# The impact of optimisation and machine learning on robotics

**Martim Brandão**
Lecturer in Robotics and Autonomous Systems
King’s College London

# Responsible Robotics and AI lab

**&**

# Safe and Trusted AI CDT

- Characterizing issues of Fairness, Accountability, Transparency, Ethics in AI
- New algorithms towards mitigating those issues

# My background: Robotics


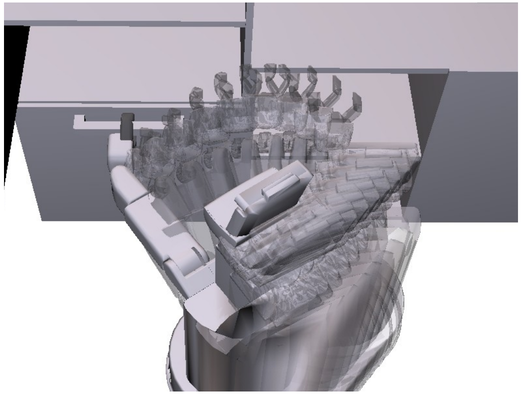

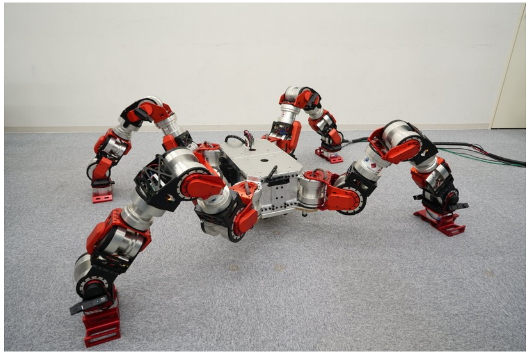

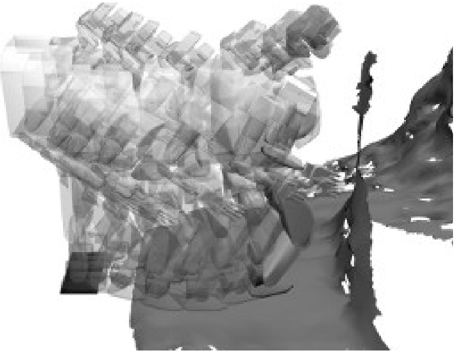

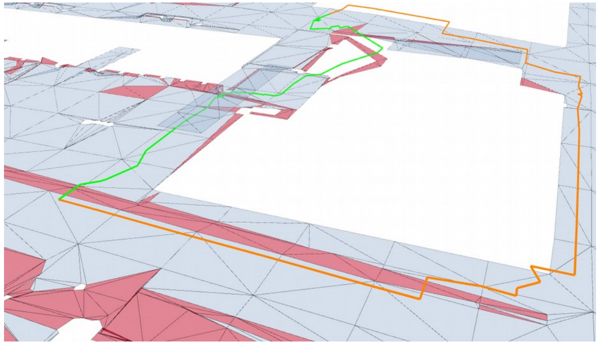

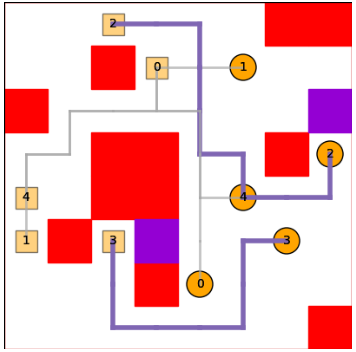

# How has modern computing/AI/ML affected robotics?


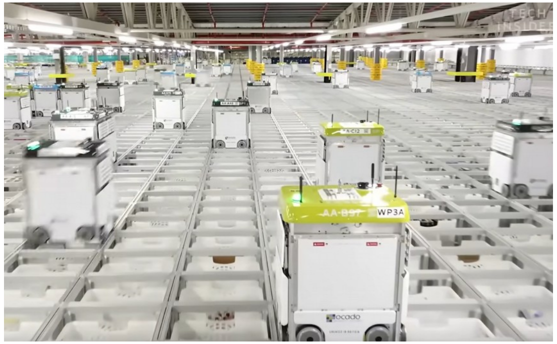


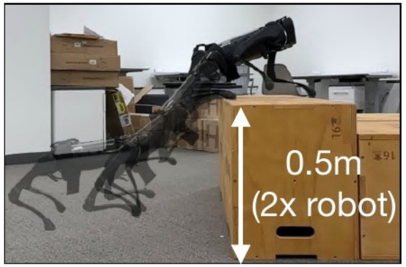


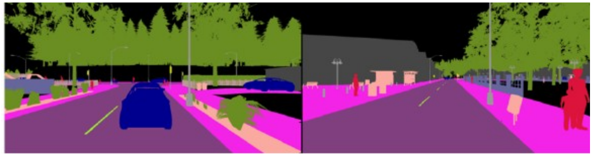

# Large-scale MAPF

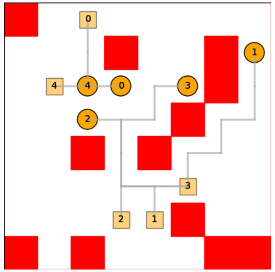


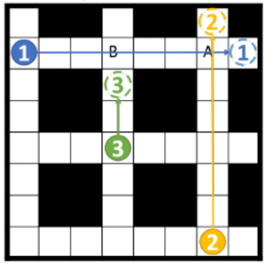


Fast sub-optimal heuristics
(conflict resolution, prioritization,
symmetry-breaking...)

