Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 950–961
https://doi.org/10.54941/ahfe1008326

Living Geometry as a Human-Centered Design Framework for Context-Aware Multi-Stage Injury Prevention

Togo Hane[1], Yuya Kawabe[1], Philip Blenwall[2], Koji Kitamura[3], Mikiko Oono[3], Tatsuhiro Yamanaka[4], and Yoshifumi Nishida[1]

[1]Institute of Science Tokyo, Tokyo, Japan
[2]Karlstad University, Karlstad, Sweden
[3]National Institute of Advanced Industrial Science and Technology, Tokyo, Japan
[4]Safe Kids Japan, Tokyo, Japan

ABSTRACT

Human-centered safety design requires more than detecting people and hazards; it must account for how risk emerges from changing relationships among human capabilities, behavior, objects, and everyday environments. This paper presents a prototype injury-prevention service based on Living Geometry, a framework that transforms physical observations into context-dependent relational structures such as reachability, attention, caregiver availability, and hazardous coupling. The framework connects RGB-D-based estimation of a person's position and gaze direction with route-aware reachability analysis and a multi-stage intervention strategy. Potentially hazardous appliances, including an electric fan and an electric pot, are evaluated in both floor-level and tabletop placements. For elevated hazards, the system considers not only direct distance but also whether an accessible route exists through surrounding furniture or structures. When an active hazardous relationship is detected, the service first notifies a caregiver, preserving human judgment as the primary response. If no response is obtained, a bear-shaped robot redirects the person's attention and trajectory. If the hazardous relationship persists, the system escalates to device-level control, such as turning off the appliance. A proof-of-concept implementation in the AIST Living Lab demonstrates that heterogeneous sensing and intervention mechanisms can be coordinated through a common relational representation. The contribution is a human-centered design framework that selects among interventions with different degrees of human involvement and technological control according to the evolving person-environment-hazard relationship.

Keywords: Living Geometry, Multi-stage intervention, Context-aware safety

INTRODUCTION

Unintentional injuries remain a major child-safety problem, and many preventive opportunities arise in and around the home (WHO and UNICEF, 2008). Many domestic injury-prevention systems detect isolated states: a person is present, an appliance is on, or an object lies within a predefined safety zone. These observations are useful, but they do

Received September 26, 2026; Revised October 30, 2026; Accepted November 4, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

not directly represent why a situation is hazardous. Risk emerges from relationships. A pot becomes hazardous when it is active and reachable by a particular person; a fan becomes actionable when the person is approaching or attending to it; and caregiver notification is useful only when a caregiver is available and can intervene in time. Human-centered safety design therefore requires a representation that connects sensing to the changing relationships among people, objects, environments, and possible interventions.

Ubiquitous and context-aware computing has developed methods for recognizing context and adapting services to it (Weiser, 1991; Dey, 2001; Satyanarayanan, 2001; Gu et al., 2005). Situation-awareness research further emphasizes interpreting environmental elements in relation to future action (Endsley, 1995). Scene-graph research has made spatial relations among people, objects, and places explicit (Rosinol et al., 2020), while embodied artificial intelligence has emphasized perception and action through physical interaction (Duan et al., 2022). However, physical grounding alone does not guarantee that a system represents what matters in everyday life. Sensor coordinates and activity labels must be connected to lived meanings such as who can reach an object, whose attention is directed toward a hazard, who can intervene, and how the situation can be changed without unnecessarily disrupting ongoing activity.

Building on the broader symbol-grounding problem (Harnad, 1990), we refer to this problem as Living Grounding: grounding measured physical states in relational structures that have meaning for everyday activity and design. Living Geometry (LG) is the computational representation used for this purpose. LG describes a living situation as a typed, dynamic relational structure and retains the geometric and contextual constraints required to recompute the structure when human capability, object state, or environmental configuration changes. Earlier work used this representation to examine perturbations in human, contextual, and environmental conditions and to re-ground selected living relationships in another home environment (Hane et al., 2026). The present study extends the framework from analysis to intervention.

