A Markov Perturbed Utility Model of Route Choice

Mogens Fosgerau[1] and Rui Yao[2]

[1]University of Copenhagen, Denmark
[2]Technion, Israel

Short summary

This paper proposes the Markov perturbed utility route choice (M-PURC) model. The model is based on a generalized entropy over link flows, constructed at each node as the perspective of a convex node perturbation with respect to the node outflow. The convex conjugate of the generalized entropy defines a surplus function whose gradient yields the optimal link flows. Lagrangian duality decomposes the surplus maximization into node-level perturbed utility problems coupled through a Bellman equation, yielding a Markovian representation of optimal behavior. The framework accommodates the recursive logit and the NGEV model as special cases. A separable specification, in which the node perturbation is the sum of link perturbations, yields a link-by-link probability-recovery formula and endogenous sparsity: dominated links receive exactly zero probability without choice set specification. We estimate the model via a Fenchel–Young loss that is convex in parameters, remains finite at corner solutions, and reduces to the recursive-logit likelihood as a special case.

Keywords: endogenous sparsity; Fenchel–Young estimation; Markov decision process; perturbed utility; recursive logit; route choice.

1 Introduction

Route choice modeling faces a trade-off between behavioral realism and computational tractability. Path-based discrete choice models require the analyst to generate a feasible set of paths, a step that is both computationally burdensome on large networks and a source of bias (Frejinger et al., 2009). The recursive logit (RL) model (Fosgerau et al., 2013), building on Dial (1971), Shen et al. (1996), and Baillon & Cominetti (2008), avoids path enumeration by formulating route choice as sequential link selection with, in effect, entropy regularization. However, entropy regularization implies that every outgoing link receives strictly positive probability, so that clearly dominated links may still be chosen with non-negligible probability on rich networks. The model can also entail substantial computational challenges on large networks (Zimmermann & Frejinger, 2020; Mai & Frejinger, 2022). The perturbed utility route choice (PURC) model of Fosgerau et al. (2022) takes a different approach. It characterizes route choice through a global convex optimization problem in which a perturbation term regularizes behavior and can induce sparsity: many links can receive exactly zero usage in optimum, without there being an *a priori* choice set. Fosgerau & Sørensen (2026) develop a convex duality theory for PURC in which optimal behavior is represented by node potentials and recovered link-by-link from conjugate perturbations. However, the PURC, in its existing forms, requires flows to be computed separately for every origin–destination pair.

This paper proposes the Markov perturbed utility route choice (M-PURC) model, which connects the recursive and the perturbed utility perspectives. The starting point is a *generalized entropy* over network link flows, constructed at each node as the perspective of a convex perturbation with respect to the node outflow. The convex conjugate of the generalized entropy defines a *surplus function* that characterizes the traveler's behavior as the solution to a global convex program. All model predictions—link flows, next-link choice probabilities, and value functions—are derived from this conjugate pair. The optimal link flows are the gradient of the surplus, and the Fenchel–Young inequality provides a natural loss function for estimation. In a companion paper, Yao & Zhang (2026) develop the dual side of the same conjugate pair, taking the surplus function as the starting point and focusing on computation and equilibrium.

The node-level perspective structure of the generalized entropy implies, via Lagrangian duality, a Markov decomposition: the global program separates into independent node-level perturbed utility problems coupled through a value function that satisfies a Bellman equation. This decomposition

holds for general convex node perturbations, extending the entropy case considered in Akamatsu (1997). The framework thus encompasses the recursive logit and the network generalized extreme value (NGEV) model of Oyama et al. (2022) as special cases corresponding to specific choices of the node perturbation.

