Simultaneous recovery of the consensus and structure of permutations

Marina Meilă<br>University of Washington<br>with Chris Meek, Microsoft Research

UW Theory Seminar 2/2/16

## The "Sushi preference" data

## $N=5000$ people ranked $n=12$ types of sushi

sake |ebi |ika |uni |tamago |kappa-maki |tekka-maki |anago |toro |maguro ebi |kappa-maki |tamago |ika |toro |maguro |tekka-maki |anago |sake |uni toro |ebi |maguro |ika |tekka-maki |uni |sake |anago |kappa-maki |tamago tekka-maki |tamago |sake |ebi |ika |kappa-maki |maguro |toro |uni |anago tamago |maguro |kappa-maki |ebi |sake |anago |uni |tekka-maki |toro |ika uni |toro |ebi |anago |maguro |tekka-maki |ika |sake |kappa-maki |tamago maguro |ika |toro |tekka-maki |ebi |uni |sake |tamago |anago |kappa-maki

Consensus Ranking Problem
Given a set of rankings $\left\{\pi_{1}, \pi_{2}, \ldots \pi_{N}\right\} \subset \mathbb{S}_{n}$ find the consensus ranking $\pi_{0}$ such that

$$
\pi_{0}=\underset{\mathbb{S}_{n}}{\operatorname{argmin}} \sum_{i=1}^{N} d\left(\pi_{i}, \pi_{0}\right)
$$

for $d=$ inversion distance / Kendall $\tau$-distance / "bubble sort" distance

## The "Sushi preference" data

## $N=5000$ people ranked $n=12$ types of sushi

sake |ebi |ika |uni |tamago |kappa-maki |tekka-maki |anago |toro |maguro ebi |kappa-maki |tamago |ika |toro |maguro |tekka-maki |anago |sake |uni toro |ebi |maguro |ika |tekka-maki |uni |sake |anago |kappa-maki |tamago tekka-maki |tamago |sake |ebi |ika |kappa-maki |maguro |toro |uni |anago tamago |maguro |kappa-maki |ebi |sake |anago |uni |tekka-maki |toro |ika uni |toro |ebi |anago |maguro |tekka-maki |ika |sake |kappa-maki |tamago maguro |ika |toro |tekka-maki |ebi |uni |sake |tamago |anago |kappa-maki

Consensus Ranking Problem
Given a set of rankings $\left\{\pi_{1}, \pi_{2}, \ldots \pi_{N}\right\} \subset \mathbb{S}_{n}$ find the consensus ranking $\pi_{0}$ such that

$$
\pi_{0}=\underset{\mathbb{S}_{n}}{\operatorname{argmin}} \sum_{i=1}^{N} d\left(\pi_{i}, \pi_{0}\right)
$$

for $d=$ inversion distance / Kendall $\tau$-distance / "bubble sort" distance

This problem is NP-hard []

## Related work

Consensus Ranking/Single parameter/Mallows model
[Cohen,S,Singer 99] CSS ALGORITHM $=$ greedy search on $Q$
improved by extracting strongly connected components
[Ailon,Newman,Charikar 05] Randomized algorithm guaranteed 11/7 factor approximation (ANC)
[Mohri, Ailon 08] linear program
[Mathieu, Schudy 07] $(1+\epsilon)$ approximation, time $\mathcal{O}\left(n^{6} / \epsilon+2^{2^{O(1 / \epsilon)}}\right)$
[Davenport,Kalagnanan 03] Heuristics based on edge-disjoint cycles used by our B\&B implementation
[Conitzer,D,K 05] Exact algorithm based on integer programming, better bounds for edge disjoint cycles (DK)
[Betzler,Brandt, 10] Exact problem reductions
[Awasthi,Blum,Sheffet,Vijayaraghavan 14]