The research question is: How can heterogeneous safety interventions be coordinated through a common representation while preserving human involvement whenever the situation permits? To address this question, we developed a prototype service that combines RGB-D sensing, gaze-direction estimation, route-aware reachability assessment, caregiver notification, robotic attention redirection, and device-level hazard control. Consistent with human-centered design and levels-of-automation research (ISO, 2019; Parasuraman et al., 2000), the framework does not prescribe one automated response. Instead, it escalates intervention only when the currently selected response fails to resolve the hazardous relationship. Figure 1 summarizes this framework.

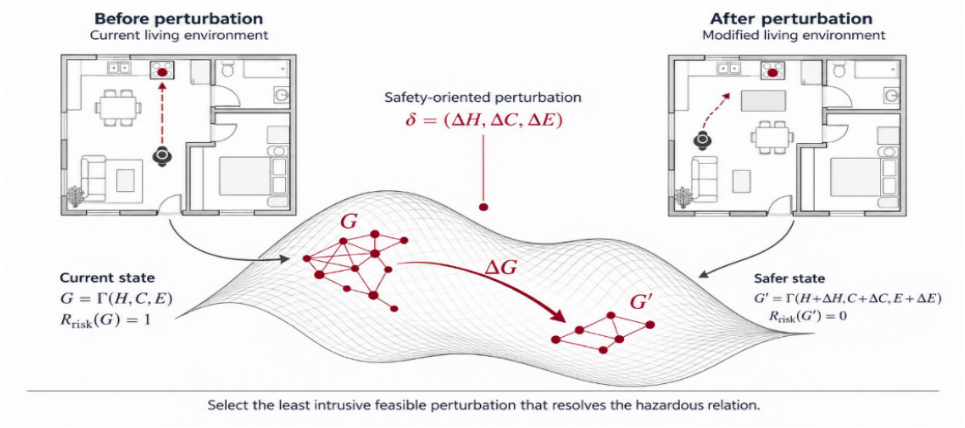


Figure 1: Living Geometry as a bridge from physical observation to context-dependent safety intervention.

LIVING GEOMETRY FOR HUMAN-CENTERED SAFETY DESIGN

Relational Representation and Living Grounding

At time t, Living Geometry is represented as a typed dynamic relational structure composed of an entity set V(t), a relation set R(t), and an attribute set A(t):

$$G(t) = \big(V(t), R(t), A(t)\big). \tag{1}$$

Entities include people, appliances, furniture, rooms, and intervention agents. Relations include located-in, near, approaches, attends-to, reachable, reachable-via-route, caregiver-of, available-for-intervention, hazard-for, and risk-active. Attributes describe states such as appliance activation, estimated human capability, object height, or response status. This representation separates the observable physical layer from the living-context layer without discarding the geometric constraints required for later computation.

Let $Y_{0:t}$ denote an observation sequence, H the person's physical and developmental conditions, C the current living context and object states, and E the environment model. A grounding operator Gamma generates Living Geometry:

$$G(t) = \Gamma\big(Y_{0:t}, H, C, E\big). \tag{2}$$

Gamma is not limited to one inference technique. In the present proof of concept, deterministic rules connect geometric features and system states to relational edges. Future implementations may use probabilistic inference or learned models. The important requirement is that the output remains interpretable as a set of context-dependent relationships that can be evaluated before and after intervention.

Figure 1 also summarizes a local sensitivity view of Living Geometry. For sufficiently small changes in human conditions H, living context C,

and environment E, the corresponding structural change can be written conceptually as

$$\Delta G \approx \frac{\partial G}{\partial H}\Delta H + \frac{\partial G}{\partial C}\Delta C + \frac{\partial G}{\partial E}\Delta E. \tag{3}$$

This expression is a conceptual local model rather than an experimentally estimated causal equation. In the present prototype, intervention acts mainly through contextual changes, including attention, caregiver response, and appliance state (ΔC); human capability (ΔH) and environmental configuration (ΔE) determine reachability and route accessibility. The safety objective is to select a feasible change that removes the risk-active relation.

