Human Preference Models: Choice models

Sanmi Koyejo

All Sections

Export PDF

Today: Choice Modeling

Tools to predict the choice behavior of a group of decision-makers in a specific choice context.

Stanford Artificial Intelligence

Application: Marketing

What features affect a car purchase?

Autoguide.com

			
Make	Toyota	Honda	Ford
Year	2013	2013	2013
Model	Camry	Accord Sdn	Fusion
Trim	4dr Sdn I4 Auto L (Natl)	4dr I4 Man LX	4dr Sdn S FWD
General Information			
MSRP	**$22,235**	**$21,680**	**$21,900**
Invoice	$20,345	$19,849	$20,422
Destination	$795.00	$790.00	$795.00
Local Dealer Pricing	Get Free Quotes	Get Free Quotes	Get Free Quotes
Fuel Economy	25 MPG city/ 35 MPG hwy	24 MPG city/ 34 MPG hwy	22 MPG city/ 34 MPG hwy
Engine	2.5L/152 Gas I4	2.4L/144 Gas I4	2.5L/152 Gas I4
Transmission	Auto, 6	Manual, 6	Auto, 6
Horsepower	178 hp @ 6000 rpm	185 hp @ 6400 rpm	170 hp @ - TBD - rpm

Application: Transportation

- How pricing affects route choice
- How much is a driver willing to pay

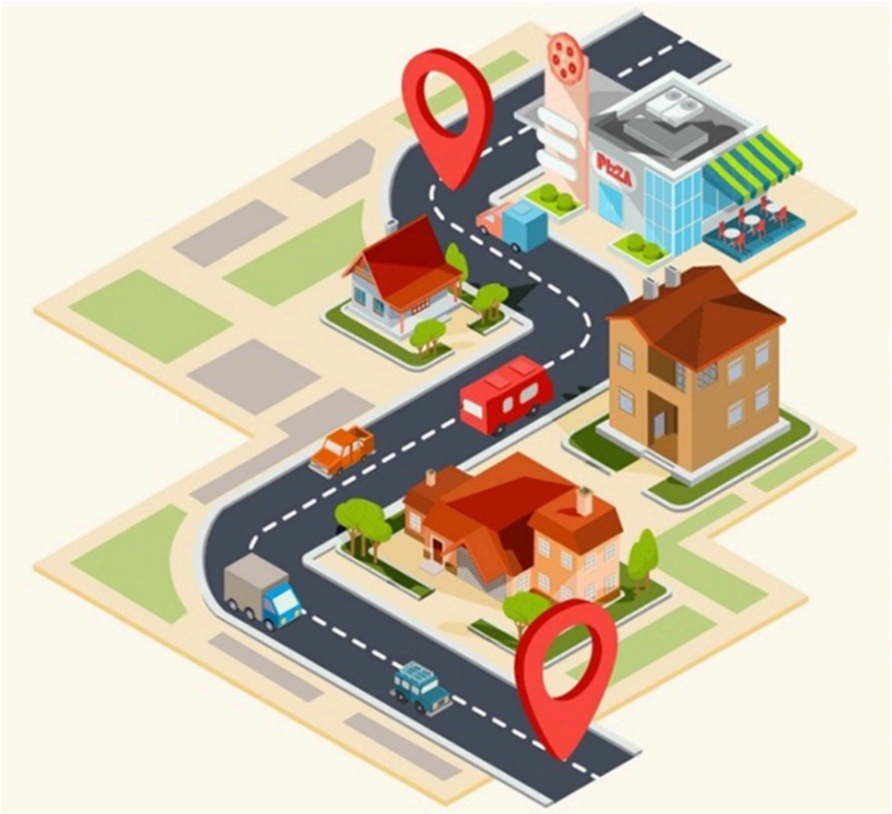

Image source: https://www.supplychain247.com/article/8_factors_to_consider_when_choosing_route_optimization_software/locus