For estimation, the Fenchel–Young loss derived from the conjugate pair resolves a fundamental difficulty in likelihood-based estimation of sparse models: because the M-PURC can assign exactly zero probability to some links, the log-likelihood diverges when an observed choice falls outside the model-implied active set. The Fenchel–Young loss remains finite and convex at such points (Blondel et al., 2020; Lin et al., 2026). We propose a single-level estimation procedure that alternates parameter gradient steps with warm-started value function updates. A key computational advantage runs through both solution and estimation algorithms: the Markov structure implies that a single value-function solve for a given destination characterizes behavior from all origins simultaneously, pooling all origins sharing a common destination into a single Bellman update.

2 Methodology

Model

Let $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ be a directed network with finite node set $\mathcal{N}$ and link set $\mathcal{L} \subseteq \mathcal{N} \times \mathcal{N}$. Fix a destination $d \in \mathcal{N}$ and, for each non-destination node i, let $\mathcal{L}(i)$ denote the set of successor nodes. We assume that the network is strongly connected. Each link (i, j) carries a utility $v_{(i,j)} \in \mathbb{R}$, parameterized in estimation as $v_{(i,j)} = \beta^\top z_{(i,j)}$, where $z_{(i,j)}$ is a vector of observable link attributes and β is the parameter vector to be estimated. A demand vector $q \in \mathbb{R}_+^{|\mathcal{N}|-1}$ specifies the number of travelers originating at each node with destination d. A link flow $y \geq 0$ is feasible if it satisfies flow conservation $Ay = q$, where A is the node-link incidence matrix (with the row for d omitted); let Y denote the set of feasible flows. The node outflow is $\mu_i = \sum_{j \in \mathcal{L}(i)} y_{(i,j)}$ and the next-link choice probabilities at node i, when $\mu_i > 0$, are $x_{(i,j)} = y_{(i,j)}/\mu_i$.

At each non-destination node i, the traveler's preferences over next-link probabilities are characterized by a *node perturbation* $\Omega_i : \mathbb{R}^{|\mathcal{L}(i)|} \to \mathbb{R} \cup \{+\infty\}$ that is closed, strictly convex on the probability simplex Δ_i, and $+\infty$ elsewhere. The node perturbation penalizes concentrated behavior: the traveler balances the utility of choosing high-value links against the perturbation cost of concentrating probability on few links. The *generalized entropy* over link flows is

$$H(y) = \sum_{i \in \mathcal{N} \setminus \{d\}} \mu_i \, \Omega_i\left(\frac{y_i}{\mu_i}\right) + \delta_Y(y), \tag{1}$$

with δ_Y the convex indicator of the feasible set and the standard perspective convention $\mu_i \Omega_i(y_i/\mu_i) := 0$ at $\mu_i = y_i = 0$. Each summand is the perspective of a convex function, hence jointly convex in (y_i, μ_i) (Boyd & Vandenberghe, 2004, Section 3.2.6). We refer all proofs to our full paper.

Proposition 1 (Convexity of the generalized entropy) *The generalized entropy H is a proper closed convex function on $\mathbb{R}^{|\mathcal{L}|}$.*

The traveler's behavior is characterized through the *surplus function*, defined as the convex conjugate of the generalized entropy:

$$H^*(v) = \sup_{y \in Y} \left\{ v^\top y - \sum_{i \neq d} \mu_i \, \Omega_i\left(\frac{y_i}{\mu_i}\right) \right\}. \tag{2}$$

The surplus $H^*(v)$ is the maximum, over feasible link flows, of total link utility minus the generalized entropy; as the conjugate of a proper closed convex function, H^* is itself proper, closed, and convex. We maintain throughout that v lies in the interior of $\operatorname{dom} H^*$, which guarantees that the supremum is attained.

