

The Fast and the Flawed: Speed vs. Reliability in Hybrid Robotic Picking Systems

Sebastian Koehler¹,
Alexander Ernst¹,
Christoph Rockenschaub²,
Domenik Kaever²

1: Karlsruhe Institute of Technology (KIT), Institute for Material Handling and Logistics (IFL)

2: Institute of Logistics Engineering, Graz University of Technology, Institute of Logistics Engineering (ITL)

ORDER picking remains one of the most labor-intensive and cost-critical processes in intralogistics, and growing workforce shortages are pushing warehouses toward partial automation. Highly automated parts-to-picker systems offer high performance but require large investments and lack flexibility, whereas traditional picker-to-parts systems remain adaptable yet labor-dependent. Hybrid picker-to-parts systems, in which mobile picking robots operate cooperatively alongside human workers, combine the strengths of both, but introduce heterogeneous resources that differ in capability, speed, and reliability, where faster robots tend to be less reliable. How these properties should be balanced to maximize system performance is still insufficiently understood. We address this gap with an analytical model that determines the net throughput of a hybrid human-robot picking system in picks and sub-orders per resource type. Starting from the workload split between robots and humans, it incorporates robotic failure probabilities, the reprocessing of failed picks by human workers, and an s-shape routing policy, and is validated against an agent-based simulation across layouts, batch sizes, picking speeds, and reliability levels. The results show that throughput is highly sensitive to robotic failures, with the largest losses occurring already at small failure probabilities, so that a slower but more reliable robot can outperform a faster but less reliable one beyond a layout-dependent threshold; travel-dominated layouts further dilute the value of faster picking. On the human side, more capable robots reduce the human's picking share, though the benefit of higher capability diminishes as the failure probability rises. The findings yield a clear design principle for cooperative hybrid picking systems: robots should be made reliable before they are made faster.

[Keywords: robotic picking, cooperative picking, heterogeneous pickers, design principles]

1 Introduction

Driven by the rapid growth of e-commerce, the efficient fulfillment of customer orders has become increasingly important. At the same time, order picking remains one of the most labor-intensive and cost-critical processes in intralogistics [1], while growing workforce shortages [2] further constrain the availability of manual labor. This creates a growing need to balance high fulfillment performance with reduced reliance on human labor, driving warehouses toward partial automation. Rather than fully automated parts-to-picker systems, which demand large investments and lack flexibility [3, 4], many operators turn to hybrid picker-to-parts systems in which mobile robots work alongside human pickers [5, 6]. Commercial examples of such autonomous picking robots include Magazino's TORU, which independently locates, grips, and transports individual items within the warehouse, operating as an autonomous picking agent rather than merely as a transport aid for human workers.

The integration of such robots into picker-to-parts systems can take different forms, which we distinguish as *collaborative* and *cooperative* picking. In collaborative picking, a robot and a human worker jointly process the same order or task in close synchronization, for example when a robot escorts a picker through the aisles and the worker places items onto it [7]. In cooperative picking, by contrast, robots and human workers operate independently and in parallel [8], each completing separate orders or tasks on their own, with their outputs potentially consolidated at a later stage. This paper focuses on the cooperative setting, in which robots such as TORU from Magazino [9] or Autopicker from Brightpick [10] pick independently alongside, but not together with, human workers in the same environment.

Introducing robots alongside human workers, however, creates a heterogeneous resource pool: robots and humans differ in *capability*, i.e., what share of SKUs or tasks a re-

source can handle; in *reliability*, i.e., how consistently they perform; and in *speed*, i.e., how fast a robot can move and retrieve articles. In practice, this creates a fundamental design trade-off. Faster robots tend to be less reliable, as they are more prone to breakdowns, mispicks, and exceptions requiring human intervention, whereas simpler, slower robots generally operate more reliably. This raises a central, practically relevant question that existing literature does not answer, and which forms the main research question of this work:

When does a fast but flawed robot outperform a slower, more reliable one, and what design principles follow for hybrid picking systems?

This paper addresses this gap with an analytical throughput model that captures the trade-off between picking speed, capability, and reliability in hybrid human-robot order fulfillment. We first review the existing literature on collaborative and cooperative human-robot order fulfillment systems and position our work within the cooperative setting. Building on this, we derive the workload and the resulting throughput capacity of each resource type, as well as of the overall system. We then introduce our design of experiments and discuss numerical results across multiple use cases, before validating our analytical model against an agent-based simulation to confirm the correctness of our results. Finally, these numerical results and the validated model are used to derive design principles and generate managerial insights for robotic resources in hybrid order fulfillment systems.