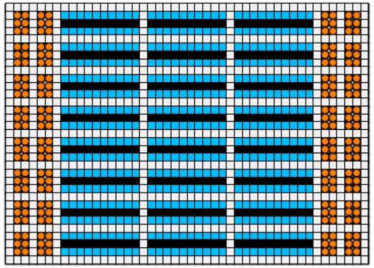

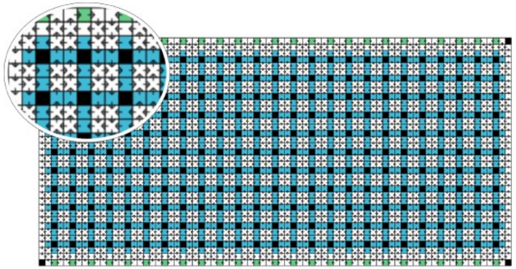

# Dynamic robot locomotion

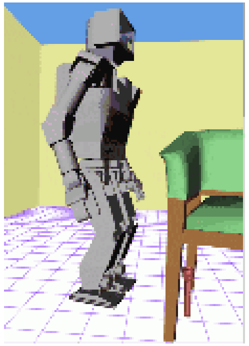

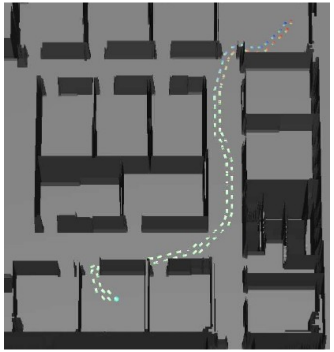

$$
\begin{aligned}
\underset{\Gamma}{\text{minimize}} \quad & \sum_{k=1}^{M} L(\mathbf{q}[k], \mathbf{v}[k], \ddot{\mathbf{r}}[k], \boldsymbol{\lambda}[k], h[k]) & \\
\text{subject to} \quad & m\ddot{\mathbf{r}}[k] = m\mathbf{g} + \sum_j \boldsymbol{\lambda}_j[k] & \text{(linear momentum)} \\
& \mathbf{k}[k] = \mathbf{A}_G^{\mathbf{k}}(\mathbf{q}[k])\mathbf{v}[k] & \text{(angular momentum)} \\
& \dot{\mathbf{k}}[k] = \sum_j (\mathbf{c}_j[k] - \mathbf{r}[k]) \times \boldsymbol{\lambda}_j & \text{(angular momentum rate)} \\
& \forall_j \;\; \boldsymbol{\lambda}_j[k] = \sum_{i=1}^{N_d} \beta_{ij}[k]\mathbf{w}_{ij} & \text{(friction)} \\
& \forall_{i,j} \;\; \beta_{ij}[k] \geq 0 & \\
& \mathbf{r}[k] = COM(\mathbf{q}[k]) & \text{(COM location)} \\
& \text{Kinematic constraints} & \\
& \text{Time integration constraints} &
\end{aligned}
$$

Full-body dynamics models,

on fast (commercial) optimization solvers

Mastalli et al 2023 inverse-dynamics MPC via nullspace resolution

Boston Dynamics

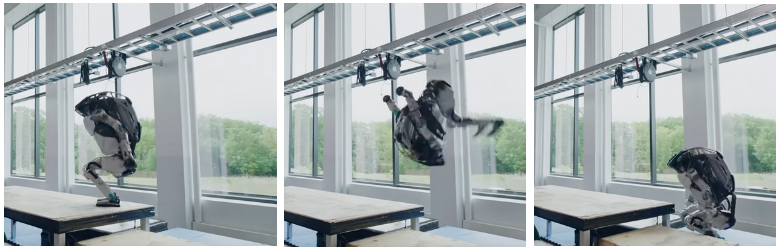

Boston Dynamics

# Robust & reactive motion

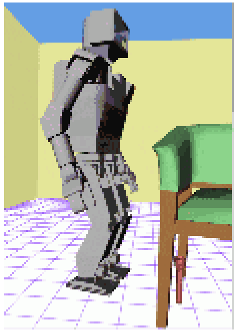

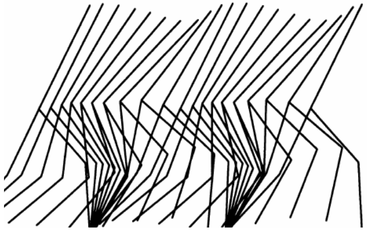

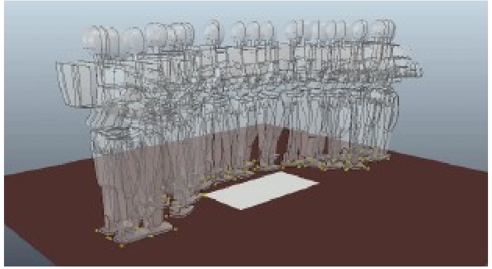

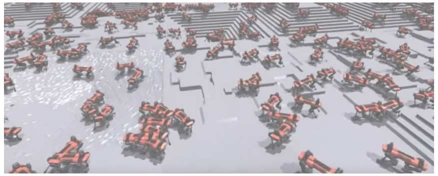

Rudin et al 2021 Learning to walk in minutes

Reinforcement Learning
on large-scale GPU-accelerated
realistic simulation platforms


