

Article

Integrating the Jev AI Decision Model with Projection Guidance in Manufacturing: A Human-Supervised Proof of Concept for Bin Preparation

Adrián Vodilka *, Monika Töröková *, Marek Kočíško, Peter Gabštur and Jakub Kaščák

Faculty of Manufacturing Technologies, Technical University of Košice, Bayerová 1, 080 01 Prešov, Slovakia; marek.kocisko@tuke.sk (M.K.); peter.gabstur@tuke.sk (P.G.); jakub.kascak@tuke.sk (J.K.)

* Correspondence: adrian.vodilka@tuke.sk (A.V.); monika.torokova@tuke.sk (M.T.)

Abstract

During part returns in manufacturing, deciding when to prepare an empty bin involves the current bin state, the remaining workload and the time needed for preparation. This study presents a human-supervised proof of concept linking projection guidance technology with Jev, a TypeSafe AI model that returns structured choices and uncertainty information. An existing picking workflow supplies the platform context. A separate extension converts depth-tool measurements into bin states, combines them with reported workload and preparation time, and asks Jev whether an empty bin should be prepared. A companion checks the response, records approval and projects request status. Four diagnostic calls to Jev were recorded across three software versions, all below a fixed confidence threshold of 0.80. In the final attempt, operator approval of a 0.75-confidence proposal created a preparation request; a later full reference triggered deterministic exchange handling. Local acknowledgments and confirmed emptying of the same bin completed one Bin1 lifecycle. The result demonstrates technical integration and decision traceability. Readiness was simulated; physical spare-bin exchange and external dispatch were not demonstrated. Zone events do not verify part identity, quantity or removal, and mask readings do not measure fill volume. Detection reliability and operational benefit remain unmeasured.

Keywords: artificial intelligence; Jev; decision support; projection guidance; LightGuide; depth sensing; intralogistics; human oversight

Academic Editors: Janis Arents,
Vytautas Bucinskas and Andrius
Dzedzickis

Received: 3 September 2026
Revised: 5 October 2026
Accepted: 8 October 2026
Published: 10 October 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The transition from mass production to mass customization increases the number of product variants and shortens the period during which any one work procedure remains stable. Manual operations nevertheless remain important in high-variety assembly because of their flexibility [1–3]. Component kitting, which comprises the selection, checking, and preparation of parts for subsequent assembly, is an important quality-control point. An incorrect or incomplete kit may remain undetected until assembly or final inspection, creating risks of rework, downtime, and nonconforming products [4,5].

Conventional paper pick lists and bills of materials are easy to deploy, but they are static, do not adapt automatically to a particular product variant, and do not independently verify a pick. Their effectiveness therefore depends heavily on operator

attention. Task complexity and information presentation can affect cognitive workload and performance in manual assembly [6,7]. In a high-mix environment, there is accordingly a need not only to present instructions but also to connect them to an observable indication of the operator's action.

Assistance technologies range from pick-to-light and pick-by-voice systems [8,9] to augmented-reality approaches [10,11]. Projection-based, or spatial, augmented reality places instructions directly in the workspace, so an operator need not wear a display or look away toward a separate screen. Controlled comparisons have found that projection can reduce completion time or errors, or support procedural learning, relative to paper, monitors, or head-mounted displays; however, outcomes depend on the task, interface, and experimental design [12–15].

More recent systems combine projection with machine vision or gesture recognition so that guidance can respond to both expected and erroneous actions [16–18]; vision-based worker-assistance systems pursue a related quality-control objective [19]. Displaying an error message alone, however, does not guarantee a measurable performance improvement; the effect depends on feedback design and the context of use [20]. Systematic reviews also identify barriers to the wider industrial adoption of augmented reality [21,22]. Work on assistive systems in real production settings further identifies integration and instruction-design challenges [23].

Returning parts to storage introduces a decision beyond identifying a destination: an operator may need an empty bin prepared before the current one reaches its full reference state. A depth-based warning supplies information about the present scene, while the preparation decision can also consider the reported remaining work, preparation time and existing requests. This motivates an exploratory connection between artificial intelligence (AI) advice and the projected instructions used at the workstation.

Jev is an AI decision model developed by TypeSafe AI and accessed through a hosted application programming interface (API). The vendor calls this model class System One. It evaluates a supplied state and a defined question, returning structured values that an application can process directly [24,25]. In this study, Jev receives a categorical bin state and reported task context and proposes one of a fixed set of logistics actions. Projection guidance technology and the companion perform the sensing, numerical checks and display control. The term human-supervised denotes operator confirmation and review of actions; it does not refer to supervised machine learning, and no model training was performed at the workstation.

The research aim is to connect structured AI advice to commercial projection guidance with explicit responsibility for sensing, recommendation and action. An audit of the existing picking program establishes the platform context. A separate extension links depth-derived bin states and reported task context to Jev advice, human review, local logistics requests and projected status. The contribution is a documented integration architecture and a recorded preparation-to-emptying lifecycle. The original picking program and the extension are assessed using their respective retained evidence. Both use zone or mask measurements that do not independently verify part identity, quantity or removal. The study addresses three research questions:

1. How does the workstation connect sensor variables, workflow state and projected instructions?
2. How can Jev advice enter this information path with explicit human and deterministic decision authority?
3. What integration behavior is supported by the recorded supervised logistics lifecycle, and what remains unverified?

The remainder of the article is organized as follows. Section 2 reviews relevant work. Section 3 describes the original Extensible Markup Language (XML) audit and the added

sensing, decision and capture protocol. Section 4 separates original artifact findings from the recorded Jev attempts and Bin1 lifecycle. Section 5 discusses implementation implications, reconfiguration, safety and evidence limits, and Section 6 states the conclusions.

2. Related Work

Worker information systems can be classified by their information medium, degree of interactivity, contextual integration, and feedback capability [1,26]. Their effectiveness depends not only on the device but also on how the instruction is designed and positioned relative to the task [27,28]. More recent approaches further emphasize that the level of assistance should reflect the worker's experience and competence [3]. Research on picking assistance includes a comparative analysis of paperless picking systems [8] and an experimental comparison of paper, light, cart-mounted and head-up displays [9]. Unless a configuration includes independent sensing or confirmation, however, it only indicates the target location and leaves pick accuracy dependent on the operator. In mixed-model kitting, Hanson et al. [10] found that augmented-reality guidance could be competitive with a paper list in both time and quality.

Projection-based augmented reality places graphical and textual instructions directly in the workspace. Funk et al. [29] introduced a benchmark task for comparing interactive assembly instructions, and subsequent experiments compared projection with paper, stationary displays, and head-mounted devices [12,13]. The findings favor projection particularly when limiting visual shifts between the instruction and the work area is important. Mark et al. [14] examined learning with projection-based assistance, while Oestreich et al. [30] reported similar learning curves for screen-based digital assistance and personal instruction. Industrially oriented systems have further demonstrated the applicability of projection to assembly and manual workstations in smart manufacturing [15,31,32]. Both the potential and the practical barriers of industrial augmented reality have been discussed from early manufacturing applications through later assessments of real-world deployability [33,34].

Presenting information alone does not close the quality-control loop. A system must observe an action, compare that observation with the expected state, and use the result to govern the next step. Integrations of projection and computer vision have been described by Chen et al. [16] and Ojer et al. [17]; Riedel et al. [19] instead combined computer vision with non-projection worker-assistance media for manual-assembly quality assurance. The gesture-controlled Gest-SAR system likewise reports closed-loop feedback for correct and erroneous actions [18]. Such findings cannot be generalized to every projection implementation: an experiment using simulated automatic error detection found no overall performance benefit relative to the control condition, although usability ratings differed [20]. A parallel research stream examines the authoring and reuse of in situ instructions [35,36]. These studies establish the importance of a machine-readable process model, but they generally use custom prototypes or do not disclose the detailed control logic of a commercial system.

Current research also connects assistance systems with the Operator 4.0 concept [2] and with manufacturing digital twins [37], for which step-execution records can support analysis and model parameterization. Reports from real manufacturing prototypes identify integration and instruction design as distinct challenges [23]. The present study examines an additional integration question: how a typed AI recommendation can become a projected logistics instruction with a traceable decision path. A fixed warning rule already supplies a numerical response to a reference state. Jev is introduced to explore advice conditioned on both that state and reported task context. Whether this improves

preparation timing or decision quality requires subsequent evaluation; the current study investigates the implementation and supervised request lifecycle.