2 Literature

Order fulfillment involving both human and robotic resources has received considerable research attention, particularly for *cooperative* settings which several simulation studies consider. [6] assess the introduction of Magazino picking robots, where orders are decomposed into robot-pickable and human-pickable subsets, reporting operational cost reductions of approximately 8%. Building on the same system, [11] use a multi-criteria assignment approach and additionally show reductions in human physical workload. [12] evaluate a cooperative robotic platform in a retail environment and find that heterogeneous systems can outperform purely manual operations under empirically measured picking times, while [13] show that heterogeneous resources may in some cases reduce throughput due to less efficient batch formation. Beyond human-robot settings, heterogeneity has also been studied among purely manual resources, for instance through carts with varying capacities [14, 15] or pickers with differing speeds and skill levels [16, 17, 18].

In parallel, a growing body of work addresses *collaborative* picking, where human and robotic resources jointly execute a single task in close synchronization, for exam-

ple through coupling strategies between AGVs/AMRs and pickers [7, 19, 20, 21, 22, 23].

Taken together, the literature demonstrates that both general aspects of resource heterogeneity, such as differing capabilities and skills, and cooperative human-robot interaction are comparatively well studied. However, only a small number of these works explicitly model heterogeneity in *hybrid* human-robot settings with analytical rigor, and those that do typically focus on cost, throughput, or workload outcomes for a specific system configuration. To the best of our knowledge, none of the existing studies derive *design principles* that relate the trade-off between resource speed, capability, and reliability to system performance. This paper addresses this gap by developing an analytical throughput model that explicitly captures this trade-off in cooperative hybrid picking systems.

3 Model

This section develops the throughput model in two steps. Section 3.1 derives the *effective workload split* between the robotic and human streams from the robot capability and its failure probability, yielding the suborder and batch-size distributions of each stream. Section 3.2 then translates these workloads into *throughput capacities*, expressing each picking-tour time as a travel and a picking component and accounting for the human reprocessing of failed robotic picks to obtain the system's *net* throughput.

3.1 Effective robot-human workload split

Let A denote the set of articles and let d_a be the expected number of picks required to satisfy the demand for article $a \in A$. The total demand over all articles is given by $D = \sum_{a \in A} d_a$. For each article, we define a binary capability parameter $\alpha_a \in \{0, 1\}$, where $\alpha_a = 1$ indicates that the article can be picked by the robot and $\alpha_a = 0$ otherwise. We assume that all picks that are technically feasible for the robot are assigned to it without rebalancing the assignment between the robotic and human streams. Consequently, the share of picks that is to be processed by robots is:

$$r_r^{\text{work}} = \frac{\sum_{a \in A} d_a \cdot \alpha_a}{D} \quad (1)$$

Robotic pick attempts may fail, so their success probabilities must be accounted for to determine the share of picks that must ultimately be completed by human workers. Let $s = (s_1, s_2, \dots, s_v)^T$ be a vector of size $k_{\max}$, where s_i is the probability that the robot's i th attempt at a pick succeeds, conditional on all preceding attempts having failed. The robot is allowed at most $k_{\max}$ attempts per pick.

Let S denote the attempt on which the robot first succeeds. Since success on attempt i requires the first $k - 1$ attempts to fail and the i th to succeed,

$$P(S = i) = \left[\prod_{j=1}^{k-1} (1 - s_j) \right] s_i. \quad (2)$$

The robot fails to complete the pick if it does not succeed within the v allowed attempts, i.e. $S > k_{max}$. This occurs only when all v attempts fail, giving the failure probability

$$p_r^{\text{fail}} = P(S > k_{max}) = \prod_{j=1}^{k_{max}} (1 - s_j). \quad (3)$$

In cooperative human robotic order fulfillment systems where robots pick independently besides human workers incoming orders need to be split between the worker and the robotic streams. Let X be the order size of an incoming customer order with a maximum value of o_{max} and an average order size of $E(X) = o_{avg}$. Therefore the probability of an order of size x is given by $P(X = x)$ which is assumed to be given from empirical observations. Incoming orders need to be split into robotic and human suborders denoted by the stochastic variables Y_r and Y_w . The suborder size probabilities of resource type i are given by:

$$P_i(Y_i = y) = \sum_{x=y}^{o_{max}} P(X = x) \binom{x}{y} (r_i^{\text{work}})^y \cdot (1 - r_i^{\text{work}})^{x-y}, \quad y = 0, 1, \dots, o_{max}. \quad (4)$$