- Most of this work based on the MinFAS view

$$
Q_{i j}>.5 \Leftrightarrow i \bullet \xrightarrow{Q_{i j}-.5} \bullet j
$$

Prune graph to a DAG removing minimum weight

## Extensions and applications to social choice

## Social choice

- Inferring rakings under partial and aggregated information [ShahJabatula08], [JabatulaFariasShah10]
- Vote elicitation under probabilistic models of choice [LuBoutillier11]
- Voting rules viewed as Maximum Likelihood [ConitzerSandholm08]
- Algorithms guaranteed to retrive certain "winners" [LinAgarwal14] "Noisy sorting"
- Using Hodge decompositions and L1, L2 distances [JiangLimYaoYe11]
- Noisy comparison [BravermanMossel08]

ML Estimation/Multiple parameters/GM model
I[VlignerVerducci 86] $\vec{\theta}$ estimation; heuristic for $\pi_{0}$

## FV algorithm/Borda rule

1. Compute $\bar{q}_{j}, j=1: n$ column sums of $Q$
2. Sort $\left(\bar{a}_{j}\right)_{j=1}^{n}$ in increasing order; $\pi_{0}$ is sorting permutation

- $\bar{q}_{j}$ are Borda counts
- FV is consistent for infinite $N$



## Generalizing consensus ranking

- Not all inversions are equally important


## Sushi preferences for uni have no consensus

sake |ebi |ika |uni |tamago |kappa-maki |tekka-maki |anago |toro |maguro ebi |kappa-maki |tamago |ika |toro |maguro |tekka-maki |anago |sake |uni toro |ebi |maguro |ika |tekka-maki |uni |sake |anago |kappa-maki |tamago tekka-maki |tamago |sake |ebi |ika |kappa-maki |maguro |toro |uni |anago tamago |maguro |kappa-maki |ebi |sake |anago |uni |tekka-maki |toro |ika uni |toro |ebi |anago |maguro |tekka-maki |ika |sake |kappa-maki |tamago maguro |ika |toro |tekka-maki |ebi |uni |sake |tamago |anago |kappa-maki

## Generalizing consensus ranking

- Not all inversions are equally important
... but there is consensus for maguro (tuna) and tekka-maki (tuna roll) sake |ebi |ika |uni |tamago |kappa-maki |tekka-maki |anago |toro |maguro ebi |kappa-maki |tamago |ika |toro |maguro |tekka-maki |anago |sake |uni toro |ebi |maguro |ika |tekka-maki |uni |sake |anago |kappa-maki |tamago tekka-maki |tamago |sake |ebi |ika |kappa-maki |maguro |toro |uni |anago tamago |maguro |kappa-maki |ebi |sake |anago |uni |tekka-maki |toro |ika uni |toro |ebi |anago |maguro |tekka-maki |ika |sake |kappa-maki |tamago maguro |ika |toro |tekka-maki |ebi |uni |sake |tamago |anago |kappa-maki


## Generalizing consensus ranking

- Not all inversions are equally important introduce importance/weight parameters $\vec{\theta}$

Irish College Admissions data
Parameters of top 10 ranks in the 33 largest clusters found


Combinatorial structure present


- described by a tree


## Recursive Inversion Models (RIM)

[Meek, M 14]

$\tau=$ tree structure
$\pi_{0}(\tau)=$ induced central ranking
$\theta_{1: n-1}=$ parameters at nodes
Recursive Inversion Models (RIM)
[Meek, M 14]

$\tau=$ tree structure
$\pi_{0}(\tau)=$ induced central ranking
$\theta_{1: n-1}=$ parameters at nodes
Inversions are penalized by $\theta_{i}$ parameters
Example: $\vec{\theta}=(0.1,1.2,0.4)$
$\operatorname{Cost}(a|b| c \mid d)=0$

## Recursive Inversion Models (RIM)

[Meek, M 14]

$\tau=$ tree structure
$\pi_{0}(\tau)=$ induced central ranking
$\theta_{1: n-1}=$ parameters at nodes
Inversions are penalized by $\theta_{i}$ parameters
Example: $\vec{\theta}=(0.1,1.2,0.4)$

$$
\begin{aligned}
\operatorname{Cost}(a|b| c \mid d) & =0 \\
\operatorname{Cost}(b|a| c \mid d) & =1.2
\end{aligned}
$$

