

CANINE: Coaching Visually Impaired Users for Interactive Navigation with a Robot Guide Dog

Cunjun Yu¹, Zishuo Wang¹, Anxing Xiao¹, Linfeng Li¹, David Hsu^{1,2}

¹School of Computing, ²Smart Systems Institute
National University of Singapore

Abstract—*Robot guide dogs offer navigation assistance that greatly expands the independent mobility of the visually impaired, but their effective use requires subtle human-robot coordination that is difficult for users to learn from generic verbal instructions. To tackle this challenge, we present CANINE, an automated coaching system that trains users for interactive navigation with a robot guide dog, through personalized, adaptive verbal feedback. CANINE decomposes a complex coordination task into sub-skills and operates at two levels. At the high level, it decides *what* to train by tracking the learner’s proficiency across sub-skills using *knowledge tracing* and prioritizing training on the weakest areas. At the low level, CANINE decides *how* to train each sub-skill by observing each human practice episode, using foundation models to infer the underlying causes of errors, and generating targeted verbal corrections adaptively. A controlled study with blindfolded participants, treated as a proxy population for quantitative evaluation, demonstrates that CANINE significantly improves both learning efficiency and final navigation performance compared to generic verbal instructions. We further validate CANINE through a retention study and an exploratory case study. The retention study shows lasting skill improvement after two weeks. The case study confirms CANINE’s effectiveness in training a visually impaired user, while revealing additional design considerations for real-world deployment. Both are well aligned with the findings of the controlled study. The video can be found on the [project page](#).*

I. INTRODUCTION

For the 2.2 billion people worldwide with visual impairments (VI) according to the WHO [63], everyday tasks like navigating unfamiliar environments require assistance that is often unavailable [38]. While robotic guide systems promise to provide on-demand support [65, 9, 6, 23, 19, 28, 31], a critical challenge remains largely unaddressed: *how can we effectively coach users to leverage these systems?* Consider a user navigating through a doorway with a robot guide dog [7] as shown in Fig. 1. This task requires precise spatial positioning relative to both the robot and the door frame, combined with exact coordination to avoid collisions, all without visual feedback. Similarly, when such robotic platforms are equipped with a manipulator arm for object handover [8, 30], users must coordinate hand positioning and grasp timing without seeing the robot’s approach. Without effective coaching, users may develop unsafe habits, and experience repeated failures that erode confidence. These scenarios exemplify a broader class of physical interaction tasks where one agent (the human) lacks key perceptual channels and must learn to coordinate with another agent (the robot) through alternate modalities.

Coaching humans to master these tasks presents two fun-




Fig. 1: CANINE. In our study, a robot guide dog coaches a visually impaired user to navigate through a doorway.

damental challenges. First, these tasks are complex processes composed of distinct phases, each requiring different coordination strategies. For example, locating a door handle requires spatial exploration, whereas passing through the door requires precise coordination. Given the heterogeneity of these phases and significant variability across users, treating the task as a single monolithic process is inefficient. Second, the user’s lack of visual perception makes providing feedback challenging. The user cannot observe the environment to correct their errors and must rely on the coaching system to provide in-situ feedback. To provide meaningful guidance, the coach must infer latent user states, such as distinguishing between spatial misalignment and cue misinterpretation, solely from noisy interaction signals. Existing approaches rely on periodic sessions with orientation and mobility (O&M) specialists [67], which are effective but require human professionals and provide limited opportunities for repetitive skill refinement.

We present CANINE, an automated coaching system that trains users for interactive navigation with robot guide dogs through personalized, adaptive verbal feedback. CANINE addresses these challenges by operating at two distinct levels, mirroring how human coaches determine *what* to practice and *how* to practice it. To enable targeted training, CANINE decomposes the overall task into *sub-skills*, modular components that can be practiced independently, e.g., open door.

At the *inter-skill level*, CANINE addresses curriculum scheduling by tracking the learner’s latent proficiency across sub-skills using *knowledge tracing* [62, 49] and dynamically

prioritizing training in the weakest areas. We formulate this as a partially observable decision problem where user proficiency cannot be directly observed but must be inferred from performance, enabling the system to maintain beliefs about user competencies and select training activities that maximize learning efficiency. At the *intra-skill level*, CANINE addresses feedback generation by observing each practice episode and providing adaptive audio feedback through a speaker at the end of each attempt. The system observes the user’s behavior, uses foundation models to infer the underlying causes of errors (e.g., whether an error stems from spatial misalignment or misunderstanding of robot cues), and generates targeted verbal corrections. To handle the complexity of modeling human behavior, we instantiate an approximate policy using a structured reasoning process: Vision-Language Models (VLMs) [42] extract behavioral states from video observations, and Large Language Models (LLMs) [5] generate actionable feedback tailored to the inferred state (e.g., “position yourself further from the robot”).

We validate CANINE through a multi-stage evaluation on doorway navigation with robot guide dogs. In a controlled proxy-user study with blindfolded sighted participants [65, 9], CANINE significantly improves learning efficiency and final performance over generic verbal instructions. We further validate CANINE through (i) a retention study showing lasting skill improvement after two weeks and (ii) an exploratory case study with a visually impaired user.

II. RELATED WORKS

Our work sits at the intersection of assistive robotics, human-robot interaction, and intelligent tutoring systems.

Physical Human-Robot Interaction. Physical Human-Robot Interaction research has explored three paradigms: *assistance*, *collaboration*, and *teaching* [24, 50]. Assistance systems infer human intent and adapt robot behavior to help users achieve immediate goals [14, 47, 64, 26, 43], such as mobility augmentation [39] or motor function restoration [10]. Collaboration systems enable humans and robots to work toward shared objectives through mutual adaptation [34], as seen in object manipulation [40] and co-manipulation tasks [48]. Teaching systems structure interactions to improve human capabilities over time [66, 60, 33, 45, 18, 56]. Recent robotic teaching work has made significant progress on key aspects of the coaching problem. Curriculum-focused methods [66, 56, 54, 16] have demonstrated effective approaches enabling adaptive skill sequencing. Feedback-focused methods [45, 55, 37] have shown how to monitor performance and generate corrective guidance. Physical assistance methods [22] have advanced techniques for modulating robot behavior during execution to support motor learning. Building on these foundations, CANINE integrates proficiency-based curriculum scheduling with diagnostic verbal feedback to address the complete coaching loop, particularly for learners who cannot visually self-correct.