Application: Energy Economics

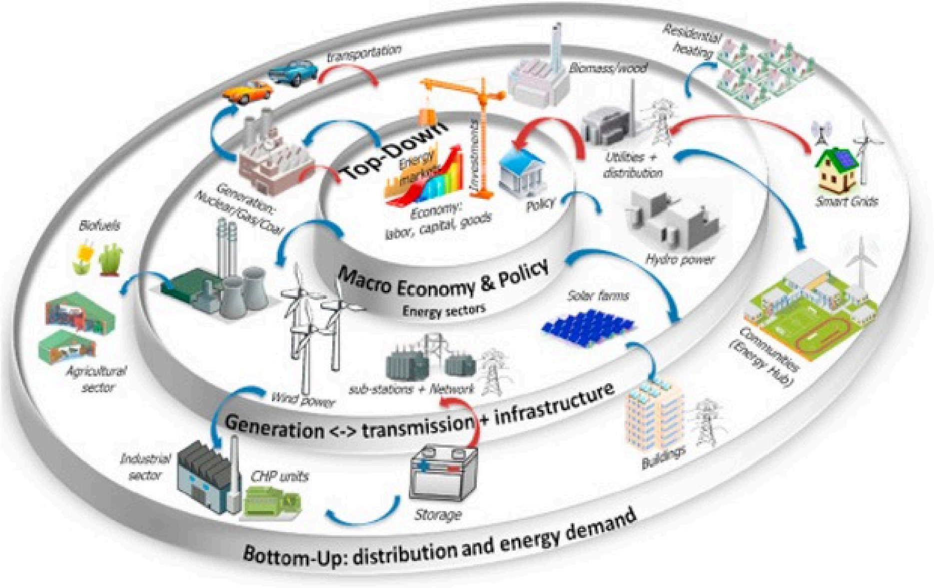

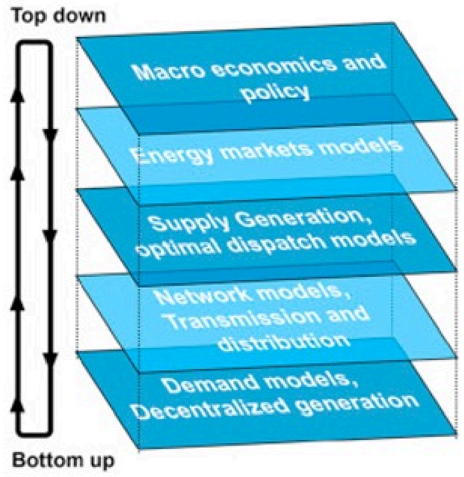

Del Granado, Pedro Crespo, Renger H. Van Nieuwkoop, Evangelos G. Kardakos, and Christian Schaffner. "Modelling the energy transition: A nexus of energy system and economic models." *Energy strategy reviews*, 20 (2018): 229-235.