From Risk Detection to Relation Transformation

A hazardous situation is represented as a coupling among multiple relations rather than as a property of an object alone. For example, an appliance may be active but not reachable; a person may be near an appliance without attending to it; or a potentially hazardous object may be reachable while a caregiver is already intervening. A risk-active relation is generated only when the relevant conditions are jointly satisfied. The objective of intervention is then expressed as transforming G(t) into a subsequent structure G(t + Δt) in which the risk-active relation no longer holds.

This formulation permits several responses to achieve the same safety objective. A caregiver may remove the hazard, a robot may redirect the person's attention or trajectory, or the system may deactivate the appliance. These actions differ in autonomy, intrusiveness, and dependence on human availability. Their common evaluation criterion is whether they resolve the hazardous relationship while minimizing unnecessary control and disruption.

CONTEXT-AWARE MULTI-STAGE INTERVENTION

Human Involvement as the First Response

Figure 2 summarizes the intervention sequence. The first intervention stage sends a notification to a caregiver's smartphone. This stage treats caregivers as active participants rather than bypassing them through immediate automation. When the caregiver acknowledges the notification and changes the situation, the system recomputes the relevant relations. If reachability, attention, or hazard activation is removed, escalation stops. This makes the intervention human-centered in an operational sense: the system reserves space for human judgment before exercising stronger technological control.

Attention Redirection as an Intermediate Response

If the caregiver does not respond within the configured interval, the system activates an attention-redirecting bear-shaped robot, drawing on the use of socially expressive robots to shape human attention and interaction (Breazeal, 2003; Goodrich and Schultz, 2007). The robot is intended to shift interest and movement away from the hazard, thereby gaining time without physically constraining the person. The robot is not treated as an isolated device; its action is evaluated through changes in the same Living Geometry

relations used for hazard detection. A successful intervention is observed when attention-to-hazard, approaches, or reachable-in-trajectory relations cease to hold before hazardous contact occurs.

Device-Level Control as the Final Response

If the hazardous relationship persists after caregiver notification and attention redirection, the system proceeds to device-level control. For an electric fan, this may involve turning the fan off. This final stage removes hazard activation rather than changing human behavior. It is deliberately placed last because direct device control can interrupt ongoing activity and may not always be desirable. The escalation order therefore reflects both urgency and the degree of system authority.

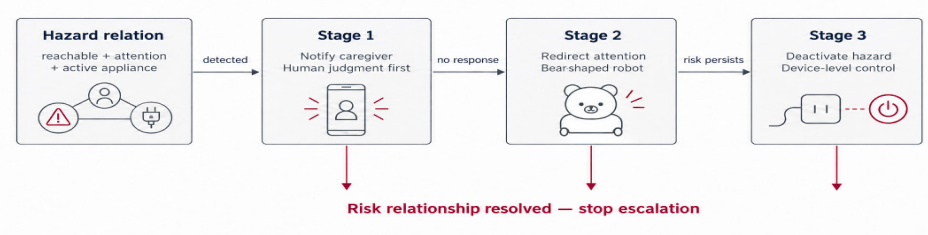


Figure 2: Multi-stage intervention logic. Escalation stops whenever the Living Geometry risk relation is resolved.

PROOF-OF-CONCEPT IMPLEMENTATION

The proof-of-concept system integrates synchronized sensing, Living Geometry-based relational reasoning, and three stages of safety intervention. Figure 3 shows MQTT-based information exchange among the sensing, Living Geometry computation, and device-control modules, with Home Assistant providing device integration.

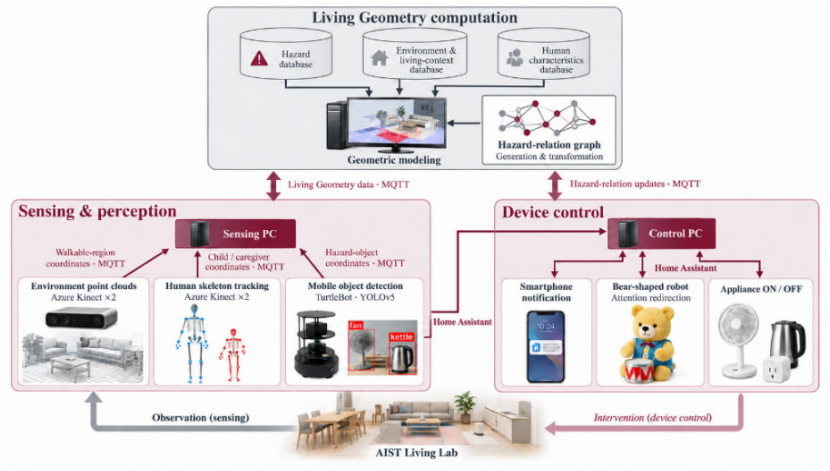


Figure 3: Living Geometry system architecture with MQTT-based messaging and Home Assistant integration.