Assistive robotics has developed navigation systems for visually impaired users [65, 9, 6]. While these systems demonstrate technical feasibility in controlled settings, they

provide no structured learning support, leaving users to acquire coordination skills through trial and error. CANINE addresses this gap by providing systematic coaching.

Intelligent Tutoring Systems. Intelligent Tutoring Systems (ITS) have a long history in education, focusing on personalized instruction and adaptive learning [44, 20, 11, 49]. However, traditional ITS primarily target cognitive skills (e.g., mathematics, language learning) [59, 36, 57] rather than embodied, physical skills. Recent work has begun to explore LLM-based coaching for physical activities, such as GPTCoach [27] which provides feedback based on self-reported performance data and sensor readings. GPTCoach takes in user self-reported structured performance metrics (e.g., exercise counts) and generates motivational feedback and workout plans. In contrast, CANINE must diagnose coordination errors from raw visual observations of human-robot interaction, requiring the system to infer what went wrong from unstructured observation alone and generate actionable corrections for spatial and temporal coordination failures that users cannot directly observe.

Foundation Models for Robotics. Recent advances in foundation models, such as LLMs and VLMs, have shown promise for robotics applications [17, 29, 3, 35, 15, 1, 51, 41, 2, 25, 61]. These models have been used for task planning, natural language instruction following, and scene understanding. However, most applications focus on enabling robots to execute tasks more effectively using language for task specification and clarification [52, 13], rather than using these models to teach humans. In contrast to prior work that uses foundation models to improve robot task execution, CANINE inverts this paradigm by using VLMs and LLMs as reasoning tools to coach humans while robots provide physical interaction.

III. COACHING NEEDS AND DESIGN PRINCIPLES

To understand the coaching needs of VI users learning to use robot guide dogs, we conducted a formative study with an expert coach (2.5 years experience with robotic guiding technology) and one VI participant (2 years experience with robot guide dogs). Fig. 2 summarizes key results.

The formative study used a semi-structured interview built around a questionnaire organized into six sections: background and prior assistive-technology training; robot-guide-dog experience; coaching and feedback preferences; learning preferences; safety and trust; and future use. We analyzed responses for recurring themes, including limited opportunities for repeated expert coaching and a preference for error-specific, actionable feedback. These themes directly informed three design choices: decomposing practice into sub-skills, generating personalized terminal feedback, and adapting the training sequence to each learner’s observed weaknesses. The full instrument is provided in Appendix P.

The core challenge: understanding $\neq$ execution. General instructions are insufficient because learners often understand the high-level goals but fail to execute the corresponding physical actions, leading to diverse error patterns. Effective coaching must therefore go beyond explaining “what” to do




Fig. 2: **Expert coach insights from formative study.** Analysis reveals common learner challenges, effective coaching strategies (timing and content), and design implications for automated coaching systems.

and adaptively diagnose “how” the execution failed, modifying feedback to correct specific physical errors.

Interactive tasks with timing constraints are most difficult.

Tasks requiring active physical coordination, such as navigating doorways, particularly under time pressure, were identified as the most challenging. We thus select doorway navigation as our focal task.

Automated coaching enables autonomy and scalability.

Both participants confirmed the necessity of dedicated coaching sessions for effective learning. In addition, visually impaired users seek autonomy and to reduce dependence on human coaches over time. While structured coaching is essential for effective learning, automating adaptive coaching would scale the process and provide independence to practice without constant supervision.

IV. PROBLEM FORMULATION

We formulate robotic assistance coaching as a two-part problem: a *target task* where humans learn to coordinate with robotic assistants, and a *teaching task* where a robot coach optimizes the training process.

Target Task: Collaborative Navigation. We model the target task as a cooperative Markov Game $\mathcal{M} = (S, A^1, A^2, T, R, \gamma)$. For doorway navigation, the state S encodes the joint physical configuration: human pose relative to the robot and door, robot pose, door state (closed/open/opening), etc. Actions comprise continuous human control inputs A^1 (body velocity, hand forces on door handle) and robot actions A^2 (base velocity). Transitions T govern the coupled non-linear dynamics, requiring tight coordination (e.g., maintaining leash contact while opening the door). The cooperative reward R incentivizes safe, efficient task completion (reaching the goal beyond the doorway) while penalizing collisions and loss of connection. The objective is for the human to learn a policy π^H that complements the fixed robot policy π^R to maximize expected return.

Teaching Task: Hierarchical Coaching. We decompose the target task into discrete sub-skills Ω (e.g., for doorway navigation: Navigate to Door, Open Door, and Enter Room) and formulate coaching as two POMDPs:

a) Inter-skill POMDP: Curriculum Sequencing:

At the episode level, the coach solves $\mathcal{M}_{inter} = (S_{skill}, A_{skill}, O_{perf}, T_{skill}, R_{learning})$. The state S_{skill} represents the learner’s latent proficiency across sub-skills (a vector of continuous proficiency values, one per sub-skill). Actions $A_{skill} = \Omega$ select the next sub-skill to practice. Observations O_{perf} are noisy performance metrics for the practiced sub-skill (e.g., completion time, success/failure).

Transitions T_{skill} model how proficiency evolves with practice. The reward $R_{learning}$ incentivizes practicing sub-skills where the learner has low proficiency, thereby maximizing overall learning progress. The coach cannot directly observe proficiency and must maintain a belief over S_{skill} based on observed performance.

b) Intra-skill POMDP: Feedback Generation: Within a chosen sub-skill ω , the coach solves $\mathcal{M}_{intra}^\omega = (S_{task}, A_{coach}, O_{intra}, T_{task}, R_{guide})$. The state S_{task} captures the physical system state during task execution (human pose, robot pose, environment state, task progress) as well as latent cognitive states (whether errors stem from spatial misalignment or cue misinterpretation). Actions A_{coach} consist of post-episode verbal feedback and robot parameter adjustments for the next attempt. Observations O_{intra} are visual observations from a chest-mounted camera plus optional user verbal feedback. Optional verbal feedback is captured through a microphone, transcribed to text, and passed to the coaching model together with the visual episode summary. The reward R_{guide} incentivizes feedback that corrects errors and ensures safety. The coach must infer the latent cognitive state from observations to provide guidance.

Exact solutions are intractable due to continuous state spaces, latent human cognition, and unknown learning dynamics. We therefore employ approximate solutions tailored to each level: Bayesian knowledge tracing with greedy action selection for the inter-skill POMDP, and a structured VLM/LLM process for the intra-skill POMDP.