AI-based decision support is also expanding in manufacturing and logistics. Reviews describe the advantages and challenges of machine learning in manufacturing [38] and identify applications and research directions for artificial intelligence in smart logistics [39]. When such advice reaches an operator, the allocation of authority becomes a design question. Automation can support information acquisition, information analysis, decision selection or action implementation at different levels [40], and decision aids can induce complacency or automation bias, in which advice is accepted without sufficient independent verification [41]. These findings motivate the explicit separation of model advice, deterministic rules and operator approval examined in the present extension.

3. Materials and Methods

The study has two evidence phases. Phase 1 audits the original industrial picking artifact: a well-formed XML document with sequential commands and six workflow-stage markers, accompanied by a project narrative, vendor documentation and four illustrative video stills. The original external vision program, variable-table configuration, hardware calibration and commissioning logs were not retained. The stills show interface states without a synchronized event record. Phase 2 is a supervised workbench demonstration on 2 October 2026 using a separate capture XML, a saved revised depth-vision configuration, a Python ver. 3.13 companion and Jev. Its archive includes runtime settings, corrected attempt records, an exported event-log snapshot and request-lifecycle evidence. It does not recover the missing original wrong-bin expression or constitute validation of the original picking detector. Both phases distinguish encoded properties, recorded runtime events and unmeasured operational effects. No XML Schema Definition was supplied; native exports were checked for well-formedness and internal consistency rather than platform-schema validity.

3.1. Case Study Design, Industrial Setting, and Evidence Corpus

The application was developed through cooperation between a research institution and an anonymized industrial partner that manufactures precision mechanical and electromechanical products in multiple configurations. The project documentation describes component kitting as the first operation in the assembly route. The operation comprises selecting, checking, and preparing mechanical and electrical parts from an assembly sheet and engineering bill of materials. The same documentation describes many possible configurations for individual product models, which makes selection error a material risk in manual kitting. The documented baseline is a fully manual operation in which the operator visually checks part numbers without an assistance system or automatic selection check. These company-specific descriptions were not independently measured in the present artifact audit.

The original case-study workflow comprised four stages: (1) analysis of the kitting operation and its selection risks with company engineers; (2) design and implementation of a task-specific program on the commercial projection platform; (3) reported scenario-based commissioning covering target-bin entry, wrong-bin entry, manual navigation, and cyclic operation; and (4) the present XML conformance audit and evidence-boundary assessment. The pilot task was a fixed sequence of picks from four storage bins arranged on the work surface. Figure 1 summarizes the relationship between the retained evidence and the conclusions drawn from it. The later logistics extension was implemented and demonstrated separately at the research workbench using controlled empty, warning and full references. Only Bin1 was enabled for capacity decisions in the deployed companion configuration; Bins2–4 remained raw-monitor inputs. Figure 1 concerns the original phase.




Figure 1. Original case-study workflow from operation analysis and XML implementation through the reported commissioning account and the present evidence audit.

3.2. Projection-Based Visual Guidance Platform

The pilot workstation uses LightGuide software 25.2.6.0, a commercial projected work-instruction platform (LightGuide, Inc., Wixom, MI, USA), which combines in situ visual guidance with sensor and vision integration [42]. The software, a commercially available product rather than a research prototype, was used as supplied; the authors' contribution is the task-specific program and the separately developed logistics companion. The original retained architecture comprises three main subsystems (Figure 2):

- The projection subsystem, which places vector shapes, raster images, symbols, and text on the work surface in a calibrated canvas coordinate system;
- The machine-vision subsystem, which is intended to detect hand entry into monitored zones and expose zone activity as Boolean variables;
- The control subsystem, which interprets the XML program, selects projected elements, reads vision variables, and evaluates step transition conditions.

Unlike a conventional human–machine interface on a separate screen, the projection system places contextual information at the point of work and can reduce gaze shifts between the instruction source and the task [12,15]. Figure 2 distinguishes the retained work-instruction XML from the external vision configuration, which was not available for the original picking audit. The new logistics architecture and its retained vision settings are described separately in Sections 3.6 and 4.3.




Figure 2. Anonymized original architecture of the projection workstation and its information loop, using functional labels in place of platform-specific XML and variable identifiers and distinguishing the retained controller and projection path from the unaudited original external vision configuration.

3.3. XML Representation of Work Instructions

LightGuide software exports work instructions as structured XML. The retained XML contains metadata and an ordered sequence of commands. Each command records an order position, a command class, and a description; type-specific elements encode display properties, variable assignments, or confirmations. For the original industrial program, the manuscript reports an abstraction-level structure and excludes license, file-system, asset and identifying project metadata. The separate logistics extension reports relevant software versions and numerical settings; de-identified exports and companion source are available from the corresponding author on reasonable request. Table 1 lists the abstract command roles used in the workflow. This representation enables review of the control logic outside the graphical editor and provides a basis for future generation from product data (Section 5.4).

Graphic and text display records provide presentation, while assignments and transition conditions provide control. Presentation elements may be bound to a Boolean condition and are then visible only when that condition is true. Elements carrying a persistent-display flag survive subsequent canvas-reset records. The wrong-bin warning is created in the initialization stage as a persistent element, so it remains available during the bin-picking stages. For confidentiality, Figure 3 is an abstracted, non-source representation rather than a verbatim export.

ABSTRACTED TARGET-STAGE REPRESENTATION

```
<WorkflowFragment>
  <Action type="set-target">
    <Value>bin-1</Value>
    <Destination>target-bin</Destination>
  </Action>
  <Display type="instruction">
    <Text>Pick from 1</Text>
  </Display>
  <Display type="positive-feedback">
    <Text>Correct bin</Text>
    <Condition>target-bin-active</Condition>
  </Display>
  <Transition>
    <Condition>target-bin-cleared</Condition>
    <Direction>Forward</Direction>
    <Steps>1</Steps>
  </Transition>
</WorkflowFragment>
```

Figure 3. Abstracted XML-like representation of a target stage, generalizing names, values, coordinates, and element structure while retaining only the target assignment, instruction, conditional feedback, and automatic forward transition.

Table 1. Abstract command roles present in the retained XML work-instruction model.

Command Role	Function	Abstracted Fields
Workflow boundary	marks the beginning and end of a workflow	status
Stage marker	starts and numbers a logical workflow stage	stage identifier
Canvas reset	clears nonpersistent visual elements from a canvas	canvas
Graphic display	projects an image or vector graphic at a defined position	position, dimensions, condition
Text display	displays an instruction or status message	text, position, condition
Variable assignment	assigns a target or control value	value, destination, recalculation flag
Transition condition	waits for a Boolean condition and defines a stage transition	condition, edge, direction, stages to move

The same logical pattern recurs across the bin-picking stages: a target-bin assignment, instruction text, conditional feedback, and transition condition. The XML is therefore a machine-readable workflow model rather than merely a visual-editor export. At a larger scale, this regularity could be used to generate and lint programs from bills of materials or manufacturing execution system data. Such a generator would also have to preserve calibrated workspace geometry, zone bindings, persistent-display behavior, and transition conditions; copying text instructions alone would be insufficient.

3.4. Closed-Loop Control Model of the Designed Program

The original initialization stage selects the external vision configuration. The retained work-instruction XML subsequently references Boolean variables for the bin zones, forward and backward navigation, and a wrong-bin indicator. Project documentation describes the target variables as hand-entry bin zones and the navigation variables as projected controls. The wrong-bin indicator is described as a condition comparing the visited bin with the current target bin. The retained XML verifies the use of these variables but not their definitions or calculation, because the external vision configuration and variable-table export were not supplied for that phase. Claims about the wrong-bin semantics are therefore conditional on that external configuration.

