

Open Access

Optimization of warehouse logistics in household woods

Anatolii Nosar

Domus Logistics Inc

Corresponding author: **Anatolii Nosar**

Received: 28, 08, 2025

Accepted: 04, 09, 2025

Published: 09, 10, 2025

Abstract

Objective: Household-wood warehouses grapple with bulky SKUs, moisture risk and same-day fulfilment, evidence on optimisation is scattered. This review quantifies operational and sustainability gains attainable through contemporary optimisation techniques.

Methods: A systematic search of Scopus, Web of Science and IEEE Xplore (2017-2025) retrieved 326 records. After PRISMA screening, fifteen empirical papers met inclusion criteria. Key performance indicators-travel time, inventory cost and cradle-to-gate CO₂ per tonne-were harmonised. Random-effects meta-analysis computed pooled Hedges g values, subgroup tests contrasted AI-driven scheduling with heuristic approaches.

Results: Combined data (n = 112 warehouse observations) reveal optimisation trims travel time by 18 % (95 % CI 14-22 %) and inventory cost by 15 % (CI 11-19 %). CO₂ intensity falls 20 % when digital picking aids shorten forklift routes. AI outperforms heuristics by eight percentage points in sites below 10 000 m², yet the edge shrinks in high-bay operations. Heterogeneity remains high (I² = 71 %), but sensitivity checks uphold central tendencies.

Conclusions: Even amid methodological noise, optimisation tools deliver repeatable gains. Practising managers should start with low-code analytics, then escalate to full automation as data maturity improves. Future researchers must capture live sensor streams to refine effect estimates and link efficiency targets with carbon agendas.

Keywords: Warehouse logistics optimization, Household wood supply chain, Order-picking efficiency, Inventory cost reduction, Industry 4.0, Sustainability metrics,

Introduction

Warehouses that trade in household woods-flat-pack furniture panels, finger-jointed boards, decorative mouldings-sit at the confluence of bulky geometry and impatient consumer demand. A single pallet of laminated particleboard may weigh less than a washing machine yet devours twice the aisle width, warping if humidity drifts by a mere three percent. Meanwhile retail platforms promise “click-to-door in thirty-six hours.” This cocktail of physical fragility and temporal pressure forces operators to orchestrate storage layout, order-picking, and replenishment with surgical care, small missteps ripple outward as stockouts, dented panels, or costly returns. Over the past decade, researchers have responded with an eclectic toolbox: genetic algorithms for route planning, lean 5S interventions, and more recently, deep-reinforcement controllers that tweak forklift dispatch in real time. The promise sounds irresistible-yet reported gains travel in discordant units, contexts, and sample sizes, leaving practitioners unsure which lever to pull first.

The academic conversation mirrors that fragmentation. Broulias, Marcoulaki, and colleagues (2005) illustrated how a serpentine picking path halved travel distance in a Greek veneered-ply facility, whereas De la Fuente et al. (2017) pursued cradle-to-gate life-cycle metrics and warned that speed-centric redesigns can backfire environmentally when idle forklifts still idle on diesel. More recently, Müller, Jaeger, and Hanewinkel (2019) positioned Industry 4.0 as the lubricant that joins disparate optimisation subfields, yet admitted empirical validation remains patchy, especially among small and mid-size warehouses that lack sensor saturation. Across these and other studies, three persistent gaps emerge. First, effect sizes fluctuate wildly, throughput reductions range from negligible to thirty percent, a spread that

likely hides contextual moderators such as floor area, SKU heterogeneity, or labour contract rigidity. Second, operational and sustainability outcomes rarely appear in the same paper, hampering integrated decision-making. Third-and most troublesome for managers-methodological transparency varies: some authors report full test-bed data, others offer only ratios, and a few rely on simulation without factory calibration. In short, literature offers sparks of insight but little consolidated fire.

This article answers that call for consolidation. It aggregates fifteen peer-reviewed investigations published between 2017 and 2025, each meeting strict inclusion criteria: empirical design, warehouse-level focus, wood or furniture domain, and quantitative key performance indicators. By translating disparate metrics into common denominators-minutes per pick, euros per cubic-metre-month, kilograms of CO₂ per tonne handled-and subjecting them to a random-effects meta-analysis, the study produces the first statistically grounded map of optimisation pay-offs in household-wood logistics. Particular attention is paid to the interaction between algorithmic sophistication and warehouse scale, an angle hinted at but not formally tested in earlier work. The synthesis deliberately excludes new field experiments, instead it mines existing data to show what is already knowable yet under-appreciated. Such an approach serves two audiences. Practitioners gain a ranked list of levers, weighted by probable impact and contextual fit. Scholars receive a clarified baseline, against which future sensor-rich case studies can benchmark rather than reinvent.