Proposition 2 (Existence and uniqueness of optimal policy) *The surplus maximization* (2) *has at least one optimum. Any two optima share the same next-link choice probabilities at every co-visited node.*

Markov decomposition

The flow-conservation constraints in (2) couple link flows across nodes. Dualizing these constraints with multipliers $V \in \mathbb{R}^{|\mathcal{N}|-1}$ (with $V_d := 0$) and substituting the change of variables $y_{(i,j)} = \mu_i x_{(i,j)}$ separates the Lagrangian across nodes: at each node, the inner maximization over $x_i \in \Delta_i$ defines the convex conjugate of the node perturbation,

$$\Omega_i^*(U_i) = \max_{x_i \in \Delta_i} \left\{ U_i^\top x_i - \Omega_i(x_i) \right\}, \tag{3}$$

where $U_{(i,j)} = v_{(i,j)} + V_j$ is the *full utility* at link (i,j)—the sum of the immediate link utility and the continuation value at the successor node. The outer maximization over $\mu_i \geq 0$ is linear and produces the Bellman inequality $V_i \geq \Omega_i^*(U_i)$. At any dual optimum these inequalities may be tightened to equalities without affecting the primal optimum, giving the *Bellman equation*

$$V_i = \Omega_i^*\big((v_{(i,j)} + V_j)_{j \in \mathcal{L}(i)}\big), \qquad V_d = 0. \tag{4}$$

Strict convexity of Ω_i makes Ω_i^* differentiable, so the optimal policy is recovered node by node as $x_i^*(U_i) = \nabla \Omega_i^*(U_i)$. The value V_i equals the maximum over randomized next-link choices of the expected one-step utility minus node perturbation plus the expected continuation value, which is the standard structure of a value function in a Markov decision process (Neu et al., 2017; Geist et al., 2019). To recover the optimal link flows from the Bellman-induced policy, we first verify that the policy is proper, meaning that the induced transition matrix has spectral radius strictly less than one.

Proposition 3 (Properness) *Any Bellman-induced policy x^* is proper: $\rho(P_{x^*}) < 1$.*

Properness implies that $(I - P_{x^*}^\top)$ is invertible, so the aggregate node outflow vector $\mu^* = (I - P_{x^*}^\top)^{-1} q$ is finite and the optimal link flows are $y^*_{(i,j)} = \mu_i^* x^*_{(i,j)}$.

Proposition 4 (Uniqueness and differentiability) *The optimal link flow $y^*(v)$ is unique. The surplus H^* is differentiable at v with $\nabla H^*(v) = y^*(v)$, and the Fenchel–Young equality $H(y^*(v)) + H^*(v) = v^\top y^*(v)$ holds.*

Proposition 5 (Jacobian of the value function) *Suppose the node operators Ω_i^* are continuously differentiable. Then the mapping $v \mapsto V^*(v)$ is continuously differentiable, with*

$$\frac{\partial (V^*)_i}{\partial v_{(k,l)}} = \left[(I - P_{x^*})^{-1} \right]_{ik} \cdot x^*_{(k,l)}. \tag{5}$$

Consequently, $H^*(v) = q^\top V^*(v)$ satisfies $\nabla_{v_{(k,l)}} H^*(v) = \mu_k^* x^*_{(k,l)} = y^*_{(k,l)}$, where $\mu^* = (I - P_{x^*})^{-\top} q$ is the vector of node outflows aggregated across origins.

The Jacobian identity reveals the computational structure of the gradient: it decomposes into aggregate node outflows μ^* obtained from a single linear solve per destination, and next-link choice probabilities x^* obtained from the node-level problems.

Two consequences of this decomposition are central. First, the optimal next-link choice probabilities at each node depend only on the current node and the destination, not on the origin: the model with general demand q is equivalent to aggregating single-origin problems, $y^*(q) = \sum_o q_o y^*(e_o)$. Second, the global problem for a given destination is solved by a single Bellman fixed-point solve; on networks with cycles we compute this fixed point by modified policy iteration, alternating one policy-improvement step $x \leftarrow \nabla \Omega^*(U(V))$ with a small number m of warm-started policy-evaluation sweeps $V \leftarrow r_x + P_x^\top V$, where $r_{x,i} = \sum_j v_{(i,j)} x_{(i,j)} - \Omega_i(x_i)$ (Puterman & Shin, 1978; Yao & Zhang, 2026).