Example: Daily activity-travel pattern of an individual

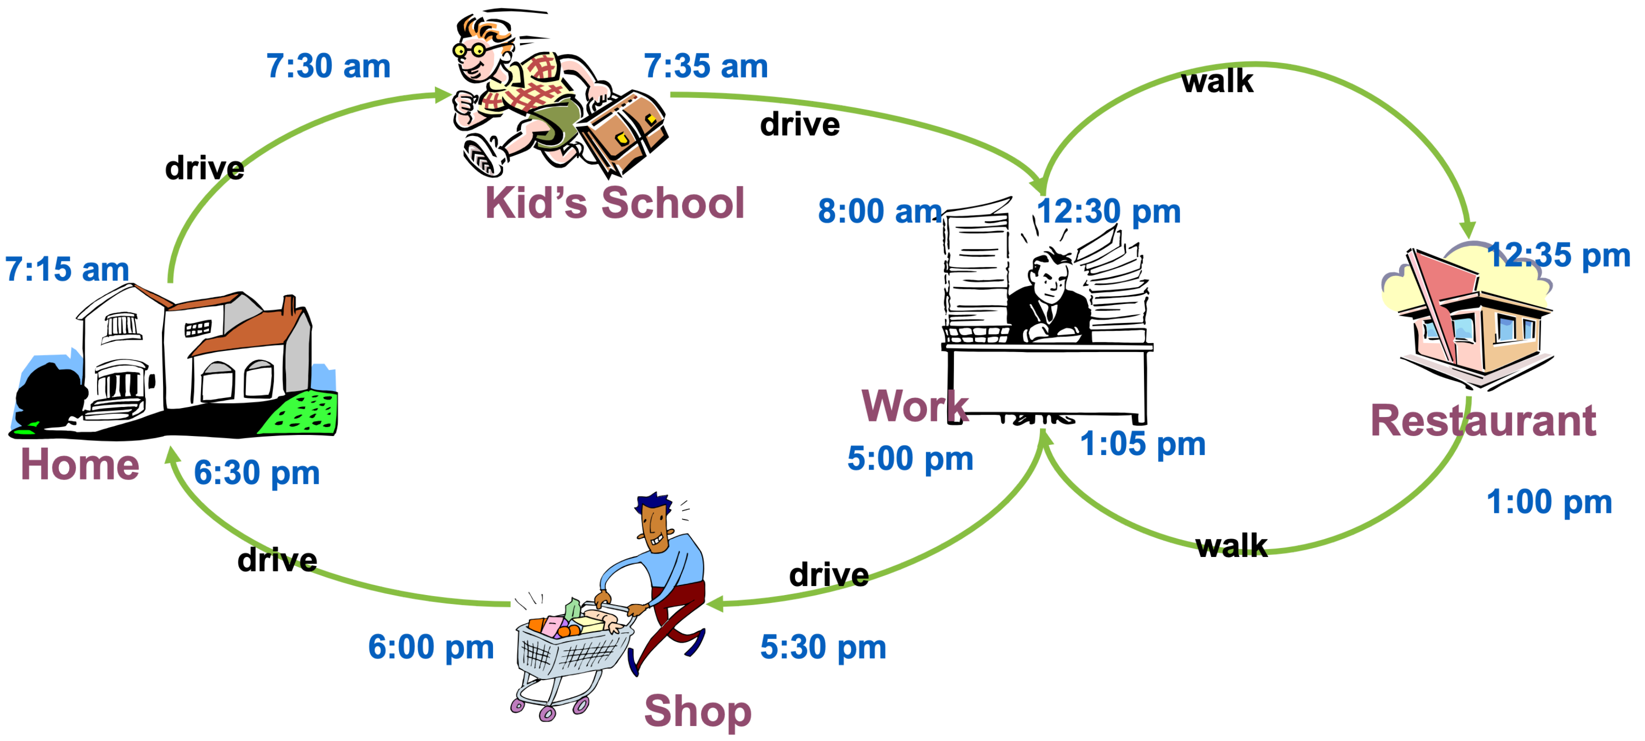

Source: Chandra Bhat, " General introduction to choice modeling"