Three objectives steer the inquiry. The first is descriptive: catalogue dominant optimisation techniques and their prevalence across geographies. The second is analytic: estimate pooled effect sizes for throughput, inventory cost, and environmental intensity, while probing heterogeneity through subgroup and meta-regression tests. The third is prescriptive: distil a pragmatic decision pathway

that aligns investment sequence with data maturity, thereby reducing the odds of over-automation or under-utilisation. Wrapped around these goals are two research questions. Which optimisation strategies deliver the most substantial, generalisable performance gains in household-wood warehouses? How do moderating factors-warehouse size, technology depth, sustainability orientation-shape those gains?

By weaving together fragmented empirical strands, the present study fills a critical void between algorithmic elegance and on-the-ground feasibility. It contends that optimisation is not a monolith but a spectrum whose sharpest edge depends on context, and it offers evidence-based guidance for selecting the right blade. In doing so, the article aspires to shift the discourse from isolated success stories toward cumulative knowledge, laying firmer footing for both industrial adoption and academic advance.

Literature review

Warehouse optimisation for household woods has travelled a winding, sometimes looping, path in the academic record. Early work treated the warehouse as a geometric puzzle, later studies reframed it as a node in a sustainability network, the most recent contributions weave cyber-physical threads through every aisle. A detailed reading of ten cornerstone papers uncovers both converging insights and stubborn blind spots-knowledge essential for any researcher who hopes to push the field beyond incremental gains.

The story opens with space and motion. In a plywood facility near Thessaloniki, Broulias, Marcoulaki, Chondrocoukis, and Laios (2005) repainted travel lanes, forcing pickers to snake through racks instead of following the habitual S-shape. Their stopwatch study, simple yet persuasive, showed a forty-plus-percent cut in walking distance. Although technology was minimal, two durable principles emerged: bulky rectangular boards impose unique aisle choreography, and low-budget layout tweaks can outperform expensive automation if tuned to material form. Those insights, though sometimes overlooked, still echo in modern heuristic design.

Physical rearrangement alone soon met its limits. As order volumes grew and sustainability entered board-room vocabulary, scholars questioned whether faster always meant better. De la Fuente, Athanassiadis, González-García, and Nordfjell (2017) stitched cradle-to-gate life-cycle metrics onto Canadian and Swedish timber warehouses. They discovered a paradox: forklifts that travelled fewer metres now idled longer to keep hydraulic systems warm, nudging net carbon upward. That finding unsettled the prevailing “speed-is-green” creed and nudged the community to elevate emissions to a first-class performance indicator.

The environmental turn deepened when Lenglet, Courtonne, and Cauria (2017) used material-flow analysis to trace French log exports. They showed that even modest shifts in warehouse dwell time ripple upstream to harvest intensity and downstream to carbon leakage across borders. Momentum gathered quickly. Müller, Jaeger, and Hanewinkel (2019) reviewed Industry 4.0 tools-sensorised pallets, pick-by-vision glasses, fleet telematics-and placed throughput, cost, and sustainability in a three-gear metaphor: turning one cog necessarily twists the others. Their conceptual map lacked broad empirical footing, yet it re-anchored the optimisation agenda around integrated performance.

Quantitative heft arrived with Santos, Carvalho, Barbosa-Póvoa, Marques, and Amorim (2019). They scanned more than a hundred forest-wood studies and discovered that fewer than ten percent shared even a single identical KPI. Such fragmentation, they argued, stifles cumulative learning and inflates the risk of cargo-cult adoption of “best practices.” Their plea for data transparency, blunt but timely, catalysed a subtle culture shift: more authors began to append raw logs, simulation parameters, or at least clear unit conversions.

Integrated modelling soon followed. Baghizadeh, Zimon, and Jum’a (2021) built a stochastic programme that merged supplier discounts, transit uncertainty, and warehouse capacity. One simulation run revealed that a tempting price cut on low-grade pine could flood racks with sluggish pallets and strangle high-margin SKUs. The lesson was brutal: procurement gains can mutate into logistical headaches if warehouse constraints remain opaque. Here, optimisation escaped the silo of internal operations and migrated into cross-functional strategy.

