# Preference Modeling 

Lirong Xia

## (1). Rensselaer

## Today's Schedule

$>$ Modeling random preferences

- Random Utility Model
$>$ Modeling preferences over lotteries
- Prospect theory


## Parametric Preference Learning



## Parametric ranking models

$>$ A statistical model has three parts

- A parameter space: $\Theta$
- A sample space: $\mathrm{S}=$ Rankings $(A)^{\mathrm{n}}$
- $A$ = the set of alternatives, $\mathrm{n}=\#$ voters
- assuming votes are i.i.d.
- A set of probability distributions over S:
$\left\{\operatorname{Pr}_{\theta}(s)\right.$ for each $s \in \operatorname{Rankings}(A)$ and $\left.\theta \in \Theta\right\}$


## Example

$>$ Condorcet's model for two alternatives
$>$ Parameter space $\Theta=\{$ 忽, $\}$
$>$ Sample space $S=\{\text { 易, }\}^{n}$
$>$ Probability distributions, i.i.d.

$$
\begin{aligned}
& \operatorname{Pr}\left(\xi^{2} \mid \xi^{2}\right) \\
& =\operatorname{Pr}(\text { Q }) \\
& =p>0.5
\end{aligned}
$$

## Mallows' model [Mallows-1957]

$>$ Fixed dispersion $\varphi<1$
$>$ Parameter space

- all full rankings over candidates

>Sample space
- i.i.d. generated full rankings
$>$ Probabilities:

$$
\operatorname{Pr}_{W}(V) \propto \varphi^{\operatorname{Kendall}(V, W)}
$$

## Example: Mallows for

$\Theta_{\operatorname{sen}}$
$\Rightarrow$ Probabilities: $Z=1+2 \varphi+2 \varphi^{2}+\varphi^{3}$


## Random utility model (RUM) [Thurstone 27]

$>$ Continuous parameters: $\Theta=\left(\theta_{1}, \ldots, \theta_{m}\right)$

- $m$ : number of alternatives
- Each alternative is modeled by a utility distribution $\mu_{i}$
- $\theta_{i}$ : a vector that parameterizes $\mu_{i}$
$>$ An agent's latent utility $U_{i}$ for alternative $c_{i}$ is generated independently according to $\mu_{i}\left(U_{i}\right)$
$>$ Agents rank alternatives according to their perceived utilities
- $\operatorname{Pr}\left(c_{2}>c_{1}>c_{3} \mid \theta_{1}, \theta_{2}, \theta_{3}\right)=\operatorname{Pr}_{U_{i} \sim \mu_{i}}\left(U_{2}>U_{1}>U_{3}\right)$



## Generating a preference-profile

$>\operatorname{Pr}\left(\right.$ Data $\left.\mid \theta_{1}, \theta_{2}, \theta_{3}\right)=\Pi_{V \in \text { Data }} \operatorname{Pr}\left(V \mid \theta_{1}, \theta_{2}, \theta_{3}\right)$


Agent 1
$P_{1}=c_{2}>c_{1}>c_{3}$

Agent $n$
$P_{n}=c_{1}>c_{2}>c_{3}$

## Plackett-Luce model

$>\mu_{i}$ 's are Gumbel distributions

- A.k.a. the Plackett-Luce (P-L) model [BM 60, Yellott 77$]$
$>$ Alternative parameterization $\lambda_{1}, \ldots, \lambda_{m}$


Pros:
$c_{1}$ is the toip perefécead tocam..., $\left.c_{m}\right\}$

- Computationally tractable
- Analytical solution to the likelihood function
- The only RUM that was known to be tractable

- Widely applied in Economics [McFadden 74], learning to rank [Liu 11], and analyzing elections [GM 06,07,08,09]
.. Cons: may not be the best model



## RUM with normal distributions

$>\mu_{i}$ 's are normal distributions

- Thurstone's Case V [Thurstone 27]


Pros:

- Intuitive
- Flexible
$\because$ Cons: believed to be computationally intractable
- No analytical solution for the likelihood function $\operatorname{Pr}(P \mid$ $\Theta$ ) is known



## Decision making

## Maximum likelihood estimators (MLE)

Model: $\mathcal{M}_{r}$

$>$ For any profile $P=\left(V_{1}, \ldots, V_{n}\right)$,

- The likelihood of $\theta$ is $L(\theta, P)=\operatorname{Pr}_{\theta}(P)=\prod_{V \in P} \operatorname{Pr}_{\theta}(V)$
- The MLE mechanism
$\operatorname{MLE}(P)=\operatorname{argmax}_{\theta} L(\theta, P)$
- Decision space = Parameter space


## Bayesian approach

$>$ Given a profile $P=\left(V_{1}, \ldots, V_{n}\right)$, and a prior distribution $\pi$ over $\Theta$
$>$ Step 1: calculate the posterior probability over $\Theta$ using Bayes' rule