Application: RL and Language

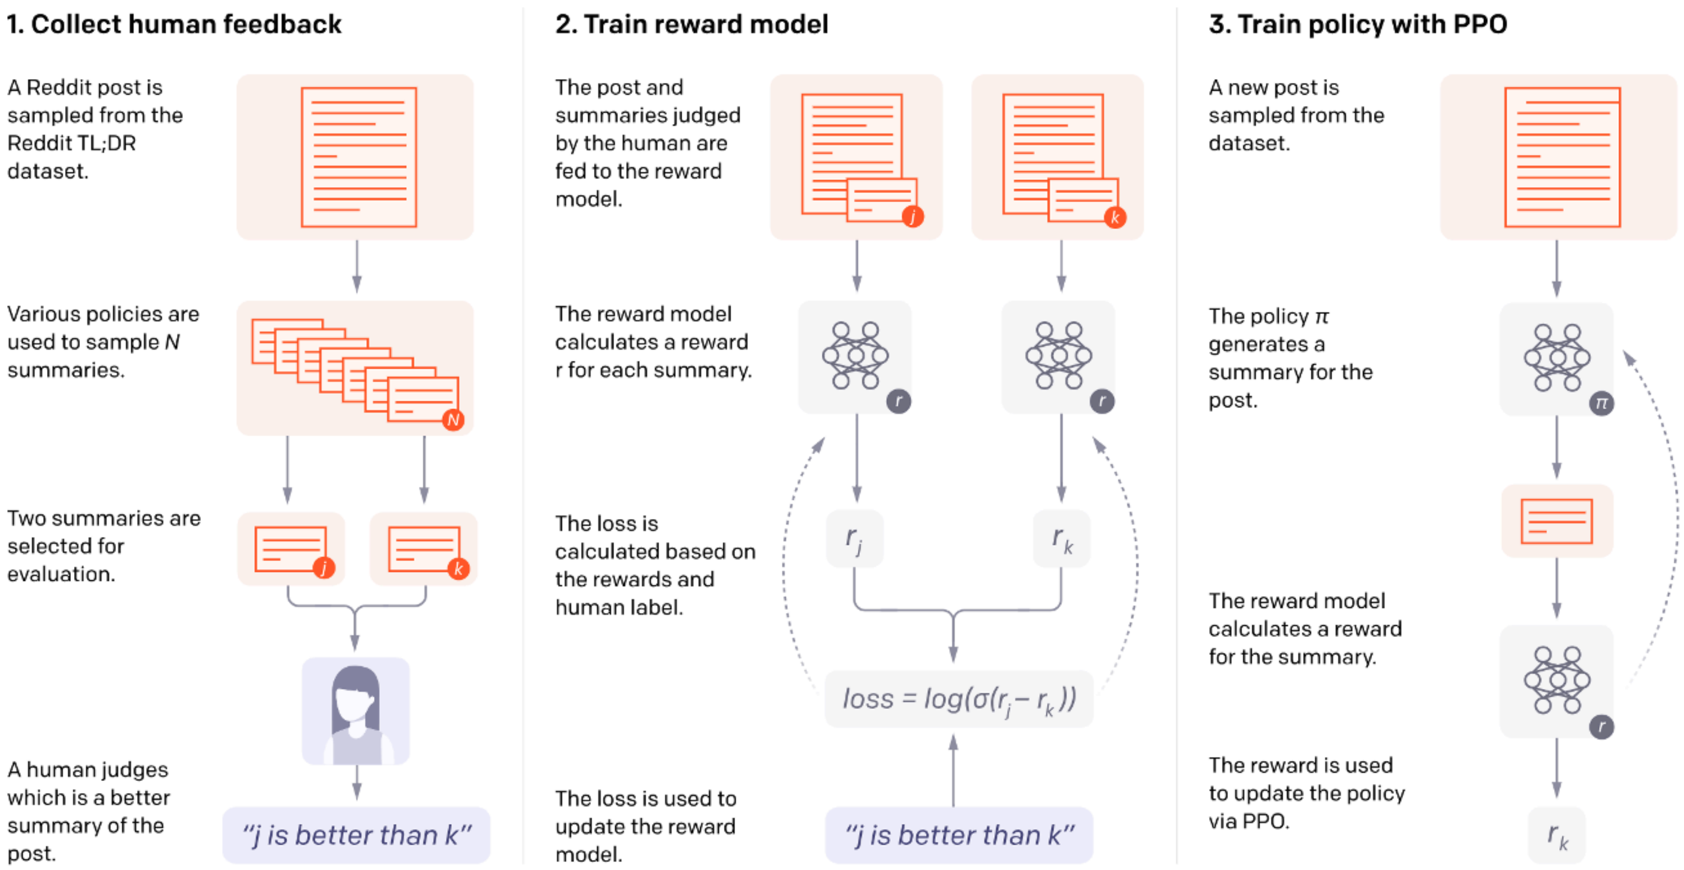

https://openai.com/research/learning-to-summarize-with-human-feedback

History

- Thurstone research into food preferences in the 1920s
- Microeconomics: Random Utility Theory (1970s)
  - McFadden: Nobel prize in 2000 for the theoretical basis for discrete choice.
- Psychology: Duncan Luce and Anthony Marley
  - Luce, R. Duncan (1959). "Conditional logit analysis of qualitative choice behavior"
- Early use in marketing
  - Predict demand for new products that are potentially expensive to produce
- Early use in transportation
  - Predict usage of transportation resources, e.g., used by McFadden to predict the demand for the Bay Area Rapid Transit (BART) before it was built

Why are we studying choice models?

- Human preferences are often gathered by asking for choices across alternatives
- Basic choice models are the workhorse for ML from preferences (Bradley-Terry, Plackett Luce)
- Our discussion will highlight some of the key assumptions, e.g., utility and rationality
  - We will cover models originally built for discrete/finite choices, which have been extended to ML applications (conditional choices)

Stanford ARTIFICIAL INTELLIGENCE

(Discrete) choice models

- Models designed to capture decision-process of individuals
- True utility is not observable, but perhaps can measure via preferences over choices
- **Main assumption**: utility (benefit, or value) that an individual derives from item A over item B is a function of the frequency that they choose item A over item B in repeated choices.
- **Useful Note**: “Utility” in choice models <=> “Reward” in RL

Modeling: Discrete choice

- Choices are collectively exhaustive, mutually exclusive, and finite

$$y_{ni} = \begin{cases} 1, & \text{if } U_{ni} > U_{nj} \; \forall j \neq i \\ 0, & \text{otherwise} \end{cases}$$

$$U_{ni} = H_{ni}(z_{ni})$$

- $z_{n,i}$ are variables describing the individual attributes and the alternative choices
- $H_{ni}(z_{ni})$ is a stochastic function, e.g., linear $H_{ni}(z_{ni}) = \beta z_{ni} + \epsilon_{ni}$, where ϵ_{ni} are unobserved individual factors