Numbers, however, never move a two-metre MDF plank, people do. Interviewing Dutch warehouse crews, Hoogstra-Klein and Meijboom (2021) found that many pickers still rely on memory maps and peer signals. Their qualitative study exposed a hidden layer of social choreography. In one site, management installed a route optimiser that instructed a senior operator to reverse sequence. The crew quietly ignored the screen, and efficiency dipped. Culture, the authors concluded, is not a backdrop but an active variable, algorithms thrive only when grafted onto existing knowledge networks.

Technology’s march continued. Hosseini and Peer (2022) surveyed optimisation across wood manufacturing and recast warehouses as the central hinge in a cyber-physical loop. Predictive maintenance on saw lines, they argued, collapses if downstream buffers lack real-time visibility, while an unclogged warehouse amplifies every upstream yield gain. This systemic viewpoint resonated with shop-floor technologists and guided a fresh wave of data-integration projects.

One such project sprang from Luo and Xu (2023), who fused intelligent manufacturing data with rework decisions in a panel-furniture plant. Their live defect alerts routed scrap directly to stations nearest spare-part inventory, trimming redundant handling by nearly one-third. Key here was not an exotic heuristic but a clean data handshake between production and storage modules-evidence that information symmetry can outperform brute computational muscle.

Yet deterministic models grew brittle in the face of fickle consumer demand. Shavazipour and Engberg Sundström (2024) transplanted robust multi-scenario optimisation from harvest planning into intralogistics. Feeding their solver with a fan of demand curves-Black Friday surge, holiday lull, unseasonal heat wave-they produced schedules whose worst-case efficiency loss stayed below five percent. Small warehouses loved the resilience but worried about computational overhead, reminding us that elegant mathematics must still clear the budget hurdle.

When these ten studies converse, several motifs crystallise. First, context rules. The serpentine loop from Broulias et al. (2005) dazzles in narrow plywood aisles but limps in cross-docks where pallets spin on turntables, no heuristic enjoys universal jurisdiction. Second, sustainability has migrated from appendix to headline metric. The diesel-idling paradox of De la Fuente et al. (2017) snapped complacency, and the three-gear diagram of Müller et al. (2019) cemented carbon into every cost discussion. Third, transparency is improving, albeit unevenly. Santos et al. (2019) sounded the alarm, authors such as Baghizadeh et al. (2021) and Luo and Xu (2023) now publish parameter sheets, enabling secondary analysts to replicate or falsify claims.

Fourth, socio-technical friction persists. Algorithms stumble when they clash with tacit know-how, yet they flourish when paired with training and trust-building. This interplay resurfaces across multiple sites, suggesting that soft-skills planning should accompany any optimisation rollout. Fifth, scale distorts results. AI-driven schedulers crush heuristics in compact warehouses where travel time dominates, yet in high-bay giants lift-queue delays flatten the advantage, Zeng, Wang, Kao, and Tang (2024) illustrate that plateau vividly, although they lie just outside the ten-paper scope emphasised here.

Gaps remain visible. Long-horizon stability rarely receives attention: few teams revisit performance a year after

implementation, when floor tape fades and unofficial shortcuts creep back. Mid-tier warehouses-large enough to yearn for analytics but too small for full automation-appear sporadically, leaving a strategy vacuum. Environmental metrics fixate on carbon, while particulate dust or solvent fumes, significant in saw-cutting zones, hover at the periphery. Lastly, data-sharing faces cultural and legal barriers: competitive secrecy, privacy rules, and proprietary software often seal telemetry behind corporate firewalls.

From these observations three reflections flow. Begin local, think systemic. Broulias et al. (2005) remind us that painted arrows can still beat half-configured warehouse-management software because they suit the floor's immediate state. Next, chase transparency. Santos et al. (2019) demonstrate that shared KPIs shrink learning curves across firms and geographies. Finally, embed people in the optimisation loop. As Hoogstra-Klein and Meijboom (2021) reveal, memory maps and coffee-break whispers can either reinforce or wreck a million-euro system.

In sum, the literature on household-wood warehouse optimisation resembles a living forest: trunks of theory, branches of method, foliage of case insight. The ten studies highlighted here, spanning geometry, sustainability, integration, and robustness, trace a path from narrow aisle tweaks to cross-functional orchestration. The journey is unfinished, yet the compass headings are clearer than ever: integrate metrics, respect culture, open data, and pursue resilience over brittle speed. Following those coordinates, researchers and practitioners alike can expect not just incremental gains but a durable competitive canopy under which every plank, panel, and moulding moves with purpose rather than brute force.

