Opportunistic Stochasticity in Shortest Path Problems:

from causal PDE-discretizations to efficient routing of autonomous vehicles

Alex Vladimirsky

Cornell University

vladimirsky@cornell.edu

August 26, 2024

Joint work with Mallory Gaspard.

An (In)famous Magic Roundabout in Swindon, UK

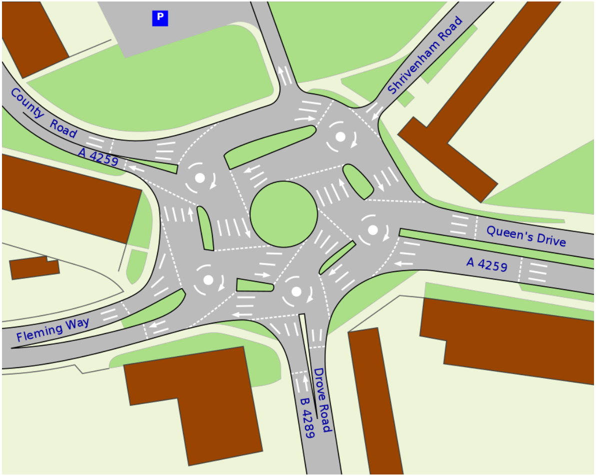

Source: Wikimedia Commons.

DETERMINISTIC "Shortest" Path Problems

- **Nodes:** $X = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M\}$
- **Bounded degree of nodes:** $|N(\boldsymbol{x}_i)| < \kappa$ for all i
- **Transition cost:** $C_{ij} \geq \delta > 0$ (assumed $+\infty$ if $\boldsymbol{x}_j \notin N(\boldsymbol{x}_i)$)
- **Exit cost:** $q(\boldsymbol{x}_i)$ for all $\boldsymbol{x}_i \in Q \subset X$

Dynamic Programming: The **value function** $U(\boldsymbol{x}_i) = U_i$ is the minimum required total-cost-to-exit starting from $\boldsymbol{x}_i$.

Bellman's Optimality Principle:

$$\begin{aligned} U_i &= \min_{\boldsymbol{x}_j \in N(\boldsymbol{x}_i)} \{C_{ij} + U_j\}, & \forall \boldsymbol{x}_i \notin Q; \\ U_i &= q(\boldsymbol{x}_i), & \forall \boldsymbol{x}_i \in Q. \end{aligned}$$

A coupled system of M non-linear equations!

Fast (Non-iterative, "Label-Setting") Methods

$$U_i = \min_{\boldsymbol{x}_j \in N(\boldsymbol{x}_i)} \{C_{ij} + U_j\}, \qquad \forall \boldsymbol{x}_i \notin Q$$

How can you de-couple a non-linear system?

Monotone Causality: Each node depends only on its "smaller" neighbors!

"If you use The Known
to tentatively compute The Still Unknown
then the smallest of The Tentatively Known
is actually Known."

Dijkstra's Method: $O(M \log M)$ complexity; uses a heap-sort.
Dial's Method: $O(M)$ complexity; uses a list of "buckets" of width δ.

General Stochastic Shortest Path (SSP) problems

- $X = \{\boldsymbol{x}_1, ..., \boldsymbol{x}_M, \boldsymbol{t} = \boldsymbol{x}_{M+1}\}$;
- $A_i = A(\boldsymbol{x}_i)$ a compact set of actions available at $\boldsymbol{x}_i$;
- choice of an action $\boldsymbol{a} \in A_i$ determines
  - the cost of the next transition $C(\boldsymbol{x_i}, \boldsymbol{a})$ and
  - the probability distribution over successor-states $p(\boldsymbol{x}_i, \boldsymbol{x}_j, \boldsymbol{a}) = p_{ij}(\boldsymbol{a}) = \mathbb{P}\left(\boldsymbol{x}_i \to \boldsymbol{x}_j \mid \text{using } \boldsymbol{a}\right)$;
- the target $\boldsymbol{t}$ is *absorbing* , i.e., $p_{tt}(\boldsymbol{a}) = 1$ and $C(\boldsymbol{t}, \boldsymbol{a}) = 0$ for $\forall \boldsymbol{a} \in A_t$.