Living Lab and Sensing Configuration

Figure 4 illustrates the Living Lab setup and representative caregiver and device-level interventions.

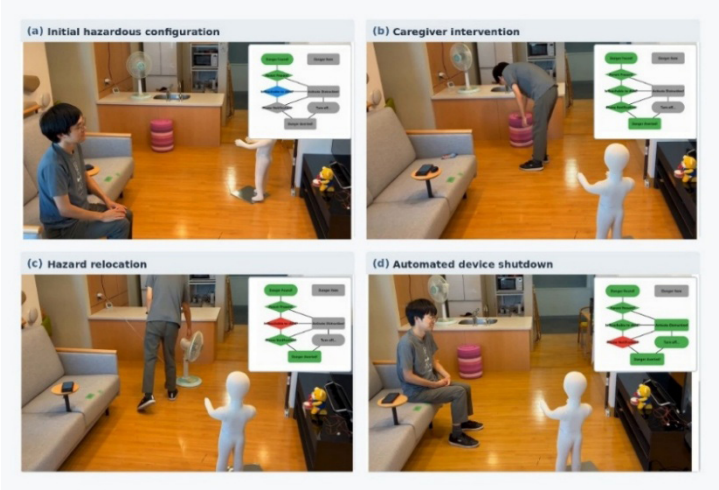


Figure 4: Real-world intervention scenarios in the AIST Living Lab: (a) initial hazardous configuration, (b) caregiver intervention, (c) relocation of the hazardous appliance, and (d) automated device shutdown. The insets indicate the corresponding decision states of the intervention system.

The proof of concept was implemented in the AIST Living Lab, a mock domestic environment containing living, dining, kitchen, bathing, care, and childcare areas. Time-synchronized RGB-D observations support environmental point-cloud measurement, human-skeleton measurement, and three-dimensional object detection (Shotton et al., 2011; Cao et al., 2017; Newcombe et al., 2011). As shown in Figure 3, two Azure Kinect sensors measure environmental point clouds, and two separate Azure Kinect sensors track human skeletons, giving four sensors for these functions. An additional Azure Kinect mounted on the TurtleBot supports mobile object detection with YOLOv5; detections are transformed into three-dimensional object coordinates. The resulting person, object, and environment states are associated with a model of rooms, furniture, support surfaces, and appliance locations, while appliance activation is obtained from the corresponding control interfaces.

Potentially hazardous appliances include an electric fan and an electric pot. They may be placed on the floor or on a table. This variation is important because height alone does not determine safety. An elevated appliance may still be accessible when surrounding furniture or structures provide a route for approach, climbing, or extended reach. Conversely, an appliance that is close in Euclidean distance may be inaccessible because no feasible route exists.

Route-Aware Reachability

The reachability module evaluates both direct access and access through intermediate environmental structures. The direct case compares the person's capability model with the position and support-surface configuration of the hazard. The route-aware case searches for a sequence of accessible environmental elements that can extend access to an elevated object. A reachable-via-route relation is generated only when the intermediate sequence satisfies the implemented geometric constraints, following the general principle of cost-constrained graph search (Hart et al., 1968). This prevents the system from equating every tabletop object with either absolute safety or absolute danger.

The system evaluated age-dependent reach and access to an appliance labeled "Kettle" through an intermediate cylindrical object. Figure 5 illustrates this model and its relation transformation. The AHFE prototype extends this principle to surrounding furniture and structures, combines it with real-time person and attention observations, and uses the resulting relations for intervention.