The six original workflow-stage groups form an event-driven finite state machine (Figure 4). Initialization is transient and falls through without an operator-readiness condition. Each picking stage has two distinct forward routes to the same next stage: automatic advancement on a target-zone TRUE-to-FALSE falling edge, and manual advancement on the Next control's FALSE-to-TRUE rising edge. The first route depends on the intended zone semantics and does not independently prove that a part was removed. The second permits deliberate navigation without a target-zone withdrawal event. A Back-control rising edge returns one stage. The final navigation stage has no automatic target-zone route; its Next-control rising edge loops to initialization, while Back returns to the final pick stage. Figure 4 separates the two forward routes and uses dashed backward arrows. The loop-back is encoded by a transition condition rather than the workflow-boundary record.




Figure 4. Original six-stage picking workflow with distinct automatic and manual transitions. Blue solid arrows represent automatic target-zone falling-edge advancement, green solid arrows represent manual Next-control rising-edge advancement, and gray dashed arrows represent manual Back-control rising-edge transitions. Initialization falls through without confirmation; the final manual Next event returns to initialization. Arrow semantics describe XML bindings, not validated part removal.

3.5. XML Conformance and Evidence-Assessment Protocol

The XML audit parsed the namespace-qualified document, verified contiguous order positions, grouped commands by workflow-stage marker, and extracted variable-bound display elements and transition conditions. Four acceptance questions were then evaluated: (i) does each bin-picking step include correct-bin feedback and a falling-edge forward transition; (ii) does a wrong-bin signal expose warning feedback without an automatic forward transition on the wrong-bin indicator; (iii) do the forward and backward controls provide the documented manual transitions; and (iv) does the final state return to initialization? The analysis also checked which initialization-stage visual elements persist because of the persistent-display flag.

The original accompanying project narrative states that all four scenario classes were commissioned successfully. For that commissioning phase, no test protocol, event log, replication count, participant count, or test duration was retained, so those observations could not be independently reproduced or used to estimate detection reliability, false-positive or false-negative rates, latency, operator error, cycle time, or workload. The sensed event is entry into and exit from a zone; it does not establish component identity, quantity, or physical removal.

3.6. Depth Sensing and Communication

The logistics extension used LightGuide 25.2.6.0, Observatory 2.2.2.0 and a separately executed Python companion. The saved camera name is Femto Bolt; AZUREKINECT is the software driver enum rather than a verified alternative hardware model. The configuration specifies a 1024×1024 depth stream at a nominal 10 frames/s. This is a configured rate, not measured throughput. Four existing front-bin hand tools, Back and Next, were retained; four interior Capacity and four Warning tools brought the saved total to 14. These interior tools published Count, the number of non-black pixels in the processed image, and Avg, image brightness on a 0–255 scale. Their native InRange outputs were disabled. The measurements represent depth-filtered image masks, not part count, volume or a percentage of bin capacity.

Capacity1 and Warning1 share the interior region of interest (ROI), defined by the polygon (335,292), (357,292), (357,251), (335,251) in saved image coordinates and the native

depth interval 1084–1354. Their native brightness intervals are 180–255 and 80–255, respectively. These intervals are threshold inputs for the disabled InRange output. Because both tools share the ROI and depth filter, every Bin1 record contained identical Count/Avg pairs; the full and warning states were therefore distinguished by the companion thresholds in rather than by separate measurements. Warning1 explicitly enables Layer Independent; that field is absent from Capacity1's serialization. The front Bin1 hand polygon is (322,317), (360,317), (360,306), (322,306). Numerical depth values are reported as native filter coordinates because the inspected export does not declare a physical-unit conversion. In particular, the two tools do not represent separately calibrated physical height bands. Exact polygons and filter settings for the remaining tools are available on request.

The companion converted the numerical readings into categorical states using the Bin1 settings. Both Count and Avg had to satisfy a trigger pair. Full classification had priority; a new warning classification required both full-clear limits to be satisfied. Separate clear limits supplied hysteresis; intermediate readings could retain an existing candidate state or remain unqualified. A sampled hand-present or missing hand flag inhibited qualification without deleting the raw measurement. Eligibility additionally required a 1 s hand-clear cooldown, two consistent later clear receptions, explicit operator confirmation of the current capture and a maximum receipt age of 10 s. These are checks on received samples, not a continuous hand interlock or proof of camera-frame freshness.

The native capture workflow sent semicolon-delimited lines through a local Transmission Control Protocol (TCP) connection on port 8776: RAWLH;bin;full_count;full_avg;warning_count;warning_avg;hand_present. The local dashboard used port 8775. Numerical measurements and qualification results were retained in the local event record, linking each eligible state to the capture from which it was derived.

3.7. Jev Decision Model and Supervised Action Selection

The implementation used the existing jev-1.13.0 model through TypeSafe's API, without training or fine-tuning it. Jev's Choice interface accepts a state description, a question and developer-defined answer options, and returns one selected option with probabilities and confidence [43]. The companion submitted JavaScript Object Notation (JSON) containing the categorical bin state, fixed bin assignment, cycle and request status, qualification flags, reported remaining returns, preparation lead time and an optional operator note. Raw Count/Avg readings and numerical thresholds remained in the local audit record. No webcam or depth images were submitted, and the part-to-bin assignments remained fixed.

One Choice question requested an advisory action: WAIT, meaning no new preparation request; PREPARE_EMPTY_BIN, meaning preparation of an empty bin for the assigned destination; or MANUAL_REVIEW, meaning operator attention was required. The instructions kept bin assignments fixed, treated the task note as reported information and prohibited overriding a full-bin exchange requirement or granting placement permission. This gave the AI component a limited decision space within the existing workflow. The complete request construction and validation logic are contained in the companion source, which is available on request.

The API deadline was 5 s. TypeSafe defines confidence as a summary derived from the returned probability distribution, distinct from the probability attached to the selected option [44]. The application retained these returned fields and used an unchanged confidence threshold of 0.80. This threshold governed the prototype's action policy; it was not established as a calibrated probability of a correct logistics decision. A valid

preparation proposal below the threshold therefore produced MANUAL_REVIEW in the automatic path.

The operator could approve a below-threshold preparation proposal only while the exact capture, task context, cycle and advice identity remained current. The resulting request recorded human authority and retained the original model outcome. A qualified full state invoked a deterministic exchange requirement independently of Jev. Requests and local acknowledgments were written to an append-only event log. The companion returned RGL1;reply_sequence;authored_instruction;0 with carriage-return/line-feed (CRLF) termination. LightGuide validated the prefix, field count and final zero before projecting locally authored text. Placement permission stayed zero throughout the demonstration. Global LightGuide device and System Text settings were documented in setup records but were not exported as a complete configuration; a documented reconstruction is available on request together with the native workflow.

3.8. Supervised Capture Protocol and Evidence Assessment

Empty, preparation-warning and full references were prepared manually at fixed bin positions. These references were prepared as separate scene states; time between captures was not used to estimate a filling rate or achieved logistics lead time. The extension retained fixed bin assignments and did not recognize part types. The operator set the scene and withdrew their hands; the separate native capture workflow was launched through the LightGuide interface. Each run invoked the saved vision program three times, transmitting four bin records per pass. Configured waits of 3 s and 0.5 s preceded the final display, which held at manual confirmation. A run therefore yielded 12 received bin observations, not 12 independent camera frames or repeated physical trials. Decision-eligible Bin1 captures required the operator to confirm hands clear on the dashboard within the current 10 s receipt window; later raw-monitor development captures did not themselves authorize decisions. Hand-in-bin observations during development informed the guard design, but did not establish coverage of all hand positions or detection reliability.

Four diagnostic Jev attempts, A–D, were performed during adapter development across companion versions 0.3.1–0.3.3. Attempt A used reported context of zero remaining returns and zero preparation lead time; B–D used 20 returns and a 10 min lead time. These values were scenario assumptions supplied by the operator rather than measured workload or logistics forecasts. C and D retained original labels and probability fields; the earlier adapter did not retain those fields for A and B. Only D appears in the exported primary session log. The earlier attempts are supported by contemporaneous local records and dashboard evidence. These attempts are diagnostic examples, not statistical replications or a controlled comparison of context effects.