Cheng et al 2023 Extreme parkour with legged robots

# Image understanding

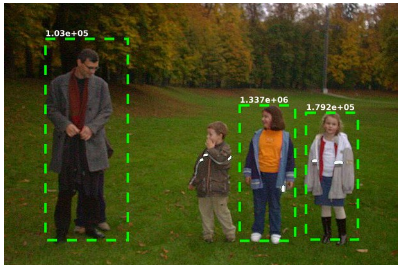


Deep Neural Networks,
large-scale real datasets and game-like realistic simulation platforms

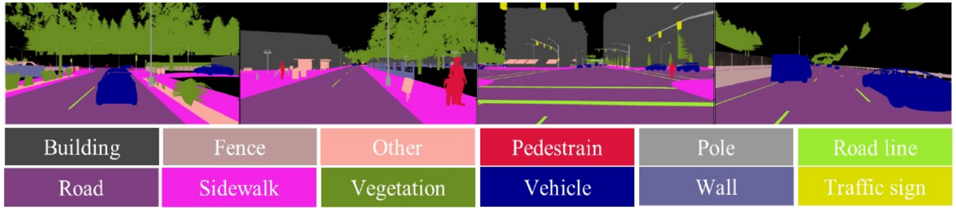


Wang et al 2022 SFNET-N

Schult et al 2023 Mask3D

# BUT NOW WE HAVE...

## Huge models

(DNNs, 3D robot/world models, full-body dynamics equations...)

## Algorithms with lots of heuristics

(Convexification, symmetry-breaking, MIP, anytime suboptimal algos...)

Debugging and accountability problems

**Need for explanation-generation algorithms ("XAI")**

# Explainability requirements for robot planning

## Explain what?

| Kinds of questions | Example questions |
|---|---|
| Plan-centered | why is this plan feasible?<br>why is this plan infeasible?<br>why is there no plan? (why did the planner fail to obtain a plan?)<br>why did agents X traverse through locations L?<br>why did agents X take paths Y, when I expected them to take path Z?<br>why did agents X take paths that are so far off of paths Y?<br>why did agents X wait/stop/get-precedence here?<br>why did agents X wait for that other agent here?<br>why did agents X wait for so long here?<br>why is agent X not at location P before time T?<br>why is agent X on a collision course?<br>why is there a deadlock between agents X?<br>why is there a congestion in area A? |
| Metric-centered | why isn't metric M higher?<br>why is this plan optimal?<br>why is this plan sub-optimal? |
| Consistency-centered | why does agent X always get precedence (over multiple problems)?<br>why does agent X agent always go through location X (over multiple problems)?<br>why is there always congestion in this area (over multiple problems)? |
| Algorithm-centered | why does heuristic H1 perform better than H2 in this problem/map?<br>why does algorithm/variant A1 perform better than A2 in this problem/map?<br>why does the algorithm take so long to find a path in this problem/map? |
| Design-centered | why do I need this many agents? |

Brandao et al, “Explainability in Multi-Agent Path/Motion Planning”, AAMAS 2022

# Explainability requirements

## Kinds of explanations

| | | | |
|---|---|---|---|
| Method-centered explanation | Because of the use of approximate collision checking<br>Because of the use of gradient methods on an unsmooth problem<br>Because the objective promotes solutions of following types<br>Because of the computation time budget used<br>Because hyper-parameter x is not y<br>Because of the algorithm's initialization scheme<br>Because the methods got stuck at an infeasible local minimum<br>Because cost weights are x not y<br>Because of chance<br>Because the algorithm did not obtain enough samples<br>Because of incorrect pruning of the search tree in state x<br>Because of the choice of planner<br>Because of a software bug | Problem-centered explanation | Because of conflicting constraints (x conflicts with y)<br>Because constraint x cannot be satisfied even by itself<br>Because of occupied space in region x (obstacle y)<br>Because of free space in region x<br>Because the robot dynamics has effect x<br>Because of the volume of robot part x<br>Because the problem has no solution<br>Because object x is not close enough<br>Because the start/waypoint/goal cannot be satisfied<br>Because the environment is cluttered<br>Because of a singularity in region x<br>Because path A is lower-cost<br>Because path A is safer/more-efficient<br>Because B is out of the workspace<br>Because path B crosses unmapped space<br>Because of a bug-trap in region x<br>Because of problem difficulty |
| Unknown | I dont know the reasons for x | Visualization-centered explanation | Visualize explored actions and their feasibility regions<br>Visualize where expected paths become infeasible<br>Visualize which part/link leads to not finding a plan and where<br>Visualize map/plan areas that are problematic/bottlenecks<br>Visualize with colours why the robot moved the way it did<br>Visualize families of good solutions coloured by performance |

Most answers to "what kind of planner-generated explanations would be **useful**?" were problem- and visualization-centered.