Fenchel–Young estimation

For perturbations that generate sparsity, the log-likelihood is undefined whenever an observed traversal falls on a link with zero model probability, since $\ln x^*_{(i,j)} = -\infty$. The conjugate pair (H, H^*) provides a loss that remains finite and convex at such corner solutions. The data consist of observed link traversals grouped by destination; for each destination d, let y_d^{obs} be the vector of aggregate observed link flows and q_d the corresponding demand vector. The *Fenchel–Young loss* for destination d is

$$\ell_d(v) = H(y_d^{\text{obs}}) + H_d^*(v) - (y_d^{\text{obs}})^\top v. \tag{6}$$

Proposition 6 (Properties of the Fenchel–Young loss) *The loss ℓ_d is non-negative, finite on the interior of* $\operatorname{dom} H_d^*$*, convex in v, zero if and only if $y_d^{\mathrm{obs}} = y_d^*(v)$, and satisfies $\nabla_v \ell_d(v) = y_d^*(v) - y_d^{\mathrm{obs}}$.*

With $v = Z^\top \beta$, these properties are inherited by the full estimation objective $J(\beta) = \sum_d \ell_d(Z^\top \beta)$, which is convex in β and has gradient

$$\nabla_\beta J(\beta) = \sum_d Z\big(y_d^*(Z^\top \beta) - y_d^{\mathrm{obs}}\big). \tag{7}$$

This is the moment condition that predicted aggregate link flows match observed aggregate link flows, weighted by link attributes.

Proposition 7 (Convexity and identification) *The objective $J : \mathbb{R}^d \to \mathbb{R}$ is convex in β. Assuming* $\operatorname{rank}(Z) = d$*, J is strictly convex if and only if for every $\beta_1 \neq \beta_2$ there exists a destination d with $y_d^*(Z^\top \beta_1) \neq y_d^*(Z^\top \beta_2)$.*

When the perturbation is the entropy with scale $l > 0$, flow-conservation telescoping shows that $-\mathrm{LL}_d(v) = (\ell_d(v) - H(y_d^{\mathrm{obs}}))/l$; the Fenchel–Young estimator and maximum likelihood then coincide, while for sparsity-generating perturbations the Fenchel–Young loss remains finite and convex where the log-likelihood is undefined.

The gradient (7) requires the Bellman fixed point only through y_d^*. Rather than solve each inner Bellman system to full precision at every parameter evaluation, as in the nested fixed-point approach of Rust (1987), we interleave a single outer gradient step with a few warm-started inner sweeps. Algorithm 1 summarizes the resulting single-level procedure; for brevity we present it for a single destination and suppress the subscript d.

Algorithm 1 Single-level inexact gradient descent for M-PURC estimation

Require: Inner depth $m \geq 1$; stepsize $\alpha > 0$; initial $\beta^{(0)}$ with $v(\beta^{(0)}) \in \operatorname{int}(\operatorname{dom} H^*)$

$\hat{V} \leftarrow V^*(\beta^{(0)})$ ▷ Warm start
for $t = 0, 1, 2, \ldots$ **do**
  $\hat{x} \leftarrow \nabla\Omega^*\big(U(\hat{V}; \beta^{(t)})\big)$; form $P_{\hat{x}}$, $r_{\hat{x}}$ ▷ Policy improvement
  for $h = 1, \ldots, m$ **do**
    $\hat{V} \leftarrow r_{\hat{x}} + P_{\hat{x}} \hat{V}$ ▷ m Bellman sweeps
    $\hat{\mu} \leftarrow q + P_{\hat{x}}^\top \hat{\mu}$ ▷ m outflow sweeps
  end for
  $\hat{g} \leftarrow Z^\top(\hat{\mu} \odot \hat{x} - y^{\mathrm{obs}})$ ▷ Inexact gradient
  $\beta^{(t+1)} \leftarrow \beta^{(t)} - \alpha\, \hat{g}$ ▷ Parameter update
end for