V. CANINE

A. Overview

CANINE is a two-level coaching system that enables visually impaired users to learn effective coordination with robot guide dogs. We decompose the overall navigation task into discrete sub-skills (e.g., for the doorway navigation task: Navigate to Door, Open Door, and Enter Room) that can be practiced independently. As the user navigates with the robot, the system observes behavior through a chest-mounted camera capturing a first-person view. The *inter-skill* module dynamically sequences the curriculum by estimating proficiency across sub-skills and selecting which sub-skill to practice next. The *intra-skill* module generates personalized feedback for each practice episode, leveraging foundation models to diagnose specific errors (e.g., distinguishing spatial misalignment from cue misinterpretation), provide corrections, and adjust robot behaviors. To minimize cognitive load during safety-critical maneuvers, feedback is delivered after each episode completes.




Fig. 3: **Overview of CANINE.** CANINE employs a two-level coaching strategy. The **inter-skill** coaching (up-right) tracks proficiency across sub-skills using knowledge tracing and selects the sub-skill to practice next. The **intra-skill** coaching (down-right) takes in video observations and generates coaching instructions for the selected sub-skill.

B. Inter-skill Coaching

The inter-skill layer uses the POMDP formulation for adaptive curriculum selection: the learner’s proficiency across sub-skills is latent, each practice episode provides a noisy performance observation, and the coach chooses the next sub-skill to practice based on the current belief.

Sub-skill Decomposition. For doorway navigation, we decompose the task into three sub-skills: (1) Navigate to Door, (2) Open Door, and (3) Enter Room. Each sub-skill ω represents a coherent behavior that can be practiced and evaluated independently by configuring the robot to start from the appropriate initial state (e.g., for “Open Door”, the robot positions the user in front of a closed door).

Knowledge Tracing. The learner’s true proficiency for each sub-skill is latent and cannot be directly observed. The coach maintains a belief b^t over proficiency, modeling each skill with a Gaussian distribution $\mathcal{N}(\mu_\omega^t, (\sigma_\omega^t)^2)$, extending Bayesian knowledge tracing [62, 49] to continuous performance metrics.

When the learner practices sub-skill ω , the coach observes a performance measurement y_t (e.g., completion time). The coach maintains a Gaussian belief over the learner’s latent expected performance, $p(\mu_\omega^t) = \mathcal{N}(\mu_\omega^t, (\sigma_\omega^t)^2)$. To model gradual skill drift, the prior variance is first inflated by adding process noise, $(\sigma_\omega^t)^2 \leftarrow (\sigma_\omega^{t-1})^2 + \sigma_{\text{proc}}^2$. The observation y_t is then incorporated via a Bayesian update under a linear-Gaussian observation model. The posterior mean μ_ω^t is mapped to a normalized proficiency score $p_\omega^t \in [0, 1]$ by linearly interpolating between novice and expert baseline performance, $p_\omega^t = \text{clip}\left(\frac{\mu_{\text{novice}} - \mu_\omega^t}{\mu_{\text{novice}} - \mu_{\text{expert}}}, 0, 1\right)$, where μ_{expert} and μ_{novice} denote expected completion times for expert and novice performance, respectively. This mapping assigns proficiency 1.0 to expert-level performance, 0.0 to novice-level performance, with linear interpolation between.

Curriculum Policy. At each step, the system selects the sub-skill with the lowest estimated proficiency for practice. This greedy policy prioritizes the learner’s weakest areas, approximating the optimal curriculum under the assumption that practicing weak skills yields the highest learning gain. This gives a concrete curriculum policy for the inter-skill decision process: the belief tracker estimates which proficiency state is weakest, and the selected sub-skill becomes the next coaching

action. The curriculum adapts dynamically as proficiency estimates evolve. Initial proficiency beliefs are personalized using each participant’s initial practice trials: the first observations initialize the participant-specific belief means, while expert and novice reference performance from a pilot study set the normalization endpoints and tracker hyperparameters. Thus, the pilot data calibrates the scale of the tracker, but each participant’s curriculum is initialized from their own measured performance in initial practice trials.

C. Intra-skill Coaching

Within a selected sub-skill, the coach faces the partial-observation problem described in Section IV: the chest-mounted camera reveals the user’s behavior and task progress, but not the underlying cause of an error. The coach must infer the latent cognitive states from observations to provide appropriate guidance. Accordingly, our intra-skill pipeline converts visual observations into symbolic state histories, summarizes the episode, diagnoses likely latent error causes, and then produces both verbal feedback and robot-parameter adjustments. For the doorway navigation task in this paper, we use terminal (post-episode) feedback to avoid interfering with time-critical safety behaviors during navigation.

Let $o_{1:T}$ denote the raw visual observation sequence (RGB images) from a chest-mounted camera during a practice episode. We use a structured four-stage process that processes visual observations, infers user states, and generates coaching actions:

$$o_{1:T} \xrightarrow{f_{\text{frame}}} h_{1:T} \xrightarrow{f_{\text{time}}} \tau \xrightarrow{f_{\text{coach}}} c \xrightarrow{f_{\text{param}}} u, \quad (1)$$

where h_t is the symbolic state representation at time t , τ is the episode summary, c is the coaching action (verbal instructions), and u are robot parameter adjustments to alter the behavior of the robot. The intermediate history $h_{1:T}$ and summary τ form the coach’s episode-level belief context, while c and u are the resulting coaching actions exposed to the user and robot. These components execute once per episode (after time T), ensuring the robot’s real-time control depends only on its low-level controller, not on LLM latency. The per-stage I/O details and prompts can be found in Appendix.

Frame-level State Extraction (f_{frame}). This module takes as input the raw visual observation sequence $o_{1:T}$ and produces




Fig. 4: **Illustration of the summarized timeline.** The timeline aggregates per-frame observations sampled every 0.5 s into a structured episode-level summary.

symbolic state representations $h_{1:T}$. Given the observation trajectory, we uniformly sample key frames and apply $h_t = f_{\text{frame}}(o_{1:T}, t)$ to each timestamp t . In practice, we capture an image every 0.5 seconds and feed it to the VLM. The output h_t encodes: user physical state (approaching, positioned, interacting), robot state, environment state (clear, cluttered, hazardous), task progress (not started, in progress, complete, failed), and safety flags (collision risk, blocked path). This structured representation extracts key information from the raw input to allow the LLM to reason about the user’s behavior explicitly using symbolic representation.