As shown in Fig. 5, the source data comprise an adult trajectory and a child-mannequin position; object states and support-mediated reachability are modeled. In $G(t)$, the toy provides an intermediate support for reaching the kettle. Switching the modeled appliance state to OFF removes the activation-dependent hazard relation while preserving reachability in $G'(t)$. These graphs illustrate the representation, rather than an additional measured intervention trial. For a hot-water appliance, residual heat and hot contents can remain hazardous after shutdown; actual hazard clearance therefore requires additional state information or access prevention.

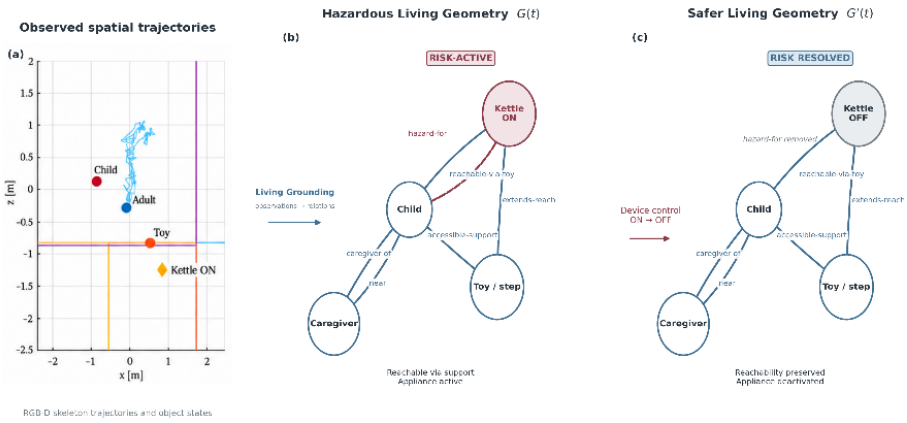


Figure 5: Measured spatial data and an illustrative Living Geometry transformation. (a) Spatial data also used in Hane et al. (2026), with configured object states. (b) Simplified risk-active graph *G*(*t*). (c) Graph *G'*(*t*) after modeled ON-to-OFF switching. "Risk resolved" refers only to the modeled activation-dependent relation, not residual thermal risk.

System Integration

As shown in Figure 3, MQTT-based messaging connects the sensing, Living Geometry computation, and control modules (OASIS, 2019), while Home Assistant integrates the intervention devices. The computation module grounds observations in G(t), updates safety-relevant relations, and selects the least escalated applicable response. The control module then coordinates caregiver notification, robotic attention redirection, and appliance power control.

DEMONSTRATION AND RESULTS

Demonstration Scenarios

Table 1 summarizes the four demonstration scenarios. The integrated system was exercised in four Living Lab scenarios that varied appliance placement, route accessibility, caregiver response, and persistence of risk. The demonstrations used a child mannequin rather than actual children. Here, "success" denotes correct system operation under the configured scenario conditions. The effectiveness of the robot in redirecting actual children's attention or behavior was not evaluated. S1 tested a directly reachable floor fan and successful caregiver intervention. S2 tested a tabletop hazard with an accessible route, no caregiver response, and successful execution of the robotic redirection sequence. S3 tested a reachable tabletop appliance for which neither caregiver notification nor robotic redirection resolved the risk, causing escalation to device control. S4 tested a tabletop appliance without an accessible route; because the appliance was not reachable under the modeled environmental constraints, no intervention was initiated.

Table 1: Outcomes of four Living Lab scenarios for route-aware risk assessment and multi-stage intervention.

ID	Hazard Configuration	Reachability	Caregiver Response	Robot Response	Device Control	Outcome
S1	Floor fan	Directly reachable	Successful	Not activated	Not activated	Risk resolved
S2	Tabletop hazard with accessible route	Route-reachable	No response	Successful redirection	Not activated	Risk resolved
S3	Tabletop appliance	Reachable	No response	Risk persisted	Activated	Hazard deactivated
S4	Tabletop appliance without an accessible route	Not reachable	Not activated	Not activated	Not activated	No intervention

Integrated Operation

The proof of concept confirmed end-to-end connectivity among real-time RGB-D sensing, gaze-direction estimation, route-aware reachability analysis, caregiver notification, robotic attention redirection, and appliance-state control.