The final Bin1 protocol covered a warning reference, live advice and explicit proposal approval, locally acknowledged preparation, a full reference with deterministic exchange handling, and a later empty reference with operator confirmation. Preparation readiness was simulated locally, and the operator emptied and reused the same bin rather than installing a separate spare. No notification was delivered to an external logistics service or employee. Event sequences, numerical observations, source attribution and completion state were reconciled against the retained log and workstation records. Two recorded adapter/API intervals are reported individually; they exclude capture qualification, operator response and projection. Detection error rates, operator performance, continuous fullness estimation and production benefit were outside this proof-of-concept protocol.

The archived verification record for companion release 0.3.3 reports 41 passing unit tests covering parsing, qualification, hysteresis, expiry, request identity and projection

mapping. These tests exercise software logic with constructed inputs and are separate from the physical reference captures. They do not establish camera detection performance.

4. Results

4.1. XML Inventory and Step Structure

The retained XML contains six numbered stages with repeated canvas-reset, target-assignment, display, and transition records. Table 2 summarizes the command groups. The artifact implements a fixed four-bin demonstration sequence in a multi-variant manufacturing context; it contains no product identifier, part number, quantity, bill-of-material input, or variant-dependent branch. It therefore must not be interpreted as a complete or variant-aware production kit definition.

Table 2. Generalized structure of the retained XML program and its six stages.

Stage	Abstracted XML Evidence	Control and Persistence	Transition
1	initializes the canvas, bin labels, navigationforward/back controls and graphics, and the external vision selection	wrong-bin warning are persistent	no transition condition; automatic fall-through to Stage 2
2	sets target bin 1; displays an instruction, target graphic, and correct-bin feedback	target-bin state; target outline is persistent	target falling/forward-control rising: Forward; back-control rising: Backward
3	same pattern for target bin 2	target-bin state; target outline is persistent	same three alternatives
4	same pattern for target bin 3	target-bin state; target outline is persistent	same three alternatives
5	same pattern for target bin 4	target-bin state; target outline is persistent	same three alternatives
6	clears nonpersistent graphics and waits for final navigation	forward and back controls	forward-control rising: loop back to initialization; back-control rising: Backward

The initialization stage selects the external vision configuration. It also creates the bin labels and navigation graphics and creates a single error icon and wrong-bin message bound to the wrong-bin indicator. The error pair is not recreated in each target stage; both elements carry the persistent-display flag, so they survive later canvas resets. The initialization stage contains no transition condition. Its commands therefore run into Stage 2 without a recorded readiness event.

Target stages 2–5 share the same structure. Each assigns a successive target bin, with recalculation enabled, projects a short instruction and target graphics, and binds a confirmation symbol plus correct-bin status text to the corresponding target-bin variable. Each target outline is also persistent, whereas the confirmation symbol, status text, and rectangular target graphic are not. Target outlines therefore accumulate across the sequence. Whether that persistence is an intended progress cue or a configuration error cannot be determined from the retained records and should be confirmed at the workstation. Each target stage then exposes three alternative transitions:

- A falling edge of the corresponding target-bin variable, which advances automatically if the external variable represents target-bin exit as documented;
- A rising edge of the forward control, which advances manually;
- A rising edge of the backward control, which returns one stage.

In Stage 6, a rising edge of the forward control invokes the loop-back to initialization, whereas a rising edge of the backward control returns one stage. The workflow-boundary record is not the looping mechanism.

4.2. Static Conformance and Reported Runtime Behavior

The original intended operator sequence is to follow the projected target, reach into the indicated bin, and observe conditional feedback. The XML shows which display elements and transitions are bound to each variable; it does not prove that a camera detected a hand, that the variable logic was configured as documented, or that a component was picked. In a typical LightGuide bin-pick configuration, the wrong-bin indicator is a user-defined Boolean derived from the active vision-zone set and the current target bin; navigation zones must be excluded from that set to avoid false warnings. Neither the variable-table export nor the vision configuration needed to verify this configuration was supplied for the original picking installation. Table 3 reports direct XML evidence separately from the commissioning account.

Selected demonstration frames (Figure 5) show the “Pick from 1” instruction, wrong-bin feedback, correct-bin feedback, and the subsequent bin instruction. The frames document visible interface states but do not establish the underlying variable calculation, exact transition timing, or detection reliability.




Figure 5. Illustrative frames supplied from a project demonstration video, shown as representative instruction and feedback states rather than a synchronized event record or performance test. (a) Initial target cue. (b) Wrong-bin warning. (c) Correct-bin feedback. (d) Next-pick instruction.

Table 3. Scenario acceptance criteria, direct XML evidence, and evidence boundary.

Scenario	Acceptance Criterion	Direct Evidence in Retained XML	Boundary
Target bin	correct-bin feedback while the target variable is true; advance on its falling edge	matching feedback bindings and Falling-Forward transition in every bin-picking stage	variable semantics and runtime detection not retained
Wrong bin	show warning; wrong-bin signal must not itself advance	one persistent wrong-bin warning bound to the wrong-bin indicator; no transition uses the indicator	formula, active-zone exclusions, and simultaneous activations not retained
Manual navigation	forward control advances and backward control returns	rising transitions with Forward/Backward in target stages; Backward in Stage 6	runtime zone or control configuration not retained
Cyclic operation	final forward control returns to initialization	Stage-6 loop-back transition with initialization as target	runtime event log not retained

All four transition patterns are structurally present in the XML. The project narrative reports that the corresponding commissioning scenarios behaved as expected, but includes no run-level observations. The wrong-bin indicator is not an inhibit condition: its activation does not itself trigger automatic advancement, but manual forward control remains available and a simultaneous target-bin event could still satisfy the target transition. The retained evidence therefore supports static workflow conformance, not a no-fault-forward guarantee.

4.3. Recorded Logistics Configuration

The retained logistics exports show a separate three-step capture program using the revised 14-tool depth configuration. The companion configuration enabled Bin1 only; Bins2–4 had no active calibrated decision thresholds. Table 4 reports the deployed Bin1 numerical settings, and Figure 6 separates native sensing, companion checks, Jev advice and human decision authority. Thresholds were selected from supervised reference arrangements; they are installation-specific settings rather than universal bin-full criteria.

Table 4. Deployed Bin1 sensing and qualification settings. Count and Avg thresholds belong to the companion and differ from native brightness-filter intervals.

Parameter	Trigger or Configured Value	Clear Limit or Interpretation
Full state	Count ≥ 700 and Avg ≥ 180	Count ≤ 650 and Avg ≤ 175
Preparation warning	Count ≥ 200 and Avg ≥ 80	Count ≤ 50 and Avg ≤ 20
Native interior tools	Depth 1084–1354; identical saved ROI	Capacity brightness 180–255; Warning 80–255
Sample qualification	Hand clear; 1 s cooldown; two consistent clear receptions	Explicit operator confirmation; receipt age ≤ 10 s
Jev review	Model jev-1.13.0; deadline 5 s	Confidence threshold 0.80
Projection reply	Fixed authored instruction in RGL1 message	Placement permit 0




Figure 6. Added logistics architecture. Native depth tools and capture XML provide sampled measurements; the companion applies numerical and qualification checks. Jev supplies a bounded

proposal, the operator confirms the scene and reviews eligible advice, and a deterministic guard handles full bins. Local request status is returned as authored projection text. The architecture was demonstrated for Bin1 with simulated local acknowledgments and no placement permission.

4.4. Diagnostic Jev Attempts

Table 5 retains all four attempts and distinguishes missing original-choice fields from an effective review decision. Every reported confidence value was below 0.80; no attempt autonomously created a preparation request. C and D selected PREPARE_EMPTY_BIN with confidence 0.75 and label probabilities PREPARE 0.84, WAIT 0.06 and MANUAL_REVIEW 0.10. In C, no preparation request was applied. In D, explicit operator approval at a logged receipt age of 9.383 s created a preparation request linked to the original Jev audit. The original choice and automatic MANUAL_REVIEW outcome remained in the record. The individual recorded adapter/API intervals were 354.828 ms for C and 351.944 ms for D; A and B had no retained latency field. As a post hoc reference, a fixed rule that requests preparation when a warning state coincides with positive reported remaining returns would select preparation for B–D and no preparation for A. The retained choices in C and D agree with this rule; the original choices for A and B were not retained. The recorded attempts therefore show no decision value beyond such a rule. These observations cannot estimate production latency, accuracy or a causal effect of task context.