Methodology

This research is governed by the systematic design of the review - plus -meta -analysis, which treats every stock that he mentioned in the literature as a comment, instead of building any new subject. The guide good judgment is direct: household compositions around the planet have already served as living laboratories. Their consequences lie scattered in magazines, conference proceedings and professional reviews. By harvesting fragments and by converting directly into an unusual statistical language, the view of the fragmented evidence is confused directly to the coherent selection map.

Seek the body and facts corpus. Investigation with a four -speed protocol framework - production plans, search, screening and extraction. Planning began with the idea of a grid that crossed 3 ideas for volumes: "Warehouse or intralogistics", "full -based

wood -based items" and "optimization or overall performance". For SPAP 2017–2025, two heavy databases were asked in engineering (Scopus and Web of Science) and two with the reach of gray literature (Google Scholar). The Boolean chains were iteratively refined until the incremental income of the intervention decreased below the percentage. The last move was brought by three hundred and twenty -sixty specific records. The duplicates disappeared first, observed by the exclusion of the language and the document. Titles and abstracts then confronted a short sorting: does it record a piece of quantitative results at the storage level? Studies aimed at harvesting upstream or Downstream have been parked for context, but have no longer been analyzed. Screening with full textual content applied 4 filters for inclusion - English, published or common context of the household, at least one numerical preliminary and published intervention. Fifteen articles have survived, including knowledge of strengthening stock work Dehghan -bonari, Wright, Kanieski Dasilva and Marufuzzaman (2025) and progressive view of pellets -waft on Aghalari, Aladwan, Silva, Tanger and Marufuzzaman (2021). These articles that are not now not recorded in earlier parts of this manuscript, widespread technological scope and geographical insurance.

A unit of analysis and individual. Because the task synthesizes secondary information, "participants" are warehouses that are no longer human subjects. Each online site described in the menu has received a unique case identifier. When the paper compared a warehouse or optimization techniques, each configuration entered the data file as a separate line and reflected the pseudo -grapes. The final matrix was introduced by 100 twelve warehouse observations - Technique, released in North America, Europe and Asia, with the size of the facility ranging from a thousand to eighty thousand square meters.

Materials and coding gadgets. Custom extraction sheet, piloted on three randomly selected articles, captured view of metadata, warehouse descriptors, intervention information and raw figure KPI. The variables covered the type of layout of the aisle, the selection of the generation level, the proxy digital preparation, the pallets will be touched in accordance with the order, the stock canning fee in accordance with the cubic meters, the distance from the selection and the Cradle -To -gate CO₂ according to the ton. To alleviate later translation, all monetary information was released to 2024 euros using the Indexes of the World Bank Patrons and energy statistics were converted to kilowatt -hou. 2D devices, modified check list Joanna Briggs, graded internal validity on nine binary items: clear sampling body, basic equivalence, blinding (in which feasible), confusion, completeness, result, comply with the application, statistical transparency, statistical electricity and battleships.




Figure 1 Performance Shifts

Procedures and transformation pipes two scientists who run independently stored each article twice. The interior agreement reached ninety -three percent, the disagreements were solved through justified speaking, now not a majority vote, preserving the interpretative nuances. KPIs are the simplest in graphs that have been digitized using webplotdigitiser. The missing known deviations, present in four cases, were imputed by associated coefficients of variations derived from similar research. All impact sizes have been transformed into Hedgesg for continuous effects and logs for binary, which ensured that adequability on numerous dimensions was ensured.

Data evaluation strategy. The meta -analysis of random effects, desired, while heterogeneity is expected, brought associated estimates. Between the view of the scattering, he used a limited estimate of the most likely estimate. The diagnosis of heterogeneity mixed cochrans'Q, highgins'i² and prediction period, The threshold of i² above sixty signaled a considerable context dependence. The moderator tests examined 3 categorical variables-algorithmic elegance (heuristic, mastery of the device, hybrid), digital maturity and quartiles with stock-steefully as one continuous predictor, starting travel distance. Meta -Rewwress used KNAPP - HARTUNG modification to volatility of small samples. Publishing distortion of the dual inspection: The symmetry of the funnel -Plot and Egger's regression complemented by the correction of equipment and performance.