If orders are processed in batches of b_i orders the probabilities of the batch size Z_i are therefore given by the discrete convolution:

$$P_{i,b_i}(Z_i = z) = \sum_{\substack{y_{i,1}, \dots, y_{i,b_i} > 0 \\ y_{i,1} + \dots + y_{i,b_i} = z}} \prod_{k=1}^{b_i} P_i(Y_{i,k}^* = y_{i,k}) \quad (5)$$

Where $Y_{i,k}^*$ denotes the k th non-empty suborder of a batch, with the normalized probabilities $P(Y_{i,k}^* = y_{i,k}) = P_i(Y_i = y_{i,k} \mid Y_i > 0)$.

3.2 Throughput capacity of the robot-human system

To determine the throughput capacity of each resource stream ($i \in I$), the average picking tour time must first be determined, which is given by:

$$t_i^{\text{pick tour}} = t_i^{\text{travel}} + t_i^{\text{pick}} \quad \forall i \in I \quad (6)$$

The first part consists of the travel time which resources need to reach all of the pick positions within the warehouse and return to the I/O point after the last pick has been done. Therefore the average time is given by:

$$t_i^{\text{travel}} = \sum_{z \in \{b_i, \dots, b_i \cdot o_{max}\}} P_{i,b_i}(Z_i = z) \cdot \frac{D_i(Z_i = z)}{v_i} \quad \forall i \in I \quad (7)$$

Where $D_i(Z_i = z)$ is the average travel distance within the warehouse. For this work we assume an s-shape routing policy where the average travel distance is given by:

$$D_i(Z_i = n) = M(d_l + l_i) \left(1 - \left(1 - \frac{1}{M} \right)^n \right) + 2d_t \left(M - \sum_{j=1}^{M-1} \left(\frac{j}{M} \right)^n \right) + 2f_{\text{last}}(n) \left(\frac{1}{2} + \frac{1}{2} \sum_{j=1}^M \binom{M}{j} \left(\frac{j}{M} \right)^n (-1)^{j-1} 2^{M-j} \right), \quad \forall i \in I. \quad (8)$$

Where M represents the number of aisles, n the number of picks, d_l the length of an aisle, and d_t the center-to-center spacing between adjacent aisles. Compared to the baseline formulation in [24], we introduce two modifications.

First, we incorporate the physical extent l_i of the picking resource of type i into each aisle traversal. This reflects the fact that the resource must fully clear an aisle before entering the next one. Although this effect is often negligible in large-scale warehouse systems, it can have a noticeable impact in more compact settings such as fast-pick zones or backroom storage areas.

Second, for cases in which the number of traversed aisles is odd, we adjust the computation of the expected position of the final pick using a return-policy correction captured by the function $f_{\text{last}}(n)$. This function is also employed in [24], although its derivation is provided in a more intuitive form in Appendix A.

The second component represents the physical picking process of products from their respective storage positions. We assume that human workers always succeed on their first attempt, whereas robots may require multiple attempts to retrieve the required article. Let N_r denote the expected number of attempts performed by the robot given by:

$$E[N_r] = \sum_{k=1}^{k_{max}} k \cdot P(S = k) + k_{max} \cdot P(S > k_{max}) \quad (9)$$

Therefore the time a robot and a human needs to retrieve an item from the shelf is defined by:

$$\begin{aligned} t_r^{\text{pick}} &= E(Z_r) \cdot t_r^{\text{retrieval}} \cdot E[N_r], \\ t_w^{\text{pick}} &= E(Z_w) \cdot t_w^{\text{retrieval}} \end{aligned} \quad (10)$$

