

Neural Dual Warm-Starts for Scalable Linear Assignment Problem

Ilay Yavlovich, Jad Agbaria, Muhamed Mhamed, Nir Weinberger, Jose Yallouz

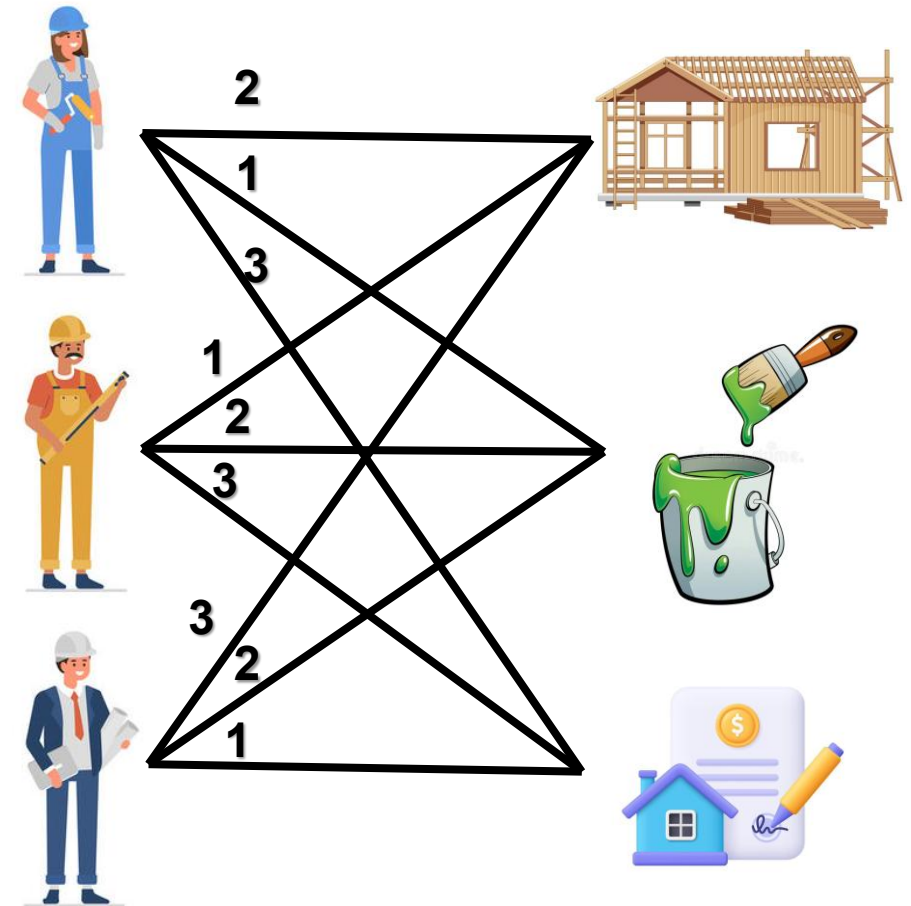


The Linear Assignment Problem

- ▶ Seeks the most efficient way to pair n workers from one set to n tasks from another to minimize the total profit of all assignments.
- ▶ Classical exact solvers scale poorly ($\mathcal{O}(n^3)$ complexity).
 - Hungarian (Kuhn 1955)
 - LAPJV (Jonker & Volgenant 1987)
- ▶ Current neural solvers either:
 - Approximate
 - Fail to scale

Workers

Tasks




The Linear Assignment Problem

- ▶ Seeks the most efficient way to pair n workers from one set to n tasks from another to minimize the total profit of all assignments.
- ▶ Classical exact solvers scale poorly ($\mathcal{O}(n^3)$ complexity).
 - Hungarian (Kuhn 1955)
 - LAPJV (Jonker & Volgenant 1987)
- ▶ Current neural solvers either:
 - Approximate
 - Fail to scale

Workers

Tasks



The Linear Assignment Problem

- ▶ Seeks the most efficient way to pair n workers from one set to n tasks from another to minimize the total profit of all assignments.
- ▶ Classical exact solvers scale poorly ($\mathcal{O}(n^3)$ complexity).
 - Hungarian (Kuhn 1955)
 - LAPJV (Jonker & Volgenant 1987)
- ▶ Current neural solvers either:
 - Approximate
 - Fail to scale



