


Epistemic Models of Voting – Part 2

Paul Gözl

Lecture 6, ORIE 4360/5360, Fall '26



Recap: Condorcet's Jury Theorem

2

Each of n voters independently votes for a with probability $1/2 < p \leq 1$ or for b with probability $1 - p$.



Nicolas de
Condorcet
(1743–1794)

Theorem [Condorcet 1785]. As $n \rightarrow \infty$, the probability that the majority is correct goes to 1.

Recap: Condorcet's Model

3

- (1) In every pairwise comparison, each voter chooses the better alternative with some fixed probability p , where $1/2 < p \leq 1$.
- (2) Each voter's judgment on every pair of alternatives is independent of her judgment on every other pair.
- (3) Each voter's judgment is independent of the other voters' judgments.
- (4) Each voter's judgment produces a ranking of the alternatives.

contradictory! ⚡

Recap: Fix 1 for Condorcet's Model

- (4) ~~Each voter's judgment produces a ranking of the alternatives.~~

Each voter independently guesses each pairwise comparison, and we collect these pairwise preferences $\succsim_i$ even if the voter is inconsistent.

(So, just for today, $\succsim_i$ can be nontransitive, e.g., $a \succsim_i b \succsim_i c \succsim_i a$. Just denote which alternative is preferred in each pair.)

Kemeny's Rule for Fix 1

For given preferences $\{ \succ_i \}_{i \in N}$ and p ,

$$\mathbb{P}[\{ \succ_i \}_{i \in N} \mid \pi] = \frac{1}{\text{some terms}} \left(\frac{1-p}{p} \right)^{\# \text{disagreements}(\pi, \{ \succ_i \}_{i \in N})}$$

don't depend on π > 0

→ MLE is π with largest $\left(\frac{1-p}{p} \right)^{\# \text{disagreements}(\pi, \{ \succ_i \}_{i \in N})}$.

→ For $\frac{1}{2} < p < 1$, this is the π with minimum $\# \text{disagreements}(\pi, \{ \succ_i \}_{i \in N})$, the **Kemeny ranking**.

Reflection on Fix 1

We saw that the Kemeny rule can be justified as the MLE under Fix 1 to Condorcet's model.

But this framing fits badly with our voting model because we allowed voters to give non-ranking preferences. Let's see Young's second approach!

Kendall Tau Distance

7

Given rankings σ, σ' over alternatives, their **Kendall Tau distance** counts how many pairs they rank in opposite order:

$$d_{KT}(\sigma, \sigma') := \left| \left\{ \{x, y\} \subseteq M \mid (x \succ_{\sigma} y \text{ and } x \prec_{\sigma'} y) \text{ or } (x \prec_{\sigma} y \text{ and } x \succ_{\sigma'} y) \right\} \right|$$

Kendall Tau Distance

$$d_{KT}(\sigma, \sigma') := \left| \left\{ \{x, y\} \subseteq M \mid (x \succ_{\sigma} y \text{ and } x \prec_{\sigma'} y) \text{ or } (x \prec_{\sigma} y \text{ and } x \succ_{\sigma'} y) \right\} \right|$$

$\sigma : c \succ_{\sigma} b \succ_{\sigma} e \succ_{\sigma} a \succ_{\sigma} d$

$\sigma' : a \succ_{\sigma'} b \succ_{\sigma'} d \succ_{\sigma'} c \succ_{\sigma'} e$



What is the Kendall Tau distance between σ and σ' ?

pollev.com/goelz

Kendall Tau Distance

$$d_{KT}(\sigma, \sigma') := \left| \left\{ \{x, y\} \subseteq M \mid (x \succ_{\sigma} y \text{ and } x \prec_{\sigma'} y) \text{ or } (x \prec_{\sigma} y \text{ and } x \succ_{\sigma'} y) \right\} \right|$$

Q: Which ranking has the highest Kendall Tau distance from $a \succ b \succ c \succ d \succ e$? Give this maximal Kendall Tau distance.

Fix 2 for Condorcet's Model

10

- (1) In every pairwise comparison, each voter chooses the better alternative with some fixed probability p , where $1/2 < p \leq 1$.
- (2) ~~Each voter's judgment on every pair of alternatives is independent of her judgment on every other pair.~~
- (3) Each voter's judgment is independent of the other voters' judgments.
- (4) Each voter's judgment produces a ranking of the alternatives.


contradictory! ⚡

Fix 2: Mallows Model

11

Independently for each voter: independently guess each pairwise comparison, get it right with probability $1/2 < p \leq 1$. If result is a ranking, done; else, start over (i.e., condition.)