Therefore, the final throughput capacity per resource type is given by:

$$\mu_i = \frac{1}{t_i^{\text{pick tour}}}, \quad i \in I. \quad (11)$$

The throughput in picks can be expressed as:

$$TH_r^{\text{picks}} = \mu_r \cdot E(Z_r) \cdot (1 - p_r^{\text{fail}}), \quad (12)$$

$$TH_w^{\text{picks}} = \mu_w \cdot E(Z_w), \quad (13)$$

and the (sub-)order throughput as

$$TH_i^{\text{order}} = \mu_i \cdot b_i, \quad i \in I. \quad (14)$$

To incorporate the robotic failures, the net throughput capacity has to be determined. While for robots $\mu_{r,\text{net}} = \mu_r$, the human throughput capacity is reduced, since the human workers have to handle the robots' failed picks. We therefore calculate the net batch size of a human worker, which incorporates the robotic failures $E(Z_{w,\text{net}})$, as

$$E(Z_{w,\text{net}}) = E(Z_w) + \frac{n_r \cdot \mu_r \cdot E(Z_r) \cdot p_r^{\text{fail}}}{n_w \cdot \mu_{w,\text{net}}} \quad (15)$$

where n_w and n_r denote the number of workers and robots in the system, respectively. We assume that the robot-to-worker ratio is chosen to balance the throughput of the two streams, such that robotic pickers and human workers perform the same number of picking tours. This is not merely a simplifying assumption but the throughput-maximizing choice for a fixed order split: since each order consists of a robotic and a human suborder that must both be completed, overall throughput is determined by the slower stream. Additional capacity in the non-bottleneck stream therefore cannot increase system throughput. A staffing ratio that left one resource type with a persistent overcapacity would therefore incur its cost without any throughput benefit; balancing the two streams is the only ratio at which no capacity is wasted. It also fixes the staffing ratio endogenously, removing n_w and n_r as free parameters and leaving throughput a function of the robot design parameters alone.

Let r_r^{work} denote the share of the total picking workload nominally assigned to the robots, so that a fraction $1 - r_r^{\text{work}}$ is handled by the human workers. Because every failed robotic pick must be reprocessed by a human, an additional fraction $r_r^{\text{work}} \cdot p_r^{\text{fail}}$ of the workload is shifted from the robots to the workers. Therefore, the effective share of the total picking workload assigned to the human workers, including the reprocessing of failed robotic picks, is:

$$r_w^{\text{work}} = (1 - r_r^{\text{work}}) + r_r^{\text{work}} \cdot p_r^{\text{fail}} \quad (16)$$

4 Design of Experiments (DoE)

The design of experiments is constructed around the central trade-off of this study: how the robot's picking speed and reliability jointly determine the throughput capacity of a hybrid system. To expose this trade-off, we vary the following factors:

- Layouts: two warehouses (Layout A with 10 aisles and Layout B with 20 aisles shown in Figure 1) with an aisle-to-aisle distance of 3 m and an aisle length of 10 m;
- Batch size: two levels, capturing the effect of order consolidation on the picking tours;
- Robot picking speed: the robotic retrieval time per article, $t_r^{\text{retrieval}}$;
- Robot reliability: the attempt success probabilities s_1 and s_2 , which determine the failure probability on the robot's first and second attempt and hence p_r^{fail} .

Speed and reliability are swept across their full ranges, since they define the trade-off under study. We additionally vary the robot capability, but each speed–reliability comparison is carried out *within* a fixed capability level, so that the speed and reliability effects are not confounded with a changing workload split. The remaining parameters are held fixed. In particular, we assume equal travel speeds for robots and human pickers, $v_r = v_w = 1$ m/s, reflecting the current state of technology: autonomous travel through the aisles is essentially solved, whereas picking remains the demanding task. The robot retrieval time is derived from the laboratory experiments in [8]. For comparison, the human retrieval time is set to the minimum observed robot retrieval time, allowing us to evaluate the case in which both resource types have the same retrieval speed. All parameter values and their levels are summarised in Table 1.




Figure 1: Schematic representation of the warehouse layout, showing the arrangement of storage racks, the input/output point, and the relevant dimensions.

Factor	Values
r_r^{work}	[30%, 50%, 70%, 90%]
b_i	[2, 3]
k_{max}	[2]
s_1, s_2	[0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9]
Layout	A (10 aisles), B (20 aisles)
o_{avg}	[10]
v_i	[1 m/s]
$t_r^{retrieval}$	[13, 15, 17, 19, 20s]
$t_w^{retrieval}$	[13s]

Table 1: Factors and their respective levels used in the DoE

5 Validation

To validate our analytical model, we compare the resulting pick tour durations $t_i^{pick\ tour}$ with average values obtained from an agent-based discrete-event simulation built in Any-Logic (Version 8.9.9) using the model parameters given in Table 2. The simulation considers a system comprising one human and one robotic picker and is terminated once both pickers have completed 200,000 suborders, providing a sufficiently large sample for comparison. Each picker independently processes randomly generated suborders as quickly as possible; subsequent order consolidation is not modelled.

