Target-Specific Pathway Redundancy and Gateway Vulnerability in Construction Human–Robot Collaboration Competency Networks

Ebenezer Olukanni, Abiola Akanmu, Ph.D., and Houtan Jebelli, Ph.D.

Abstract— **Construction robots can reduce workers' exposure to hazardous tasks, but human–robot collaboration (HRC) can also create new safety hazards when workers operate, monitor, or work near robotic systems. Effective preparation, therefore, depends not only on whether relevant competencies are included in training but also on whether important operational and safety-related capabilities are supported through independent competency pathways or depend on narrow gateways. Existing construction HRC competency research identifies competencies and their dependencies but does not quantify this target-specific vulnerability. This study analyzes theory-driven and expert-derived HRC competency networks containing 78 and 99 intra-module relationships, respectively, among 50 unique competencies. Six operational and safety-relevant targets were evaluated using node- and edge-disjoint connectivity, single-competency and relationship removal, and exhaustive multiple-competency removal. Removing Human–robot interface from the expert network left only one of eight baseline upstream competencies with access to Spatial awareness, while removing Teamwork eliminated all theory-network access to Robotic system reasoning. The expert network showed greater redundancy for Regulation standard compliance, but the larger expert edge set did not consistently correspond to lower vulnerability across targets. This study contributes a target-specific structural diagnostic that extends HRC competency modeling beyond dependency identification to distinguish broad connectivity from independently supported target access and identify gateways requiring focused assessment and empirical validation.**

Keywords— Human–robot collaboration, construction robotics, competency networks, pathway vulnerability

I. INTRODUCTION

Construction remains a high-risk industry, with 1,034 fatal work injuries reported in the United States in 2024 [1]. Robots can reduce worker exposure to dangerous, repetitive, and physically demanding activities, but their introduction can also create new interaction hazards, including struck, caught-in, compressed, and crushed events [2]. As robots increasingly operate near workers on dynamic construction sites, workforce preparation must address the competencies needed to operate, supervise, interpret, and safely interact with these systems [3, 4]. Previous studies have identified competencies involving robot applications, sensors, robot control, task planning, human–robot interfaces, safety standards, communication, technical proficiency, problem-solving, decision-making, spatial awareness, and safety awareness [5-7].

Identifying these competencies and their dependencies does not establish that the resulting training pathways are structurally dependable. A competency network may show several routes to Spatial awareness, Safety management, or Robot operation and maintenance, while most ultimately pass through the same supporting competency. If that gateway is omitted from instruction, not assessed, or not mastered, much of the modeled pathway to the target can be lost. Thus, the presence of multiple pathways can create an appearance of strong preparation even when those pathways share the same point of dependence. This does not mean that structural gateway failure predicts injury; it identifies potential weaknesses in competency models intended to support preparation for worker–robot interactions where inadequate readiness can have serious consequences. Existing construction HRC competency research identifies what workers may need and how competencies may be related [5-8], but it has not examined whether target access remains available when individual competencies or relationships are removed. Three properties distinguish this gap: coverage, the number of upstream competencies that can reach a target; pathway redundancy, the availability of structurally independent routes; and gateway vulnerability, the loss of baseline upstream access following removal of a competency or relationship. Broad coverage can therefore coexist with low redundancy and high vulnerability.

Network-removal studies show that disruption depends on the structural position of nodes and links [9-11], while Menger [12]'s theorem relates network separation to the number of internally disjoint paths. These concepts enable a target-specific question that conventional competency-network summaries cannot answer: if one supporting competency or dependency becomes unavailable, how much modeled access to an important HRC capability remains? This study addresses that question using independently developed theory-driven and expert-derived construction HRC competency networks. Six targets represent operational readiness (Robot operation and maintenance; Robotic system reasoning), interaction awareness (Spatial awareness), and safety and regulatory readiness (Regulation standard compliance; Safety management; Safety awareness). The targets have operational or safety relevance based on their definitions; they are not treated as validated predictors of incidents or safe performance. The study addresses three research questions (RQs):

- RQ1: Which competencies and dependency relationships create the greatest structural vulnerability in pathways to selected operational

*Research supported by the National Science Foundation.

Ebenezer Olukanni, Ph.D. Candidate at the Myers Lawson School of Construction, Virginia Tech, Blacksburg, VA 24060 USA (e-mail: ooebenezer@vt.edu).