A function $\mu : X \mapsto \left(\bigcup_{i=1}^{M} A_i \right)$ is a *stationary policy* if $\mu(\boldsymbol{x}_i) \in A_i$ for all $\boldsymbol{x}_i \in X$.
Starting from $\boldsymbol{x}_i$, the expected *cumulative* cost of using μ is $\mathcal{J}(\boldsymbol{x}_i, \mu)$.

The value function $U_i = U(\boldsymbol{x}_i) = \inf_{\mu} \mathcal{J}(\boldsymbol{x}_i, \mu)$.

A policy μ_* is *optimal* if $U(\boldsymbol{x}_i) = \mathcal{J}(\boldsymbol{x}_i, \mu_*)$ for all $\boldsymbol{x}_i \in X$.

SSP: Dynamic Programming and Value Iterations

Optimality conditions:

$$\begin{aligned} U_t &= 0, \\ U_i &= \min_{\boldsymbol{a} \in A_i} \left\{ C(\boldsymbol{x}_i, \boldsymbol{a}) + \sum_{j=1}^{M+1} p_{ij}(\boldsymbol{a}) U_j \right\}, \qquad \text{for } \forall \boldsymbol{x}_i \in X \backslash \{\boldsymbol{t}\}. \end{aligned}$$

$\Psi : \mathbb{R}^M \mapsto \mathbb{R}^M$ is defined componentwise: $(\Psi W)_i = \min_{\boldsymbol{a} \in A_i} \left\{ C(\boldsymbol{x}_i, \boldsymbol{a}) + \sum_{j=1}^{M+1} p_{ij}(\boldsymbol{a}) W_j \right\}$

and $U = \begin{bmatrix} U_1 \\ \vdots \\ U_M \end{bmatrix}$ is a fixed point of Ψ.

Value iterations: $W^{n+1} := \Psi W^n$ starting from an initial guess $W^0 \in \mathbb{R}^M$.

SSP: Dynamic Programming and Value Iterations

Optimality conditions:

$$
\begin{aligned}
U_t &= 0, \\
U_i &= \min_{\boldsymbol{a} \in A_i} \left\{ C(\boldsymbol{x}_i, \boldsymbol{a}) + \sum_{j=1}^{M+1} p_{ij}(\boldsymbol{a}) U_j \right\}, \qquad \text{for } \forall \boldsymbol{x}_i \in X \backslash \{\boldsymbol{t}\}.
\end{aligned}
$$

$\Psi : \mathbb{R}^M \mapsto \mathbb{R}^M$ is defined componentwise: $(\Psi W)_i = \min_{\boldsymbol{a} \in A_i} \left\{ C(\boldsymbol{x}_i, \boldsymbol{a}) + \sum_{j=1}^{M+1} p_{ij}(\boldsymbol{a}) W_j \right\}$

and $U = \begin{bmatrix} U_1 \\ \vdots \\ U_M \end{bmatrix}$ is a fixed point of Ψ.

Value iterations: $W^{n+1} := \Psi W^n$ starting from an initial guess $W^0 \in \mathbb{R}^M$.

[Bertsekas & Tsitsiklis; 1991] : In general, Ψ isn't a contraction, but the convergence is guaranteed for any W^0 provided:

- **(A0)** All $C(\boldsymbol{x}_i, \boldsymbol{a})$ are lower-semicontinuous and all $p_{ij}(\boldsymbol{a})$ are continuous.
- **(A1)** There exists at least one *proper policy* (i.e., a policy, which reaches the target $\boldsymbol{t}$ with probability 1 regardless of the initial state $\boldsymbol{x} \in X$).
- **(A2)** Every improper policy μ will have cost $\mathcal{J}(\boldsymbol{x}, \mu) = +\infty$ for at least one $\boldsymbol{x} \in X$.

Value Iterations vs. Label-setting

An SSP is *causal* if only finitely many value iterations are needed.
A *dependency digraph* G_μ defined for every stationary policy μ.

Bertsekas: the SSP is causal if $\exists$ an *optimal* policy μ_* such that G_{μ_*} is *acyclic*.

Still requires $O(M^2)$ operations! But Dijkstra-like and Dial-like methods need only $O(M \log M)$ and $O(M)$ operations respectively. Can they be used instead?

Value Iterations vs. Label-setting

An SSP is *causal* if only finitely many value iterations are needed.
A *dependency digraph* G_μ defined for every stationary policy μ.

Bertsekas: the SSP is causal if $\exists$ an *optimal* policy μ_* such that G_{μ_*} is *acyclic*.