Potential interactions between pickers, such as interference while traversing aisles or simultaneous access to the same storage location, are not represented in the simulation. Failures of the robotic picker are incorporated by collecting all unsuccessful robotic picks and assigning them to the human picker's subsequent pick tour.

Factor	Values
r_r^{work}	[30%, 50%, 70%]
b_i	[2]
k_{max}	[2]
s_1, s_2	[1, 0.9, 0.8, 0.7, 0.6, 0.5]
Layout	A (10 aisles)
o_{avg}	[10]
v_i	[1 m/s]
$t_r^{retrieval}$	[13, 15, 17s]
$t_w^{retrieval}$	[13s]

Table 2: Model parameters applied for validation

The validation results are shown in Figures 2 and 3. For the robotic pick time (Figure 2), only negligible deviations are observed across all tested settings. For the human pick time (Figure 3), however, the deviations originate from the simplification in Equation (16), which overestimates the number of failures for low robotic capabilities and underestimates the number of failed picks for high robotic capabilities. These deviations nevertheless remain within an acceptable range and, because the subsequent analysis reports throughput under a balanced staffing ratio, they do not affect the reported findings.




Figure 2: Validation of the robotic pick tour duration against the agent-based simulation for different capability levels.




Figure 3: Validation of the human pick tour duration against the agent-based simulation for different capability levels.

6 Numerical results

We split our numerical results into two parts. First we discuss the influence of different failure probabilities on the robotic throughput in Section 6.1. After doing so we evaluate the impact on the human workload in Section 6.2.


6.1 Robotic throughput

First, we evaluate the net throughput in picks for both resource types. The results for the base case with $r_r^{work} = 0.5$ are shown in Figure 4 for the two layout alternatives at a batch size of $b_i = 2$. Across all tested settings, an increase in the failure probability causes a disproportionate loss in throughput, revealing a pronounced non-linear sensitivity to robotic failures. The largest marginal losses oc-


cur when moving from an almost perfectly reliable robot to comparatively small positive failure probabilities, whereas the curves flatten at higher failure levels. As a consequence, a slower but more reliable robot can outperform a faster one with a higher failure probability up to a certain point. The intersections between configurations can be interpreted as speed–reliability equivalence points. Figure 4(a) shows that the throughput of a perfectly reliable robot with a picking time of 20 s can be approximately matched by the considerably faster 15 s configuration, even when the latter exhibits an effective failure probability of roughly 0.2–0.25. Beyond this point, however, the additional picking speed can no longer compensate for the losses caused by unsuccessful robotic picks.

The layout alters the relative value of robotic picking speed, as shown in Figure 4(b). In the smaller layout, the throughput difference between the fastest and slowest robot configurations is more pronounced than in the larger layout. As travel distances grow, travel time accounts for a larger share of the total picking-tour duration, so differences in the retrieval time become comparatively less influential. Consequently, investments in faster robotic picking mechanisms yield greater throughput benefits in compact picking areas, whereas their effect becomes increasingly diluted in travel-dominated systems.

Batch size, layout, and capability all shift the throughput level while leaving the shape of the curves unchanged. Increasing the batch size raises throughput by roughly 10% across the base case; the larger layout lowers the attainable throughput relative to the smaller one; and higher capabilities raise it (Figure 5), with the gap between pick times staying relatively constant. The capability effect follows from the number of picks per tour: a higher capability places more picks on each tour, reducing the travel distance per pick and completing the tours more efficiently. In every case the response to the failure probability is the same, so we omit the remaining DOE cases for reasons of space.





(a) Layout A



(b) Layout B

Figure 4: Net throughput versus failure rate for Layout A and B, at $b_r = 2$ and $r_r^{work} = 0.5$.



(a) Layout A, $b_r = 2$, $r_r = 0.3$ (b) Layout A, $b_r = 2$, $r_r = 0.7$ Figure 5: Net throughput versus failure rate for Layout A, at $b_r \in \{2, 3\}$ for $r_r = 0.3$ and $r_r = 0.7$.