Say $\pi : a \succ b \succ c \succ d$,
and we're determining σ_i .




Fix 2: Mallows Model

12

Independently for each voter: independently guess each pairwise comparison, get it right with probability $1/2 < p \leq 1$. If result is a ranking, done; else, start over (i.e., condition.)

Say $\pi : a \succ b \succ c \succ d$,
and we're determining σ_i .

Result: $b \succ_{\sigma_i} c \succ_{\sigma_i} a \succ_{\sigma_i} d$.



A Ranking's Likelihood in Mallows

Let π be true ranking. How likely does ranking σ_i result from Mallows model?

$$\mathbb{P}[\sigma_i \mid \pi] = \frac{\text{get } \sigma_i \text{ without restarting}}{\text{flipping coins yields valid ranking}}$$

A Ranking's Likelihood in Mallows

14

Let π be true ranking. How likely does ranking σ_i result from Mallows model?

$$\mathbb{P}[\sigma_i \mid \pi] = \frac{\begin{array}{c} = \binom{m}{2} - d_{KT}(\pi, \sigma_i) \\ \downarrow \\ p^{\text{\#agree}} \cdot (1 - p)^{\text{\#disagree}} \end{array}}{\text{flipping coins yields valid ranking}}$$

A Ranking's Likelihood in Mallows

Let π be true ranking. How likely does ranking σ_i result from Mallows model?

$$\mathbb{P}[\sigma_i \mid \pi] = \frac{p^{\binom{m}{2} - d_{KT}(\pi, \sigma_i)} \cdot (1 - p)^{d_{KT}(\pi, \sigma_i)}}{\text{flipping coins yields valid ranking}}$$

A Ranking's Likelihood in Mallows

16

Let π be true ranking. How likely does ranking σ_i result from Mallows model?

$$\mathbb{P}[\sigma_i \mid \pi] = \frac{p^{\binom{m}{2} - d_{KT}(\pi, \sigma_i)} \cdot (1 - p)^{d_{KT}(\pi, \sigma_i)}}{\sum_{\sigma'_i \in \mathcal{L}} p^{\binom{m}{2} - d_{KT}(\pi, \sigma'_i)} \cdot (1 - p)^{d_{KT}(\pi, \sigma'_i)}}$$

by symmetry, does not depend on π !

A Ranking's Likelihood in Mallows

17

Let π be true ranking. How likely does ranking σ_i result from Mallows model?

$$\mathbb{P}[\sigma_i \mid \pi] = \frac{p^{\binom{m}{2} - d_{KT}(\pi, \sigma_i)} \cdot (1 - p)^{d_{KT}(\pi, \sigma_i)} \cdot p^{-\binom{m}{2}}}{\text{some terms} \cdot p^{-\binom{m}{2}}}$$

don't depend on π > 0

A Ranking's Likelihood in Mallows

18

Let π be true ranking. How likely does ranking σ_i result from Mallows model?

$$\mathbb{P}[\sigma_i \mid \pi] = \frac{\left(\frac{1-p}{p}\right)^{d_{KT}(\pi, \sigma_i)}}{\text{some terms}}$$

don't depend on π > 0 doesn't depend on π > 0

$\cdot p^{-\binom{m}{2}}$

A Profile's Likelihood in Mallows

19

What is the likelihood of observing an entire profile $\sigma = (\sigma_1, \dots, \sigma_n)$?

$$\mathbb{P}[\sigma \mid \pi] = \prod_{i=1}^n \frac{\left(\frac{1-p}{p}\right)^{d_{KT}(\pi, \sigma_i)}}{\text{some terms}} = \frac{\left(\frac{1-p}{p}\right)^{\sum_{i=1}^n d_{KT}(\pi, \sigma_i)}}{\text{some terms}}$$

don't depend on π > 0 don't depend on π > 0

MLE of Mallows Model

$$\mathbb{P}[\boldsymbol{\sigma} \mid \pi] \propto \left(\frac{1-p}{p} \right)^{\sum_{i=1}^n d_{KT}(\pi, \sigma_i)}.$$

Monotone decreasing in $\sum_{i=1}^n d_{KT}(\pi, \sigma_i)$ for $1/2 < p < 1$.