Still requires $O(M^2)$ operations! But Dijkstra-like and Dial-like methods need only $O(M \log M)$ and $O(M)$ operations respectively. Can they be used instead?

Bertsekas: An optimal policy μ_* is *consistently improving* if

$$p_{ij}(\mu_*(\boldsymbol{x}_i)) > 0 \qquad \Longrightarrow \qquad U_i > U_j.$$

Existence of such μ_* $\Longrightarrow$ applicability of a Dijkstra-like method.

AV: Given $\delta > 0$, an optimal policy μ_* is *consistently δ-improving* if

$$p_{ij}(\mu_*(\boldsymbol{x}_i)) > 0 \qquad \Longrightarrow \qquad U_i \geq U_j + \delta.$$

Existence of such μ_* $\Longrightarrow$ applicability of a Dial-like method with bin-width δ.

But all these sufficient conditions are *implicit*... Can we do any better?

What makes an SSP "Opportunistically" Stochastic?

Definition (OSSP:)

We will refer to an SSP as *Opportunistically Stochastic* (OSSP) if

$$\exists \boldsymbol{a} \in A_i \text{ s.t. } p_{ij}(\boldsymbol{a}) > 0 \qquad \Longrightarrow \qquad \exists \tilde{\boldsymbol{a}} \in A_i \text{ s.t. } p_{ij}(\tilde{\boldsymbol{a}}) = 1$$

holds for all i and j.

Every stochastically realizable path is also deterministically realizable.

But stochastic actions might be still advantageous to reduce the cost!

Example: when driving on a highway, I might be able to guarantee a successful lane change if I slow down enough. But is it always worth it?

Shorthand action-focused notation

Focusing on any specific action $\boldsymbol{a} \in A_i$, we define

- A set of possible successor nodes
  $\mathcal{I}(\boldsymbol{a}) \;=\; \{\boldsymbol{x} \in X \;|\; p(\boldsymbol{x}_i, \boldsymbol{x}, \boldsymbol{a}) > 0\}\,.$
- The number of possible successor nodes $m = |\mathcal{I}(\boldsymbol{a})|$
  and their enumeration $\mathcal{I}(\boldsymbol{a}) = \{\boldsymbol{z}_1, \ldots, \boldsymbol{z}_m\}$.
- The probabilities of transition $\xi_j = p(\boldsymbol{x}_i, \boldsymbol{z}_j, \boldsymbol{a})$. Using these, $\boldsymbol{a}$ can be identified with a point $(\xi_1, ..., \xi_m)$ in a probability simplex Ξ_m.
- The costs C_j corresponding to deterministic $\boldsymbol{x}_i \to \boldsymbol{z}_j$ transitions.
  (Since this is an OSSP, such deterministic actions are available.)

If $m = 1$, this $\boldsymbol{a}$ is deterministic itself and $C_1 = C(\boldsymbol{x}_i, \boldsymbol{a})$.

These m, $\boldsymbol{z}_j$, ξ_j,and C_js are always understood to be $\boldsymbol{a}$-specific.

Monotone Causality of OSSPs

Theorem

Suppose there exists a $\delta \geq 0$ such that,
for all $\boldsymbol{x}_i \neq \boldsymbol{t}$, $\boldsymbol{a} \in A_i$, and every $r \in \{1, ..., m\}$,

$$C(\boldsymbol{x}_i, \boldsymbol{a}) \;\geq\; \sum_{j=1, j\neq r}^{m} \xi_j C_j + \xi_r \delta.$$

If these conditions are satisfied, this OSSP is monotone (δ-)causal.

This criterion is easy to check for any m, and it is "sharp" for $m = 2$.

(A sharp criterion for $m > 2$ is also available, but it is more complicated. See the paper.)

Unlike in prior work on MC for SSPs **(Vladimirsky, 2008)**, this does not assume anything about convexity or smoothness of $C(\boldsymbol{x}_i, \boldsymbol{a})$.