## Recursive Inversion Models (RIM)

[Meek, M 14]

$\tau=$ tree structure
$\pi_{0}(\tau)=$ induced central ranking
$\theta_{1: n-1}=$ parameters at nodes
Inversions are penalized by $\theta_{i}$ parameters
Example: $\vec{\theta}=(0.1,1.2,0.4)$
$\operatorname{Cost}(a|b| c \mid d)=0$
$\operatorname{Cost}(b|a| c \mid d)=1.2$
$\operatorname{Cost}(c|b| a \mid d)=1.2+$

## Recursive Inversion Models (RIM)

[Meek, M 14]

$\tau=$ tree structure
$\pi_{0}(\tau)=$ induced central ranking
$\theta_{1: n-1}=$ parameters at nodes
Inversions are penalized by $\theta_{i}$ parameters
Example: $\vec{\theta}=(0.1,1.2,0.4)$
$\operatorname{Cost}(a|b| c \mid d)=0$
$\operatorname{Cost}(b|a| c \mid d)=1.2$
$\operatorname{Cost}(c|b| a \mid d)=1.2+$

## Recursive Inversion Models (RIM)

[Meek, M 14]

$\tau=$ tree structure
$\pi_{0}(\tau)=$ induced central ranking
$\theta_{1: n-1}=$ parameters at nodes
Inversions are penalized by $\theta_{i}$ parameters
Example: $\vec{\theta}=(0.1,1.2,0.4)$

$$
\begin{aligned}
\operatorname{Cost}(a|b| c \mid d) & =0 \\
\operatorname{Cost}(b|a| c \mid d) & =1.2 \\
\operatorname{Cost}(c|b| a \mid d) & =1.2+
\end{aligned}
$$

RIM distribution $P_{\tau, \vec{\theta}}$
Let $v_{i}=$ number of inversions of $\pi$ at node $i$

$$
P_{\tau, \vec{\theta}}(\pi) \propto \prod_{i \in \text { nodes }} \exp \left(-\theta_{i} v_{i}\right)
$$

$$
\begin{aligned}
P(a|b| c \mid d) & \propto e^{0} \\
P(b|a| c \mid d) & \propto e^{-1.2} \\
P(c|b| a \mid d) & \propto e^{-1.2-}
\end{aligned}
$$

## Recursive Inversion Models (RIM)

[Meek, M 14]

$\tau=$ tree structure
$\pi_{0}(\tau)=$ induced central ranking
$\theta_{1: n-1}=$ parameters at nodes
Inversions are penalized by $\theta_{i}$ parameters
Example: $\vec{\theta}=(0.1,1.2,0.4)$

$$
\begin{aligned}
\operatorname{Cost}(a|b| c \mid d) & =0 \\
\operatorname{Cost}(b|a| c \mid d) & =1.2 \\
\operatorname{Cost}(c|b| a \mid d) & =1.2+
\end{aligned}
$$

$$
\begin{aligned}
P(a|b| c \mid d) & \propto e^{0} \\
P(b|a| c \mid d) & \propto e^{-1.2} \\
P(c|b| a \mid d) & \propto e^{-1.2-}
\end{aligned}
$$

## RIM distribution $P_{\tau, \vec{\theta}}$

Let $v_{i}=$ number of inversions of $\pi$ at node $i$

$$
P_{\tau, \vec{\theta}}(\pi) \propto \prod_{i \in \text { nodes }} \exp \left(-\theta_{i} v_{i}\right)
$$

Normalization constant

$$
Z(\tau, \theta)=\prod_{i \in \text { nodes }} G\left(L_{i}, R_{i}, \exp \left(-\theta_{i}\right)\right)
$$

$$
\text { with } G(L, R, q)=\frac{(q)_{L+R}}{(q)_{L}(q)_{L}}, \quad(q)_{n}=\prod_{i=1}^{n}\left(1-q^{i}\right)
$$