Implications of the choice model

- Only the utility differences matter

$$\begin{aligned} P_{ni} &= Pr(y_{ni} = 1) \\ &= Pr(U_{ni} > U_{nj}, \forall j \neq i) \\ &= Pr(U_{ni} - U_{nj} > 0, \forall j \neq i) \end{aligned}$$

- Note that utility here is scale-free
  - May be invariant to monotonic transformations
  - Ok within a single context, but will need to normalize for comparing across datasets
  - Common approach: normalize scale by standardizing the variance

Stanford Artificial Intelligence

Example: Binary choice with individual attributes

- Benefit of action depends on s_n = individual characteristics

$$\begin{cases} U_n = \beta s_n + \epsilon_n \\ y_n = \begin{cases} 1 & U_n > 0 \\ 0 & U_n \leq 0 \end{cases} \end{cases} \Rightarrow \quad P_{n1} = \frac{1}{1 + \exp(-\beta s_n)}$$

- $\epsilon \sim$ Logistic
- Replacing $\epsilon \sim$ Standard Normal gives the probit model

$$P_{n1} = \Phi(\beta s_n)$$

- Where $\Phi(.)$ is the normal CDF

Example: Utility is linear function of variables that vary over alternatives (Bradley-Terry Model)

- The utility of each alternative depends on the attributes of the alternatives (which may include individual attributes)
- Unobserved terms are assumed to have an extreme value distribution

$$\begin{cases} U_{n1} = \beta z_{n1} + \epsilon_{n1} \\ U_{n2} = \beta z_{n2} + \epsilon_{n2} \\ \epsilon_{n1}, \epsilon_{n2} \sim \text{iid extreme value} \end{cases} \quad \Rightarrow \quad P_{n1} = \frac{\exp(\beta z_{n1})}{\exp(\beta z_{n1}) + \exp(\beta z_{n2})}$$

- Equivalently $P_{n1} = \frac{1}{1+\exp(-\beta(z_{n1} - z_{n2}))}$
- Can replace noise with Standard Normal $P_{n1} = \Phi(\beta(z_{n1} - z_{n2}))$

Example: Utility for each alternative depends on attributes of that alternative

- Unobserved terms are assumed to have an extreme value distribution
- With J alternatives

$$\begin{cases} U_{ni} = \beta z_{ni} + \epsilon_{ni} \\ \epsilon_{ni} \sim \text{iid extreme value} \end{cases} \quad \Rightarrow \quad P_{ni} = \frac{\exp(\beta z_{ni})}{\sum_{j=1}^{J} \exp(\beta z_{nj})}$$

- Compare to standard model for multiclass classification (multiclass logistic)
- Can also replace noise model with Gaussians

Capturing correlations across alternatives

- All the prior models use the logistic model which does not capture correlations in noise.
- This can be fixed using a joint distribution over the noise e.g.,

$$\begin{cases} U_{ni} = \beta z_{ni} + \epsilon_{ni} \\ \epsilon_n \equiv (\epsilon_{n1}, \cdots, \epsilon_{nJ}) \sim N(0, \Omega) \end{cases}$$

Estimation

- **Linear case**: maximum likelihood estimators
  - Logistic model: use (binary or multinomial) logistic regression
  - Gaussian Model: use probit regression
- **More complex function classes**: use standard ML fitting tools for (regularized) maximum likelihood, e.g., stochastic gradient descent (SGD)
- **Standard tradeoffs**, e.g., bias-variance tradeoff
  - More complex utility models generally require more data
  - Most ML applications pool the model across individuals, individual differences may matter (more on this in future class)

Stanford Artificial Intelligence

What of measuring ordered preferences?