6.2 Human throughput

The human throughput exhibits a counterintuitive response to robotic failures: as the robot becomes less reliable, the human completes more picks per hour but her suborders. The two measures move in opposite directions, and distinguishing them is necessary for a correct assessment of the human workload.


Figure 6 reports the human pick throughput, which increases with the failure probability at every capability level. Each unsuccessful robotic pick is reassigned to the human and appended to a picking tour that would have been travelled regardless, reducing the travel distance per pick and raising the number of items retrieved per unit of travel. Across capability levels, a more capable robot lowers the human pick throughput, as it retains the denser picks and leaves the human with sparser ones, thereby increasing the travel distance per pick. The capability curves converge as the failure probability rises, reflecting the diminishing

marginal benefit of appending a further pick to an already long tour.


(a) Layout A, $b_w = 2$ (b) Layout B, $b_w = 3$ Figure 6: Human pick throughput versus the effective robot failure probability, with one curve per robot capability level, for Layout A ($b_w = 2$) and Layout B ($b_w = 3$).

Measured in completed suborders, the relationship inverts, as shown in Figure 7. The additional picks are reassigned robotic failures rather than new customer suborders; the human therefore retrieves more items but completes no additional orders, while the lengthened tours increase the time required per suborder. The rise in pick throughput at higher failure probabilities consequently coincides with a decline in suborder throughput, as the per-pick efficiency gain is absorbed by the reprocessing of robotic failures rather than by the completion of new work. In line with this, a more capable robot yields a higher human suborder throughput, since the human retains fewer picks per suborder and completes shorter tours. As for the pick throughput, an increasing failure probability compresses the differences

between capability levels, rendering the choice of capability progressively less consequential at high failure rates.



(a) Layout A, batch size 2



(b) Layout B, batch size 3

Figure 7: Human suborder throughput versus the effective robot failure probability, with one curve per robot capability level, for Layout A ($b_w = 2$) and Layout B ($b_w = 3$).

7 Insights

This work quantifies the effect of hybrid human–robot picking on net throughput, measured in picks and suborders per resource type and hour, through an analytical model that relates a robot’s capability and speed to system performance. The analysis yields a clear design principle: robots should be made reliable before they are made faster. Because technological development inherently trades reliability against speed, this ordering carries direct practical relevance. Reliability turns out to matter more than speed. Once the failure probability passes a threshold that depends on the layout, even the fastest robot we tested is overtaken by a much

slower but more reliable one, showing that a large speed advantage cannot offset even a modest loss in reliability.

On the human side, more capable robots absorb a larger share of the picks and leave the worker with fewer picks per suborder. This lowers the human pick throughput but raises the suborder throughput, as the shorter tours allow more suborders to be completed per hour. The benefit of higher capability nevertheless diminishes as the failure probability increases.

Several extensions remain open. Applying the model to further warehouse settings would clarify the influence of different order-size distributions and routing policies, and the share-level approximation of the human picking-tour duration could be replaced by the exact tour-level formulation to improve accuracy. A natural further step is an economic analysis: since gains in both reliability and speed come at a cost, quantifying the associated trade-offs would extend the design principles derived here toward investment decisions.

8 Acknowledgement

This research was funded by the Deutsche Forschungsgemeinschaft (DFG - German Research Foundation) - 1574 - 471687386.

9 AI Statement

During the preparation of this work, the authors used OpenAI’s ChatGPT and Anthropic’s Claude for editorial purposes and programming support. After using the tools, the authors reviewed and edited all code and content as needed and take full responsibility for the content of the published article.

REFERENCES

- [1] R. De Koster, T. Le-Duc, and K. J. Roodbergen, “Design and control of warehouse order picking: A literature review,” *European journal of operational research*, vol. 182, no. 2, pp. 481–501, 2007.
- [2] H. Diefenbach, E. H. Grosse, and C. H. Glock, “Human-and-cost-centric storage assignment optimization in picker-to-parts warehouses,” *European Journal of Operational Research*, vol. 315, no. 3, pp. 1049–1068, 2024.
- [3] S. Kumar, J.-B. Sheu, and T. Kundu, “Planning a parts-to-picker order picking system with consideration of the impact of perceived workload,” *Transportation Research Part E: Logistics and Transportation Review*, vol. 173, p. 103088, 2023.