At each outer iteration, a policy-improvement step computes the current policy $\hat{x}$ from the current value estimate, followed by m warm-started Bellman and outflow sweeps that update $\hat{V}$ and $\hat{\mu}$; the inexact gradient $\hat{g}$ is then formed from the residual $\hat{\mu} \odot \hat{x} - y^{\mathrm{obs}}$ and used to update β. Carrying $(\hat{V}, \hat{\mu})$ across outer iterations avoids any full inner solve after initialization. Each destination contributes independently to J and its gradient, so destinations can be parallelized and the Markov structure pools all origins sharing a common destination into a single Bellman update. The algorithm is related to the nested pseudo-likelihood family of Aguirregabiria & Mira (2002); Kasahara & Shimotsu (2012) but replaces both the full inner solve and the inner parameter optimization by a single warm-started step; because J is convex in β with smooth gradient through the continuously differentiable value function, the single-loop iterates converge to a global minimizer under the standard analysis of convex optimization with inexact gradients (Devolder et al., 2014).

Separable perturbations and endogenous sparsity

A natural specialization builds Ω_i from independent link perturbations:

$$\Omega_i(x_i) = \sum_{j \in \mathcal{L}(i)} h_{(i,j)}(x_{(i,j)}) + \delta_{\Delta_i}(x_i),$$

where each h is a closed, strictly convex function on $\mathbb{R}_+$ with $h(\xi) = +\infty$ for $\xi < 0$ and $h(0) = 0$. Dualizing the simplex constraint at node i with multiplier λ_i (the node potential) makes the node-level problem separable across links: at an interior optimum, $x^*_{(i,j)} = \nabla h^*_{(i,j)}(U_{(i,j)} - \lambda_i)$, where ∇h^*

is the gradient of the conjugate link perturbation. The potential λ_i^* is determined endogenously by probability conservation,

$$\sum_{j \in \mathcal{L}(i)} \nabla h^*_{(i,j)}(U_{(i,j)} - \lambda_i^*) = 1, \tag{8}$$

a one-dimensional monotone equation with a unique solution. When h satisfies $h'(0^+) = 0$, the conjugate gradient ∇h^* vanishes on $(-\infty, 0]$; link (i, j) then receives positive probability if and only if $U_{(i,j)} > \lambda_i^*$, so the active set at each node is determined endogenously without a prespecified choice set.

The entropy perturbation $h(\xi) = \xi \ln \xi$ recovers the recursive logit $x^*_{(i,j)} \propto \exp(U_{(i,j)})$, with every link active. The quadratic perturbation $h_{(i,j)}(\xi) = \frac{1}{2} l_{(i,j)} \xi^2$ gives the closed-form sparse recovery formula

$$x^*_{(i,j)} = \frac{\max\{0,\, U_{(i,j)} - \lambda_i^*\}}{l_{(i,j)}}, \qquad \sum_{j \in \mathcal{L}(i)} \frac{\max\{0,\, U_{(i,j)} - \lambda_i^*\}}{l_{(i,j)}} = 1, \tag{9}$$

which is structurally related to the sparsemax operator of Martins & Astudillo (2016).