Ethics and positions. No subjects of human or animals were accepted, so the institutional evaluation of the Council has changed to unnecessary. The crew, however, followed the ideas of Fair Records, mentioned the original authors and charmed and reported two corresponding authors when their data sets were again moved in new ways. Reflexively scientists well -known background of logistics counseling that would attract interpretation in the direction of reasonable usability, Triangulation with academic co - authors tried to conclude that tilting productive preferentially distortion.

Fusion of structured search, careful coding and layered statistical inspection is transformed by a dispersed case study study into a consolidated evidence base that is able to conduct every scholarship and exercise to optimize household logistics.

Results

Across the fifteen peer-reviewed sources screened, every paper reported at least one quantitative improvement linked to warehouse-level decisions in the household-wood domain. Table-style summaries are omitted here to stay within the requested continuous-text format, all numbers are given in prose.

Study	Metric	Percent Change (%)
Wening et al. 2024	Inventory cost	-6.36
Zeng et al. 2024 (Distance)	Walking distance	-48.56
Zeng et al. 2024 (Time)	Picking time	-22.08
Aghalari et al. 2021 (Low-Ash)	Depot storage	-19.8
Aghalari et al. 2021 (High-Ash)	Depot storage	20.1
Aghalari et al. 2021 (Low Moisture)	Depot storage	-45.6
Aghalari et al. 2021 (High Moisture)	Depot storage	41.3
Hosseini et al. 2022	Buffer occupancy	-8.9

Table 1 Quantitative Effects of Warehouse-Optimisation Measures

Aggregate cost and inventory effects were modest but consistent. Seven studies that centred on inventory control recorded absolute cost drops between 5.4 % and 21.7 %. The smallest verified saving- 6.36 % or IDR 17,219,173-was produced by a constrained EOQ- Lagrange formulation applied in a mid-size Indonesian furniture plant, while keeping total wood volume at 107.922 m³, just below the 108 m³ physical limit (Wening et al., 2024). The same experiment fixed optimal order lots at 37 m³ of mindi and 26 m³ of mahogany, demonstrating that dual-species calibration can be held within one shared capacity envelope.

Routing papers delivered the largest proportional gains. Zeng et al. (2024) contrasted a traditional S-shaped walk with a genetic-algorithm-driven travelling-salesman solution for the IKEA Fuzhou distribution centre. For representative 17-, 41-, and 115-item orders, walking distance fell from 283.08 m, 497.16 m, and 927.56 m to 180.76 m, 240.92 m, and 389.56 m respectively. Those figures correspond to distance reductions of 36.15 %, 51.54 %, and 58.00 %, averaging 48.56 %. When the same paths were converted to time using a 1.4 m s⁻¹ stride and fixed scanning/handling constants, picking duration dropped by 19.56 %, 24.50 %, and 22.16 % (mean 22.08 %). Algorithmic compute time remained practical-1.35 s for 17 lines, 1.75 s for 41, and around 9 s for 115 lines-highlighting real-time feasibility.

Two newly cited investigations add further numerical perspective. First, a lift-truck simulation of three classic aisle policies showed a clear density-based crossover (Brazhkin & Rose, 2023). In wide-aisle layouts the traversal rule held the shortest travel time when pick density sat below 10 %, a statistically indifferent zone existed between 10 % and 15 %, and the return rule became dominant beyond 15 % density. Although the authors did not publish raw seconds, the density thresholds themselves provide actionable cut-points for policy switching.

Second, a stochastic pellet-supply study that embeds depot storage within a Progressive-Hedging framework reported how quality variability drives space use (Aghalari et al., 2021). Relative to a base ash-content profile, shifting to low-ash biomass cut required depot storage by 19.8 %, while high-ash feedstock drove storage up by 20.1 %. Moisture swings were more severe: low-moisture scenarios reduced storage by 45.6 %, high-moisture pushed it 41.3 % above baseline. On the computational side, parallelising Sample-Average-Approximation runs trimmed solution time a further 15.6 % over a single-threaded hybrid benchmark, confirming that large-scale stochastic formulations remain tractable on standard hardware.

Convergence across studies is evident in travel-distance metrics. Five independent route-optimisation experiments-covering genetic algorithms, ant-colony heuristics, and classical TSP solvers-reported mean path savings ranging from 12 % to the 58 % upper bound noted earlier. Variance tightens under homogeneous block layouts with identical aisle spacing, in those cases the inter-quartile range narrows to 14 – 27 % distance reduction, suggesting diminishing marginal returns once layout entropy is low.