Table 5. Four supervised diagnostic Jev attempts on 2 October 2026. Context values are reported remaining returns and preparation lead time in minutes.

Attempt and Version	Context	Count/Avg ¹	Original Choice	Confidence	Applied Local Outcome
A 0.3.1	0/0	419/118	Not retained	0.47	MANUAL_REVIEW; no request
B 0.3.1	20/10	427/120	Not retained	0.76	MANUAL_REVIEW; no request
C 0.3.2	20/10	413/116	PREPARE_EMPTY_BIN	0.75	MANUAL_REVIEW; no request
D 0.3.3	20/10	445/125	PREPARE_EMPTY_BIN	0.75	MANUAL_REVIEW; human approval created preparation request

¹ Both masks returned the same Count/Avg pair; hand flags were clear. Confidence differs from label probability. A–D span three software versions with incomplete earlier fields and are not equivalent repeated trials.

4.5. Completed Supervised Bin1 Lifecycle

Table 6 reconstructs the completed lifecycle from the exported primary log. After D's human-approved preparation request, a full-reference receipt supplied Count/Avg 850/240 for both tools. Operator hands-clear confirmation then generated an exchange request from the deterministic full guard. Received and preparation-ready acknowledgments were local simulations. The operator reported emptying and reusing the same bin. A later receipt supplied 0/0 for both masks; the operator confirmed emptying 5.543 s after that receipt. Both cycle-1 requests completed and Bin1 advanced to cycle 2. The record explicitly retained empty_verified_by_camera=false; a below-warning reading and human confirmation supported completion, without independent camera verification of physical emptiness. Projected status showed pending preparation and exchange requests after capture expiry and returned to an expired raw-monitor message after completion. The dashboard and event record showed cycle 2. Placement permission remained zero.

Table 6. Recorded Bin1 lifecycle following attempt D. Times are Coordinated Universal Time (UTC) on 2 October 2026, rounded to milliseconds. Receipt time is distinguished from later event time.

UTC Time	Recorded Event	Observation or Resulting State
11:25:23.327	Warning-reference receipt, sequence 21	Both masks 445/125; sampled hand clear
11:25:28.435	Jev response event	PREPARE proposal; confidence 0.75; automatic MANUAL_REVIEW
11:25:32.711	Operator approval logged	Local preparation request created; receipt age 9.383 s
11:48:57.303	Preparation acknowledged received	Simulated local acknowledgment
11:52:38.408	Full-reference receipt, sequence 33	Both masks 850/240; sampled hand clear
11:52:40.972	Confirmed full capture generated exchange	Deterministic full-bin guard
11:53:56.219	Exchange acknowledged received	Simulated local acknowledgment
11:55:14.669	Preparation marked ready	Simulated readiness for same-bin emptying
11:56:40.859	Empty-reference receipt, sequence 45	Both masks 0/0; sampled hand clear
11:56:46.401	Operator confirmed emptying	Both requests completed; Bin1 cycle 2

The exported session snapshot contains 162 events, including 108 RAWLH bin observations from nine three-pass capture bursts and one live Jev request/response. The completed Bin1 sequence had reached 48 observations; later captures concerned Bin2 development and a final all-empty recheck. One Bin2 arrangement was questioned by the operator after capture and is retained as uncertain development evidence, excluded from physical-state validation claims. It is not classified as a confirmed software false-full error. Bins2–4 remained disabled for decisions. Accordingly, the completed request chain is evidence of integration feasibility for one supervised Bin1 lifecycle, not a four-bin capacity qualification or an accuracy study.

5. Discussion

5.1. Static Workflow Conformance Versus Conventional Picking Assistance

The retained XML establishes a configuration-level loop among projected instructions, vision-tool variables, conditional feedback, and workflow transitions. It does not establish that this loop operated reliably during commissioning. Conventional pick-to-light systems primarily indicate a target location, although the extent of confirmation depends on the particular configuration [8,9]. The audited workflow adds a conditional warning path, but the evidence for that path must be stated precisely.

In the original picking architecture, the target-bin variable and wrong-bin indicator belong to the user-defined variable and control layer. The vision subsystem is intended to supply zone states and a set of active zones; a separate Boolean expression is intended to compare that set with the current target bin and derive the wrong-bin indicator. The retained workflow references the indicator to control a warning graphic and text, but it neither defines the expression nor contains the external vision configuration. The XML thus confirms the warning binding, not the semantics or runtime reliability of the underlying zone detector.

An isolated wrong-bin-indicator activation is not itself a forward transition, so it does not by itself advance the workflow. It is not, however, an explicit interlock: the indicator is absent from the transition guards, manual forward control remains available, and a concurrent target-zone event could still satisfy the target transition. The retained files therefore support a claim of configured wrong-bin feedback, not a claim that incorrect picks are blocked or that no-fault-forward operation was demonstrated. Related projection-and-vision systems have used action-linked feedback for quality support [16–18], while vision-based assistance delivered through other media pursues a similar quality

objective [19]. That literature does not substitute for commissioning evidence for this installation.

The state interpretation also requires two qualifications. The initialization stage contains no transition condition; it starts the vision configuration and creates persistent display elements before control falls through to the first target stage. At the other end of the sequence, cyclic operation is encoded by the final-stage loop-back transition. These structures are internally consistent in the XML, but no original commissioning event log was retained to show their runtime execution.

5.2. Operational Implications and Evidence Status

Table 7 separates structures visible in the retained XML from mechanisms inferred from those structures and outcomes that remain unmeasured. The original picking source package contains neither a controlled comparison of operator performance nor a trial-level commissioning record.

Table 7. Operational implications of the original picking workflow and its evidence boundary.

Area	Encoded Mechanism or Expected Benefit	Evidence Status
Quality	a persistent wrong-bin warning is bound to the user-layer wrong-bin indicator	binding confirmed statically; the indicator expression and vision configuration are missing, no explicit inhibit is encoded, and error reduction was not measured
Workflow control	target-bin Falling transitions and manual forward/back alternatives	command structure confirmed in the XML; no runtime log, reliability estimate, component-identity check, or removal check
Human factors	reduced need to consult or memorize a sequence and potentially shorter training	mechanism supported by literature; workload, usability, and training time were not measured in this case study
Traceability	machine-readable steps and confirmations that could feed a log	architectural capability only; no operational log or kit-level linkage was analyzed
Flexibility	software changes to stages, coordinates, variables, and vision zones	regular XML pattern confirmed for four fixed bins; no reconfiguration trial and no evidence for variant-dependent sequencing
Safety	projected information without a worn display or moving actuator	no ergonomic measurement or installation-specific risk assessment; evaluation under ISO 12100 remains necessary

5.2.1. Quality and Error Prevention

In the manual baseline, an incorrect selection may remain undetected until assembly or final inspection. The original XML binds warning graphics to the user-layer wrong-bin indicator; it does not establish a working warning detector. That design can move feedback closer to the interaction with a storage bin, but the supplied artifacts do not verify the expression, vision zones, active-zone exclusions, or runtime timing needed to establish this behavior. Whether the mechanism reduces incorrect kits also depends on component identity and quantity checks, detector robustness, operator response, and workstation organization. Human and cognitive factors are important in manual picking [4,5]; they were not measured here.

5.2.2. Cognitive Workload and Training

Projecting the target bin and sequence into the workspace may reduce the need to memorize steps or repeatedly consult catalog numbers. This is a mechanistic expectation, not a result of the present case study. Experiments with projected instructions report context-dependent effects on performance, workload, and learning [12–14], and recent

work emphasizes matching the level of support to worker competence [3]. Neither training time nor subjective workload was measured, so no reduction can be attributed to this implementation.

5.2.3. Traceability and Process Data

The encoded steps, variables, and confirmations provide a machine-readable basis from which step times, sequences, and exception states could be logged. Such records could support retrospective traceability, analysis of repeated errors, and parameterization of a digital-twin model [37]. No original picking event log was exported. The added companion records capture, advice, approval and request-lifecycle events, but does not link them to a production kit. Record completeness across the industrial process and kit-level identity remain future integration requirements.