3 Results and discussion

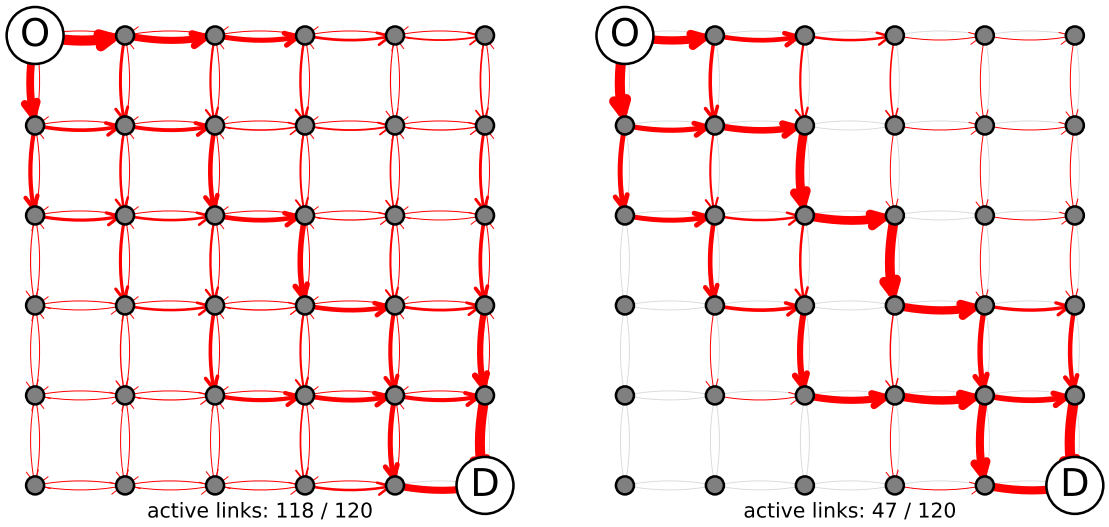
We illustrate the M-PURC model and the Fenchel–Young estimator on a 6×6 bidirectional grid with 36 nodes and $|\mathcal{L}| = 120$ directed links, with origin at the top-left corner and destination at the bottom-right corner. Links carry $d = 3$ observed attributes $z_{(i,j)} \in \mathbb{R}^3_+$: a travel-time component proportional to link length plus two synthetic covariates drawn independently from the uniform distribution on $[0, 1]$ and fixed throughout the experiments. Deterministic link utilities are $v_{(i,j)} = -z_{(i,j)}^\top \beta^*$ with ground truth $\beta^* = (1.0, 0.5, 0.3)$, and link perturbations have unit weight $l_{(i,j)} = 1$. The grid is strongly connected with cycles in every direction.

Predicted link flows: entropy versus quadratic

We solve the M-PURC prediction problem at $\beta = \beta^*$ with one unit of demand injected at the origin, under the entropy perturbation $h(\xi) = \xi \ln \xi$, which corresponds to the recursive logit and yields interior full-support link flows, and under the quadratic perturbation $h(\xi) = \frac{1}{2}\xi^2$, which induces corner solutions and endogenous exclusion of links from the active set. Figure 1 plots the resulting predicted link flows; arrow widths are proportional to link flow, and dashed background edges are inactive. Two features are visible. First, the entropy-based equilibrium activates every link of the grid (except two outgoing links from the destination), distributing flow over detours and backward links. Second, the quadratic perturbation concentrates flow on a small subset of links close to the south-east direction from origin to destination; many links, including most backward and laterally-redundant ones, carry exactly zero flow.

Simulation experiment

We now assess the proposed Fenchel–Young estimator under both the entropy perturbation (recursive logit) and the quadratic perturbation, using the same grid, the same feature draws, and the same true parameter. For each sample size $N \in \{500, 1000, 2000, 5000\}$ we simulate N independent origin-to-destination paths by random walk under the model-implied next-link choice probabilities. Each of $B = 100$ replications uses a fresh draw of paths and runs the single-level estimator with $m = 5$ inner Bellman and outflow sweeps per outer step. Table 1 reports, per parameter and separately for each model, the empirical mean of the estimates, the empirical standard deviation in parentheses, the joint root-mean-square error, and the rescaled $\sqrt{N}$-RMSE.



(a) Recursive logit (b) Quadratic perturbation

Figure 1: Predicted link flows on the 6×6 bidirectional grid. Red arrow width is proportional to flow; dashed background edges are inactive.

Table 1: Monte Carlo simulation
empirical mean estimates with empirical standard deviations in parentheses below.