- Example: On a 1-5 scale where 1 means disagree completely and 5 means agree completely, how much do you agree with the following statement: “I am enjoying this class so far”
- Use ordinal regression, e.g.,

$$U_n = H_n(z_n) \qquad y_n = \begin{cases} 1, & \text{if } U_n < a \\ 2, & \text{if } a < U_n < b \\ 3, & \text{if } b < U_n < c \\ 4, & \text{if } c < U_n < d \\ 5, & \text{if } U_n > d \end{cases}$$

- For some real numbers a, b, c, d (parameters)

Ordered Logit

- For linear utility: $U_n = \beta z_n + \epsilon$, $\epsilon \sim$ Logistic

$$Pr(\text{choosing } 1) = Pr(U_n < a) = Pr(\epsilon < a - \beta z_n) = \frac{1}{1+\exp(-(a-\beta z_n))}$$

$$Pr(\text{choosing } 2) = Pr(a < U_n < b) = Pr(a - \beta z_n < \epsilon < b - \beta z_n)$$
$$= \frac{1}{1+\exp(-(b-\beta z_n))} - \frac{1}{1+\exp(-(a-\beta z_n))}$$

$$\cdots$$

$$Pr(\text{choosing } 5) = Pr(U_n > d) = Pr(\epsilon > d - \beta z_n) = 1 - \frac{1}{1+\exp(-(d-\beta z_n))}$$

- Can also replace with Gaussian for ordered probit regression

Plackett-Luce Model

- Ranking models the **sequence of choices** (Plackett and Luce in 1970s)
- Probability of choice 1, 2, …, J is

$$Pr(\text{ranking } 1, 2, \ldots, J) = \frac{\exp(\beta z_1)}{\sum_{j=1}^{J} \exp(\beta z_{nj})} \cdot \frac{\exp(\beta z_2)}{\sum_{j=2}^{J} \exp(\beta z_{nj})} \cdots \frac{\exp(\beta z_{J-1})}{\sum_{j=J-1}^{J} \exp(\beta z_{nj})}$$

- PL is common in biomedical literature
- aka rank ordered logit (econometrics ~1980s), or exploded logit model
- All the extensions mentioned also apply (nonlinear utility, correlated noise, etc.)

Modeling and estimation summary

- Choose the utility model, i.e., how the attributes and alternatives define the utility e.g., linear function of attributes with logistic noise
- Choose the response/observation model, e.g., binary, multiple choice, ordered choice.
- Fit the model using (regularized) maximum likelihood

Aside: "Revealed preference" vs "stated preference"

- **Revealed preference:** Use observed data about the choices to estimate value ascribed to items.
  - Generally offline observational data about real choices
- **Stated Preference:** Use the choices made by individuals under experimental conditions to estimate these values
  - Generally online experimental data (can include controlled experiments)
- Revealed preference is considered a "real" choice, so can be more accurate
  - In simulated situations, participants may not respond well to hypotheticals
  - OTOH: observed data may not cover the space, hence the appeal of experiments

Exercise (inclass): choice model for class(es)

- **"Should you take CS 329H or not?"**
  - What are the attributes/features (describe what to measure about a class)?
  - What utility model?
  - What is the observation/response model?
  - Revealed preference (observed choices) or stated preference (hypothetical)?
- **"Should you take CS 329H or CS 221 or CS 229?"**
  - What are the attributes/features?
  - What utility model?
  - What is the observation/response model?
  - Revealed preference or stated preference?

Stanford Artificial Intelligence

Exercise (inclass): choice model for language

Design a choice model to evaluate the quality of a language model?

- What utility model?
    - What are the attributes/features?
- What is the observation/response model?
- Revealed preference or stated preference?
- Who should you query?
    - Individual or pooled responses: why or why not?
- What are some pro/cons of your design?

Stanford
ARTIFICIAL
INTELLIGENCE

References

- Train, K. (1986). *Qualitative Choice Analysis: Theory, Econometrics, and an Application to Automobile Demand.* MIT Press. ISBN 9780262200554. Chapter 8.
- McFadden, D.; Train, K. (2000). "Mixed MNL Models for Discrete Response" (PDF). *Journal of Applied Econometrics.* 15 (5): 447–470.
- Luce, R. D. (1959). *Individual Choice Behavior: A Theoretical Analysis.* Wiley.
- Additional:
  - Ben-Akiva, M.; Lerman, S. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand.* Transportation Studies. Massachusetts: MIT Press.
  - Park, Byeong U.; Simar, Léopold; Zelenyuk, Valentin (2017). "Nonparametric estimation of dynamic discrete choice models for time series data" (PDF). *Computational Statistics & Data Analysis.* 108: 97–120. doi:10.1016/j.csda.2016.10.024.
  - Rafailov, Rafael, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. "Direct preference optimization: Your language model is secretly a reward model." arXiv preprint arXiv:2305.18290 (2023).