Abiola Akanmu, Ph.D., Professor, Myers Lawson School of Construction, Virginia Tech, Blacksburg, VA 24060 USA (e-mail: abiola@vt.edu). *Corresponding Author*

Houtan Jebelli, Ph.D., Assistant Professor, University of Illinois Urbana-Champaign, IL 61801 (e-mail: hjebelli@illinois.edu).

and safety-relevant HRC competencies?

- RQ2: How do theory-driven and expert-derived networks differ in pathway redundancy and vulnerability to competency and dependency removal?
- RQ3: How consistently does the expert network's larger edge set correspond to greater target-specific pathway redundancy and lower vulnerability?

The contribution of this study is a target-specific structural diagnostic that combines independent-path connectivity with controlled node, edge, and multiple-competency removal. It extends HRC competency analysis from identifying whether dependencies exist to determining whether access to an intended capability is independently supported or concentrated through structurally consequential gateways.

II. Methods

A. Competency Networks and Targets

The analysis used two directed construction HRC competency networks containing 50 unique competencies organized into seven instructional modules (Table 1). The theory-driven network retained 78 intra-module relationships above the 92nd-percentile semantic-similarity threshold, directed using RBT progression, K→S→A ordering, and deterministic tie-breaking, with weakest-edge cycle removal; the expert-derived network independently retained 99 dependency relationships achieving ≥70% expert agreement.

TABLE I. HRC TRAINING MODULES AND COMPETENCIES

Module	HRC Competencies
M1 – Foundations of Construction Robotics and HRC	Types of robots (K1); Construction robot applications (K2); Robot anatomy and technical specifications (K3); HRC-related fields (K9); ROS (K20); Robotics-related jobs (K24); Robot operation and maintenance (S11)
M2 – Human–Robot Interfaces and Communication	Teamwork (A1); Communication (A2); Spatial awareness (A10); Cultural and social awareness (A11); Robotic system reasoning (A13); Sensors (K4); Immersive virtual environments (K10); Communication modes and technologies (K11); Human–robot interface (K12); Effective communication (S1); Human–robot interface proficiency (S8)
M3 – HRC Task Planning, Programming, and Control	Critical thinking (A9); Task planning (K5); Robot control system (K13); System integration (K14); Programming (K15); Task planning proficiency (S2); Technical skill (S5); Programming proficiency (S6)
M4 – Modeling, Simulation, and Systems Thinking	Problem-solving (A4); Attention to detail (A6); Analytical aptitude (A7); Modeling and simulation (K16); Computational design (K19); Simulation and modeling (S10); Systems thinking (S12)
M5 – Data Analytics, Machine Learning, and Robot Learning	Continuous learning (A3); Data analytics and machine learning (K17); Robot learning methods (K18); Data analytics and management (S7); Application of machine learning algorithms (S9)
M6 – Safety, Ethics, Regulation, and Standards	Decision-making (A8); Safety awareness (A12); HRC ethics and regulation (K6); HRC safety and standards (K7); Regulation standards compliance (S3); Safety management (S4)
M7 – Evaluation, Performance, ROI, and Deployment	Adaptability (A5); Robot acceptance and operational reliability (A14); HRC evaluation (K8); Cost and ROI in robotics (K21); Robotic project management (K22); Robotics sustainability (K23)

For evidence source $s \in \{T, E\}$ (1)

$$G_s = (V_s, E_s), \qquad (2)$$

where T is the theory-driven network, E denotes an expert-derived network, V_s represents competencies and E_s represents directed dependency relationships. Problem-solving occurred in Modules 4 and 5 and was represented separately in each module context, producing 51 contextualized nodes. Only intra-module relationships were analyzed. Theory weights represented semantic similarity, while expert weights represented consensus; neither was interpreted as a probability of learning, mastery, or failure. The six targets were: Robot operation and maintenance (S11) in M1; Spatial awareness (A10) and Robotic system reasoning (A13) in M2; Regulation standards compliance (S3); Safety management (S4); and Safety awareness (A12) in M6 (Figure 1). Figure 1 shows the theory-driven and expert-derived network structures for the three modules containing the six selected target competencies, with identical node positions used within each module to facilitate structural comparison.

Figure 1. Theory-driven and expert-derived competency network structures for the modules containing the six operational and safety-relevant target competencies.