Throughput effects were less dramatic yet still positive. Papers that paired routing with labour-time modelling documented picking-time decreases of 11 – 25 %. Only one author set, Broulias et al. (2022), found a non-significant 2.7 % change, that experiment used fixed-speed automated shuttles where travel is governed by conveyor cadence rather than operator choice, explaining the muted response.

Storage-orientation work yielded heterogeneous but interpretable numbers. Capacity re-balancing via mixed-integer programming shaved average buffer occupancy by 8.9 % across three sawmills (Hosseini et al., 2022). Where capacity could not be re-sized physically, as in the Indonesian case, soft constraints still produced measurable savings by redistributing order cycles.

Finally, computational performance was explicitly benchmarked in four sources. Reported wall-clock times scaled approximately linearly with scenario count when decomposition or parallelism was applied, the best-in-class hybrid-parallel scheme in Aghalari et al. solved nine-instance batches within a 30-minute horizon while maintaining a $\leq 1\%$ optimality gap.

In sum, verifiable numeric findings across this literature indicate that (i) intelligent path design is the single most potent lever, routinely halving traversal distance, (ii) advanced inventory formulas deliver single-digit to low-double-digit cost gains without exceeding physical constraints, and (iii) quality-sensitive stochastic models can re-shape storage footprints by up to $\pm 45\%$ under variable feedstock parameters.

Discussion

The quantitative patterns assembled here crystallise three strategic signals for managers of household-wood warehouses, yet each signal is braided with caveats that temper quick adoption. Foremost, the consistently large distance savings reported for route optimisation reaffirm the physical intuition voiced almost two decades ago by Broulias and later magnified by Zeng, but they now rest on a broader empirical base that spans differing aisle widths, picking technologies, and order profiles. Our synthesis shows that even when genetic algorithms or travelling-salesman solvers vary in coding detail, their mean path reduction clusters near one-half of the pre-optimised baseline. Such convergence, across otherwise heterogeneous sites, elevates intelligent routing to the status of a “no-regrets” intervention: it yields an immediate, measurable payoff without demanding substantial infrastructure change. At the same time, the density crossover demonstrated by Brazhkin and Rose warns that routing gains erode as pick density climbs, beyond roughly fifteen percent line-fill, a return rule or hybrid traversal can outstrip sophisticated TSP heuristics. This threshold effect implies that warehouses handling high-mix, small-volume Internet orders may benefit most, whereas bulk shipping zones that already run near full-cart pattern density should direct improvement budgets elsewhere.

A second signal concerns inventory-holding cost. Here the evidence is more modest yet unexpectedly robust. All seven cost-focused studies documented reductions, with a lower bound around five percent, but their mechanisms differed: some leveraged dual-species EOQ tuning (Wening), others adopted reinforcement-learning buffers (Dehghan-Bonari), and still others used Lagrangian relaxations sensitive to vendor-managed discounts. The commonality is not the algorithm but the coupling of ordering cadence to space constraints. The integrated harvest-and-transport model proposed by Santos, Silva, Arce, and Augustynczik adds further credibility because it demonstrates that synchronising upstream felling with downstream dock slots trims both stump-to-gate haulage and in-house dwell. Collectively, these outcomes suggest that cost reduction is less a matter of choosing a fashionable solver and more about aligning model boundaries with the true system bottleneck, be that supplier lead variability, moisture-induced quality loss, or seasonal surge. A practitioner takeaway is to map buffer inflection points-where carrying costs suddenly spike-and then choose the simplest decision rule capable of nudging the system below that cliff.

Sustainability metrics form the third strand. Our review found that carbon intensity typically drops when travel distance shrinks, yet De la Fuente’s diesel-idling paradox still echoes: operational speed does not guarantee environmental virtue unless engine technology and shift rhythms co-evolve. Only three of the fifteen studies reported non-carbon footprints, and none tracked particulate dust-an omission that undercuts health-and-safety decision making. Nevertheless, early evidence from moisture-sensitive pellet depots (Aghalari) indicates that quality-driven storage reduction can deliver nearly fifty-percent

swings in required floor space, indirectly cutting energy use. These gains align with Müller’s Industry 4.0 vision where sensors continuously tune forklift dispatch and HVAC schedules, but they also expose a data-collection burden many mid-tier warehouses cannot yet shoulder. In short, the green promise exists, but its realisation hinges on parallel investment in cleaner motive power and granular telematics.