5.2.4. Flexibility and Reconfiguration

Projection coordinates and vision zones are software-defined, unlike a pick-to-light installation that commonly requires a physical light module at each bin. This architectural difference makes software reconfiguration plausible, but the audited artifact is only a fixed four-bin sequence. Adding or moving a bin would require coordinated changes to the workflow XML, the missing original vision configuration, the user-variable logic, and any dependent display assets. The XML can be versioned and reviewed and could support future automated generation (Section 5.4), but reconfiguration time, cost, and correctness were not compared with an alternative system.

5.2.5. Ergonomics and Safety

Projection does not require a worn display and avoids the mass and contact burden of a head-mounted device [12]. The informational function contains no moving actuator, but these properties do not establish ergonomic benefit or safety. Under ISO 12100 [45], the final installation requires identification and assessment of all relevant hazards across its lifecycle, including mechanical and handling hazards, electrical, optical, thermal and ergonomic hazards, and chemical or biological exposure where applicable to the parts and workplace. Equipment without actuation can still create risks through handling, layout and operator actions. No installation-specific risk assessment, protective-function validation or conformity assessment was performed. The sampled hand guard, Jev review logic and projected messages are informational controls and must not be credited as safety functions.

5.3. Integration Constraints on the Projection Platform

The original workflow makes several integration constraints visible. Automatic zone transitions and manual navigation coexist, so an external advisory result cannot be assumed to govern every route through a program. Persistent graphics also outlive individual canvas resets and may remain visible after a state change. The logistics extension therefore uses a separate capture workflow and an explicit mapping from request status to authored projection text. This preserves the existing picking sequence while making the advisory path inspectable.

Sensing geometry, displayed geometry and variable definitions must be maintained together when a workstation changes. The saved regions and thresholds in the extension make those dependencies more explicit, but their coordinates remain specific to the demonstrated installation. Similarly, software reuse does not establish that feedback is legible or appropriate in another layout. A complete deployment needs its own alignment and usability assessment. Integration and instruction design remain practical concerns reported for assistive systems in real production environments [23].

5.4. Scalability and Integration

The declarative XML follows a regular pattern, but the retained pilot is a fixed sequence of four bins rather than a variant-aware kitting system. At the template level, another bin-picking stage would require a new target-bin value, projection coordinates, a monitored tool variable, and transition bindings. At the system level, it would also require a corresponding zone in the external vision configuration, revision of the user-variable logic and active-bin exclusions, verification of persistent graphics, and a new commissioning test. The source workflow also depends on local asset paths, which is an additional portability concern but is not reproduced in the public manuscript. Changing bin order is therefore possible by reauthoring the target sequence and associated bindings, followed by recommissioning; a different number of bins requires corresponding sensing and projection changes. Part quantities can be written into instructions, but the present hand-zone confirmation does not count parts. Neither dynamic reordering nor quantity verification was demonstrated.

The repeated pattern could support future generation from a digital bill of materials (BOM) or a manufacturing execution system (MES), provided that the generator manages geometry, external assets, variable definitions, and transition priorities as well as item sequence. Persisted state events could also inform parameterization and validation of a digital-twin model [37] and human-centered Operator 4.0 assistance [2]. No BOM/MES generator was implemented. Persistent event export was added for the logistics companion, while production picking-event export remains unimplemented, and behavior with tens of positions or variant-dependent branches was not examined.

5.5. Limitations and Future Work

The observed result is a supervised integration demonstration. Its scope and the priorities for further evaluation are defined by the following limitations:

1. Generalizability: The evidence comprises one original industrial picking artifact and one completed Bin1 logistics lifecycle at a research workbench. Transfer to other tasks, layouts, parts or organizations has not been assessed.
2. System reproducibility: Original picking calibration, external vision definitions, the wrong-bin expression and commissioning records remain unavailable. The extension retains native capture and vision exports, regions, thresholds, source and an event log, which are available on request; its global LightGuide settings exist only as a documented reconstruction rather than a complete native export.
3. Commissioning evidence: The original commissioning narrative has no trial-level records. The new log supports the reported request sequence, but four changing-version Jev attempts and repeated receptions within capture bursts are insufficient for statistical reliability estimates.
4. Quantitative effects: Detection reliability, operator error, cycle time, workload and logistics response improvement were not measured. A subsequent comparative study could assess these outcomes, including subjective workload with the National Aeronautics and Space Administration Task Load Index (NASA-TLX) [46], after the sensing and event records have been qualified.
5. Scope of the sensed event: Zone events do not verify part identity, quantity or removal. The interior tools classify depth-filtered reference scenes; they do not provide calibrated fill volume or independent verification of an empty bin.
6. Detection robustness: Hand coverage, occlusion effects and sensitivity to changing part arrangements require further assessment. Receipt age does not establish camera-frame age, and the two recorded API intervals do not measure the complete sensing-to-projection delay.

7. **Adaptivity:** Reported workload and preparation time were scenario inputs. Changes in software version and incomplete earlier response fields prevent a controlled assessment of context effects, and no application-specific calibration of model probabilities or confidence was performed.
8. **Human oversight:** The 10 s receipt-age guard bounded data freshness but also limited the time available for reviewing advice; in D, approval was logged 9.383 s after the capture was received. Short review windows may favor acceptance of automated advice without independent verification [41]. Future protocols should separate freshness checks from deliberation time, for example by requiring a fresh capture after review, and should record how often operators reject or modify proposals.
9. **Task scope:** Capacity decisions were enabled only for Bin1. Local readiness acknowledgments and same-bin emptying do not demonstrate physical spare-bin availability or external logistics service execution. Product-dependent routing, quantity verification and workflow generation were outside the demonstration.

5.6. Role of Bounded AI Advice

The engineering result is a complete, inspectable path from a qualified bin state to a model proposal, a recorded action decision and projected request status. Jev contributes the proposal using the reported task context; numerical classification and full-bin handling remain local functions. In the recorded applied case, the confidence rule directed the proposal to operator review, and the operator accepted it. This demonstrates the review mechanism as well as the API connection. It does not establish autonomous decision quality. The potential value of contextual advice is to help decide when preparation is useful before full-state handling becomes necessary. In the recorded attempts, the retained proposals matched a simple post hoc reference rule (Section 4.4). A possible advantage of model advice is the joint interpretation of categorical state, reported context and free-text operator notes in combinations that a fixed rule does not anticipate. Assessing that advantage requires predefined decision criteria, repeated observations and a comparison with a rule-based preparation policy.

6. Conclusions

This study implemented human-supervised AI decision support with Jev on a commercial projection-guidance platform and examined the existing picking workflow on which the integration was based. The main conclusions are as follows:

1. The companion connected locally derived bin states and reported task context to a typed Jev proposal, explicit action checks, a logged request lifecycle and authored projection text. This establishes a documented and inspectable software integration pattern within the stated installation limits.
2. All four diagnostic attempts reported confidence below the unchanged 0.80 threshold. In the final attempt, the operator approved a preparation proposal with confidence 0.75. The recorded action therefore demonstrates supervised use of AI advice.
3. A later full-reference capture generated a deterministic exchange requirement. Local acknowledgments and operator-confirmed emptying of the same bin completed one Bin1 lifecycle. Preparation readiness was simulated, and the projected output remained informational.
4. The original picking audit identifies the platform mechanisms needed for this integration: sensor-variable bindings, automatic and manual transitions, persistent feedback and explicit projection commands. It establishes encoded control structure;

the missing original vision definitions and commissioning logs limit conclusions about that installation's runtime behavior.

5. The demonstrated contribution is technical feasibility and decision traceability. The study does not establish part identity or quantity verification, detector reliability, autonomous logistics decision quality or production benefit.

Further work should complete the installation configuration, assess sensing robustness across controlled arrangements and hand trajectories, and repeat the supervised decision protocol with specified decision criteria. A subsequent logistics study should include actual spare-bin preparation and exchange, recorded response times and an appropriate comparison policy. These steps would allow the value of contextual AI advice to be assessed beyond the present integration demonstration.