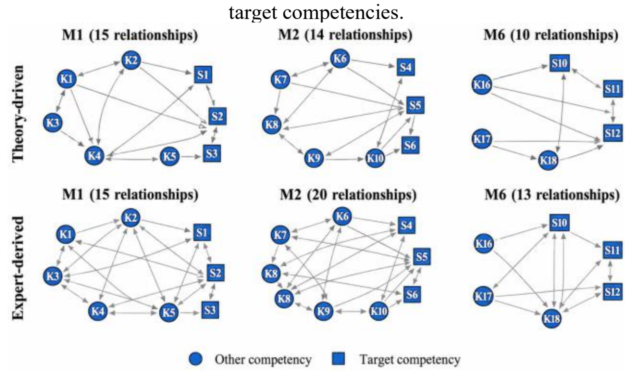

The competency labels and codes correspond to the finalized HRC competency framework. For target t, the baseline upstream set was

$$U_s(t) = \{u \in V \setminus \{t\} : u \rightsquigarrow t\}, \qquad (3)$$

where $u \rightsquigarrow t$ indicates that at least one directed path exists from competency u to target t. The baseline upstream set was fixed before any removal and remained the reference set for all loss and retention calculations.

B. Dependency Construction and Comparison

For each upstream competency u, local node connectivity $\kappa_s(u,t)$ was defined as the maximum number of internally node-disjoint directed paths from u to t. Local edge connectivity was defined similarly using edge-disjoint paths.

Mean target-specific node connectivity was defined as

$$\bar{\kappa}_s(t) = \frac{1}{[U_s(t)]} \sum_{u \in U_s(t)} \kappa_s(u,t). \qquad (4)$$

The proportion of upstream competencies with at least two internally node-disjoint paths was also calculated. Each eligible nontarget competency x in the target's module was then removed individually. If $R_s^x(t)$ denotes the original upstream competencies retaining access after removal, the loss fraction was

$$L_s^x(t) = 1 - \frac{[R_s^x(t)]}{[U_s(t)]} \quad (5)$$

When belonged to the baseline upstream set, its deliberate removal counted as loss, together with any additional upstream competencies that subsequently lost access. The largest value of $L_s^x(t)$ identified the target's most consequential competency gateway.

Each dependency relationship was also removed individually. For edge e,

$$L_s^e(t) = 1 - \frac{[R_s^e(t)]}{[U_s(t)]} \quad (6)$$

where $R_s^e(t)$ denotes baseline upstream competencies still able to reach the target after removal of e. Node and edge removals were evaluated separately because vulnerability can arise from a gateway competency, a critical relationship, or several relationships converging on a common gateway.

C. *Exact Multiple-Removal Robustness*

Single-removal analysis identifies immediate gateways but does not show how quickly access deteriorates when several competencies become unavailable. For a removal set $X \subseteq U_s(t)$, total target retention was

$$\rho_s(X,t) = \frac{[\{u \in U_s(t) \setminus X : u \rightsquigarrow t \text{ in } G_s - X\}]}{[U_s(t)]} \quad (7)$$

For exactly k removals, every possible subset satisfying $[X] = k$ was evaluated. Exact average and worst-case retention were calculated for each removal size. No random subset sampling or approximation was used. The analysis is deterministic. Node removal is a structural stress-test abstraction. Omission, nonassessment, and nonmastery are pedagogically distinct, but each represents a condition in which the competency cannot be assumed to provide reliable downstream support; removal does not model learning or failure probability.

III. Results and Discussion

A. *Target-Specific Competency and Relationship Gateways*

Structural vulnerability varied substantially across targets. Table 2 identifies the most consequential competency gateway for each target, while Figure 2 compares the corresponding maximum loss of baseline upstream access across the theory- and expert-derived networks.

TABLE II. HRC Training Modules and Competencies

Target	Theory gateway	Loss	Expert gateway	Loss
S11	K3	33.3%	K1	50.0%
A10	S8	40.0%	K12	87.5%
A13	A1	100.0%	K12	83.3%
S3	K7	66.7%	K7	50.0%
S4	K7	50.0%	A12	75.0%
A12	K7	40.0%	K7 and K6	50.0%

Figure 2. Maximum fraction of baseline upstream competencies losing access after the most consequential single-competency removal.