Structure $\tau$ known as Riffle Independence model [Huang, Guestrin 12]

The RIM is a general flexible model


- any tree structure
- any parameters (but $\theta_{j} \geq 0$ suffices)
- includes the Mallows and Generalized Mallows models



## Max Likelihood Estimation for RIM

[M,Meek 14]

- Problem Given permutations $\pi_{1}, \ldots \pi_{N}$, infer $\tau, \theta$


## Max Likelihood Estimation for RIM

[M,Meek 14]

- Problem Given permutations $\pi_{1}, \ldots \pi_{N}$, infer $\tau, \theta$
- Identifiability of $\theta$
- reorder to obtain cannonical representation, with $\theta_{i} \geq 0$ for all $i \in$ nodes
- given $\tau, \theta_{i}$ can be estimated by convex univariate minimization



## Max Likelihood Estimation for RIM

[M,Meek 14]

- Problem Given permutations $\pi_{1}, \ldots \pi_{N}$, infer $\tau, \theta$
- Identifiability of $\theta$
- reorder to obtain cannonical representation, with $\theta_{i} \geq 0$ for all $i \in$ nodes
- given $\tau, \theta_{i}$ can be estimated by convex univariate minimization

- Identifiability of $\tau$

Theorem[M, Meek 14] A model $\tau, \theta$ is identifiable iff

1. $\theta_{i}>0$ for all $i \in$ nodes
2. $\theta_{i} \neq \theta_{p a(i)}$ for all $i \in \operatorname{nodes}(p a(i)$ is the parent of node $i$ in $\tau$ )

## Max Likelihood Estimation for RIM

[M,Meek 14]

- Problem Given permutations $\pi_{1}, \ldots \pi_{N}$, infer $\tau, \theta$
- Identifiability of $\theta$
- reorder to obtain cannonical representation, with $\theta_{i} \geq 0$ for all $i \in$ nodes
- given $\tau, \theta_{i}$ can be estimated by convex univariate minimization

- Identifiability of $\tau$

Theorem[M, Meek 14] A model $\tau, \theta$ is identifiable iff

1. $\theta_{i}>0$ for all $i \in$ nodes
2. $\theta_{i} \neq \theta_{p a(i)}$ for all $i \in \operatorname{nodes}(p a(i)$ is the parent of node $i$ in $\tau$ )

- Hardness of $\tau$ estimation
- Estimating $\pi_{0}$ is NP-hard [Duchi, Mackey, Jordan 13]
- Estimating $\tau$ structure given $\pi_{0}$ is tractable


## Sufficient statistics



$$
Q(d|a| b \mid c)=\left\lvert\, \begin{array}{cccc|c}
a & b & c & d & \\
\hline- & 1 & 0 & 0 & a \\
0 & - & 1 & 0 & b \\
0 & 0 & - & 0 & c \\
1 & 1 & 1 & - & d \\
\hline
\end{array}\right.
$$

## Sufficient statistics



$$
Q(d|a| b \mid c)=\begin{array}{|cccc|c}
a & b & c & d & \\
\hline- & 1 & 1 & 0 & a \\
0 & - & 1 & 0 & b \\
\cline { 1 - 2 } & 0 & - & 0 & c \\
1 & 1 & 1 & - & d \\
\hline
\end{array}
$$

$$
\operatorname{Cost}(d|a| b \mid c)=0.1 \times 2+1.2 \times 0+0.4 \times 1
$$

Max Likelihood Estimation algorithm(s)


Estimating $\tau$ given $\pi_{0}$ is tractable

Max Likelihood Fstimation algorithm(s)


- Estimating $\tau$ given $\pi_{0}$ is tractable
- by Dynamic Programming (DP) algorithm, similar to Matrix Chain Multiplication, Inside(-Outside) algorithm $\mathcal{O}\left(n^{4}\right)$
- contains $\theta_{j}$ estimation at each DP "partial solution"
- Estimating $\pi_{0}$ : Stochastic local search over $\pi_{0}$ space, similar to Simulated Annealing