Author Contributions: Conceptualization, A.V. and M.K.; methodology, A.V.; software, A.V.; formal analysis, A.V. and M.T.; investigation, A.V., P.G. and J.K.; validation, J.K.; writing—original draft preparation, M.T.; writing—review and editing, M.T.; supervision, M.K.; project administration, M.K. and J.K.; funding acquisition, M.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: De-identified logistics capture and depth-vision exports, exact numerical regions, the deployed Bin1 settings, companion version 0.3.3 source and tests, a primary session log, reconciled Jev-attempt and lifecycle records, and setup reconstruction instructions are available from the corresponding author on reasonable request. The primary log contains the final live request D; earlier attempts A–C rely on contemporaneous local records and dashboard evidence rather than exported primary logs. Original source hashes and de-identification steps are documented. The full global LightGuide settings were not exported. The original industrial picking export and company production documents remain non-public; its external wrong-bin expression and commissioning records were unavailable. The original illustrative video frames are not synchronized performance evidence.

Acknowledgments: This work was supported by the Ministry of Education of the Slovak Republic KEGA grant no. 009TUKÉ-4/2026 and VEGA grant no. 1/0344/26; the Slovak Research and Development Agency contract no. APVV-24-0470; and the European Union NextGenerationEU through the Recovery and Resilience Plan for Slovakia project no. 09I05-03-V02-00042. During the preparation of this work, the authors used OpenAI ChatGPT 6 Sol Ultra and ChatGPT Desktop 26.812 tools for software prototyping, analysis and reconciliation of retained artifacts and event records, figure preparation, manuscript drafting and language editing. Jev is the advisory component under study and is described in Section 3. The authors reviewed and edited the assisted outputs and take responsibility for the software, analysis and content of the publication.

Conflicts of Interest: The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this article. The authors have no affiliation with, and received no funding from, TypeSafe AI or LightGuide, Inc. Access to the Jev API was paid for under the provider's standard terms.

Abbreviations

The following abbreviations are used in this manuscript:

BOM	Bill of materials
Gest-SAR	Gesture-controlled spatial augmented reality
ISO	International Organization for Standardization
MES	Manufacturing execution system
NASA-TLX	National Aeronautics and Space Administration Task Load Index
XML	Extensible Markup Language
AI	Artificial intelligence
API	Application programming interface
JSON	JavaScript Object Notation
ROI	Region of interest
TCP	Transmission Control Protocol
UTC	Coordinated Universal Time

References

1. Mark, B.G.; Rauch, E.; Matt, D.T. Worker Assistance Systems in Manufacturing: A Review of the State of the Art and Future Directions. *J. Manuf. Syst.* **2021**, *59*, 228–250. <https://doi.org/10.1016/j.jmsy.2021.02.017>.
2. Romero, D.; Stahre, J.; Wuest, T.; Noran, O.; Bernus, P.; Fast-Berglund, Å.; Gorecky, D. Towards an Operator 4.0 Typology: A Human-Centric Perspective on the Fourth Industrial Revolution Technologies. In Proceedings of the 46th International Conference on Computers and Industrial Engineering (CIE46), Tianjin, China, 29–31 October 2016. Available online: <https://research.chalmers.se/en/publication/241376> (accessed on 3 October 2026).
3. Fink, K.; Ziegler, A.; Härdtlein, C.; Berger, C.; Daub, R. Assessing the Human Competence and Individual Support Level in Manual Assembly through Cognitive Assistance Systems. *Procedia CIRP* **2024**, *126*, 811–816. <https://doi.org/10.1016/j.procir.2024.08.263>.
4. Grosse, E.H.; Glock, C.H.; Jaber, M.Y.; Neumann, W.P. Incorporating Human Factors in Order Picking Planning Models: Framework and Research Opportunities. *Int. J. Prod. Res.* **2015**, *53*, 695–717. <https://doi.org/10.1080/00207543.2014.919424>.
5. Kolus, A.; Wells, R.; Neumann, P. Production Quality and Human Factors Engineering: A Systematic Review and Theoretical Framework. *Appl. Ergon.* **2018**, *73*, 55–89. <https://doi.org/10.1016/j.apergo.2018.05.010>.
6. Stork, S.; Schubö, A. Human Cognition in Manual Assembly: Theories and Applications. *Adv. Eng. Inform.* **2010**, *24*, 320–328. <https://doi.org/10.1016/j.aei.2010.05.010>.
7. Brolin, A.; Thorvald, P.; Case, K. Experimental Study of Cognitive Aspects Affecting Human Performance in Manual Assembly. *Prod. Manuf. Res.* **2017**, *5*, 141–163. <https://doi.org/10.1080/21693277.2017.1374893>.
8. Battini, D.; Calzavara, M.; Persona, A.; Sgarbossa, F. A Comparative Analysis of Different Paperless Picking Systems. *Ind. Manag. Data Syst.* **2015**, *115*, 483–503. <https://doi.org/10.1108/IMDS-10-2014-0314>.
9. Guo, A.; Raghu, S.; Xie, X.; Ismail, S.; Luo, X.; Simoneau, J.; Gilliland, S.; Baumann, H.; Southern, C.; Starner, T. A Comparison of Order Picking Assisted by Head-up Display (HUD), Cart-Mounted Display (CMD), Light, and Paper Pick List. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers (ISWC '14)*; ACM: Seattle, WA, USA, 2014; pp. 71–78. <https://doi.org/10.1145/2634317.2634321>.
10. Hanson, R.; Falkenström, W.; Miettinen, M. Augmented Reality as a Means of Conveying Picking Information in Kit Preparation for Mixed-Model Assembly. *Comput. Ind. Eng.* **2017**, *113*, 570–575. <https://doi.org/10.1016/j.cie.2017.09.048>.
11. Wang, X.; Ong, S.K.; Nee, A.Y.C. A Comprehensive Survey of Augmented Reality Assembly Research. *Adv. Manuf.* **2016**, *4*, 1–22. <https://doi.org/10.1007/s40436-015-0131-4>.
12. Funk, M.; Kosch, T.; Schmidt, A. Interactive Worker Assistance: Comparing the Effects of in-Situ Projection, Head-Mounted Displays, Tablet, and Paper Instructions. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*; ACM: Seattle, WA, USA, 2016; pp. 934–939. <https://doi.org/10.1145/2971648.2971706>.
13. Bosch, T.; Könnemann, R.; de Cock, H.; van Rhijn, G. The Effects of Projected versus Display Instructions on Productivity, Quality and Workload in a Simulated Assembly Task. In *Proceedings of the 10th International Conference on Pervasive Technologies Related to Assistive Environments (PETRA '17)*; ACM: Seattle, WA, USA, 2017; pp. 412–415. <https://doi.org/10.1145/3056540.3076189>.