Geometric interpretation of (δ-)MC criterion when $m = 2$:

$$\mathcal{I}(\boldsymbol{a}) = \{\boldsymbol{z}_1, \boldsymbol{z}_2\}; \quad \xi_1 = p \in (0,1); \quad \xi_2 = (1-p); \quad K(p) = C(\boldsymbol{x}_i, \boldsymbol{a}).$$

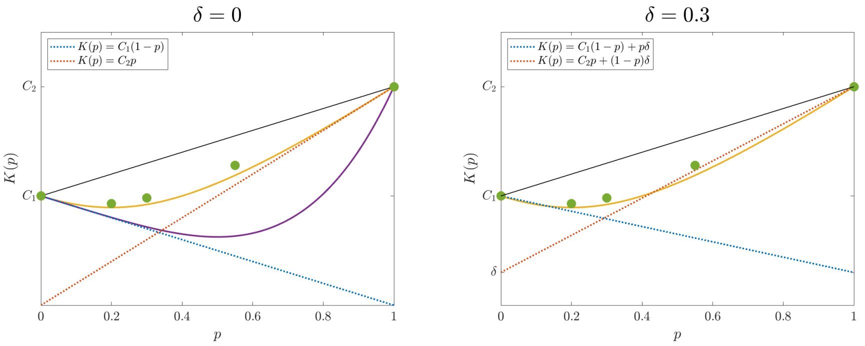

OSSP is monotone δ-causal if all points $(p, K(p))$ are on or above the dotted restriction lines.

Examples above: Orange and green graphs are MC, but purple is not. Green is also δ-MC for $\delta = 0.3$, but orange is not.

A new OSSP based AV-routing framework

- Deterministic transitions to continue in the same lane, with traffic dependent costs.
- Stochastic transitions to attempt lane changes ($m = 2$ possible outcomes).

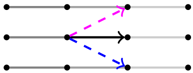

- (Infinitely) many lane change actions available to reflect different urgency levels, interpreted as probability of success $p \in [0, 1]$.
- "Urgency" translates into willingness to alter velocity; so, the cost $K(p)$ is monotone increasing.
- Easy to find suitable cost models that ensure MC; e.g., $K(p) = \beta p^2 + \gamma$, with $\beta, \gamma > 0$ determined by traffic patterns.

Subject to U.S. Provisional Patent 10471-02-US.

Significantly extends a previous SSP-routing approach (Jones, Haas-Heger, and van den Berg; 2022).

A simplified Magic Roundabout (MR)

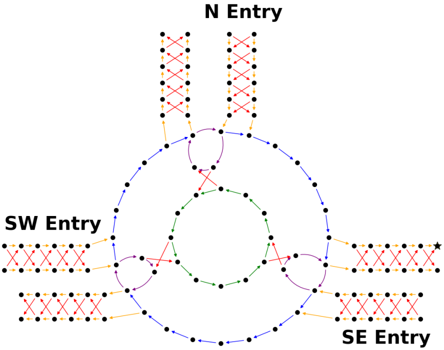

- Trying to reach ★ (on SE Exit), minimizing the expected travel time.
- When approaching MR, your lane determines the initial direction of travel (clockwise or counterclockwise).
- Which one is quicker/easier depends on the traffic distribution.
- The success of lane-change attempts is uncertain, but you can influence it (e.g., by slowing down).

How urgently should you try to switch lanes while approaching MR?