Brandao et al, "How experts explain motion planner output", ROMAN 2021

# Kinds of explanations

[multi-agent setting]

**Table 3: Kinds of explanations for MAPF/MRMP**
**(I=Initial examples from interviewees; R=Refinements by questionnaire participants)**

| Category | Sub-category | Example explanations | I | R |
|---|---|---|---|---|
| Plan-based | Agent interference | because that would delay agents X | | * |
| | | because that would create conflicts with agents X | | * |
| | | because of a deadlock between agents X | * | |
| | Event chains | because that would affect agent X, which would affect agent Y | | * |
| | | because of event X at time T that propagated | * | |
| | Landmarks | because a large number of agents have to go through area A | * | |
| Problem-based | Problem properties | because the environment is not well-formed/well-structured | * | |
| | Metric | because that would be worse according to the metric | * | * |
| | | because that would require change X to the metric | | * |
| | | because the metric would be the same | | * |
| | Constraints | because agents X have higher priority (imposed by the problem) | * | * |
| | | because that would require an extra constraint/precondition X | * | * |
| | | because of the kinodynamic constraints of agents X | | * |
| | Map | because of the costs assigned to edges/regions X | * | |
| | | because object O is at location X | * | |
| | | because of obstacles X | * | * |
| | | because of the size of the map | | * |
| | | because of obstacle density (at location X) | | * |
| | | because of a choke point at location X | * | * |
| | | because that would require a change to the map | * | * |
| | Agents | because there are too many/few agents | * | * |
| | | because that would require X more/less agents | * | * |
| | | because that would require agents to have priorities X | * | |
| | Example | for the same reason as in this smaller problem | * | |
| Algorithm-based | Planner properties | because the planner does not provide safety guarantees | * | |
| | | because the planner is incomplete | * | |
| | | because the planner is sub-optimal | * | |
| | | because heuristic H is inadmissible | * | |
| | | because the planner is better on small/large maps | | * |
| | | because the planner cannot handle a large number of agents well | | * |
| | | because the planner explores movement in direction X first | | * |
| | | because heuristic prefers moving agents close to their destination first | | * |
| | Planner decisions | because of an incorrect state expansion/prune | * | |
| | | because of planner-added constraint X (e.g. PBS, MAPF/C, k-delay MAPF) | * | * |
| | | because of the priority ordering of the agents (imposed by planner) | * | * |
| | Planner parameters | because algorithm A1 was used instead of A2 | | * |
| | | because heuristic H1 was used instead of H2 | | * |
| | | because hyperparameter X was equal/lower/higher than Y | | * |
| | Information | because of incorrect information (decentralized algorithms) | | * |
| | | because of insufficient information (decentralized algorithms) | | * |
| | | because the algorithm was not trained on similar examples (learning algorithms) | | * |
| | Bug | because of a bug in function/class/file/line X | * | * |
| Execution-based | Execution gap | because agents X stopped due to technical failure | * | |
| | | because agents X are moving slower/faster than expected | | * |

Brandao et al, “Explainability in Multi-Agent Path/Motion Planning”, AAMAS 2022

# Explainability use-cases

## In motion planning

- **Warehouse automation**
  - Large-scale (1000s agents), dynamic, anytime sub-optimal planners.
  - Explanations could **help developers** tune cost functions, **onsite engineers** add constraints/params on the fly, **layout designers** improve layout
- **Computer games**
  - Large-scale, dynamic, sub-optimal planners, no safety guarantees.
  - Explanations could **help game designers** understand reasons for undesired behavior, know what to change to improve game.
- **Mining**
  - Continuous, kinodynamic, uncertainty-aware, distributed planners.
  - Explanations could **help managers** understand why metrics are being optimized.
- **All applications**
  - Explanations could **help developers and researchers** inspect models and algorithms, find bugs, understand why some heuristic/parameters/variants work better than others.

Brandao et al, “Explainability in Multi-Agent Path/Motion Planning”, AAMAS 2022

# Explainable **path planning**

Problem-centred explanations

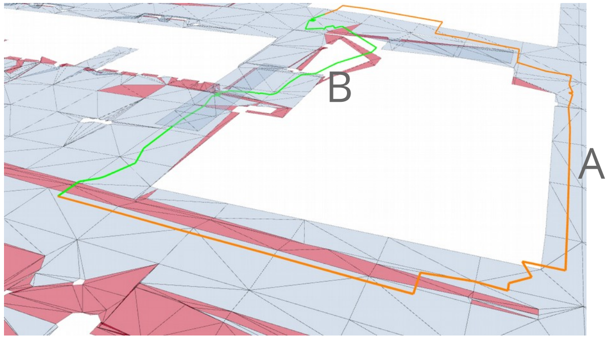


Why plan A, instead of B which I expected?