Timeline Summarization (f_{time}). This module takes as input the frame-level states $h_{1:T}$ and produces a compact episode summary τ . While the frame sequence captures observable physical states, latent cognitive components (error causes, mental model) remain hidden. Applying $\tau = f_{\text{time}}(h_{1:T})$ produces a structured summary containing: (1) episode statistics (duration, terminal state, success/failure), (2) performance metrics (efficiency, safety scores), and (3) a discrete event sequence with timestamps (task start, progress milestones, failures, safety violations). One example can be found in Fig. 4. This makes behavioral patterns such as hesitation, repeated errors, or unsafe behaviors explicit and observable.

Feedback Generation (f_{coach}). This module takes as input the episode summary τ and optional user verbal feedback, and produces a coaching action c containing an error diagnosis and a verbal instruction. We use an LLM to infer the likely latent cognitive state (i.e., the error cause) from the observable evidence: it analyzes the summary and user feedback to diagnose the source of errors (e.g., spatial misalignment, cue misinterpretation), then generates the coaching action. User feedback helps disambiguate the visual summary, and the diagnosis schema remains open-ended rather than limited to the example error causes. The action includes: an error diagnosis and an actionable verbal instruction tailored to the inferred error cause (e.g., “position yourself closer to the door before grasping the handle”).

Robot Adaptation (f_{param}). This module takes as input the coaching action c (which contains the error diagnosis and verbal instruction) and produces robot behavior adjustments u . Based on the diagnosed error cause, this module generates specific parameter updates (e.g., if the error is spatial misalignment due to positioning too far, adjust the robot’s stopping position to a spot closer to the door or change the facing angle to make it easier for the human to grasp the handle). These




Fig. 5: **Robot guide dog hardware setup.** Unitree Go2 quadruped robot equipped with Jetson AGX Orin for onboard computation, Hesai JT16 LiDAR for navigation, and chest-mounted camera on the user for first-person view analysis.

parameter adjustments implement scaffolding: the robot adapts task difficulty to match the learner’s proficiency, making the task easier when errors are detected and gradually increasing difficulty as proficiency improves.

To reduce hallucinations, every stage is prompted to emit a constrained JSON object matching a predefined schema; outputs are parsed and validated before being passed downstream, with invalid or missing fields rejected and regenerated or replaced by conservative fallbacks. The frame stage restricts open-ended visual reasoning to explicit state fields, while later stages reason over structured summaries rather than raw free-form video descriptions.

We instantiate f_{frame} with a VLM that processes images, while f_{time} , f_{coach} , and f_{param} use LLM with symbolic representations. In practice, we use GPT-5.1 for all of them. We use these models zero-shot with fixed prompts and JSON schemas, without fine-tuning on our study data.

Example. Consider an episode where the user struggles to open the door efficiently. The VLM (f_{frame}) extracts frame-level states: “t=0s: user approaching, door closed”, “t=5s: user positioned far away from the door, hand searching”, “t=12s: user grasped handle, door opening”, “t=18s: door fully opened”. The timeline summarizer (f_{time}) aggregates this into: “Duration: 18s, Outcome: Success, Timeline: t=0s to 5s, the user hesitated at the door for 3 seconds, t=5s to 12s, the user grasped the handle and opened the door, ...”. The coaching policy (f_{coach}) identifies the inefficiency cause (spatial misalignment: positioned too far) and generates: “Position yourself closer to the door before grasping the handle—about arm’s length away”. Finally, f_{param} determines the appropriate robot adjustment: move stopping position closer to the door.

D. Hardware Platform

The robot guide dog system (shown in Fig. 5) consists of a Unitree Go2 quadruped robot equipped with a Jetson AGX Orin for onboard computation and a Hesai JT16 LiDAR for environment sensing and obstacle detection. The robot provides autonomous navigation using a planner based on Xiao et al. [65] that guides users to specified goals while avoiding obstacles. Users interact with the robot through a rigid leash attachment. A DJI Action 6 camera mounted on the user’s chest captures a first-person view for VLM-based behavior




Fig. 6: **Adaptive curriculum sequencing achieves faster learning.** Violin plots show the distribution of teaching actions needed to reach mastery for different numbers of sub-skills. Adaptive curriculum sequencing achieves faster learning.

analysis, enabling the coaching system to observe the current state. A microphone records the user feedback. Generated instructions are delivered through a speaker.

VI. EXPERIMENT

We structure our experiment around the following research questions.

- RQ1 (Component Feasibility): Are components of CANINE technically feasible for coaching? (Section VI-A)
- RQ2 (Human Subject Evaluation): Does CANINE achieve better objective learning outcomes and subjective user experience? (Section VI-B)
- RQ3 (Learning Retention): Do skills learned with CANINE persist over time? (Section VI-C)
- RQ4 (Case Study with a Visually Impaired User): What are the feasibility insights and additional design considerations with a visually impaired user? (Section VI-D)

We address these questions across four studies. Study 1 validates the system’s technical components, showing that our adaptive curriculum and decomposed reasoning pipeline outperform baselines in simulation and human judgment benchmarks. Study 2 demonstrates that CANINE significantly improves navigation performance and user perceived usefulness compared to general instruction in a controlled user study. Study 3 confirms that these learned skills persist after two weeks. Finally, Study 4 provides a case study with a visually impaired user, delivering design insights for deployment and understanding of extensibility to a handover task.

A. Study 1: Feasibility Evaluation

To address RQ1, we evaluated the technical feasibility of CANINE’s key components.

1) *Simulated User Evaluation*: To evaluate adaptive curriculum sequencing, we developed a simulated learner following [12, 62, 46, 32] capturing key learning dynamics (Power Law of Practice, slips/guesses, forgetting, fatigue; full details in Appendix A). This allows us to compare curriculum policies across many learner trajectories before deploying them in the physical study. We compared three curriculum strategies




Fig. 7: **Frame analysis accuracy comparison.** Preference rates of VLM-generated descriptions versus human annotations across different language models (GPT-5.1, Claude-4.5-Sonnet, Gemini-3.0-Pro and Qwen3-VL).

TABLE I: **Human evaluation on coaching feedback quality.** Preference rate indicates percentage of decisive preference. Other metrics are mean ratings on 1-5 Likert scales (higher is better). Latency is average time to generate feedback.