When Congestion is Heaviest Around INNER Roundabout

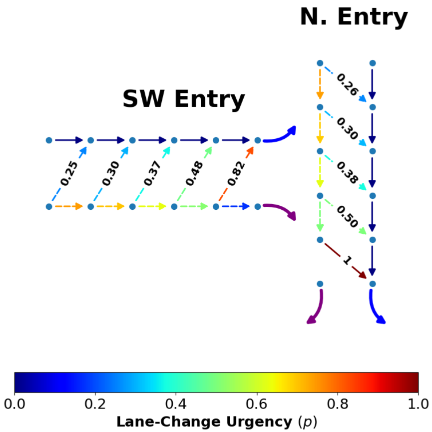

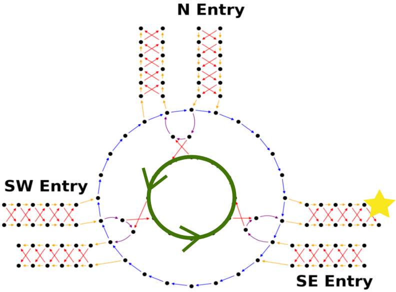

When Congestion is Heaviest Around OUTER Roundabout

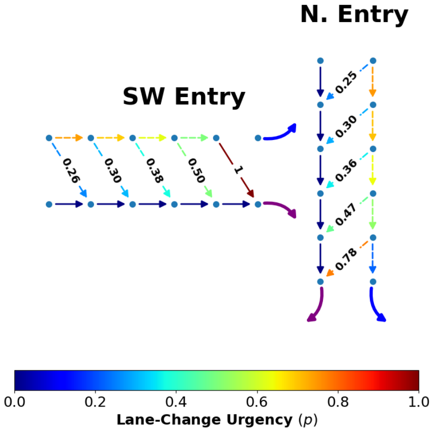

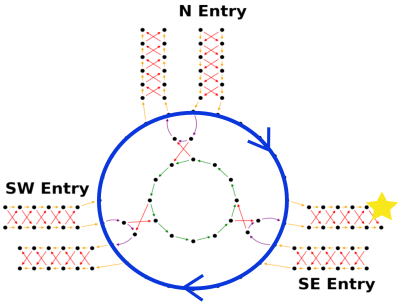

Controlled system:

$$\begin{cases} \boldsymbol{y}'(t) = \boldsymbol{v}(\boldsymbol{y}(t), \boldsymbol{a}(t)), & \text{velocity } \boldsymbol{v} : \Omega \times A \longmapsto \mathbb{R}^d; \\ \boldsymbol{y}(0) = \boldsymbol{x}, & \boldsymbol{x} \in \Omega \subset \mathbb{R}^d. \end{cases}$$

Time-to-destination $T_{\boldsymbol{a}(\cdot), \boldsymbol{x}} = \min \left\{ t \in \mathbb{R}_{+,0} \,\middle|\, \boldsymbol{y}(t) \in Q \subset \partial\Omega \right\}.$

Value function $u(\boldsymbol{x}) = \inf\limits_{\boldsymbol{a}(\cdot)} T_{\boldsymbol{a}(\cdot), \boldsymbol{x}}.$

Viscosity solution of a Hamilton-Jacobi-Bellman PDE:

$$\begin{aligned} \min_{\boldsymbol{a} \in A} \{ \nabla u(\boldsymbol{x}) \cdot \boldsymbol{v}(\boldsymbol{x}, \boldsymbol{a}) + 1 \} \;=\; 0, \qquad & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = 0, \qquad & \boldsymbol{x} \in Q. \end{aligned}$$

Controlled system:

$$\begin{cases} \boldsymbol{y}'(t) = \boldsymbol{v}(\boldsymbol{y}(t), \boldsymbol{a}(t)), & \text{velocity } \boldsymbol{v} : \Omega \times A \longmapsto \mathbb{R}^d; \\ \boldsymbol{y}(0) = \boldsymbol{x}, & \boldsymbol{x} \in \Omega \subset \mathbb{R}^d. \end{cases}$$

Time-to-destination $T_{\boldsymbol{a}(\cdot), \boldsymbol{x}} = \min \{ t \in \mathbb{R}_{+,0} \mid \boldsymbol{y}(t) \in Q \subset \partial\Omega \}$.

Value function $u(\boldsymbol{x}) = \inf_{\boldsymbol{a}(\cdot)} T_{\boldsymbol{a}(\cdot), \boldsymbol{x}}$.

Viscosity solution of a Hamilton-Jacobi-Bellman PDE:

$$\begin{aligned} \min_{\boldsymbol{a} \in A} \{ \nabla u(\boldsymbol{x}) \cdot \boldsymbol{v}(\boldsymbol{x}, \boldsymbol{a}) + 1 \} &= 0, & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) &= 0, & \boldsymbol{x} \in Q. \end{aligned}$$

Geometric dynamics: velocity $\boldsymbol{v} = f(\boldsymbol{x}, \boldsymbol{a})\boldsymbol{a}$ with the speed f and controls = directions of motion (i.e., $A = S^1$).

The isotropic case: direction-independent speed (i.e., $f(\boldsymbol{x}, \boldsymbol{a}) = f(\boldsymbol{x})$) results in a much simpler Eikonal PDE: $\|\nabla u(\boldsymbol{x})\| f(\boldsymbol{x}) = 1$.

But why should control-theorists care about SSPs?

SSPs are useful in approximating continuous optimal control; see, e.g. **[Kushner, 1977]**.