In S1, caregiver action resolved the risk and the system did not activate the robot or device control. In S2, the absence of a caregiver response triggered the robot; the system successfully completed the configured robotic redirection sequence and stopped escalation without activating device control. In S3, the risk persisted after robotic intervention, and appliance deactivation removed the active hazard. These cases demonstrate that escalation stopped at different stages according to the observed relational outcome.

S4 demonstrated the distinction between an elevated hazard and an accessible hazard. The tabletop appliance was not judged solely from its height or straight-line distance from the person. Because no accessible route existed under the modeled environmental constraints, the system generated no active reachability relation and initiated no intervention. Comparison of S2 and S4 shows that surrounding furniture and structures can change risk assessment even when the hazard category and elevated placement are similar. This relation-based assessment can reduce unnecessary alerts and interventions while retaining sensitivity to indirect access routes.

These results establish technical feasibility, not comparative effectiveness. The present study does not claim that the robot reliably redirects every user, that one escalation interval is optimal, or that the staged strategy reduces injury incidence. Those questions require repeated trials with defined participants, intervention timing measures, behavioral outcomes, and acceptability assessment. The current contribution is the integrated design framework and its operational proof of concept.

DISCUSSION

Contribution to Human-Centered Design

The principal contribution is not a new sensor, gaze estimator, robot, or smart appliance in isolation. It is a design framework that organizes these components around the changing relationship between a person and an everyday hazard. This changes the design question from 'Which device should respond?' to 'Which relationship must change, and which available response can change it with the least necessary escalation?' The relational formulation also makes the intervention rationale inspectable: the system can indicate whether it acted because the hazard became reachable, attention shifted toward it, the caregiver was unavailable, or a previous response failed.

The ordering of interventions embodies a specific human-centered principle. Human participation is retained when time and availability permit; behavioral redirection is used as an intermediate, potentially less disruptive response; and direct technological control is reserved for persistent risk. This allocation of agency follows established concerns in human-centered design and human-AI interaction, including visibility of system status, appropriate control, and graceful escalation (Norman, 2013; Amershi et al., 2019). Human-centeredness here does not imply that participatory design or user acceptability evaluation has already been completed. Rather, it refers to the explicit representation of human capability, attention, availability, and action in safety decisions.

Relationship to the Previous Study

The previous study introduced Living Geometry as a bridge between real-world sensing and simulation (Hane et al., 2026). It demonstrated that relations could change under perturbations in human condition, context, and environment and that selected relations could be preserved when re-grounded in another home. The present paper uses the same grounding concept but addresses a different design function: selecting and coordinating interventions in an evolving situation. Retargeting and relation-preserving re-grounding are therefore treated as foundations for future deployment across homes rather than as the main experimental result.

Limitations and Future Work

The current evaluation is limited to a Living Lab proof of concept and a small set of configured scenarios. Reachability and route accessibility are based on explicit geometric rules and do not yet model the full diversity of posture, climbing strategy, support use, or individual physical capability. Gaze direction estimated from RGB-D data is an indicator of visual orientation and should not be interpreted as a complete measure of attention. The intervention logic also uses configured thresholds and escalation intervals rather than parameters learned from longitudinal household data.

Future work will quantify detection latency, intervention lead time, caregiver response time, robot-induced attention change, unnecessary intervention, activity disruption, and the probability of resolving the hazardous relationship at each stage. User studies should examine acceptability and perceived control among caregivers and household members. The framework should also be tested across multiple home layouts by using relation-preserving re-grounding to identify where a strategy remains valid and where environmental geometry produces a breaking point. Such evaluation would connect the intervention-oriented AHFE prototype with the broader goal of making living contexts computationally transformable and testable.

CONCLUSION

This paper presented Living Geometry as a human-centered design framework for context-aware multi-stage injury prevention. The framework grounds RGB-D observations, gaze direction, environment geometry, and appliance state in relations describing reachability, attention, caregiver availability, hazardous coupling, and intervention outcomes. It evaluates hazards on the floor and on tables, including access through intermediate environmental routes. When risk becomes active, the system prioritizes caregiver intervention, proceeds to robotic attention redirection if necessary, and finally deactivates the appliance if risk persists. A Living Lab proof of concept demonstrated that these heterogeneous mechanisms can be integrated through a common relational representation. The work establishes a basis for designing safety services that adapt the degree of automation to the evolving situation while preserving human involvement whenever possible.