	Entropy (recursive logit)				Quadratic (sparsemax)			
N	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	RMSE $\sqrt{N}$-RMSE	$\hat{\beta}_0$	$\hat{\beta}_1$	$\hat{\beta}_2$	RMSE $\sqrt{N}$-RMSE
500	0.998 (0.054)	0.508 (0.041)	0.294 (0.049)	0.084 1.874	0.990 (0.052)	0.496 (0.032)	0.303 (0.028)	0.068 1.514
1000	0.995 (0.044)	0.510 (0.032)	0.303 (0.036)	0.066 2.078	0.993 (0.036)	0.496 (0.024)	0.302 (0.018)	0.048 1.508
2000	0.994 (0.031)	0.507 (0.022)	0.305 (0.025)	0.046 2.075	0.996 (0.025)	0.498 (0.017)	0.302 (0.013)	0.033 1.476
5000	0.996 (0.020)	0.501 (0.014)	0.305 (0.016)	0.030 2.086	0.999 (0.014)	0.499 (0.011)	0.300 (0.007)	0.019 1.358

Two patterns emerge and hold for both perturbations. First, the empirical bias is small at every sample size, and all three coordinates of β are recovered to within one-percent accuracy at $N = 5000$. Second, the joint RMSE decreases monotonically with sample size while the $\sqrt{N}$-RMSE stays essentially flat, suggesting $\sqrt{N}$-consistency. The quadratic specification is only piecewise smooth in β—the active set at each node is locally constant and the gradient changes at boundary transitions—yet parameter recovery is empirically as fast as in the smooth entropy case.

4 Conclusions

The M-PURC model connects two previously separate approaches to route choice—the recursive logit and the perturbed utility framework—through a generalized entropy constructed from node-level perspective functions. At the general level, Lagrangian duality yields a Markov decomposition and a Bellman equation for any convex node perturbation, encompassing the recursive logit and the NGEV model as special cases. The separable specification adds a link-by-link probability-recovery formula and endogenous sparsity via convex conjugate structure: links whose full utility falls below a node-specific threshold receive exactly zero probability, without requiring a prespecified choice set. The Fenchel–Young estimator complements the model by providing a loss that is well-defined at corner solutions and convex in the link utilities, while reducing to standard maximum

likelihood for the recursive logit. The Markov structure implies that a single value-function solve per destination characterizes behavior from all origins simultaneously, making both solution and estimation scalable to large networks.

References

Aguirregabiria, V., & Mira, P. (2002, July). Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models. *Econometrica*, *70*(4), 1519–1543. Retrieved 2012-01-25, from http://onlinelibrary.wiley.com.globalproxy.cvt.dk/doi/10.1111/1468-0262.00340/abstract (arXiv: 1011.1669v3 ISBN: 0012-9682) doi: 10.1111/1468-0262.00340

Akamatsu, T. (1997). Decomposition of path choice entropy in general transport networks. *Transportation Science*, *31*(4), 349-362. Retrieved from https://doi.org/10.1287/trsc.31.4.349 doi: 10.1287/trsc.31.4.349

Baillon, J. B., & Cominetti, R. (2008, January). Markovian traffic equilibrium. *Mathematical Programming*, *111*(1-2), 33–56. doi: 10.1007/s10107-006-0076-2

Blondel, M., Martins, A. F. T., & Niculae, V. (2020). Learning with Fenchel-Young losses. *Journal of Machine Learning Research*, *21*(35), 1–69. Retrieved 2026-01-07, from http://jmlr.org/papers/v21/19-021.html

Boyd, S., & Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.

Devolder, O., Glineur, F., & Nesterov, Y. (2014). First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, *146*(1), 37–75.