FAST



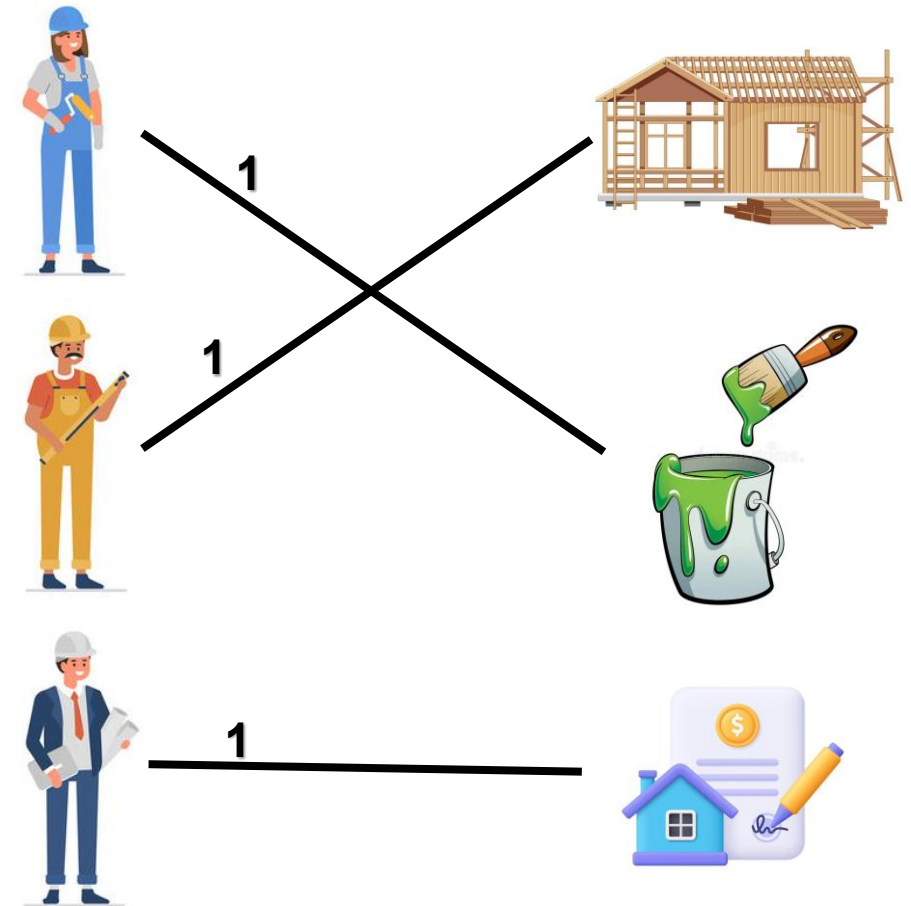
EXACT



SCALABLE

Workers

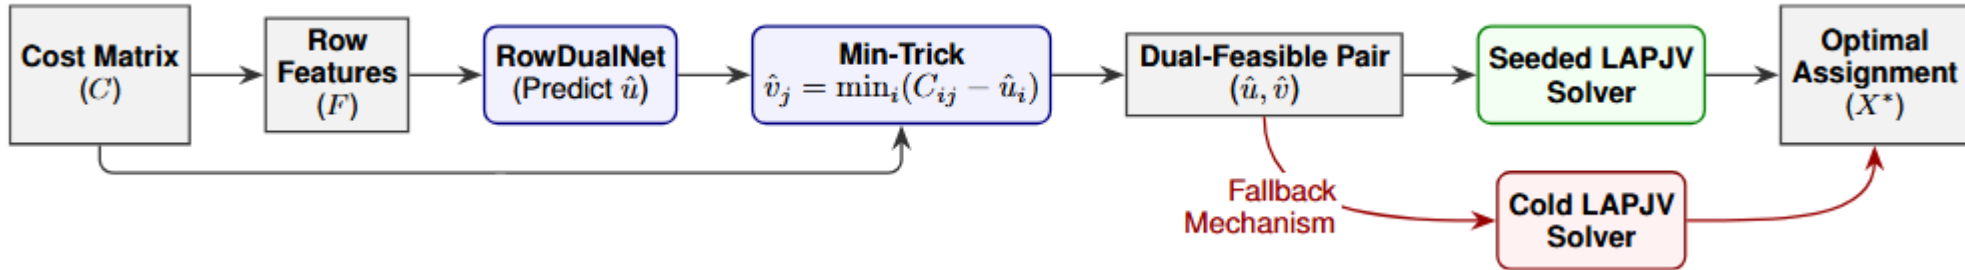
Tasks



LAPJV (Jonker & Volgenant 1987)