ACKNOWLEDGMENT

This work was supported in part by the BRIDGE program of the Cabinet Office of Japan, R7-H05, under research and development activities related to strengthening AI safety in industrial fields. OpenAI's ChatGPT was used to assist with English-language editing and figure preparation; all technical content was reviewed by the authors.

REFERENCES

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., Teevan, J., Kikin-Gil, R. and Horvitz, E. (2019). Guidelines for human-AI interaction. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Paper 3, 1–13.

Breazeal, C. (2003). Toward sociable robots. Robotics and Autonomous Systems, 42(3–4), 167–175.

Cao, Z., Simon, T., Wei, S.-E. and Sheikh, Y. (2017). Realtime multi-person 2D pose estimation using part affinity fields. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 7291–7299.

Dey, A.K. (2001). Understanding and using context. Personal and Ubiquitous Computing, 5(1), 4–7.

Duan, J., Yu, S., Tan, H.L., Zhu, H. and Tan, C. (2022). A survey of embodied AI: From simulators to research tasks. IEEE Transactions on Emerging Topics in Computational Intelligence, 6(2), 230–244.

Endsley, M.R. (1995). Toward a theory of situation awareness in dynamic systems. Human Factors, 37(1), 32–64.

Goodrich, M.A. and Schultz, A.C. (2007). Human-robot interaction: A survey. Foundations and Trends in Human-Computer Interaction, 1(3), 203–275.

Gu, T., Pung, H.K. and Zhang, D.Q. (2005). A service-oriented middleware for building context-aware services. Journal of Network and Computer Applications, 28(1), 1–18.

Hane, T., Kawabe, Y., Blenwall, P., Kitamura, K., Oono, M. and Nishida, Y. (2026). Bridging real-world sensing and simulation through Living Geometry: A framework for living-context grounding, perturbation, and retargeting. Proceedings of SI2026, forthcoming (in Japanese).

Harnad, S. (1990). The symbol grounding problem. Physica D: Nonlinear Phenomena, 42(1–3), 335–346.

Hart, P.E., Nilsson, N.J. and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. IEEE Transactions on Systems Science and Cybernetics, 4(2), 100–107.

ISO (2019). ISO 9241-210:2019 Ergonomics of human-system interaction - Part 210: Human-centred design for interactive systems. International Organization for Standardization, Geneva.

Newcombe, R.A., Izadi, S., Hilliges, O., Molyneaux, D., Kim, D., Davison, A.J., Kohli, P., Shotton, J., Hodges, S. and Fitzgibbon, A. (2011). KinectFusion: Real-time dense surface mapping and tracking. Proceedings of the 10th IEEE International Symposium on Mixed and Augmented Reality, 127–136.

Norman, D.A. (2013). The Design of Everyday Things: Revised and Expanded Edition. Basic Books, New York.

OASIS (2019). MQTT Version 5.0. OASIS Standard, 7 March 2019.

Parasuraman, R., Sheridan, T.B. and Wickens, C.D. (2000). A model for types and levels of human interaction with automation. IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans, 30(3), 286–297.

Rosinol, A., Gupta, A., Abate, M., Shi, J. and Carlone, L. (2020). 3D dynamic scene graphs: Actionable spatial perception with places, objects, and humans. Robotics: Science and Systems XVI.

Satyanarayanan, M. (2001). Pervasive computing: Vision and challenges. IEEE Personal Communications, 8(4), 10–17.

Shotton, J., Fitzgibbon, A., Cook, M., Sharp, T., Finocchio, M., Moore, R., Kipman, A. and Blake, A. (2011). Real-time human pose recognition in parts from single depth images. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 1297–1304.

Weiser, M. (1991). The computer for the 21st century. Scientific American, 265(3), 94–104.

WHO and UNICEF (2008). World Report on Child Injury Prevention. World Health Organization, Geneva.