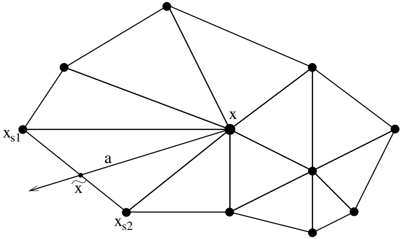

- $\tilde{\boldsymbol{x}} = \xi_1 \boldsymbol{x}_{\boldsymbol{s},1} + \xi_2 \boldsymbol{x}_{\boldsymbol{s},2}$
- $D(\xi) = \|\tilde{\boldsymbol{x}} - \boldsymbol{x}\| = \|(\xi_1 \boldsymbol{x}_{\boldsymbol{s},1} + \xi_2 \boldsymbol{x}_{\boldsymbol{s},2}) - \boldsymbol{x}\|$.
- $\boldsymbol{a} = \boldsymbol{a}_\xi = \frac{\tilde{\boldsymbol{x}} - \boldsymbol{x}}{D(\xi)}$.

$$\begin{aligned} V_s(\boldsymbol{x}) &= \min_{\xi \in \Xi_2} \left\{ \frac{D(\xi)}{f(\boldsymbol{x}, \boldsymbol{a}_\xi)} + \xi_1 U(\boldsymbol{x}_{\boldsymbol{s},1}) + \xi_2 U(\boldsymbol{x}_{\boldsymbol{s},2}) \right\}; \\ U(\boldsymbol{x}) &= \min_{s \in S(\boldsymbol{x})} V_s(\boldsymbol{x}); \qquad \forall \boldsymbol{x} \in X \cap \Omega. \end{aligned}$$

- $S(\boldsymbol{x})$ is the set of adjacent simplexes and $C^s(\boldsymbol{x}, \xi) = D(\xi)/f(\boldsymbol{x}, \boldsymbol{a}_\xi)$.
- $\mathrm{U}(\boldsymbol{x}) = 0$ for all $\boldsymbol{x} \in X \cap \partial\Omega$.

Two simple stencils in $\mathbb{R}^2$:

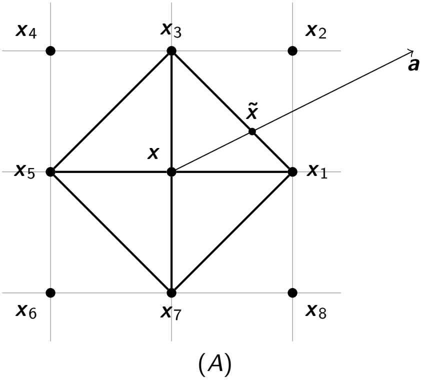

For Eikonal PDEs (the isotropic case, $f(\boldsymbol{x}, \boldsymbol{a}) = f(\boldsymbol{x})$):
Tsitsiklis (1995) showed that semi-Lagrangian discretizations are (MC) on both stencils **and** (B) is also "δ-causal" for $\delta = \frac{h}{f_{max}\sqrt{2}}$.

On stencil (A), Tsitsiklis' first algorithm is **equivalent** to Sethian's Fast Marching Method (1996).

The "monotone ordering" decoupling does not work here:
characteristics and gradient lines do not have to be the same.
Nor do they have to lie in the same simplex!

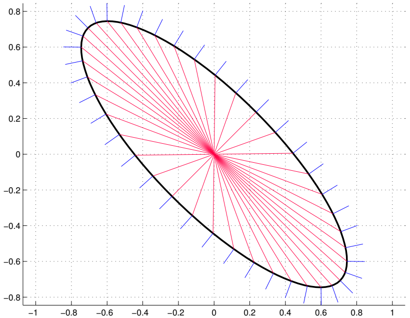

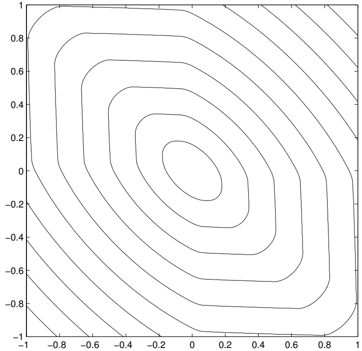

characteristic for $\boldsymbol{x}$ lies in the simplex $\boldsymbol{x}\boldsymbol{x}_1\boldsymbol{x}_2$ $\quad \not\Rightarrow \quad$ $u(\boldsymbol{x}) > \max\{u(\boldsymbol{x}_1), u(\boldsymbol{x}_2)\}$

Label-setting methods for HJB Equations

If the problem is **isotropic** (i.e. $f(\boldsymbol{x}, \boldsymbol{a}) = f(\boldsymbol{x})$), the same monotone de-coupling works:
"If you use The Known to tentatively compute The Still Unknown, then the smallest of The Tentatively Known is actually Known."

