

Dynamic Warehouse Picking Optimization

Powered by Reinforcement Learning

Industry Context:

High-throughput warehouses face challenges from fluctuating demand, diverse SKU orders, labor shortages, and stringent service-level agreements (SLAs). Key issues include inefficient picker travel, suboptimal batching and slotting, aisle congestion, and frequent priority changes.

Given that picking can account for approximately 50% of warehouse operating costs, even small improvements significantly enhance unit economics and service reliability.

The Challenge

Rule-based heuristics and static optimizers struggle to adapt to real-world changes, such as rush orders, stockouts, layout modifications, or inconsistent work-in-process. They fail to holistically consider factors such as batching, routing, congestion, and priorities, which results in extended cycle times, missed deadlines, and increased overtime.

As the unpredictability of orders and service expectations grows, the discrepancy between planned operations and actual factory floor conditions becomes increasingly expensive and operationally untenable.

Our Edge:

By reframing picking as a sequential decision problem, we train policies within a digital twin of your warehouse. Our approach moves beyond rigid, hard-coded rules, instead learning policies that offer significant advantages:

- **Joint Optimization:** We optimize batching, picker routing, and release timing simultaneously to reduce travel time and minimize aisle conflicts.
- **Adaptive Flexibility:** Our policies adapt to various layouts and constraints, including zones, one-way aisles, no-go areas, dangerous goods handling, and shift patterns, without requiring expensive re-engineering.
- **Disruption Resilience:** The system reacts to disruptions like rush orders, stockouts, or equipment downtime by re-sequencing and re-routing decisions in near real-time.
- **Business Alignment:** Policies are aligned with your specific business objectives, such as SLA adherence, makespan, labor hours, and priority orders, through customizable objective functions.
- **Safe Deployment:** We ensure safe deployment through a phased approach: simulation, followed by shadow mode, and then a gradual rollout. Improvements are delivered via periodic policy retraining, avoiding the risks associated with "live learning."

Expected added value:

In digital-twin pilots and shadow-mode trials with a global industrial provider operating a high-mix, time-sensitive network, our approach is designed to deliver:

- **Higher on-time fulfillment** through priority-aware batching and dynamic re-sequencing (shadow-mode evidence available under NDA).
- **Shorter pick cycles & fewer steps** by congestion-aware routing and zone-smart batching, improving throughput without extra labor.
- **Scalable performance** across small and large layouts; robust under seasonality and promotion spikes (no rule rewrites).
- **Faster change incorporation** (new constraints/objectives modeled in the simulator, policies retrained within days rather than months).

Get in Touch

Explore how AI is reshaping the backbone of warehouse operations and discover how an optimized process validated in a digital twin would impact your SLA compliance and labor profile.

hello@instadeep.com