- ▶ Relies on finding row and column potentials (u, v) for the cost matrix.
- ▶ Two phases:
 - **Initialization phase:** Initial dual feasible solution
 - **Augmentation phase:** Shortest-path search to update u and v until every worker has a zero-cost task

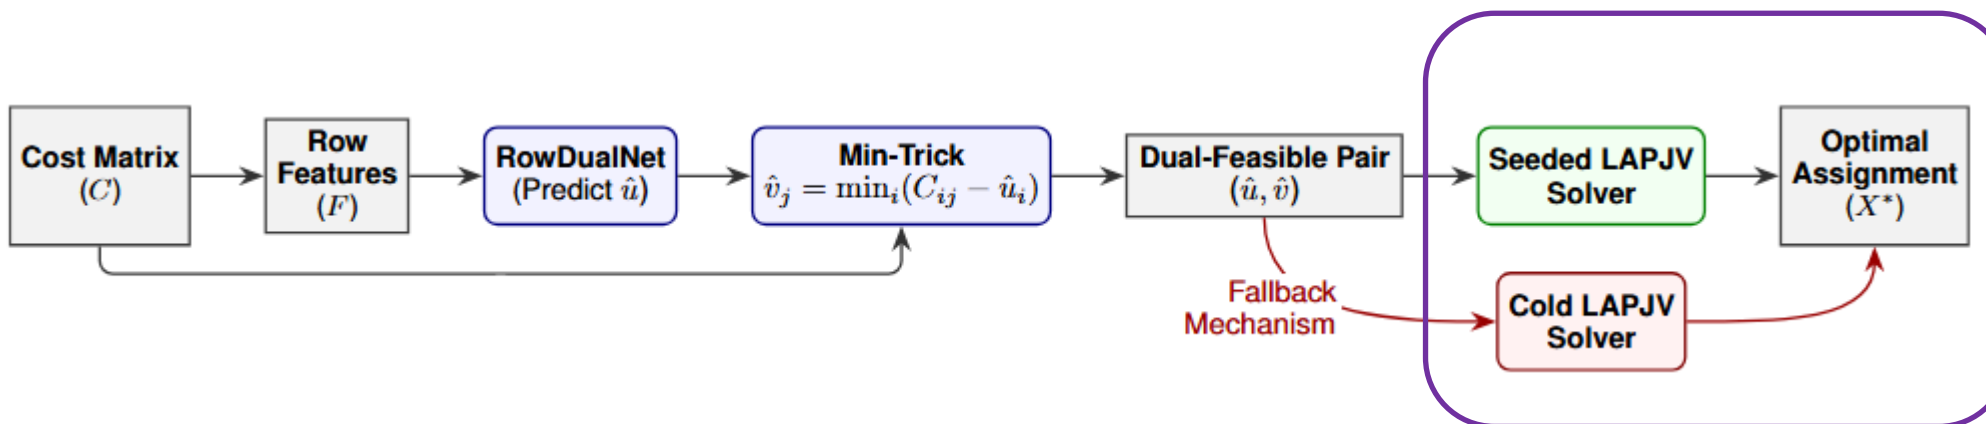
Accelerating via Warm Start



- ▶ **Dual Warm Starting:** Instead of predicting the assignment we predict the dual variables u (row potentials).
- ▶ **Primal-Dual theory:** High-quality duals prune the search space, fewer augmentation steps.




Accelerating via Warm Start



- ▶ **Dual Warm Starting:** Instead of predicting the assignment we predict the dual variables u (row potentials).
- ▶ **Primal-Dual theory:** High-quality duals prune the search space, fewer augmentation steps.