- **Dijkstra-like:** (Tsitsiklis, 1995); (Sethian, 1996); (Kimmel & Sethian, 1998); (Sethian, 1999); (Sethian & AV, 2000); (Potter & Cameron, 2019 & 2021).
- **Dial-like:** (Tsitsiklis, 1995); (Kim et al., 2000); (AV, 2008).

For **anisotropic** HJB equations, "local" stencils need not be causal.
But extended stencils can be used to restore MC!

(Sethian & AV, 2001 & 2003); (AV, 2008); (Alton & Mitchell, 2012); (Cameron, 2012); (Mirebeau, 2014); (Dahiya & Cameron, 2018); (Desquilbet et al., 2021).

Previous criteria for checking whether a stencil is causal for a particular anisotropic problem were analytic & somewhat cumbersome.

Our δ-MC OSSP criteria yield a simple/geometric interpretation

and identify **all** anisotropic problems compatible with a specific stencil.

For which speed profiles is your chosen stencil (δ-)MC?

In $\mathbb{R}^2$, a simple geometric answer based on our OSSP MC criterion!

1. For each stencil-represented direction $\boldsymbol{z}_i$,
draw the corresponding velocity vector $\boldsymbol{v}_i$.
2. Form parallelograms based on pairs
of velocity vectors from each simplex.
3. A union of these parallelograms defines a "sunflower".
4. If the speed profile $\mathcal{V}_f(\boldsymbol{x}) = \left\{ f(\boldsymbol{x}, \boldsymbol{a})\boldsymbol{a} \mid \boldsymbol{a} \in S^1 \right\}$ is fully contained
in the sunflower drawn at that gridpoint for each $\boldsymbol{x} \in X$,
then the stencil is MC.

For $\delta > 0$, the δ-MC condition is the same, but parallelograms are replaced by smaller quadrilaterals, with one δ-dependent vertex in each.

5 MC and 1 non-MC stencil/speed profile combinations

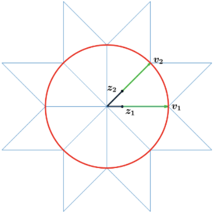

(A)

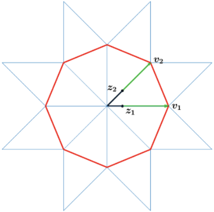

(B)

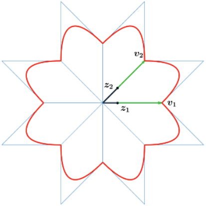

(C)

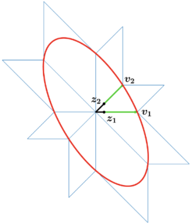

(D)

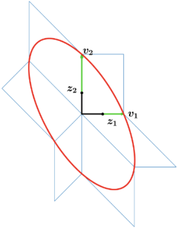

(E)

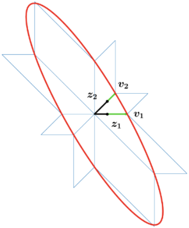

(F)

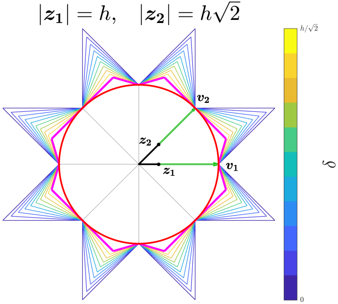

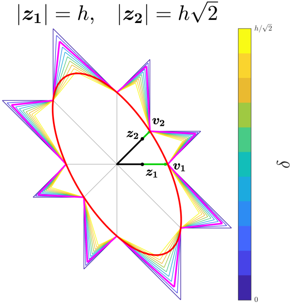

Each "sunflower" color corresponds to a specific $\delta > 0$.
Magenta indicates the largest δ that works for shown profiles.
The bigger is δ, the faster Dial's method will generally be.