Metric	VLM	Structured VLM	Ours
Overall Preference Rate (%) ↑	16.3	32.5	51.2
Helpfulness ↑	2.99	3.49	3.66
Factuality ↑	3.26	3.50	3.59
Specificity ↑	3.18	3.95	3.91
Latency (s) ↓	44.36	42.15	20.62

across 500 diverse learner profiles: **Uniform** (uniform sub-skill selection), **Greedy** (use the recent trial performance only to select sub-skill), and **Adaptive** (ours, maintains a cumulative proficiency model). As shown in Fig. 6, adaptive curriculum sequencing achieved faster learning compared to baselines.

2) *Frame Description Quality*: We evaluated whether VLMs can accurately extract high-quality state descriptions from video frames by comparing VLM predictions against human annotations on 100 frames. Accurate frame descriptions are important because they are the evidence the coach uses to summarize an episode and diagnose causes of errors. Independent human judges chose their preferred description through pairwise comparisons (Fig. 7) for the same frame. Given the evaluation results, we chose GPT-5.1 for our experiments.

3) *Comparison of Different VLM Pipelines*: While directly processing videos with VLMs could simplify the architecture, it leads to prolonged waiting times that disrupt interactive coaching and fails on longer episodes due to context length constraints. To validate our decomposed approach, we compared it against two alternatives on 10 navigation videos: (1) **VLM**: direct video prompting, (2) **Structured VLM**: full video input with structured prompt similar to CANINE, and (3) **Ours**: frame-level decomposition. This comparison evaluates whether the staged representation used by CANINE preserves enough episode context to generate useful coaching while reducing latency. Eight human judges evaluated feedback quality through overall preference rankings, with additional Likert-scale ratings on three supporting dimensions: *Helpfulness* (overall utility for learning), *Factuality* (accuracy of observations), and *Specificity* (level of actionable detail) (de-




Fig. 8: **Navigation task setup and experimental environment.** We show the three sub-skills of door navigation: Navigate to Door, Open Door, and Enter Room. **Top:** human behavior. **Bottom:** visualization of the navigation map.

tails in Appendix D. We report four metrics in Table I: Overall Preference Rate is the percentage of decisive judge preferences won by each method; Helpfulness, Factualty, and Specificity are mean 1–5 Likert ratings, where higher is better; and Latency is the average wall-clock time in seconds from video input to generated feedback on the same hardware and API endpoints. Instructions generated by CANINE were preferred in 51.2% of comparisons while achieving faster generation as shown in Table I. Having validated the technical components, we next evaluate the complete system with human participants. The mean end-to-end latency of our decomposed pipeline was 20.62 s, which supports terminal feedback between practice episodes; the full latency breakdown is provided in the supplementary material.

B. Study 2: Human Subject Evaluation

To address RQ2, we evaluated CANINE through a two-stage approach: a controlled between-subject study with blindfolded sighted participants as a proxy followed by an exploratory case study with a visually impaired participant (Section VI-D).

Task Description. Participants navigated a doorway using the robot guide dog (Fig. 8). The task requires mastering the three sequential sub-skills defined in Section V: (1) Navigate to Door, (2) Open Door, and (3) Enter Room.

Study Design. We recruited 22 blindfolded sighted participants ($M = 27, SD = 3.66$, 5 females): 2 for pilot, 20 for the main between-subjects study comparing CANINE against a baseline providing task-level verbal instructions without sub-skill decomposition, personalized feedback, or adaptive sequencing (full details in Appendix G). We use blindfolded participants to obtain controlled quantitative evidence and interpret this population as a proxy rather than as a substitute for statistically powered evaluation with VI users.

1) *Pilot Study:* We conducted a pilot study with 2 participants to establish optimal performance, identify common failure modes, and refine the experimental protocol.

2) *Main Study:* Following the pilot study, we recruited 20 blindfolded sighted participants for the main study, randomly assigned to one of the two teaching methods (10 participants per group). All participants gave written informed consent and were compensated for their time. None of the participants had




Fig. 9: **Objective performance.** a) Performance improvement across sub-skills. b) Final completion time in evaluation trials.




Fig. 10: **Subjective evaluation results.** Participants rated CANINE as having significantly higher Perceived Usefulness (PU) compared to general instruction.

prior experience using the robot guide dog. The study was IRB-approved.

Experimental Protocol. Each participant completed a single session lasting approximately 40 minutes. The session consisted of five phases: Introduction (5 min), Practice (2 trials, 7 min), Training (11 min), Evaluation (2 trials, 7 min), and Survey (10 min). During the Training phase, the baseline group performed 3 full end-to-end trials. The CANINE group practiced up to 6 sub-skill training episodes in total selected by the adaptive curriculum, with the session ending early if the system determined all sub-skills were mastered.

Evaluation Metrics. Objective: Time to complete route (in seconds), averaged across the two evaluation trials. Subjective: NASA Task Load Index (NASA-TLX) [21], System Usability Scale (SUS) [4], Customized Questionnaires for Perceived Usefulness (PU) and Perceived Ease of Use (PEU) (full questionnaire in Appendix D).

3) *Results on Objective Performance:* Coaching significantly improves learning rates and final performance for interaction-intensive sub-skills As shown in Fig. 9 (all comparisons use Welch’s t-test; all times in seconds), for the “Navigate to Door” sub-skill, no significant difference was observed. This is primarily because the task involves straight-forward following of the robot’s guidance without requiring complex coordination or spatial reasoning. Thus, we focus our analysis on the interactive parts of the task. Coaching is particularly effective for these coordination tasks. For the “Open Door” task, the experimental group showed significantly greater improvement from pre-training to post-training (improvement: $M = 8.64, SD = 0.99$ vs. baseline $M = 4.53, SD = 1.00$), $t(18) = -2.78, p = .012, g = -1.19$, resulting in faster final completion times ($M = 11.07, SD = 0.89$) compared to baseline ($M = 14.54, SD = 1.23$), $t(18) = 2.12, p = .045, g = 0.93$. Similarly, for “Enter Room,” the CANINE group achieved greater improvement ($M = 6.96, SD = 1.05$ vs. $M = 2.27, SD = 1.37$), $t(18) =$




Fig. 11: **Qualitative examples of generated instructions.** Top-left shows the image captured by the chest-mount camera.




Fig. 12: **User ratings of CANINE components.** Using a 1 - 7 Likert scale (1 = not helpful, 7 = very helpful).