Dial, R. B. (1971). A probabilistic multipath traffic assignment model which obviates path enumeration. *Transportation Research*, *5*(2), 83-111. Retrieved from https://www.sciencedirect.com/science/article/pii/0041164771900128 doi: https://doi.org/10.1016/0041-1647(71)90012-8

Fosgerau, M., Frejinger, E., & Karlstrom, A. (2013). A link based network route choice model with unrestricted choice set. *Transportation Research Part B: Methodological*, *56*, 70–80. doi: 10.1016/j.trb.2013.07.012

Fosgerau, M., Paulsen, M., & Rasmussen, T. K. (2022). A perturbed utility route choice model. *Transportation Research Part C*, *136*(103514). doi: https://doi.org/10.1016/j.trc.2021.103514

Fosgerau, M., & Sørensen, J. R.-V. (2026). *Convex Duality in Perturbed Utility Route Choice* [Working paper].

Frejinger, E., Bierlaire, M., & Ben-Akiva, M. (2009, December). Sampling of alternatives for route choice modeling. *Transportation Research Part B: Methodological*, *43*(10), 984–994. Retrieved 2015-04-30, from http://www.sciencedirect.com/science/article/pii/S0191261509000381 (ISBN: 0191-2615) doi: 10.1016/j.trb.2009.03.001

Geist, M., Scherrer, B., & Pietquin, O. (2019). *A theory of regularized markov decision processes.* Retrieved from https://arxiv.org/abs/1901.11275

Kasahara, H., & Shimotsu, K. (2012, September). Sequential Estimation of Structural Models With a Fixed Point Constraint. *Econometrica*, *80*(5), 2303–2319. Retrieved 2015-06-10, from http://dx.doi.org/10.3982/ECTA8291 (ISBN: 0012-9682) doi: 10.3982/ECTA8291

Lin, X., Yin, Y., & Liu, T. (2026). *Fenchel-young estimators of perturbed utility models.* Retrieved from https://arxiv.org/abs/2602.21376

Mai, T., & Frejinger, E. (2022). Undiscounted recursive path choice models: Convergence properties and algorithms. *Transportation Science*, *56*(6), 1469-1482. Retrieved from https://doi.org/10.1287/trsc.2022.1145 doi: 10.1287/trsc.2022.1145

Martins, A. F. T., & Astudillo, R. F. (2016). *From softmax to sparsemax: A sparse model of attention and multi-label classification.* Retrieved from https://arxiv.org/abs/1602.02068

Neu, G., Jonsson, A., & Gómez, V. (2017). *A unified view of entropy-regularized markov decision processes.* Retrieved from https://arxiv.org/abs/1705.07798

Oyama, Y., Hara, Y., & Akamatsu, T. (2022, January). Markovian traffic equilibrium assignment based on network generalized extreme value model. *Transportation Research Part B: Methodological*, *155*, 135–159. Retrieved 2025-12-14, from https://www.sciencedirect.com/science/article/pii/S0191261521001934 doi: 10.1016/j.trb.2021.10.013

Puterman, M. L., & Shin, M. C. (1978). Modified policy iteration algorithms for discounted markov decision problems. *Management Science*, *24*(11), 1127-1137. Retrieved from https://doi.org/10.1287/mnsc.24.11.1127 doi: 10.1287/mnsc.24.11.1127

Rust, J. (1987). Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher. *Econometrica*, *55*(5), pp.999–1033. Retrieved from http://www.jstor.org/stable/1911259 (arXiv: 1011.1669v3 ISBN: DR000152 00129682 DI952691 95P03192) doi: 10.2307/1911259

Shen, W., Zhang, H. M., & Akamatsu, T. (1996, October). Cyclic flows, Markov process and stochastic traffic assignment. *Transportation Research Record*, *30*(530), 1–34. Retrieved from http://www.sciencedirect.com/science/article/pii/0191261596000033 (ISBN: 0191-2615) doi: 10.1016/0191-2615(96)00003-3

Yao, R., & Zhang, K. (2026). *Perturbed utility markovian traffic equilibrium: theory and computation.*

Zimmermann, M., & Frejinger, E. (2020, June). A tutorial on recursive models for analyzing and predicting path choice behavior. *EURO Journal on Transportation and Logistics*, *9*(2), 100004. (arXiv:1905.00883 [stat]) doi: 10.1016/j.ejtl.2020.100004