*Terrain-MILP:*

$$
\begin{aligned}
\min_{l',\pi,\lambda} \quad & ||l' - l||_1 \\
\text{s.t.} \quad & \textstyle\sum_i A_{ij}\pi_i = \sum_k d_j c_k l'_{k,r(j)} \quad \forall_{j:x'_j=1} \\
& \textstyle\sum_i A_{ij}\pi_i + \lambda_j = \sum_k d_j c_k l'_{k,r(j)} \quad \forall_{j:x'_j=0} \\
& \pi \in \mathbb{R}^{|V|} \\
& \lambda \in \mathbb{R}^{|E|} \\
& \lambda_j \geq 0 \quad \forall_{j:x'_j=0} \\
& l' \in \{0,1\}^{K|V^r|} \\
& \textstyle\sum_k l'_{k,i} = 1 \quad \forall_i,
\end{aligned}
$$

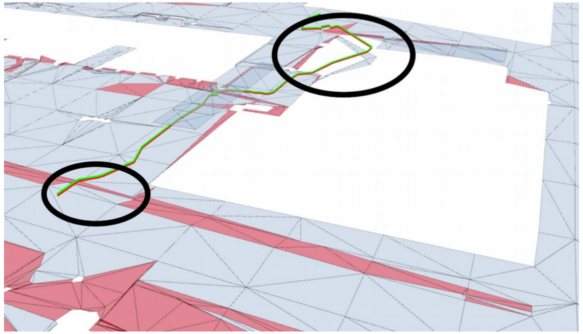

“Because these areas are high-cost”

*COT-MILP:*

$$
\begin{aligned}
\min_{c',\pi,\lambda} \quad & ||c' - c||_1 \\
\text{s.t.} \quad & \textstyle\sum_i A_{ij}\pi_i = \sum_k d_j c'_k l_{k,r(j)} \quad \forall_{j:x'_j=1} \\
& \textstyle\sum_i A_{ij}\pi_i + \lambda_j = \sum_k d_j c'_k l_{k,r(j)} \quad \forall_{j:x'_j=0} \\
& \pi \in \mathbb{R}^{|V|} \\
& \lambda \in \mathbb{R}^{|E|} \\
& \lambda_j \geq 0 \quad \forall_{j:x'_j=0} \\
& c' \in \mathbb{R}^K_+,
\end{aligned}
$$

“Because blue areas are too low cost compared to red”

Brandao, Coles, Magazzeni, “Explaining Path Plan Optimality…”, ICAPS 2021.