$-2.61, p = .019, g = -1.11$, leading to faster completion times ($M = 11.36, SD = 0.52$ vs. $M = 14.42, SD = 0.76$), $t(18) = 3.00, p = .006, g = 1.35$.

4) *Results on Subjective Experience:* Participants rated the coaching system as more useful than baseline instruction, valuing its targeted feedback and adaptive structure. As shown in Fig. 10, participants rated CANINE as having significantly higher Perceived Usefulness ($M = 13.10, SD = 1.29$ vs. $M = 10.80, SD = 2.82$), $p = .036, g = 1.05$. General usability metrics (SUS, Workload) showed no significant differences; we hypothesize that these general measures primarily capture participants' experience with the overall robot guide dog system rather than the coaching system specifically.

Thus, we examined which coaching features participants found most helpful (Fig. 12). After training, participants rated each CANINE component on a 1-7 Likert scale, where 1 indicated "not helpful" and 7 indicated "very helpful." The most appreciated features were: (1) dedicated sub-skill training, (2) personalized feedback, and (3) adaptive curriculum. They reported that they learned skills (e.g., "Try to align with the robot dog direction during movement.") from personalized feedback. This pattern is consistent with the component studies: the simulated-learner results suggest that adapting practice to estimated proficiency improves curriculum selection, while the pipeline comparison suggests that the decomposed feedback generator can produce useful episode-level coaching. In the full system, participants most clearly

noticed the human-facing parts of this design: targeted sub-skill practice, personalized feedback, and adaptive sequencing. Robot-side adaptation acted mainly as conservative scaffolding through task-difficulty adjustments, which may explain why blindfolded participants found it less salient than the feedback and curriculum components.

Following the main experiment, we conducted a post-hoc comparison with the same 10 participants from the CANINE group, where a human expert coached them on the same task. When asked for their preference, 70% of participants favored the human expert, citing two key advantages: (1) the ability to provide physical corrections (e.g., manually adjusting body posture), and (2) the delivery of concurrent, real-time feedback that addresses questions and offers instructive guidance during execution. This preference also clarifies the role of our terminal-feedback design: it supports reflection between attempts, but does not correct body pose or timing while the user is acting. Conversely, two participants preferred the robot coach, valuing its consistent, standardized feedback and the sense of autonomy and privacy it offered. One participant remained neutral. Qualitative examples can be found in Fig. 11.

C. Study 3: Retention Analysis

To address RQ3, we conducted follow-up sessions with a subset of participants ($N=5$) from the CANINE group in Study 2, two weeks after their training. This within-subject analysis provides preliminary evidence of whether skills learned with CANINE persist over time. Because only five participants returned, we interpret the retention study as preliminary longitudinal evidence rather than a definitive estimate of long-term retention. Participants completed both the same routes from the evaluation phase of Study 2 and novel routes with different layouts to assess temporal retention.

1) *Results:* Skills learned with coaching persist over two weeks, with complex coordination skills showing continued improvement and simpler skills remaining stable.

As illustrated in Fig. 13, performance at the two-week retention session varied by sub-skill. Participants demonstrated




Fig. 13: **Skill retention.** Bars show the mean retention change for each skill in second.

substantial improvement in the “Open Door” task. In contrast, “Enter Room” and “Navigate to Door” showed minor performance degradation, indicating these skills remained relatively stable. Interestingly, 60% of participants identified “Open Door” as the most challenging aspect of the task. We hypothesize that the observed improvement in this difficult sub-skill reflects the incubation effect [53], where complex motor skills continue to consolidate during rest periods as the brain integrates and refines learned patterns.

D. Study 4: Case Study with Visually Impaired User

To address RQ4, we conducted an exploratory case study with a totally blind participant (Male, 70, with no residual vision) to identify design considerations and evaluate system usability for navigation and extensibility to robotic handover tasks (details in Appendix L and Appendix M).

1) *Convergence with Study 2:* Training with CANINE yielded clear learning progression, with the participant’s positive feedback (“The feedback is very accurate.”). The participant completed the same navigation coaching-session structure as the controlled-study participants and improved “Open Door” completion time from 22.4 s to 9.1 s, a trend consistent with Study 2. This session followed the same structure as Study 2, with two initial practice trials, up to six coached sub-skill episodes, and two evaluation trials. Preferences aligned remarkably with Study 2: the participant independently ranked sub-skill practice, personalized feedback, and adaptive curriculum as the top features, and also confirmed “Open Door” as the most challenging sub-skill.

2) *VI-Specific Designs:* However, the study revealed critical distinctions. Unlike blindfolded proxies who may rely on residual visual-spatial mental models, the VI participant emphasized a need for *multimodal grounding*, specifically haptic cues to supplement verbal feedback and provide spatial information that language alone cannot convey. In addition, the VI user requested greater control over the coaching system, specifically the ability to pause, replay, and skip feedback to manage cognitive load.

3) *Extensibility to Robotic Handover:* We applied CANINE to a robotic handover task [58, 8, 30] as shown in Fig. 14. In this task, a robotic arm hands an object to a human, requiring human hand positioning and grasp timing without vision. We use this study as an instantiation showing that the coaching pipeline can be configured for another interaction task. Cross-task comparison revealed task-dependent coaching needs: while terminal feedback succeeded for navigation, the participant required *concurrent* (real-time) feedback for handover to adjust hand position during execution. This suggests




Fig. 14: **Experimental setup for the handover task.** Human practices receiving objects from a robot manipulator.

coaching systems must adapt not only feedback content but also its timing based on task dynamics. We treat this handover result as an extensibility demonstration rather than a full evaluation: it shows that the architecture can be instantiated for a different interactive assistance task, but does not establish task-general effectiveness without a larger controlled study.

VII. CONCLUSION

We presented CANINE, an automated coaching system designed to teach visually impaired users how to navigate with a robot guide dog. CANINE dynamically sequences sub-skills and generates personalized, context-aware feedback. By leveraging foundation models and knowledge tracing, our system delivers a dynamic curriculum and actionable guidance. Evaluations with simulated agents and human participants demonstrate that CANINE accelerates human learning in complex coordination tasks compared to standard instruction, with skill retention persisting over two weeks.