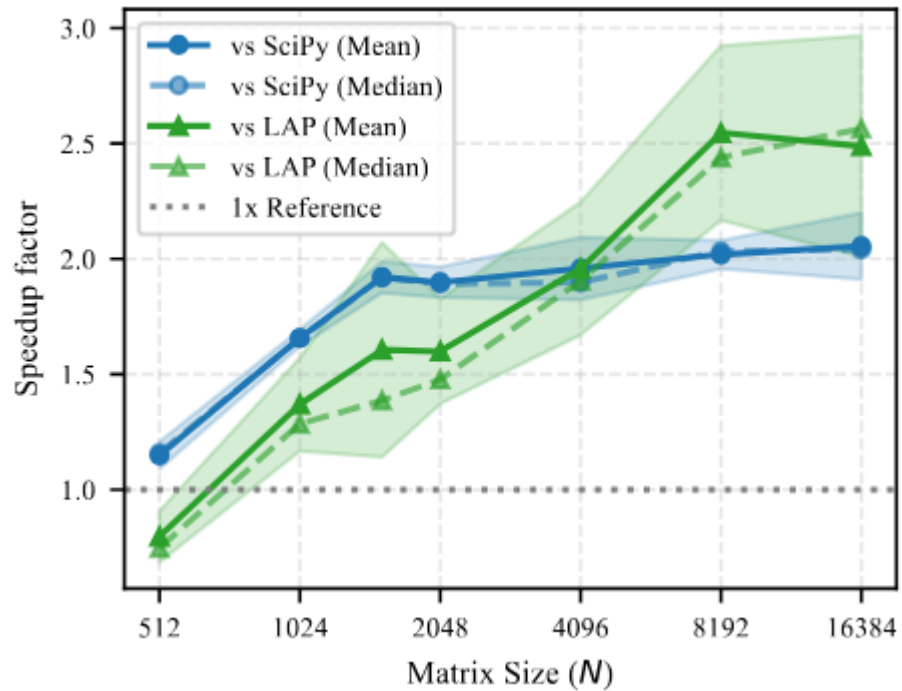
Main Model Features



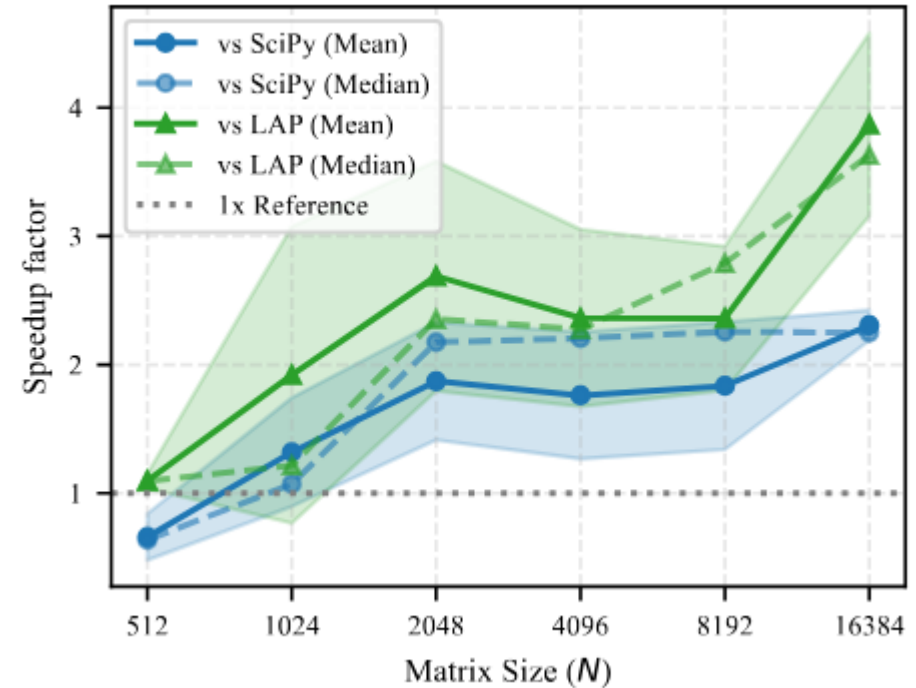
- ▶ **Feature extraction:** Reduce matrix size to $D \cdot N$. Only looking at important features (distribution, ambiguity, competition, etc.)
- ▶ **RowDualNet:** Row-independent DNN that treats the problem as a set-processing task to capture inter-agent competition.
- ▶ **Feasibility via the “Min-Trick”:** Constructively derive column potentials to ensure no violation of the dual feasibility constraint