# Explainable **road navigation**

Problem-centred explanations

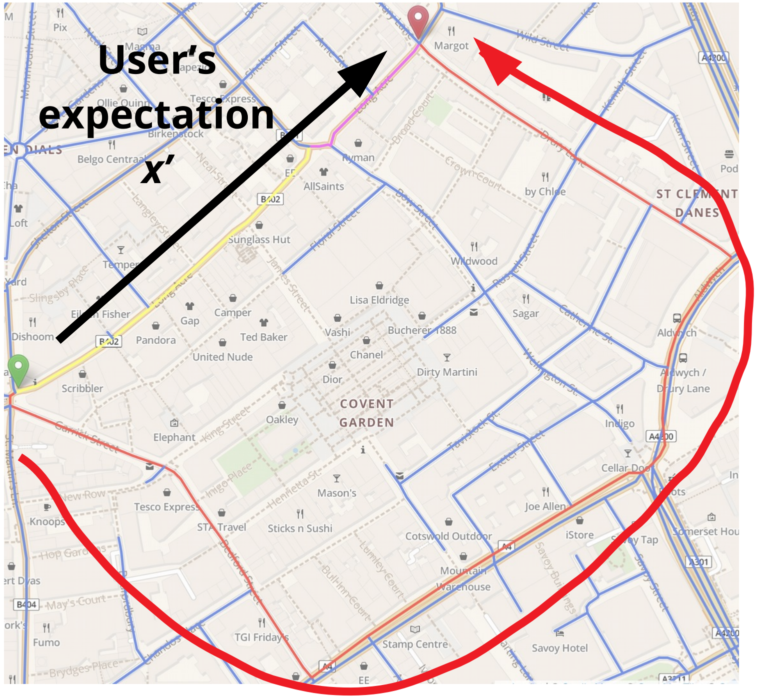


**Shortest Path**

***x=argmin w.x***

**x'**

$$
\begin{aligned}
\underset{\boldsymbol{n}',\boldsymbol{c}',\boldsymbol{b}^{s'},\boldsymbol{b}^{m'},\boldsymbol{\pi},\boldsymbol{\lambda}}{\text{minimize}} \quad & \|\boldsymbol{n}'-\boldsymbol{n}\|_1+\|\boldsymbol{c}'-\boldsymbol{c}\|_1+\|\boldsymbol{b}^{s'}\|_1+\|\boldsymbol{b}^{m'}\|_1 && \text{(5a)}\\
\text{subject to} \quad & && \\
& \textstyle\sum_i A_{ij}\pi_i = s'_j l_j + n'_j M + c'_j M, \quad \forall_{j:x'_j=1 \,\wedge\, m_j/s_j<r}, && \text{(5b)}\\
& \textstyle\sum_i A_{ij}\pi_i = m'_j l_j + n'_j M + c'_j M, \quad \forall_{j:x'_j=1 \,\wedge\, m_j/s_j\geqslant r}, && \text{(5c)}\\
& \textstyle\sum_i A_{ij}\pi_i + \lambda_j = s'_j l_j + n'_j M + c'_j M, \quad \forall_{j:x'_j=0}, && \text{(5d)}\\
& \lambda_j \geqslant 0, \quad \forall_{j:x'_j=0}, && \text{(5e)}\\
& n'_j = n_j, \; c'_j = c_j, \quad \forall_{j:x^*_j=1 \,\wedge\, x'_j=0}, && \text{(5f)}\\
& n'_j = 0, \; c'_j = 0, \quad \forall_{j:x'_j=1}, && \text{(5g)}\\
& \boldsymbol{n}' \in \{0,1\}^{|E|}, \; \boldsymbol{c}' \in \{0,1\}^{|E|}, \; \boldsymbol{b}^{s'} \in \{0,1\}^{|E|}, && \\
& \boldsymbol{b}^{m'} \in \{0,1\}^{|E|}, \; \boldsymbol{\lambda} \in \mathbb{R}^{|E|}, \; \boldsymbol{\pi} \in \mathbb{R}^{|V|}, && \text{(5h)}
\end{aligned}
$$

**$n', c'$, $b^{s'}$, $b^{m'}$**

**"The desired path is not optimal because Endell Street is currently closed; and Long Acre is a one-way road."**

$c'_{EndellStreet}=0$ ; $n'_{LongAcre}=0$

(i.e. these changes make x' optimal)

Alsheeb, Brandao, "Towards Explainable Road Navigation Systems", ITSC2023.

# Explainable **multi-agent planning**

Problem-centred explanations

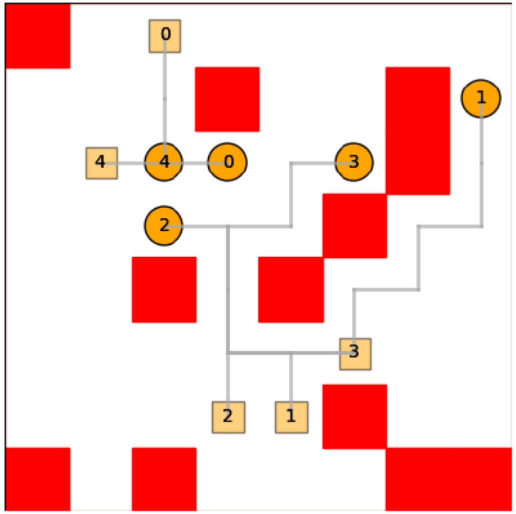


Original plan π

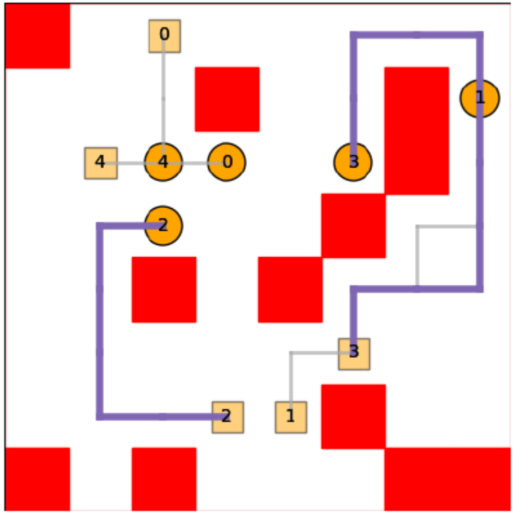


Expected/desired plan π’

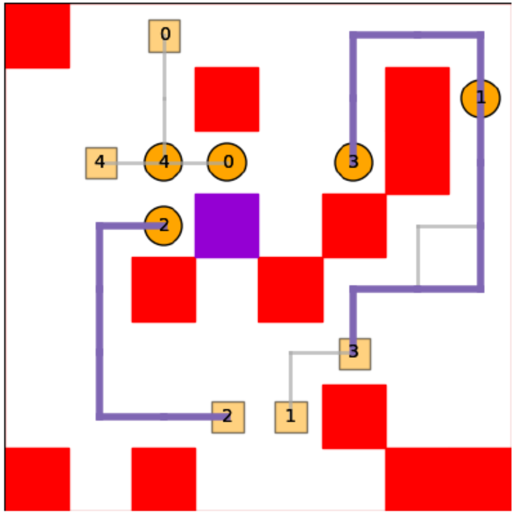


New map where π’ optimal

**“why not π’ instead of π?”**

**→ “because the cells in purple are free**

**(but if they were occupied then π’ would be optimal)”**

[XAIP2022]