Our current approach relies on expert-defined reasoning schemas to ensure safety and interpretability, a necessary trade-off that limits fully open-ended adaptation. Furthermore, while coaching is automated, the system relies on expert-defined skill decomposition, and physical safety and task resetting still require human oversight. The current study has four limitations: terminal feedback avoids cognitive overload during safety-critical navigation but cannot support in-the-moment correction; the chest-mounted camera can suffer from occlusion, blur, or viewpoint limitations; robot-side adaptation was evaluated as conservative scaffolding rather than through a full factorial ablation; and the visually impaired case study plus five-participant retention study provide preliminary evidence rather than population-level validation. Longer deployments are needed to study mutual adaptation over time. Future work will focus on longitudinal deployment with the visually impaired community to validate real-world efficacy, developing adaptive feedback mechanisms that transition from terminal to concurrent guidance based on task demands, using multi-camera sensing or explicit body-orientation guidance, automating skill discovery to reduce reliance on expert-defined sub-skill decomposition, conducting ablations of system components, including human-selected curriculum baselines, and extending CANINE to more complex interactive navigation tasks such as elevator use, escalators, and navigation in crowded environments.

ACKNOWLEDGMENTS

This research is supported in part by the National Research Foundation, Singapore under its AI Singapore Programme under AISG Award AISG2-PhD-2022-01-036[T] by Ministry of Education Academic Research Fund (grant no. T1-251RES2406), and by Ministry of Digital Information under its AIVP program. We thank the anonymous reviewers for their valuable feedback, and our participants for their time and insights. We also thank Shaojun Cai for his support and discussions that helped shape this work.

REFERENCES

- [1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do As I Can, Not As I Say: Grounding language in robotic affordances. In *Conference on Robot Learning (CoRL)*, 2022.
- [2] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0 : A vision-language-action flow model for general robot control. In *Robotics: Science and Systems (RSS)*, 2025.
- [3] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, et al. RT-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning (CoRL)*, 2023.
- [4] John Brooke. SUS: A “quick and dirty” usability scale. In *Usability Evaluation in Industry*, pages 189–194. Taylor & Francis, 1996.
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [6] Shaojun Cai, Ashwin Ram, Zhengtai Gou, Mohd Alqama Wasim Shaikh, Yu-An Chen, Yingjia Wan, Kotaro Hara, Shengdong Zhao, and David Hsu. Navigating real-world challenges: A quadruped robot guiding system for visually impaired people in diverse environments. In *CHI Conference on Human Factors in Computing Systems (CHI)*, 2024.
- [7] Shaojun Cai, Nuwan Janaka, Ashwin Ram, Janidu Shehan, Yingjia Wan, Kotaro Hara, and David Hsu. Navigation beyond wayfinding: Robots collaborating with visually impaired users for environmental interactions. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2026.
- [8] Wesley P Chan, Chris AC Parker, HF Machiel Van der Loos, and Elizabeth A Croft. A human-inspired object handover controller. *The International Journal of Robotics Research (IJRR)*, 32(8):971–983, 2013.
- [9] Yanbo Chen, Zhengzhe Xu, Zhuozhu Jian, Gengpan Tang, Liyunong Yang, Anxing Xiao, Xueqian Wang, and Bin Liang. Quadruped guidance robot for the visually impaired: A comfort-based approach. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2023.
- [10] Yu Chen, Shu Miao, Jing Ye, Gong Chen, Jianghua Cheng, Ketao Du, and Xiang Li. Upper-limb rehabilitation with a dual-mode individualized exoskeleton robot: A generative-model-based solution. *The International Journal of Robotics Research (IJRR)*, 2025.
- [11] Benjamin Clement, Didier Roy, Pierre-Yves Oudeyer, and Manuel Lopes. Multi-armed bandits for intelligent tutoring systems. *Journal of Educational Data Mining*, 7(2):20–48, 2015.
- [12] Albert T. Corbett and John R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction (UMUI)*, 4(4):253–278, 1994.
- [13] Yuchen Cui, Siddharth Karamcheti, Raj Palleti, Nidhya Shivakumar, Percy Liang, and Dorsa Sadigh. No, to the Right: Online language corrections for robotic manipulation via shared autonomy. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*. Association for Computing Machinery, 2023.
- [14] Anca D Dragan and Siddhartha S Srinivasa. A policy-blending formalism for shared control. *The International Journal of Robotics Research (IJRR)*, 32(7):790–805, 2013.
- [15] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. PaLM-E: An embodied multimodal language model. In *International Conference on Machine Learning (ICML)*, 2023.
- [16] Maximilian Du, Jennifer Grannen, Shuran Song, and Dorsa Sadigh. Will people enjoy a robot trainer? a case study with snooPie the pacerbot. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2026.
- [17] Roya Firoozi, Johnathan Tucker, Stephen Tian, Anirudha Majumdar, Jiankai Sun, Weiyu Liu, Yuke Zhu, Shuran Song, Ashish Kapoor, Karol Hausman, et al. Foundation models in robotics: Applications, challenges, and the future. *The International Journal of Robotics Research (IJRR)*, 44(5):701–739, 2025.
- [18] Deepak Gopinath, Xiongyi Cui, Jonathan DeCastro, Emily Sumner, Jean Costa, Hiroshi Yasuda, Allison Morgan, Laporsha Dees, Sheryl Chau, John Leonard, et al. Computational teaching for driving via multi-task imitation learning. *IEEE International Conference on*

Robotics and Automation (ICRA), 2025.