$$\hat{v}_j = \min_i (C_{ij} - \hat{u}_i)$$

Synthetic Experimental Results

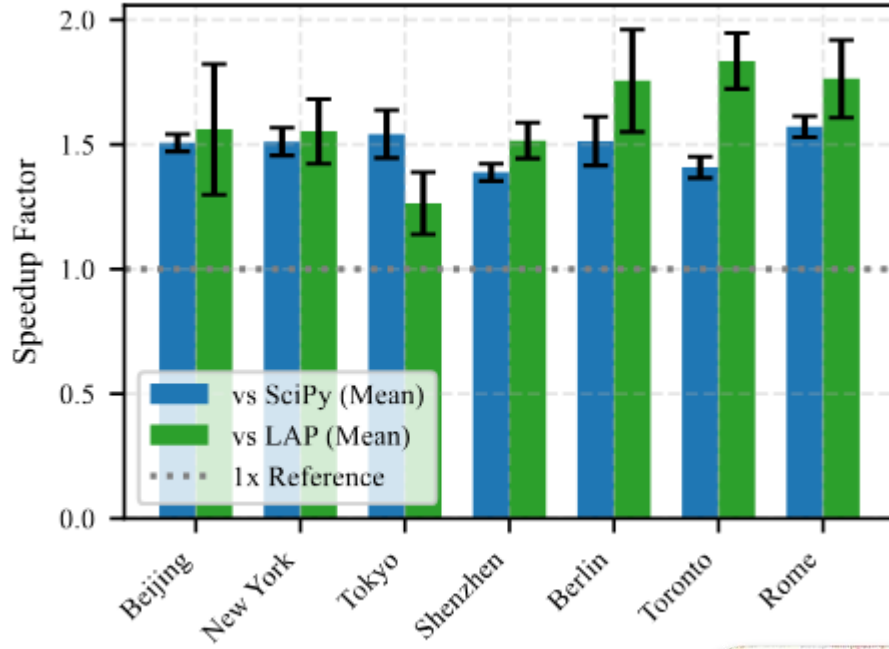


► Dense models

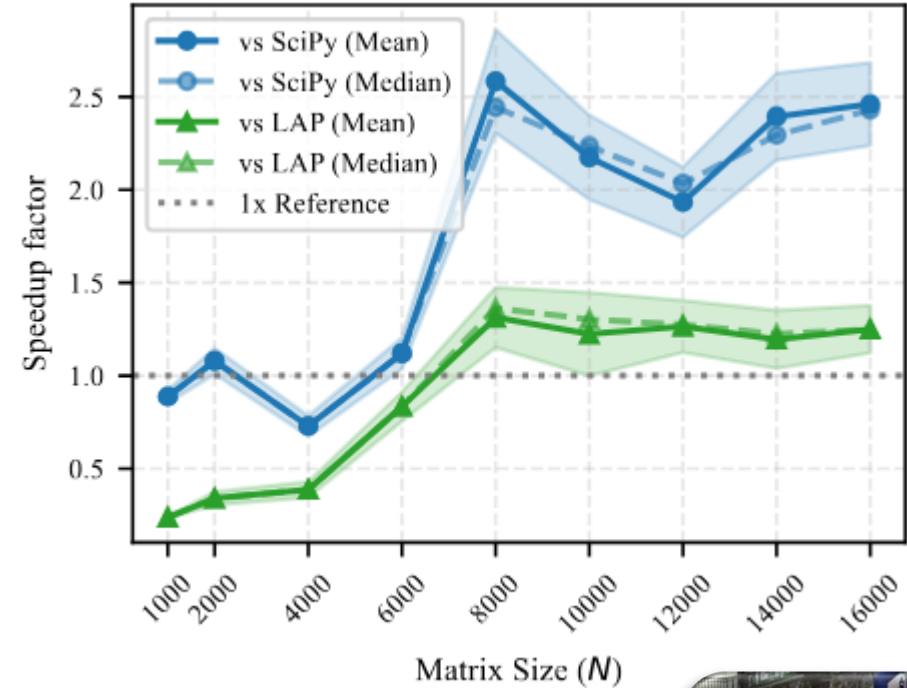
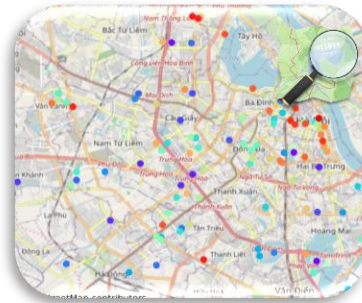


► Block-Structure models


Real-World Experimental Results



Open-Street Map



Multiple Object Tracking



Conclusion

- ▶ Introduced a solution to the LAP via warm-starting the LAPJV algorithm
 - Scaling to massive cost matrices (2k to 16k)
 - Guarantee mathematically feasible duals, eliminating correction phases
 - End-to-end acceleration
 - Preserving exact optimality