# Explainable **motion planning**

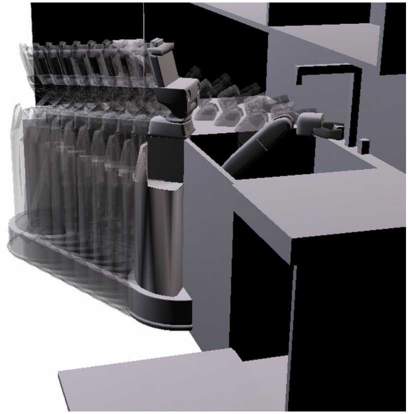

Why did you fail?

## **Problem-centred explanation**

**Algorithm 1** (constraint-based failure explanation):
for each $\mathcal{H}_i$ in $\mathcal{P}(\mathcal{H})$:
    $\xi_i = \text{argmin}_\xi\ C[\xi] + \mathbf{1}_{H_{\text{target}} \in \mathcal{H}_i} \alpha C_{\text{target}}[\xi]$ s.t. $\mathcal{H}_i$
$k = \text{argmin}_i\ C[\xi_i] - \beta|\mathcal{H}_i| + \gamma C_{\text{target}}[\xi_i]$ s.t. $H_{\text{collision}}$
return Message("Not all constraints could be satisfied. The problem would be feasible if constraints $\mathcal{H} \setminus \mathcal{H}_k$ were dropped and the target was $C_{\text{target}}[\xi_k]$ meters away from the original.")

Because constraints "target" and "collision" conflict with each other. The problem would be feasible if the target was 0.15m away.

## **Algorithm-centred explanation**

**Algorithm 2** (initialization-based failure explanation):
for $i = 1, ..., N_{\text{max}}$:
    Pick random initialization and use it below
    $\xi_i = \text{argmin}_\xi\ C[\xi]$ s.t. $\mathcal{H}$
    if $\xi_i$:
        return Message("The initialization was in the basin of attraction of an infeasible local minimum. The planner would succeed with initialization $\xi_i$.")
    else:
        return Message("Unfeasible or hard problem for random initialization.")

Because the initialization was in the basin of attraction of an infeasible local minimum. The planner would have succeeded with a random initialization policy.

Brandao et al, "Towards providing explanations for robot motion planning", ICRA 2021

# NOW WE HAVE...

Huge models

(DNNs, 3D robot/world models, full-body dynamics equations...)

Algorithms with lots of heuristics

(Convexification, symmetry-breaking, MIP, anytime suboptimal algos...)