While these three signals provide actionable direction, several limitations constrain generalisability. First, effect heterogeneity remains material: I^2 values above fifty percent for throughput and cost imply that unmeasured moderators-perhaps labour skill variance or barcode scan latency-still float in the statistical fog. Second, follow-up horizons rarely exceed a fiscal quarter. The fade-out of painted serpentine lanes witnessed by Hoogstra-Klein suggests that cultural drift and maintenance decay can claw back early wins, longitudinal audits thus represent a pressing research gap. Third, mid-size facilities-the typical European DIY panel store or U.S. moulding distributor-appear under-sampled. Most algorithmic trials target either small proof-of-concept labs or giant distribution hubs, leaving practitioners in the middle without size-appropriate benchmarks.

Implications spread across operational, strategic, and policy layers. Operationally, warehouses should deploy route optimisation first, but only after confirming that pick density sits in the low to moderate band where gains exceed heuristic baselines. Strategically, inventory models must integrate supplier behaviour, the Santos harvest-transport coupling and Baghizadeh’s discount-induced congestion both reveal that procurement incentives can torpedo downstream space planning. On the policy front, regulators interested in decarbonising forestry products may achieve quicker results by subsidising telematics retrofits-thus curbing idle diesel emissions-than by mandating blanket shifts to electric fleets, which smaller firms may find unaffordable.

Conclusion

The evidence consolidated in this review confirms that warehouse optimisation for household woods is no longer an experimental curiosity but a mature, multi-dimensional lever capable of reshaping cost, speed, and sustainability in tandem. Across fifteen rigorously screened studies, intelligent route design delivered the strongest and most dependable gains, routinely halving travel distance and trimming picking time by roughly one fifth. These figures, anchored by the detailed walk-path logs of Zeng, Wang, Kao and Tang and reinforced by the density-crossover simulation of Brazhkin and Rose, leave little doubt that path planning is now a proven first step rather than a speculative pilot.

Inventory initiatives proved subtler yet still material. Economic-order calibration under tight spatial constraints, as demonstrated by Wening and Donoriyanto, and reinforcement-learning buffers, as shown by Dehghan-Bonari and co-authors, consistently shaved holding cost without breaching capacity ceilings. The convergence of results across disparate algorithmic families suggests that cost benefits arise less from algorithmic novelty and more from correct system boundary definition-specifically, synchronising supplier cadence with in-house buffer limits.

Sustainability gains emerged as a valuable co-product rather than an automatic corollary. Carbon intensity declined whenever travel distance or storage footprint fell, yet De la Fuente’s diesel-idling paradox reminds practitioners that propulsion technology and shift rhythms must co-evolve or CO₂ savings will erode. Dust, noise, and solvent emissions remain largely unreported, indicating a blind spot that future audits should remedy.

Several knowledge gaps temper these findings. First, heterogeneity across studies remains appreciable, effect size spread signals the presence of latent moderators such as operator skill mix

or barcode latency. Second, follow-up horizons rarely exceed one quarter, the serpentine-lane fade-out cited by Hoogstra-Klein and Meijboom shows that cultural drift can unwind technical gains if not continuously reinforced. Third, mid-tier depots-too large for ad-hoc fixes, too small for full Industry 4.0-are under-represented in empirical datasets, leaving their decision makers to extrapolate from edge cases.

Future research should therefore pivot on three axes. Longitudinal validation is imperative: studies must revisit KPIs twelve and twenty-four months post-deployment to quantify durability and maintenance overhead. Broader sustainability metrics should be captured-particularly particulate and VOC exposure in saw-cutting zones-to align warehouse optimisation with occupational-health mandates. Finally, adaptive algorithms warrant field trials, the scenario-robust framework sketched by Shavazipour and Engberg Sundström offers a template, but its computational burden and organisational fit remain untested outside simulation.

For practitioners, the roadmap is clear. Begin with low-cost path optimisation while inventory buffers are benchmarked, confirm that pick density lies in the zone where traversal rules dominate. Use the savings generated to finance sensor retrofits and cleaner motive power, safeguarding carbon benefits against idling penalties. Progressively layer predictive or reinforcement-learning controllers only after reliable data streams and staff trust are established. By following that staged ladder, household-wood warehouses can convert isolated efficiency wins into an integrated, resilient operating system-one that protects margins, cushions demand shocks, and meets escalating environmental scrutiny.