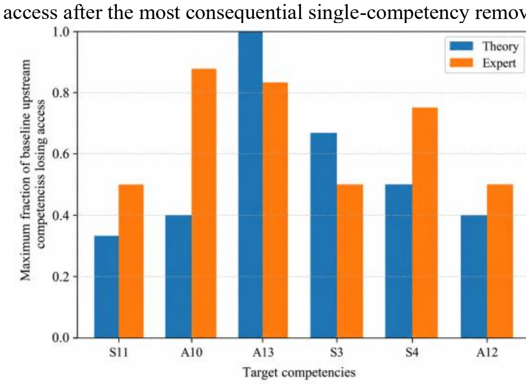

The most severe theory-network vulnerability occurred for A13. A1 eliminated access from all four baseline upstream competencies, and removal of A1→A13 produced the same 100% loss. A1, therefore, functions as a single point of failure in the modeled theory structure. In the expert network, removing K12 left only one of eight baseline upstream competencies with access to A10; removing K11→K12 produced a 75% loss. As Fig. 1 shows, vulnerability was not uniformly greater in either network; the relative severity depended on the target. Thus, RQ1 identifies target-specific structural gateways that warrant attention in training and assessment. K12's gateway role suggests assessing HRI competence before A10-focused activities and testing alternative routes to A10, while A1's role for A13 represents a proposed prerequisite requiring learner-based validation rather than a fixed sequence.

B. *Pathway Redundancy and Removal Vulnerability*

Spatial awareness showed broad expert coverage but weak independent redundancy. The expert network contained eight upstream competencies compared with five in theory, yet mean node connectivity was lower (1.125 versus 1.40), and only 12.5% of expert upstream competencies had at least two node-disjoint routes compared with 40% in theory. Much of the expert access, therefore, converged through K12.

S3 showed the opposite pattern. Theory had mean node connectivity of 1.0 and no upstream competency with two node-disjoint routes. Expert mean connectivity was 1.75; 75% of upstream competencies had at least two independent routes; maximum node loss decreased from 66.7% to 50%; and maximum edge loss was 25%. For this target, the additional expert relationships created independent alternatives around individual gateways. S4 showed why mean connectivity alone is insufficient: both networks had mean node connectivity of 1.5, yet their maximum node losses were 50% for theory and 75% for expert. Theory also showed greater mean node connectivity for A12 (2.0 versus 1.25) and Robot operation and maintenance (2.33 versus 1.50). RQ2 therefore shows that coverage, redundancy, and vulnerability capture different properties: more competencies can reach a target without having more independent ways of reaching it.

C. Effect of the Larger Expert Edge Set

The expert network contained 99 relationships compared with 78 in theory, but the direction of the differences in redundancy and vulnerability varied by target. S3 gained independent expert routes and lower vulnerability, while the additional expert routes to A10 largely converged through K12. To examine whether these differences persisted beyond a single removal, Figure 3 compares exact worst-case retention for three representative patterns: A10, where the expert structure was more vulnerable; A13, where both structures were highly vulnerable; and S3, where the expert structure provided stronger redundancy.

Figure 3. Exact worst-case target-pathway retention under multiple upstream competency removals for three representative targets.

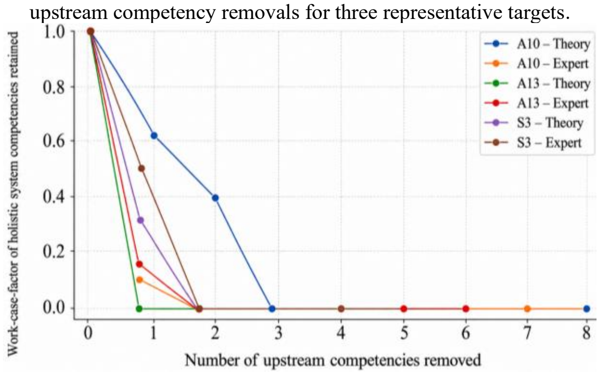

For A10, worst-case retention after one removal was 60% in theory and 12.5% in expert. After two removals, theory retained 40% under its worst case, while expert could be reduced to zero. For A13, one theory removal eliminated all upstream access; the expert structure reached zero under the worst two-removal combination. For S3, worst-case one-removal retention was 33.3% in theory and 50% in expert.

RQ3, therefore, shows that edge count alone is not evidence of stronger target support. Additional relationships reduce vulnerability when they create independent routes around gateways; when they converge on the same gateway, they can increase coverage without providing comparable redundancy. For workforce preparation, gateway status should therefore be treated as a signal for direct assessment and empirical validation, not as evidence that a competency necessarily causes readiness or should automatically receive more instructional time.