**Debugging and accountability problems**

**Also: algorithmic bias problems**

**In robotics: biased models → biased physical safety**

# Bias in pedestrian detection?

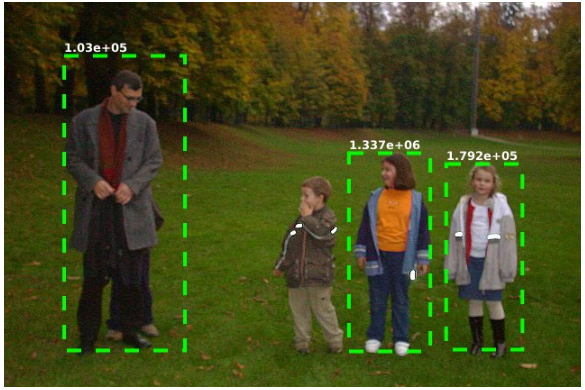



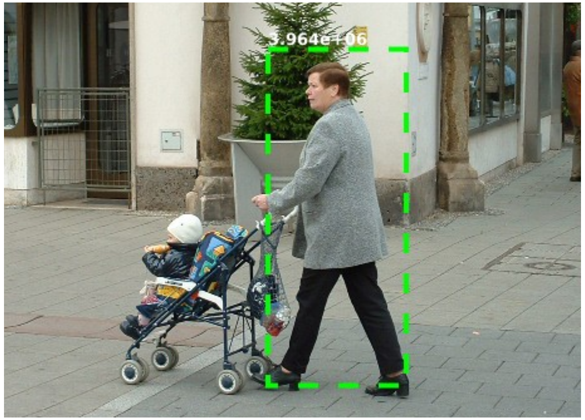


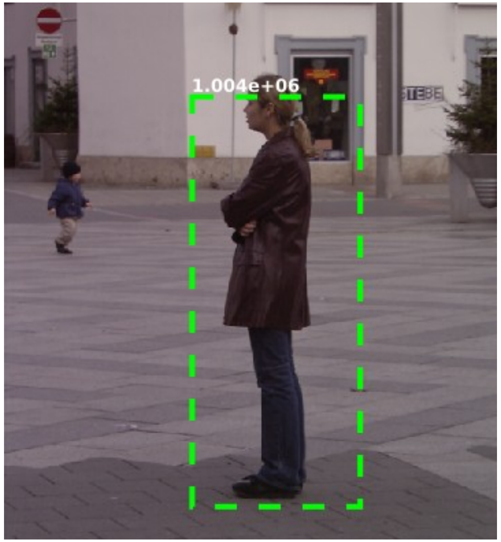


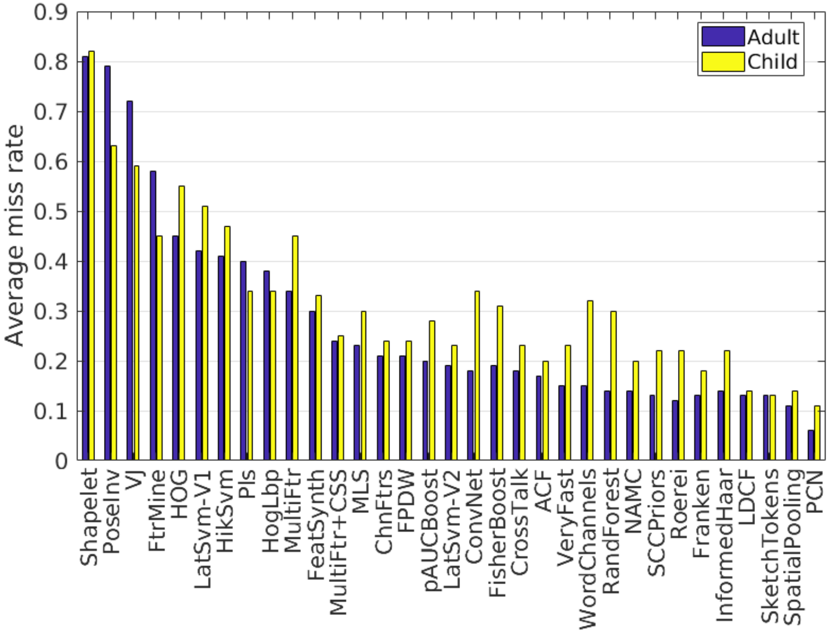


- **All top-24** methods have **higher miss rates on children**.
- Best methods almost **2x miss rates** on children vs adults.
- **Physical safety** differences

Brandao, “Age and gender bias in pedestrian detection algorithms”, CVPR-FATECV 2019 [ACMMM Ambassador Award]

# Pedestrian detection trends

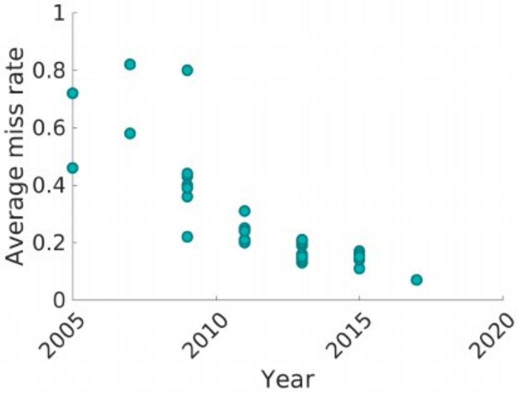


Miss rates going down

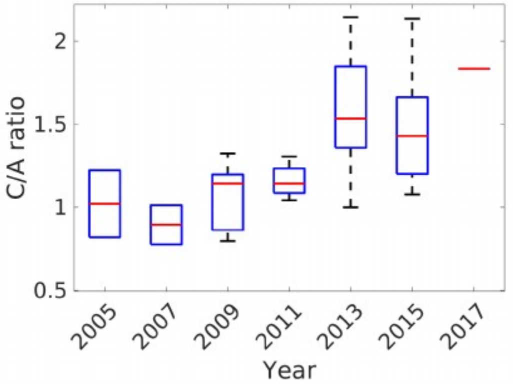


Disparity going up

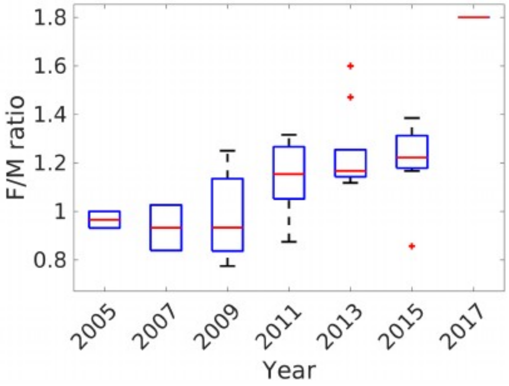


Disparity going up

Increase in **average performance** + Increase in **performance gap** (majority-minority)

Need bias mitigation methods, overfitting avoidance

Brandao, “Age and gender bias in pedestrian detection algorithms”, CVPR-FATECV 2019 [ACMMM Ambassador Award]

# Conclusion

- **Robots have become:** more dynamic, robust, adaptive, large-scale interaction-ready, world-understanding
- **Thanks to:** large physics models, fast optimization solvers, heuristics, large simulated worlds, large neural networks
- **With disadvantages:** difficulties in debugging, predictability, explainability; growing issues of bias on minorities

Still some work to do...

**Thank you**