1. Sample $\pi_{0}{ }^{\text {new }}$ from proposal distribution current $P_{\tau, \theta}$
2. Given $\pi_{0}{ }^{\text {new }}$, find $\tau^{\text {opt }}, \theta^{\text {opt }}$ by Dynamic Programming
3. Bring to cannonical form $\Rightarrow \tau^{\text {new }}, \theta^{\text {new }} \succeq 0$
4. Compute log-likelihood score, accept/reject like in Metropolis-Hastings, return to step 1

## Experiments - Sushi preferences data



## Data

$N=5000$ permutations of $n=10$ items Compared with:
alph $\pi_{0}$ fixed, $\tau, \theta \mid \pi_{0}$ optimize
GM fixed $\tau$, optimize $\pi_{0}, \theta$
HG fixed $\tau$ from [Huang, Guestrin,12], optimize $\theta$
SA Simulated Annealing

Test set log-likelihood w.r.t SA


## Partial rankings

"Sushi preference" data $n=12$
types of sushi
"My top 3 preferences are ika, maguro, tekka, in this order"
"I like uni least of all"
"I prefer fish to non-fish"


Three good things about the RIM

- RIM is a general model (includes Mallows, generalized Mallows)
- likelihood $P(\pi \mid \tau(\vec{\theta}))$ factors according to tree (and partition function $Z$ tractable)
- RIM has sufficient statistics


## Partial rankings

"Sushi preference" data $n=12$
types of sushi
ika|maguro|tekka|\{all other types\} \{all but ebi\}|ebi


Partial rankings
"Sushi preference" data $n=12$
types of sushi
ika|maguro|tekka|\{all other types\} \{all but ebi\}|ebi


Partial ranking $\sigma$ [Huang \& al, 10] $\sigma=\left(E_{1}\left|E_{2}\right| \ldots \mid E_{K}\right)$ with

- $E_{1} \cup E_{2} \cup \ldots E_{K}=$ set of items
- shape $\left(n_{1}, \ldots n_{K}\right)$, $n_{k}=\left|E_{k}\right|, \sum n_{k}=n$

Partial rankings
"Sushi preference" data $n=12$ types of sushi ika|maguro|tekka|\{all other types\} \{all but ebi\}|ebi


Partial ranking $\sigma$ [Huang \& al, 10]

$$
\sigma=\left(E_{1}\left|E_{2}\right| \ldots \mid E_{K}\right) \text { with }
$$

- $E_{1} \cup E_{2} \cup \ldots E_{K}=$ set of items
- shape $\left(n_{1}, \ldots n_{K}\right)$, $n_{k}=\left|E_{k}\right|, \sum n_{k}=n$

Three good things about the RIM

- RIM is a general model (includes Mallows, generalized Mallows)
- likelihood $P(\pi \mid \tau(\vec{\theta}))$ factors according to tree ? YES [Huang et al, 10]
- RIM has sufficient statistics ? NO

Inferences with partial rankings in the RIM. Are they tractable?
The meaning of "tractable"

- Estimation of $\pi_{0}$ for RIM is intractable in the worst case
- We define tractable as $\mathcal{O}(N$ poly $(n)) \times$ time (memory) for complete data

Inferences with partial rankings in the RIM. Are they tractable?
The meaning of "tractable"

- Estimation of $\pi_{0}$ for RIM is intractable in the worst case
- We define tractable as $\mathcal{O}(N$ poly $(n)) \times$ time (memory) for complete data

Main technical difficulty

- marginal probability of a partial ranking $\sigma$

$$
P(\sigma \mid \tau(\vec{\theta}))=\sum_{\pi \sim \sigma} P(\pi \mid \tau(\vec{\theta}))
$$

where linear extension $\{\pi \sim \sigma\}$ of $\sigma$ can have exponential size