- $\operatorname{Pr}(\theta \mid P) \propto \pi(\theta) \operatorname{Pr}_{\theta}(P)$
$>$ Step 2: make a decision based on the posterior distribution
- Maximum a posteriori (MAP) estimation
- $\operatorname{MAP}(P)=\operatorname{argmax}_{\theta} \operatorname{Pr}(\theta \mid P)$
- Technically equivalent to MLE when $\pi$ is uniform
- $\Theta=\{$, Example
- $\mathrm{S}=\left\{\mathrm{g}_{\mathrm{g}}^{0},\right\}^{n}$
- Probability distributions:

$$
\begin{aligned}
& \operatorname{Pr}\left(\sum^{Q}\right) \\
= & \operatorname{Pr}(\mathrm{E}) \\
= & 0.6
\end{aligned}
$$

- Data $P=\left\{10 @\right.$, Q $\left._{6}\right\}$
- MLE

$$
\begin{aligned}
& -L(O)=\operatorname{Pr}_{0}(O)^{6} \operatorname{Pr}_{0}(M)^{4}=0.6^{10} 0.4^{8} \\
& -L(M)=\operatorname{Pr}_{M}(O)^{6} \operatorname{Pr}_{M}(M)^{4}=0.4^{10} 0.6^{8} \\
& -L(O)>L(M), O \text { wins }
\end{aligned}
$$

- MAP: prior O:0.2, M:0.8
$-\operatorname{Pr}(\mathrm{O} \mid \mathrm{P}) \propto 0.2 \mathrm{~L}(\mathrm{O})=0.2 \times 0.6^{10} 0.4^{8}$
$-\operatorname{Pr}(\mathrm{M} \mid \mathrm{P}) \propto 0.8 \mathrm{~L}(\mathrm{M})=0.8 \times 0.4^{10} 0.6^{8}$
$-\operatorname{Pr}(\mathrm{M} \mid \mathrm{P})>\operatorname{Pr}(\mathrm{O} \mid \mathrm{P}), \mathrm{M}$ wins


## Decision making under uncertainty

- You have a biased coin: head w/p $p$
- You observe 10 heads, 4 tails
- Do you think the next two tosses will be two heads in a row?
> MLE-based approach
- there is an unknown but fixed ground truth
- $p=10 / 14=0.714$
- $\operatorname{Pr}(2$ heads $\mid p=0.714)$
$=(0.714)^{2}=0.51>0.5$
- Yes!
- Bayesian
- the ground truth is captured by a belief distribution
- Compute $\operatorname{Pr}(p \mid$ Data $)$ assuming uniform prior
- Compute $\operatorname{Pr}(2$ heads $\mid$ Data $)=0.485<0$ . 5
- No!


## Prospect Theory: Motivating Example

$>$ Treat lung cancer with Radiation or Surgery

|  | Radiation | Surgery |
| :---: | :---: | :---: |
| Q1 | $100 \%$ immediately survive <br> $22 \% ~ 5-y e a r ~ s u r v i v e ~$ | $90 \%$ immediately survive <br> $34 \%$ 5-year survive |

$>$ Q1: 18\% choose Radiation
$>$ Q2: 49\% choose Radiation

## More Thoughts

$>$ Framing Effect

- The baseline/starting point matters
- Q1: starting at "the patient dies"
- Q2: starting at "the patient survives"
$>$ Evaluation
- subjective value (utility)
- perceptual likelihood (e.g. people tend to overweight low probability events)


## Prospect Theory


> Framing Phase (modeling the options)

- Choose a reference point to model the options
- $\mathrm{O}=\left\{\mathrm{o}_{1}, \ldots, \mathrm{o}_{\mathrm{k}}\right\}$
$>$ Evaluation Phase (modeling the preferences)
- a value function v: $\mathrm{O} \rightarrow \mathrm{R}$
- a probability weighting function $\pi:[0,1] \rightarrow R$
$>$ For any lottery $L=\left(p_{1}, \ldots, p_{k}\right) \in \operatorname{Lot}(O)$

$$
\mathrm{V}(\mathrm{~L})=\sum \pi(\mathrm{pi}) \mathrm{v}\left(\mathrm{o}_{\mathrm{i}}\right)
$$

## Example: Insurance

$>$ potential loss of $\$ 1000$ @1\%
$>$ Insurance fee $\$ 15$
$>$ Q1 (reference point: current wealth)

- Buy: Pay $\$ 15$ for sure. $V=v(-15)$
- No: \$0@99\% + \$-1000@1\%.V = $\quad$ (.99)v(0) + $\quad$ (.01)v($1000)=\pi(.01) v(-1000)$
$>$ Q2 (reference point: current wealth-1000)
- Buy: $\$ 985$ for sure. V = v(985)
- No: \$1000@99\% + \$0@1\%. V = $\quad$ (.99) v(1000) + $\pi(.01) v(0)=\pi(.99) v(1000)$