- [4] E. H. Grosse, C. H. Glock, and W. P. Neumann, "Human factors in order picking: a content analysis of the literature," *International journal of production research*, vol. 55, no. 5, pp. 1260–1276, 2017.
- [5] N. Boysen, R. De Koster, and F. Weidinger, "Warehousing in the e-commerce era: A survey," *European Journal of Operational Research*, vol. 277, no. 2, pp. 396–411, 2019.
- [6] M. Zhang, S. Winkelhaus, and E. H. Grosse, "Evaluation of human workload in a hybrid order picking system," *IFAC-PapersOnLine*, vol. 54, no. 1, pp. 458–463, 2021.
- [7] M. Ghorashi Khalilabadi, D. Roy, and R. de Koster, "Picking the best bot: Collaboration strategies for humans and bots in order pick systems with traveling salesman problem routing," *Transportation Science*, 2025.
- [8] J. Bhowmick, J. Pazour, and I. Dayarian, "Store fulfillment with autonomous mobile robots and in-store customers," *IIE Transactions*, vol. 58, no. 3, pp. 272–289, 2026.
- [9] MAGAZINO. (2026) Autonomer schuhlogistik- und kommissionier-roboter toru. MAGAZINO. Accessed: 2026-02-02. [Online]. Available: <https://www.magazino.eu/schuhlogistik/>
- [10] Brightpick. (2026) Solution — brightpick. Brightpick. Accessed: 2026-02-02. [Online]. Available: <https://brightpick.ai/solution/>
- [11] V. Vijayakumar and A. Sobhani, "Performance optimisation of pick and transport robot in a picker to parts order picking system: a human-centric approach," *International Journal of Production Research*, vol. 61, no. 22, pp. 7791–7808, 2023.
- [12] J. Bhowmick, S. Köhler, G. Arndt, G. Fischer, M. Padhy, K. Furmans, and J. Pazour, "Assessing economic and operational feasibility of a designed and lab demonstrated robotic platform for omnichannel logistics," *Computers & Industrial Engineering*, p. 111304, 2025.
- [13] A. Dehghan, M. Cevik, and M. Bodur, "Dynamic task allocation in intelligent warehouses with hybrid workforce of automated guided vehicles and human pickers," *ACM Transactions on Evolutionary Learning*, 2025.
- [14] J. A. Cano, P. Cortés, E. A. Campo, and A. A. Correa-Espinal, "Multi-objective grouping genetic algorithm for the joint order batching, batch assignment, and sequencing problem," in *Sustainable Logistics Systems Using AI-based Meta-Heuristics Approaches*. Routledge, 2023, pp. 39–55.
- [15] S. Wagner and L. Mönch, "A variable neighborhood search approach to solve the order batching problem with heterogeneous pick devices," *European journal of operational research*, vol. 304, no. 2, pp. 461–475, 2023.
- [16] S. Mou, "Integrated order picking and multi-skilled picker scheduling in omni-channel retail stores," *Mathematics*, vol. 10, no. 9, p. 1484, 2022.
- [17] —, "In-store order fulfilment in omni-channel supermarkets with heterogeneous workforce: A bi-objective optimisation approach," *Computers & Industrial Engineering*, vol. 171, p. 108516, 2022.
- [18] M. Matusiak, R. De Koster, and J. Saarinen, "Utilizing individual picker skills to improve order batching in a warehouse," *European Journal of Operational Research*, vol. 263, no. 3, pp. 888–899, 2017.
- [19] S. Srinivas and S. Yu, "Collaborative order picking with multiple pickers and robots: Integrated approach for order batching, sequencing and picker-robot routing," *International Journal of Production Economics*, vol. 254, p. 108634, 2022.
- [20] M. Löffler, N. Boysen, and M. Schneider, "Human-robot cooperation: Coordinating autonomous mobile robots and human order pickers," *Transportation science*, vol. 57, no. 4, pp. 979–998, 2023.
- [21] I. Žulj, H. Salewski, D. Goeke, and M. Schneider, "Order batching and batch sequencing in an amr-assisted picker-to-parts system," *European Journal of Operational Research*, vol. 298, no. 1, pp. 182–201, 2022.
- [22] K. Azadeh, D. Roy, R. de Koster, and S. M. G. Khalilabadi, "Zoning strategies for human–robot collaborative picking," *Decision Sciences*, vol. 56, no. 1, pp. 50–70, 2025.
- [23] K. J. Min, B. D. Chung, and G. S. Cho, "Agent-based modelling for human–robot collaborative order picking system considering workers' performance," *International Journal of Production Research*, pp. 1–24, 2025.
- [24] K. Dmytrów, "Analytical and simulation determination of order picking time in a low storage warehouse for shared storage systems," *Operations Research and Decisions*, vol. 32, no. 2, pp. 34–51, 2022.