## Contributions

1. for marginal probability $P(\sigma \mid \tau(\vec{\theta}))$

- exact formula and polynomial algorithm
- proved algorithm no more than 2 Nn more costly than for complete permutations (and sometimes much faster)

2. for pairwise marginals $E\left[Q_{a b}\right]=\operatorname{Pr}[a$ precedes $b \mid \sigma, \tau(\vec{\theta})]$

- exact recursive (polynomial) algorithm
- proved algorithm no more costly than for complete permutations

3. for parameter $\vec{\theta}$ estimation (Maximum Likelihood)

- convex univariate minimization algorithm for each $\theta i$
- proved algorithm is $\mathcal{O}(\mathrm{Nn})$ more costly than for complete permutations

4. for structure search (Maximum Likelihood) previous work

- complete data: local (simulated annealing) search algorithm with exact, tractable steps [Meek M 14]
- partial rankings: EM algorithm with approximate (or exponential) E step [Huang \& al 10]
our contributions
- new "E step" based on completing the pairwise marginals $E\left[Q_{a b}\right]$
- algorithms above can use the completed pairwise marginals as if they were complete data

Computing the marginal probability $P(\sigma \mid \tau, \vec{\theta})$


RIM probability for complete data $P(\pi \mid \tau, \vec{\theta})$

$$
\begin{aligned}
P(a|b| c \mid d) & \propto e^{0} \\
P(b|a| c \mid d) & \propto e^{-1.2} \\
P(c|b| a \mid d) & \propto e^{-1.2-2 \times 0.1}
\end{aligned}
$$

$$
\begin{aligned}
P_{\tau, \vec{\theta}}(\pi) & =\prod_{i \in \text { nodes }} \frac{e^{-\theta_{i} v_{i}}}{G_{L_{i}, R_{i}}\left(\exp \left(-\theta_{i}\right)\right)} \\
\text { with } \quad G_{L, R}(q) & =\frac{(q)_{L+R}}{(q)_{L}(q)_{R}}, \quad(q)_{n}=\prod_{i=1}^{n}\left(1-q^{i}\right) .
\end{aligned}
$$

RIM probability for partial ranking $\sigma$ [ $M$, Meek in prep]

$$
P_{\tau, \vec{\theta}}(\sigma)=\prod_{i \in \text { nodes }}(\text { factor at node i) }
$$

Marginal $P(\pi \mid \tau, \vec{\theta})$ for partial ranking $\sigma$


Sufficient to consider root node Complete ranking $\pi=(c|a| b \mid d)$

$$
\text { factor }=\frac{e^{-2 \theta}}{G_{2,2}\left(e^{-\theta}\right)}
$$

Partial ranking $\sigma=(c \mid\{a, b, d\})$
factor $=\frac{e^{-2 \theta} G_{0,1}\left(e^{-\theta}\right) G_{2,1}\left(e^{-\theta}\right)}{G_{2,2}\left(e^{-\theta}\right)}$

Marginal $P(\pi \mid \tau, \vec{\theta})$ for partial ranking $\sigma$


Sufficient to consider root node
Complete ranking $\pi=(c|a| b \mid d) \quad$ Partial ranking $\sigma=(c \mid\{a, b, d\})$

$$
\text { factor }=\frac{e^{-2 \theta}}{G_{2,2}\left(e^{-\theta}\right)}
$$

$$
\text { factor }=\frac{e^{-2 \theta} G_{0,1}\left(e^{-\theta}\right) G_{2,1}\left(e^{-\theta}\right)}{G_{2,2}\left(e^{-\theta}\right)}
$$

