A neighborhood search heuristic for the p-median problem with continuous demand

Shane Auerbach and Rebekah Dix

Lyft and Wisconsin, respectively

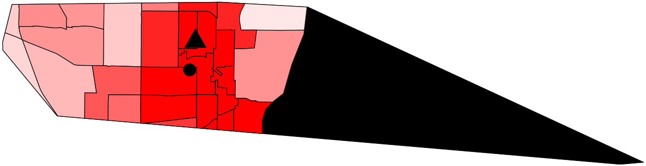

The Voronoi cell of Super Fresh Market in Waukegan, IL

The standard p-median problem (Hakimi, 1964)

Essentially a weighted (Fermat-)Weber problem.

Inputs and decision variable:

$$\begin{aligned} \mathcal{N} &= \text{a set of nodes in a network} \\ w_m &= \text{demand at node } m \in \mathcal{N} \\ d(m,n) &= \text{distance between node } m \in \mathcal{N} \text{ and node } n \in \mathcal{N} \\ p &= \text{number of facilities to locate} \end{aligned}$$

$$s = (s_1, s_2, \ldots, s_p) = \text{choice of } p \text{ nodes at which facilities are to be located}$$

Problem:

$$\underset{s}{\text{Minimize}}\ z = \sum_{m \in \mathcal{N}} w_m \cdot \min_k d(m, s_k)$$

The p-median problem with continuous demand

Inputs and decision variable:

$$
\begin{aligned}
X \subseteq \mathbb{R}^2 &\text{ is a subset of Euclidean space} \\
f : X \to \mathbb{R}_{\geq 0} &\text{ is an integrable consumer density function} \\
d(m, n) &= \text{distance between } m, n \in X \\
p &= \text{number of facilities to locate}
\end{aligned}
$$

$$
s = (s_1, s_2, \ldots, s_p) = \text{choice of } p \text{ locations for facilities (each } s_k \in X)
$$

Problem:

$$
\operatorname*{Minimize}_{s} \; z = \int_X f(x) \cdot \min_k d(x, s_k) \, \mathrm{d}x
$$

Why continuous demand?

Two settings with continuous demand:

- **Uncertainty:** Ride-sharing service wants to place its idle drivers so as to minimize wait-times for passengers.
  - See Auerbach and Dix (2018).
- **Spatially aggregated data:** Census data tells us which tract respondents are in, but not precise location.
  - We work with a census application in this paper.

Why not discretize?

- Introduces error that is hard to quantify. See Francis et al. (2009).
  - Tradeoff between computability and accuracy.
  - With exogenous spatial discretization, sparsity issues.

Neighborhood search on a network (Maranzana, 1963)

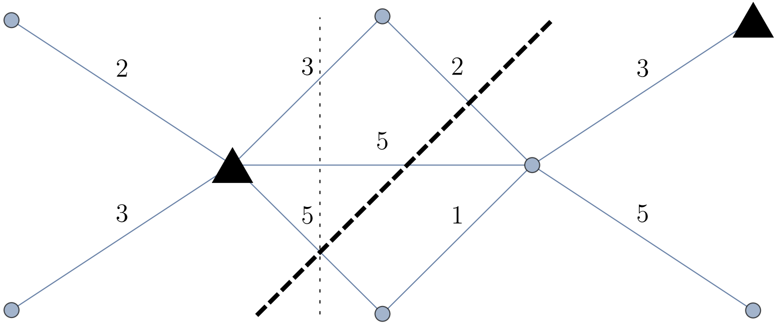

Neighborhood search on the plane

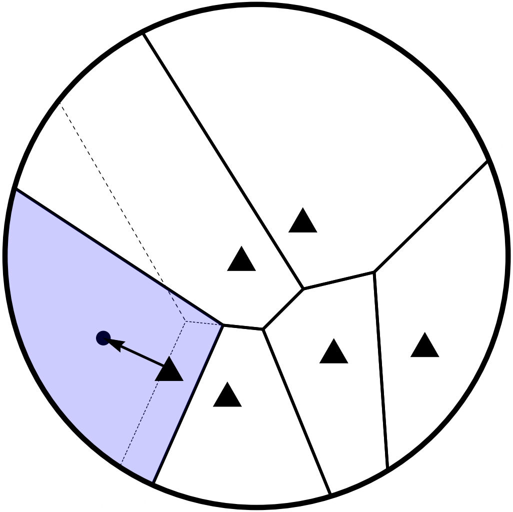

Heuristic

Procedure 1 Simple neighborhood search heuristic (SNS)

function SNS(X, f, s, tolerance) ▷ s is an arbitrary initial allocation
    complete $= 0$
    while complete $= 0$ **do**
        for $i = 1$ to I **do**
            Compute $\text{Vor}(s)$
            $s_i' = \arg\min_{y \in V_i(s)} \int_{V_i(s)} f(x) \cdot d(x, y) \, \mathrm{d}x$ ▷ Find optimal adjustment
            $\Delta_i = d(s_i', s_i)$ ▷ Measure adjustment size
            $s_i = s_i'$ ▷ Implement adjustment
        end for
        if $\max_i \Delta_i <$ tolerance **then**
            complete $= 1$, ▷ Terminate when adjustment sizes are below tolerance
        end if
    end while
end function