Appendix A: S-Shape Travel time estimation

Consider a warehouse with M aisles. A batch consists of n picks, where each pick is independently assigned to an aisle with equal probability $1/M$. After n picks, the system is described by the state (k, h, y) , where:

- k is the number of distinct aisles that are visited,
- $h \in \{0, \dots, M - 1\}$ is the index of the highest-numbered visited aisle,
- y is the number of picks assigned to aisle h .

Let $P_n(k, h, y)$ denote the probability of being in state (k, h, y) after n picks. From a state (k, h, y) , the next pick leads to the following transitions:

$$\begin{aligned} (k, h, y) &\rightarrow (k, h, y + 1) && \text{probability } \frac{1}{M}, \\ (k, h, y) &\rightarrow (k, h, y) && \text{with probability } \frac{k-1}{M}, \quad k > 1, \\ (k, h, y) &\rightarrow (k+1, h, y) && \text{probability } \frac{h-(k-1)}{M}, \\ (k, h, y) &\rightarrow (k+1, j, 1) && \text{probability } \frac{1}{M}, \\ &&& j = h+1, \dots, M-1 \end{aligned}$$

The initial distribution after the first pick is

$$P_1(1, j, 1) = \frac{1}{M}, \quad j = 0, \dots, M-1$$

Let $f_{\max}(L, y)$ denote the expected maximum pick depth in an aisle of length L with y uniformly distributed picks:

$$f_{\max}(L, y) = \sum_{i=1}^L i \left[\left(\frac{i}{L} \right)^y - \left(\frac{i-1}{L} \right)^y \right]$$

For a given state (k, h, y) , the longitudinal travel distance $D_1(k, y)$ is defined as:

$$D_1(k, y) = \begin{cases} k \cdot (L + l_i), & \text{if } k \text{ is even,} \\ (k-1) \cdot (L + l_i) + 2 \cdot f_{\max}(L, y), & \text{if } k \text{ is odd.} \end{cases}$$

The expected longitudinal distance is therefore:

$$E[D_1(n)] = \sum_{k, h, y} P_n(k, h, y) D_1(k, y)$$

And the expected distance in the last aisle if k is odd:

$$f_{last}(m) = \sum_{k, h, y} P_n(k, h, y \mid k \text{ odd}) \cdot f_{\max}(L, y)$$

Institute of Technology (KIT).

Address: Institut für Fördertechnik und Logistiksysteme (IFL), Karlsruher Institut für Technologie (KIT), Gotthard-Franz-Straße 8, Geb. 50.38, 76131 Karlsruhe, Germany,
E-Mail: alexander.ernst@kit.edu

Dipl.-Ing. Christoph Rockenschaub, Research Associate at the Institute of Logistics Engineering (ITL), Graz University of Technology (TU Graz). Address: Institut für Technische Logistik (ITL), Technische Universität Graz, Inffeldgasse 25/E, 8010 Graz, Austria,
E-Mail: c.rockenschaub@tugraz.at

Univ.-Prof. Dr.-Ing. Domenik Kaever, M.Sc., Full Professor and Head of the Institute of Logistics Engineering (ITL), Graz University of Technology (TU Graz). Address: Institut für Technische Logistik (ITL), Technische Universität Graz, Inffeldgasse 25/E, 8010 Graz, Austria,
E-Mail: domenik.kaever@tugraz.at

Sebastian Köhler, M.Sc., Research Associate at the Institute for Material Handling and Logistics (IFL), Karlsruhe Institute of Technology (KIT).

Address: Institut für Fördertechnik und Logistiksysteme (IFL), Karlsruher Institut für Technologie (KIT), Gotthard-Franz-Straße 8, Geb. 50.38, 76131 Karlsruhe, Germany,
E-Mail: sebastian.koehler@kit.edu

Alexander Ernst, M.Sc., Research Associate at the Institute for Material Handling and Logistics (IFL), Karlsruhe