In general, at some internal node where

- set $\mathcal{L}$ is merged with set $\mathcal{R}$
- partial ranking $\sigma$ restricted to $\mathcal{L} \cup \mathcal{R}$ is $E_{1}\left|E_{2}\right| \ldots \mid E_{K}$ with $E_{k}=L_{k} \cup R_{k}$, $L_{k} \subseteq \mathcal{L}, r_{k} \subseteq \mathcal{R}$
- factor of $P(\sigma \mid \tau(\vec{\theta}))$ at this node is

$$
g\left(l_{1: K}, r_{1: K}, \theta\right)=\frac{e^{-\theta v} G_{l_{1}, r_{1}}\left(e^{-\theta}\right) G_{l_{2}, r_{2}}\left(e^{-\theta}\right) \ldots G_{K_{K}, r_{K}}\left(e^{-\theta}\right)}{G_{|\mathcal{L}|,|\mathcal{R}|}\left(e^{-\theta}\right)}
$$

where $v=\#$ inversions in $\sigma$ at node $\leq \#$ inversions in $\pi \sim \sigma$
$\underset{\sim}{\mathbb{N}}$ Marginal $P(\pi \mid \tau, \vec{\theta})$ - how much extra computation?
How many additional factors?
Rem $1 G_{0, r}=G_{l, 0}=1$

Marginal $P(\pi \mid \tau, \vec{\theta})$ - how much extra computation?
How many additional factors?
Rem $1 G_{0, r}=G_{l, 0}=1$
Rem 2 at each node, at least one of $L_{k}, R_{k}$ decreases (and their initial sum is $n$ )

- Hence, no more than $n-1$ extra factors (but sometimes much fewer)

Marginal $P(\pi \mid \tau, \vec{\theta})$ - how much extra computation?
How many additional factors?
Rem $1 G_{0, r}=G_{l, 0}=1$
Rem 2 at each node, at least one of $L_{k}, R_{k}$ decreases (and their initial sum is $n$ )

- Hence, no more than $n-1$ extra factors (but sometimes much fewer)
- Example top- $t$ rankings $\sigma=$ (ika|maguro|sake|\{everything else\}) $P(\sigma \mid \tau, \vec{\theta})$ has at most $t-1$ non-trivial factors

Marginal $P(\pi \mid \tau, \vec{\theta})$ - how much extra computation?
How many additional factors?
Rem $1 G_{0, r}=G_{l, 0}=1$
Rem 2 at each node, at least one of $L_{k}, R_{k}$ decreases (and their initial sum is $n$ )

- Hence, no more than $n-1$ extra factors (but sometimes much fewer)
- Example top- $t$ rankings $\sigma=(i k a \mid$ maguro|sake $\mid\{$ everything else $\}) P(\sigma \mid \tau, \vec{\theta})$ has at most $t-1$ non-trivial factors
How much additional computation?
- $G_{L, R}$ is computed recursively over $I=0, \ldots L, r=1, \ldots R$
- Hence, all $G_{l, r}(\theta)$ in numerator are cached while computing the denominator
- Overhead for whole sample of size $N$ is no more than $n N$ lookups+multiplications
- For comparison, for a complete whole sample
- computation of sufficient statistics is $\mathcal{O}\left(n^{2} N\right)$
- computation of $Z$ given $\vec{\theta}$ is $\mathcal{O}\left(n^{2} \log n\right)$

Independence properties


- define $Q_{a b}=1$ iff a precedes b
- $Q_{a b} \perp Q_{c d}$ whenever path $(a, b) \cap \operatorname{path}(c, d)=\emptyset$

Independence properties


- define $Q_{a b}=1$ iff a precedes $b$
- $Q_{a b} \perp Q_{c d}$ whenever $\operatorname{path}(a, b) \cap \operatorname{path}(c, d)=\emptyset$
- Indepence checking can reveal the "branching structure" (but not $\pi_{0}$ )
- In progress: combine independence tests with local search to estimate $\tau$


## Conclusion: No need to compromise!

Goals of inference in models on permutations

- Flexible w.r.t observation model (i.e. input data)
- partial rankings, pairwise observations
- Flexible w.r.t generative model
- RIMs are a class of flexible, identifyable, intepretable models
- Exact and tractable algorithms, closed form expression