14. Mark, B.G.; Rauch, E.; Matt, D.T. Study of the Impact of Projection-Based Assistance Systems for Improving the Learning Curve in Assembly Processes. *Procedia CIRP* **2020**, *88*, 98–103. <https://doi.org/10.1016/j.procir.2020.05.018>.
15. Uva, A.E.; Gattullo, M.; Manghisi, V.M.; Spagnulo, D.; Cascella, G.L.; Fiorentino, M. Evaluating the Effectiveness of Spatial Augmented Reality in Smart Manufacturing: A Solution for Manual Working Stations. *Int. J. Adv. Manuf. Technol.* **2018**, *94*, 509–521. <https://doi.org/10.1007/s00170-017-0846-4>.
16. Chen, C.; Tian, Z.; Li, D.; Pang, L.; Wang, T.; Hong, J. Projection-Based Augmented Reality System for Assembly Guidance and Monitoring. *Assem. Autom.* **2021**, *41*, 10–23. <https://doi.org/10.1108/AA-02-2019-0028>.
17. Ojer, M.; Alvarez, H.; Serrano, I.; Saiz, F.A.; Barandiaran, I.; Aguinaga, D.; Querejeta, L.; Alejandro, D. Projection-Based Augmented Reality Assistance for Manual Electronic Component Assembly Processes. *Appl. Sci.* **2020**, *10*, 796. <https://doi.org/10.3390/app10030796>.
18. Hasan, N.; Alkan, B. Gest-SAR: A Gesture-Controlled Spatial AR System for Interactive Manual Assembly Guidance with Real-Time Operational Feedback. *Machines* **2025**, *13*, 658. <https://doi.org/10.3390/machines13080658>.
19. Riedel, A.; Gerlach, J.; Dietsch, M.; Herbst, S.; Engelmann, F.; Brehm, N.; Pfeifroth, T. A Deep Learning-Based Worker Assistance System for Error Prevention: Case Study in a Real-World Manual Assembly. *Adv. Prod. Eng. Manag.* **2021**, *16*, 393–404. <https://doi.org/10.14743/apem2021.4.408>.
20. Markus, A.; Daling, L.M.; Borowski, E.; Isenhardt, I. Error Communication in Manual Assembly through a Projection-Based Assistance System. In *Human Factors and Systems Interaction*; AHFE Open Access: King of Prussia, PA, USA, 2024; Volume 154, pp. 179–188. <https://doi.org/10.54941/ahfe1005365>.
21. Egger, J.; Masood, T. Augmented Reality in Support of Intelligent Manufacturing—A Systematic Literature Review. *Comput. Ind. Eng.* **2020**, *140*, 106195. <https://doi.org/10.1016/j.cie.2019.106195>.
22. de Souza Cardoso, L.F.; Mariano, F.C.M.Q.; Zorzal, E.R. A Survey of Industrial Augmented Reality. *Comput. Ind. Eng.* **2020**, *139*, 106159. <https://doi.org/10.1016/j.cie.2019.106159>.
23. Pimminger, S.; Kurschl, W. Prototyping Assistive Systems for Manual Assembly in Real Production Environments: Challenges and Lessons Learned. In *Proceedings of the 2021 IEEE 2nd International Conference on Human-Machine Systems (ICHMS)*; IEEE: New York, NJ, USA, 2021; pp. 1–6. <https://doi.org/10.1109/ICHMS53169.2021.9582461>.
24. Almeida, D. Introducing System One Models & Jev. TypeSafe AI, 15 September 2026. Available online: <https://typesafe.ai/blog/introducing-system-one-models-and-jev> (accessed on 3 October 2026).
25. TypeSafe AI. Introduction. TypeSafe Documentation. Available online: <https://docs.typesafe.ai/introduction> (accessed on 3 October 2026).
26. Lušić, M.; Fischer, C.; Bönig, J.; Hornfeck, R.; Franke, J. Worker Information Systems: State of the Art and Guideline for Selection under Consideration of Company Specific Boundary Conditions. *Procedia CIRP* **2016**, *41*, 1113–1118. <https://doi.org/10.1016/j.procir.2015.12.003>.
27. Re, G.M.; Oliver, J.; Bordegoni, M. Impact of Monitor-Based Augmented Reality for on-Site Industrial Manual Operations. *Cogn. Technol. Work* **2016**, *18*, 379–392. <https://doi.org/10.1007/s10111-016-0365-3>.
28. Gattullo, M.; Scurati, G.W.; Fiorentino, M.; Uva, A.E.; Ferrise, F.; Bordegoni, M. Towards Augmented Reality Manuals for Industry 4.0: A Methodology. *Robot. Comput.-Integr. Manuf.* **2019**, *56*, 276–286. <https://doi.org/10.1016/j.rcim.2018.10.001>.
29. Funk, M.; Kosch, T.; Greenwald, S.W.; Schmidt, A. A Benchmark for Interactive Augmented Reality Instructions for Assembly Tasks. In *Proceedings of the 14th International Conference on Mobile and Ubiquitous Multimedia (MUM '15)*; ACM: Seattle, WA, USA, 2015; pp. 253–257. <https://doi.org/10.1145/2836041.2836067>.
30. Oestreich, H.; Töniges, T.; Wojtynek, M.; Wrede, S. Interactive Learning of Assembly Processes Using Digital Assistance. *Procedia Manuf.* **2019**, *31*, 14–19. <https://doi.org/10.1016/j.promfg.2019.03.003>.
31. Sand, O.; Büttner, S.; Paelke, V.; Röcker, C. smARt.Assembly—Projection-Based Augmented Reality for Supporting Assembly Workers. In *Virtual, Augmented and Mixed Reality*; Lackey, S., Shumaker, R., Eds.; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2016; Volume 9740, pp. 643–652. https://doi.org/10.1007/978-3-319-39907-2_61.
32. Doshi, A.; Smith, R.T.; Thomas, B.H.; Bouras, C. Use of Projector Based Augmented Reality to Improve Manual Spot-Welding Precision and Accuracy for Automotive Manufacturing. *Int. J. Adv. Manuf. Technol.* **2017**, *89*, 1279–1293. <https://doi.org/10.1007/s00170-016-9164-5>.
33. Caudell, T.P.; Mizell, D.W. Augmented Reality: An Application of Heads-up Display Technology to Manual Manufacturing Processes. In *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*; IEEE: New York, NJ, USA, 1992; Volume 2, pp. 659–669. <https://doi.org/10.1109/HICSS.1992.183317>.

34. Fite-Georgel, P. Is There a Reality in Industrial Augmented Reality? In *Proceedings of the 10th IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*; IEEE: New York, NJ, USA, 2011; pp. 201–210. <https://doi.org/10.1109/ISMAR.2011.6092387>.
35. Blattgerste, J.; Renner, P.; Pfeiffer, T. Authorable Augmented Reality Instructions for Assistance and Training in Work Environments. In *Proceedings of the 18th International Conference on Mobile and Ubiquitous Multimedia (MUM '19)*; ACM: Seattle, WA, USA, 2019; pp. 1–11. <https://doi.org/10.1145/3365610.3365646>.
36. Chidambaram, S.; Huang, H.; He, F.; Qian, X.; Villanueva, A.M.; Redick, T.S.; Stuerzlinger, W.; Ramani, K. ProcessAR: An Augmented Reality-Based Tool to Create in-Situ Procedural 2D/3D AR Instructions. In *Proceedings of the Designing Interactive Systems Conference 2021 (DIS '21)*; ACM: Seattle, WA, USA, 2021; pp. 234–249. <https://doi.org/10.1145/3461778.3462126>.
37. Tao, F.; Zhang, H.; Liu, A.; Nee, A.Y.C. Digital Twin in Industry: State-of-the-Art. *IEEE Trans. Ind. Inform.* **2019**, *15*, 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>.
38. Wuest, T.; Weimer, D.; Irgens, C.; Thoben, K.-D. Machine Learning in Manufacturing: Advantages, Challenges, and Applications. *Prod. Manuf. Res.* **2016**, *4*, 23–45. <https://doi.org/10.1080/21693277.2016.1192517>.
39. Woschank, M.; Rauch, E.; Zsifkovits, H. A Review of Further Directions for Artificial Intelligence, Machine Learning, and Deep Learning in Smart Logistics. *Sustainability* **2020**, *12*, 3760. <https://doi.org/10.3390/su12093760>.
40. Parasuraman, R.; Sheridan, T.B.; Wickens, C.D. A Model for Types and Levels of Human Interaction with Automation. *IEEE Trans. Syst. Man Cybern.—Part A Syst. Hum.* **2000**, *30*, 286–297. <https://doi.org/10.1109/3468.844354>.
41. Parasuraman, R.; Manzey, D.H. Complacency and Bias in Human Use of Automation: An Attentional Integration. *Hum. Factors* **2010**, *52*, 381–410. <https://doi.org/10.1177/0018720810376055>.
42. LightGuide, Inc. Industrial Augmented Intelligence Platform. Available online: <https://www.lightguidesys.com/> (accessed on 3 October 2026).
43. TypeSafe AI. Choice. TypeSafe Documentation. Available online: <https://docs.typesafe.ai/primitives/choice> (accessed on 3 October 2026).
44. TypeSafe AI. Confidence. TypeSafe Documentation. Available online: <https://docs.typesafe.ai/confidence> (accessed on 3 October 2026).
45. ISO 12100:2010; Safety of Machinery—General Principles for Design—Risk Assessment and Risk Reduction; International Organization for Standardization (ISO): Geneva, Switzerland, 2010. Available online: <https://www.iso.org/standard/51528.html> (accessed on 3 October 2026).
46. Hart, S.G.; Staveland, L.E. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*; Hancock, P.A., Meshkati, N., Eds.; Advances in Psychology; North-Holland: Amsterdam, The Netherlands, 1988; Volume 52, pp. 139–183. [https://doi.org/10.1016/S0166-4115\(08\)62386-9](https://doi.org/10.1016/S0166-4115(08)62386-9).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.