IV. Conclusion

This study examined target-specific pathway redundancy and gateway vulnerability across six operational and safety-relevant construction HRC competencies. The results show that broad connectivity does not necessarily indicate independently supported pathways. The expert network provided broader access to Spatial awareness, yet removal of Human–robot interface left only one of eight baseline upstream competencies connected. Robotic system reasoning was highly vulnerable in both networks, while the expert network showed stronger redundancy for Regulation standard compliance. The theory network provided stronger pathway support for Safety awareness and Robot operation and maintenance. Overall, vulnerability depended on how pathways were organized around specific gateways, not on which network contained more relationships.

The study contributes both methodologically and substantively. Methodologically, it extends construction HRC competency analysis from identifying dependencies to evaluating target-specific pathway independence and removal vulnerability. Substantively, it demonstrates that greater connectivity cannot be interpreted as stronger workforce preparation unless additional relationships create independent routes around critical gateways.

The findings are structural and do not establish causal learning relationships, injury risk, or safe performance. The analysis is limited to intra-module relationships, six selected targets, and one expert-elicitation process. Its value is in identifying the competency relationships for which stronger empirical evidence is most needed. These gateways can be prioritized for direct assessment, alternative instructional pathways, and learner-based validation before competency networks inform training or readiness decisions in worker–robot environments where inadequate preparation can have serious safety consequences.

References

[1] U. S. Bureau of Labor Statistics, "Fatal occupational injuries by industry and event or exposure, all United States, 2024," *Census of Fatal Occupational Injuries*, U.S. Department of Labor, Bureau of Labor Statistics, 2026.

[2] S. a. H. National Institute for Occupational, “Transforming Construction: Automation and Robotics for a Safer Future,” *NIOSH Science Bulletin*, 2024–11–12, 2024.

[3] X. Wang, C.-J. Liang, C. C. Menassa, and V. R. Kamat, “Interactive and immersive process-level digital twin for collaborative human–robot construction work,” *Journal of Computing in Civil Engineering,* vol. 35, no. 6, pp. 04021023, 2021.

[4] M. Zhang, R. Xu, H. Wu, J. Pan, and X. Luo, “Human–robot collaboration for on-site construction,” *Automation in Construction,* vol. 150, pp. 104812, 2023/06/01/, 2023.

[5] E. Olukanni, A. Akanmu, and H. Jebelli, “Systematic Literature Review of Competencies for Human–Robot Collaboration in Construction: A Task-Technology Fit Perspective,” *Journal of Computing in Civil Engineering,* vol. 40, no. 1, pp. 03125001, 2026.

[6] E. Olukanni, A. Akanmu, H. Jebelli, A. Ammar, and A. Okunola, “Competencies for human–robot collaboration in the construction industry: perspectives from academia and industry,” *Construction Robotics,* vol. 9, no. 2, pp. 15, 2025/06/16, 2025.

[7] E. Olukanni, A. Akanmu, H. Jebelli, and S. Terreno, “Competencies for human-robot collaboration in the construction industry–academia’s perspective,” *International Journal of Construction Education and Research,* vol. 20, no. 4, pp. 444–463, 2024.

[8] E. Olukanni, A. Akanmu, and H. Jebelli, “Industry perception of competencies for human—robot collaboration in the construction industry: A Delphi study,” *Frontiers of Engineering Management,* vol. 12, no. 4, pp. 854–879, 2025/12/01, 2025.

[9] R. Albert, H. Jeong, and A.-L. Barabási, “Error and attack tolerance of complex networks,” *nature,* vol. 406, no. 6794, pp. 378–382, 2000.

[10] H. Cetinay, K. Devriendt, and P. Van Mieghem, “Nodal vulnerability to targeted attacks in power grids,” *Applied network science,* vol. 3, no. 1, pp. 34, 2018.

[11] M. Bellingeri, D. Bevacqua, F. Scotognella, R. Alfieri, and D. Cassi, “A comparative analysis of link removal strategies in real complex weighted networks,” *Scientific reports,* vol. 10, no. 1, pp. 3911, 2020.

[12] K. Menger, “Zur allgemeinen kurventheorie,” *Fundamenta mathematicae,* vol. 10, pp. 96–115, 1927.