A sharp (δ-)MC criterion for OSSPs

Assuming $\boldsymbol{a} \in A_i$ is not deterministic (i.e., $m > 1$) and choosing any specific $r \in \{1, ..., m\}$, we define $\boldsymbol{\gamma}_r = (\gamma_{r,1}, ..., \gamma_{r,m})$ to be an oblique (proportional) projection of $\boldsymbol{\xi}$ as follows

$$\gamma_{r,j} = \begin{cases} 0, & \text{if } j = r; \\ \xi_j/(1-\xi_r), & \text{otherwise.} \end{cases}$$

Theorem

Suppose there exists a $\delta \geq 0$ such that,
for all $\boldsymbol{x}_i \neq \mathbf{t}$, $\boldsymbol{a} \in A_i$,

- *if $\boldsymbol{a}$ is deterministic, then $C(\boldsymbol{x}_i, \boldsymbol{a}) \geq \delta$;*
- *if $\boldsymbol{a}$ is not deterministic, then*

$$C(\boldsymbol{x}_i, \boldsymbol{a}) \geq (1-\xi_r)\check{C}(\boldsymbol{\gamma}_r) + \xi_r \delta, \qquad \forall r \in \{1, ..., m = |\mathcal{I}(\boldsymbol{a})|\}.$$

If these conditions are satisfied, this OSSP is monotone causal and Dijkstra's method is applicable. If $\delta > 0$, the OSSP is monotone δ-causal and Dial's method with buckets of width δ is also applicable.

Sharp for any m. Equivalent to our previous criterion for $m = 2$.

Example: MC criterion in $\mathbb{R}^3$

Question: Suppose we can move with unit speed in each coordinate plane. How anisotropic can the full 3D speed profile be if we want Dijkstra's method to work on a Cartesian grid with a standard 6-point stencil?

Example: MC criterion in $\mathbb{R}^3$

Question: Suppose we can move with unit speed in each coordinate plane. How anisotropic can the full 3D speed profile be if we want Dijkstra's method to work on a Cartesian grid with a standard 6-point stencil?

Answer: Our sharp MC criterion guarantees that Dijkstra's will solve the HJB-discretization correctly as long as the speed profile $\mathcal{V}_f$ is contained in a **tri-cylinder**:

$$\mathcal{V}_f \subset \left\{ (v_1, v_2, v_3) \in \mathbb{R}^3 \mid \max\left(v_1^2 + v_2^2,\, v_2^2 + v_3^2,\, v_1^2 + v_3^2, \right) \leq 1 \right\}.$$

Example: MC criterion in $\mathbb{R}^3$

Question: Suppose we can move with unit speed in each coordinate plane. How anisotropic can the full 3D speed profile be if we want Dijkstra's method to work on a Cartesian grid with a standard 6-point stencil?

Answer: Our sharp MC criterion guarantees that Dijkstra's will solve the HJB-discretization correctly as long as the speed profile $\mathcal{V}_f$ is contained in a **tri-cylinder**:

$$\mathcal{V}_f \ \subset \left\{ (v_1, v_2, v_3) \in \mathbb{R}^3 \ \middle| \ \max\left(v_1^2 + v_2^2,\ v_2^2 + v_3^2,\ v_1^2 + v_3^2, \right) \ \leq \ 1 \right\}.$$

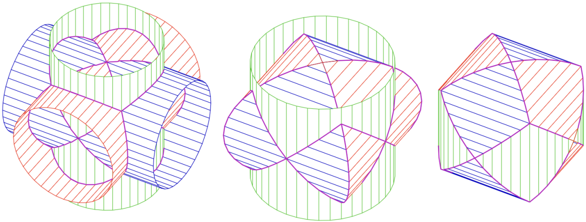

Source: Wikimedia Commons. Created by Ag2gaeh;

Conclusions

- SSPs are useful models for discrete dynamic programming and discretization of HJB PDEs, but can be computationally costly.
- OSSPs are an important subclass, for which the applicability of label-setting algorithms is easy to verify a priori.
- Such Monotone (δ)-Causal OSSPs are much more practical, allowing for frequent online replanning in dynamic environments.
- Strategic-Tactical Plans based on MC OSSPs provide an efficient routing approach for autonomous vehicles, capturing the inherent uncertainty of lane-change maneuvers and modeling a spectrum of "urgency levels" in implementing them.

Details: M. Gaspard and A. Vladimirsky, "Monotone Causality in Opportunistically Stochastic Shortest Path Problems".
Submitted to *Mathematics of Operations Research*.
https://arxiv.org/abs/2310.14121