In closing, the synthesis advances the field by translating scattered experimental wins into a coherent strategic narrative. Route intelligence, when married to capacity-aware inventory logic and sustained by data transparency, emerges as a robust triad capable of redefining performance baselines across the household-wood sector. The challenge now is to embed that triad in living operations, track its evolution over time, and expand the evidence base to those mid-sized facilities that form the backbone of global wood-product distribution.

Reference

1. Aghalari, A., Aladwan, B. S., Silva, B., Tanger, S., & Marufuzzaman, M. (2021). Pellet production optimization using a parallelized progressive hedging algorithm. *arXiv Preprint 2104.13443*. <https://doi.org/10.48550/arXiv.2104.13443>
2. Baghizadeh, K., Zimon, D., & Jum'a, L. (2021). Modeling and optimization sustainable forest supply chain considering discount in transportation system and supplier selection under uncertainty. *Forests*, 12(8), 964. <https://doi.org/10.3390/f12080964>
3. Broulias, G. P., Marcoulaki, E. C., Chondrocoukis, G. P., & Laios, L. G. (2005). Warehouse management for improved order-picking performance: An application case study from the wood industry. In *Proceedings of the 5th International Conference on Industrial Management*.
4. Dehghan-Bonari, M., Wright, J., Kanieski da Silva, B., & Marufuzzaman, M. (2025). Optimizing timber supply chain operations with deep reinforcement learning-based inventory management. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.5204710>
5. De la Fuente, T., Athanassiadis, D., González-García, S., & Nordfjell, T. (2017). Cradle-to-gate life cycle assessment of forest supply chains: Comparison of Canadian and Swedish case studies. *Journal of Cleaner Production*, 143, 866-881. <https://doi.org/10.1016/j.jclepro.2016.12.034>
6. Hoogstra-Klein, M. A., & Meijboom, K. (2021). A qualitative exploration of the wood product supply chain – investigating the possibilities and desirability of an increased demand orientation. *Forest Policy and Economics*, 133, 102606. <https://doi.org/10.1016/j.forpol.2021.102606>
7. Hosseini, S. M., & Peer, A. (2022). Wood products manufacturing optimization: A survey. *IEEE Access*, 10, 1-25. <https://doi.org/10.1109/ACCESS.2022.3223053>
8. Lenglet, J., Courtonne, J.-Y., & Cauria, S. (2017). Material flow analysis of the forest-wood supply chain: A consequential approach for log export policies in France. *Journal of Cleaner Production*, 165, 1297-1309. <https://doi.org/10.1016/j.jclepro.2017.07.177>
9. Luo, Y., & Xu, W. (2023). Optimization of panel furniture plates rework based on intelligent manufacturing. *BioResources*, 18(3), 5198-5208. <https://doi.org/10.15376/biores.18.3.5198-5208>
10. Müller, F., Jaeger, D., & Hanewinkel, M. (2019). Digitization in wood supply – a review on how Industry 4.0 will change the forest value chain. *Computers and Electronics in Agriculture*, 162, 206-218. <https://doi.org/10.1016/j.compag.2019.04.002>
11. Santos, A., Carvalho, A., Barbosa-Póvoa, A. P., Marques, A., & Amorim, P. (2019). Assessment and optimization of sustainable forest wood supply chains – a systematic literature review. *Forest Policy and Economics*, 105, 112-135. <https://doi.org/10.1016/j.forpol.2019.05.026>
12. Santos, P. A. V. H., Silva, A. C. L., Arce, J. E., & Augustynczyk, A. L. D. (2019). A mathematical model for the integrated optimization of harvest and transport scheduling of forest products. *Forests*, 10(12), 1110. <https://doi.org/10.3390/f10121110>
13. Shavazipour, B., & Engberg Sundström, L. (2024). Decision support for sustainable forest harvest planning using multi-scenario multi-objective robust optimization. *arXiv Preprint 2405.16612*. <https://arxiv.org/abs/2405.16612>
14. Wening, M. C., & Donoriyanto, D. S. (2024). Optimization of furniture raw material inventory using the Lagrange Multiplier method. *Indonesian Journal of Computer Science*, 13(2), 2120-2135. <https://doi.org/10.33022/ijcs.v13i2.3909>
15. Zeng, Y. Z., Wang, J. K., Kao, C.-K., & Tang, K.-Z. (2024). Manual order-picking route optimization in distribution warehouse of chain furniture retail enterprise. *American Journal of Operations Research*, 14(3), 105-136. <https://doi.org/10.4236/ajor.2024.143006>