Theorem: Procedure 1 converges to an approximate local minimum of the *p*-median problem with continuous demand.

Each adjustment improves the objective (proof sketch)

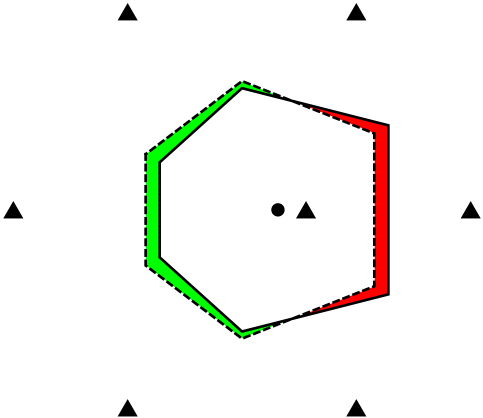

Similar logic Lloyd's algorithm (Lloyd, 1982)!

Challenges and solutions

Finding the geometric median of a weighted polygon is hard.

- Draw points randomly from the polygon based on f and find the geometric median of the set of points.

The set of local minima can be large and diverse.

- Start from a variety of initializations and take best local minimum.
  - Use a greedy algorithm to generate a good initialization.
- Include local search in hybrid heuristic with global adjustment.

Heuristic can be computationally demanding.

- Compute adjustments for all features before recomputing the Voronoi cells, making the heuristic *embarrassingly parallelizable*.

Application: Measuring spatial inefficiency

Spatial inefficiency in consumer travel times

- Suppose spatial inefficiency is proportional to the average Euclidean distance between a customer and her nearest firm/feature.
- $\bar{d}(s)$ and $\bar{d}(s^*)$ are the average distances under the observed and optimized allocations, respectively.
- $\xi(s) = (\bar{d}(s) - \bar{d}(s^*))/\bar{d}(s^*)$, i.e. a % difference in averages.

We evaluate ξ for supermarkets, hospital, and fire stations.

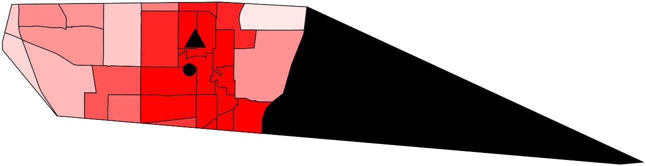

The Voronoi cell of Super Fresh Market in Waukegan, IL

Supermarkets, fire stations, and hospitals in cities

To work in 2D and account for customer density directly, we look at allocations of firms/features in cities.

Summary statistics on regions and features therein

Region	Sqmi	Pop. (M)	Tracts	S	H	F
(1) Atlanta City (AC)	163	0.47	130	19	4	38
(2) Atlanta MSA (AM)	8835	5.53	951	174	37	502
(3) Chicago City (CC)	227	2.78	803	76	25	96
(4) Chicago MSA (CM)	6304	9.56	2210	259	80	759
(5) LA City (LC)	536	4.44	1109	100	29	137
(6) LA MSA (LM)	4754	13.14	2923	277	86	535

S, **H**, and **F** are numbers of supermarkets, hospitals and fire stations, respectively.

Results

Region	Supermarkets		Hospitals		Fire Stations	
	$\bar{d}(s^{\text{act}})$	$\bar{d}(s^{*})$	$\bar{d}(s^{\text{act}})$	$\bar{d}(s^{*})$	$\bar{d}(s^{\text{act}})$	$\bar{d}(s^{*})$
(1) AC	1.90	1.01	3.47	2.24	0.86	0.71
(2) AM	3.59	2.01	5.13	4.24	1.47	1.18
(3) CC	1.09	0.60	1.49	1.06	0.61	0.54
(4) CM	2.40	1.39	3.06	2.34	0.90	0.75
(5) LC	1.43	0.73	2.06	1.36	0.75	0.63
(6) LM	1.77	1.00	2.30	1.77	0.82	0.68

Region	$\xi(s^{\text{act}})$		
	S	**H**	**F**
(1) Atlanta City	87	55	22
(2) Atlanta MSA	79	21	25
(3) Chicago City	82	41	14
(4) Chicago MSA	73	30	21
(5) LA City	96	51	19
(6) LA MSA	78	30	20

Spatial inefficiency in consumer travel costs is significantly larger for supermarket allocations, those generated by a competitive mechanism.