Hence, MLE again minimizes disagreements:

$$f_{\text{Kemeny}}(\boldsymbol{\sigma}) := \operatorname{argmin}_{\pi'} \sum_{i=1}^n d_{KT}(\pi', \sigma_i).$$

MLE of Mallows Model

Theorem [Young 1988]. π is the MLE for a preference profile σ under the Mallows model if and only if it is a Kemeny ranking.

Proof. We just did it 😊

Compute the Kemeny Ranking!

22

8x	9x	2x	1x	3x	2x
a	b	b	c	d	d
c	c	d	d	a	b
d	a	a	a	b	a
b	d	c	b	c	c

Kemeny score of π : $\sum_{i=1}^n d_{KT}(\pi, \sigma_i)$



Q: On the profile above, what is the lowest Kemeny score you can find a ranking π for?

pollev.com/goelz

Some Tricks for Computing Kemeny

23

$\sum_{i=1}^n d_{KT}(\pi, \sigma_i) = \sum_{x \succ_{\pi} y} \left| \{i \in N \mid x \prec_{\sigma_i} y\} \right|$, so
can consider again a matrix of head-to-head
election votes.

8x	9x	2x	1x	3x	2x		a	b	c	d
a	b	b	c	d	d	a	-	12	15	17
c	c	d	d	a	b	b	13	-	16	11
d	a	a	a	b	a	c	10	9	-	18
b	d	c	b	c	c	d	8	14	7	-

$$\min \left\{ \left| \{i \in N \mid x \succ_{\sigma_i} y\} \right|, \left| \{i \in N \mid x \prec_{\sigma_i} y\} \right| \right\}$$

disagreements, can subtract them out.


At least $12+10+9+8+11+7=57$ disagreements.

Some Tricks for Computing Kemeny

25


Represent as a **margin graph**. Edge weights give additional disagreements π pays for reversing.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>a</i>	-	0	5	9
<i>b</i>	1	-	7	0
<i>c</i>	0	0	-	11
<i>d</i>	0	3	0	-



$\rightarrow = \succ_{\text{maj}}$

Any ranking π must reverse at least one edge in any cycle.




Any π must have ≥ 3 more disagreements.

Some Tricks for Computing Kemeny

26

Represent as a **margin graph**. Edge weights give additional disagreements π pays for reversing.

	<i>a</i>	<i>b</i>	<i>c</i>	<i>d</i>
<i>a</i>	-	0	5	9
<i>b</i>	1	-	7	0
<i>c</i>	0	0	-	11
<i>d</i>	0	3	0	-



If we reverse $d \rightarrow b$ edge, get a ranking $(b \succ a \succ c \succ d)$, with only 3 further disagreements.

Kemeny score = $57 + 3 = 60$.

Lowest possible, so a Kemeny ranking.

$\rightarrow = \succ_{\text{maj}}$

Kemeny Is Condorcet Consistent

27

Can use this machinery to prove:

Theorem. If there is a Condorcet winner, some Condorcet winner is ranked first by Kemeny.

Kemeny Is Condorcet Consistent

28

Theorem. If there is a Condorcet winner, some Condorcet winner is ranked first by Kemeny.


Proof. Only prove for odd n for simplicity. For contradiction, let a be Condorcet winner and b be higher in Kemeny ranking.

In margin graph, a only has outgoing edges (Condorcet winner). $\rightarrow = \succ_{\text{maj}}$
Consider ordering according to Kemeny ranking, reverses some majority edges (↻). Edge from b to a .
 $\rightarrow = \succ_{\text{Kemeny}}$

Kemeny edges among $M \setminus \{a\}$ acyclic.

If we make all of a 's edges outgoing, still acyclic and have a ranking, reverse subset of edges, better score.

Better Kemeny score, contradiction.




That Being Said...

Computing the Kemeny ranking is again “NP-hard”, which means that we don’t think there are any efficient algorithms for finding it.

Up Ahead

30

In the next lectures, we will bring the epistemic approach to liquid democracy...



and to non-voting applications including chess and LLM training.



Bibliography

Elkind & Slinko: *Rationalizations of Voting Rules*, which is Chapter 8 in Brandt et al. (2016): *Handbook of computational social choice*.

Young: *Condorcet's Theory of Voting*. American Political Science Review (1988).