis and Stability

For each pair of configurations (f_t, f^*), the ***transition matrix*** T specifies the probability $T(f_t, f^*)$ of our Markov chain accepting f^* as the next configuration, given f_t is the current one. To study how small perturbations in T affect long-term behavior, let's examine the k-step difference between the original T and its perturbed version T_ε. Using a standard telescoping identity, we can get the expansion below:

$$T^k - T_\varepsilon^k = \sum_{i=0}^{k-1} T^{k-1-i}(T - T_\varepsilon)T_\varepsilon^i,$$

$$\|AB\| \le \|A\|\|B\|, \qquad \text{(triangle inequality)}$$

$$\|T^k - T_\varepsilon^k\| \le \sum_{i=0}^{k-1} \|T^{k-1-i}\|\|T - T_\varepsilon\|\|T_\varepsilon^i\|.$$

Since both T and T_ε are transition matrices (so their rows sum to 1), $\|T^k\| \le 1$ and $\|T_\varepsilon^k\| \le 1$ hold for all $k \ge 0$. Therefore,

$$\|T^k - T_\varepsilon^k\| \le \sum_{i=0}^{k-1} 1 \cdot \|T - T_\varepsilon\| \cdot 1.$$

$$\boxed{\|T^k - T_\varepsilon^k\| \le k\|T - T_\varepsilon\|.}$$

This bound shows that ***small perturbations to the transition probabilities have, at worst, a linear effect in the long run.***

Future Work

- Generalize algorithms to different graph structures
- Improve/get tighter bounds on stability
- Start writing up the paper!

References

[1] Thomas C. Schelling. "Dynamic models of segregation". In: ***The Journal of Mathematical Sociology*** 1.2 (1971).

[2] Adam Douglas Henry et al. "Emergence of segregation in evolving social networks". In: ***Proceedings of the National Academy of Sciences*** 108.21 (May 2011).

[3] Olle Häggström. ***Finite Markov Chains and Algorithmic Applications***. London Mathematical Society Student Texts. Cambridge: Cambridge University Press, 2002. ISBN: 978-0-521-89001-4.