Visualizing supermarkets in Chicago MSA (1)

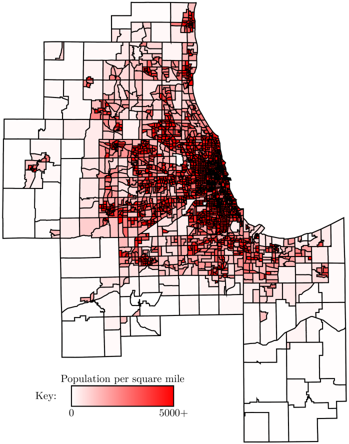

Census tracts and pop. density

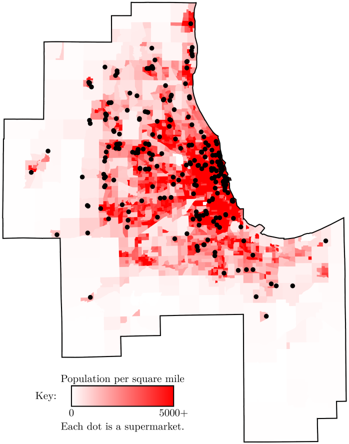

Supermarkets and pop. density

Supermarkets follow population density roughly, but cluster.

Visualizing supermarkets in Chicago MSA (2)

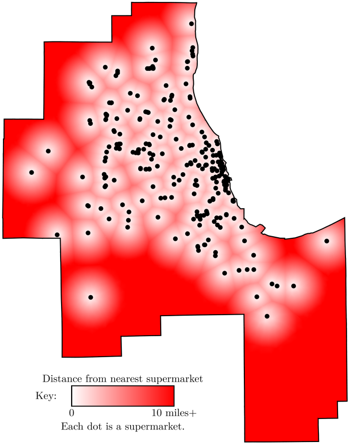

Distance to supermarket in s^{act}

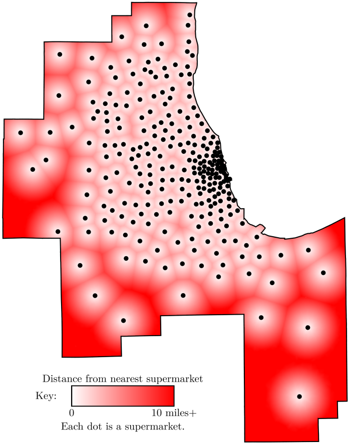

Distance to supermarket in s^*

s^* declusters supermarkets and moves them outwards.

Visualizing supermarkets in Chicago MSA (3)

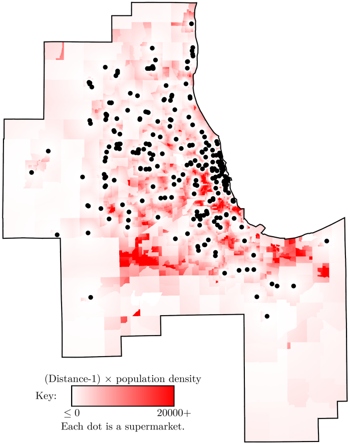

Supermarket scarcity in s^{act}

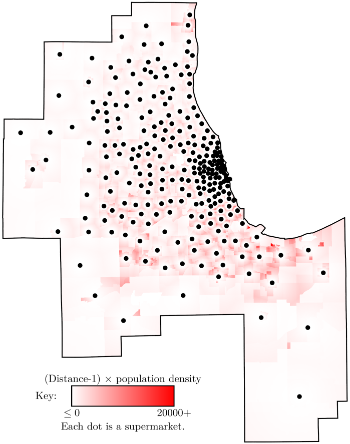

Supermarket scarcity in s^*

s^* significantly reduces spatial scarcity.

Conclusion and extensions

Conclusions:

- Use local neighborhood search (in a hybrid heuristic) for the p-median problem with continuous demand.
- Does competition yield good spatial allocations?
  - Supermarkets don't look great.
  - In our ridesharing work (Auerbach and Dix, 2018), results of competition look more positive.

Extensions:

- What about when facilities have capacities?
- A dynamic model for facility repositioning.
  - Then may be relevant to scooter or AV allocation.