- [19] João Guerreiro, Daisuke Sato, Saki Asakawa, Huixu Dong, Kris M Kitani, and Chieko Asakawa. Cabot: Designing and evaluating an autonomous navigation robot for blind people. In *ACM ASSETS 2019 (ACM SIGACCESS Conference on Computers and Accessibility)*, 2019.
- [20] Ronald K. Hambleton and Hariharan Swaminathan. *Item Response Theory: Principles and Applications*. Evaluation in Education and Human Services. 1985.
- [21] S. G. Hart and Lowell E. Staveland. Development of nasa-tlx (task load index): Results of empirical and theoretical research. *Advances in Psychology*, 52:139–183, 1988.
- [22] Zhimin Hou, Cunjun Yu, David Hsu, and Haoyong Yu. Teachingbot: Robot teacher for human handwriting. *IEEE Robotics and Automation Letters (RA-L)*, 11(3): 2610–2617, 2026.
- [23] Hochul Hwang, Soowan Yang, Jahir S Monon, Nicholas A Giudice, Sunghoon I Lee, Joydeep Biswas, and Donghyun Kim. Guidenav: User-informed development of a vision-only robotic navigation assistant for blind travelers. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2026.
- [24] Shuhei Ikemoto, Heni Ben Amor, Takashi Minato, Bernhard Jung, and Hiroshi Ishiguro. Physical human-robot interaction: Mutual learning and adaptation. *IEEE Robotics & Automation Magazine (RAM)*, 19(4):24–35, 2012.
- [25] Bahar Irfan, Sanna-Mari Kuoppamäki, and Gabriel Skantze. Between reality and delusion: Challenges of applying large language models to companion robots for open-domain dialogues with older adults. *Autonomous Robots*, 49(1):9, 2025.
- [26] Rajat Kumar Jenamani, Tom Silver, Ben Dodson, Shiqin Tong, Anthony Song, Yuting Yang, Ziang Liu, Benjamin Howe, Aimee Whitneck, and Tapomayukh Bhattacharjee. FEAST: A flexible mealtime-assistance system towards in-the-wild personalization. In *Robotics: Science and Systems (RSS)*, 2025.
- [27] Matthew Jörke, Shardul Sapkota, Lyndsea Warkenthien, Niklas Vainio, Paul Schmiedmayer, Emma Brunskill, and James A. Landay. Gptcoach: Towards llm-based physical activity coaching. In *CHI Conference on Human Factors in Computing Systems (CHI)*, CHI '25, 2025.
- [28] Rie Kamikubo, Seita Kayukawa, Yuka Kaniwa, Allan Wang, Hernisa Kacorri, Hironobu Takagi, and Chieko Asakawa. Beyond omakase: Designing shared control for navigation robots with blind people. In *CHI Conference on Human Factors in Computing Systems (CHI)*, 2025.
- [29] Callie Y Kim, Christine P Lee, and Bilge Mutlu. Understanding large-language model (llm)-powered human-robot interaction. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2024.
- [30] Andras Kupcsik, David Hsu, and Wee Sun Lee. Learning dynamic robot-to-human object handover from human feedback. In Antonio Bicchi and Wolfram Burgard, editors, *Robotics Research*, volume 2 of *Springer Proceedings in Advanced Robotics*, pages 161–176. Springer, Cham, 2018.
- [31] Masaki Kuribayashi, Tatsuya Ishihara, Daisuke Sato, Jayakorn Vongkulbhisal, Karnik Ram, Seita Kayukawa, Hironobu Takagi, Shigeo Morishima, and Chieko Asakawa. Pathfinder: Designing a map-less navigation system for blind people in unfamiliar buildings. In *CHI Conference on Human Factors in Computing Systems (CHI)*, 2023.
- [32] Andrew S. Lan, Andrew E. Waters, Christoph Studer, and Richard G. Baraniuk. Sparse factor analysis for learning and content analytics. *Journal of Machine Learning Research (JMLR)*, 15(1):1959–2008, 2014.
- [33] Chenzui Li, Xi Wu, Tao Teng, Sylvain Calinon, and Fei Chen. Towards robo-coach: Robot interactive stiffness/position adaptation for human strength and conditioning training. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [34] Yanan Li, Keng Peng Tee, Wei Liang Chan, Rui Yan, Yuanwei Chua, and Dilip Kumar Limbu. Continuous role adaptation for human–robot shared control. *IEEE Transactions on Robotics (T-RO)*, 31(3):672–681, 2015.
- [35] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
- [36] Yu Lu, Yang Pian, Penghe Chen, Qinggang Meng, and Yunbo Cao. RadarMath: An intelligent tutoring system for math education. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [37] Mariah Lynn Schrum, Srijan Srivatsa, Laporsha Dees, Evelyn Dixon, Patricio Reyes Gomez, Deepak Gopinath, Emily Sarah Sumner, Guy Rosman, and Tiffany L. Chen. Skill Modulates Coaching Language in Embodied Motor Learning. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA)*, 2026.
- [38] Mina Mizukoshi, Mana Kondo, and Toru Nakamura. Evaluation of the potential suitability of guide dog candidates by continuous observation during training. *Journal of Veterinary Behavior*, 3(5):193–198, 2008.
- [39] Dean D Molinaro, Keaton L Scherpereel, Ethan B Schonhaut, Georgios Evangelopoulos, Max K Shepherd, and Aaron J Young. Task-agnostic exoskeleton control via biological joint moment estimation. *Nature*, 635(8038): 337–344, 2024.
- [40] Stefanos Nikolaidis, Anton Kuznetsov, David Hsu, and Siddhartha Srinivasa. Formalizing human-robot mutual adaptation via a bounded memory based model. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2016.
- [41] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An

- open-source generalist robot policy. In *Robotics: Science and Systems (RSS)*, 2024.
- [42] OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [43] Akhil Padmanabha, Jessie Yuan, Tanisha Mehta, Rajat Kumar Jenamani, Eric Hu, Victoria de León, Anthony Wertz, Janavi Gupta, Ben Dodson, Yunting Yan, Carmel Majidi, Tapomayukh Bhattacharjee, and Zackory Erickson. WAFFLE: A wearable approach to bite timing estimation in robot-assisted feeding. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2026.
- [44] Joseph Psotka, Leonard Daniel Massey, and Sharon A. Mutter, editors. *Intelligent Tutoring Systems: Lessons Learned*. 1988.
- [45] Peizhu Qian, Filip Bajraktari, Carlos Quintero-Peña, Qingxi Meng, Shannan K. Hamlin, Lydia E. Kavradi, and Vaibhav Unhelkar. ASTRID: A robotic tutor for nurse training to reduce healthcare-associated infections. In *Robotics: Science and Systems (RSS)*, 2025.
- [46] Anna N. Rafferty, Emma Brunskill, Thomas L. Griffiths, and Patrick Shafto. Faster teaching by pomdp planning. In *Artificial Intelligence in Education*, 2011.
- [47] Siddharth Reddy, Anca D. Dragan, and Sergey Levine. Shared autonomy via deep reinforcement learning. In *Robotics: Science and Systems (RSS)*, 2018.
- [48] Leonel Rozo, Sylvain Calinon, Darwin G Caldwell, Pablo Jimenez, and Carme Torras. Learning physical collaborative robot behaviors from human demonstrations. *IEEE Transactions on Robotics (T-RO)*, 32(3):513–527, 2016.
- [49] Thorsten Schodde, Kirsten Bergmann, and Stefan Kopp. Adaptive robot language tutoring based on bayesian knowledge tracing and predictive decision-making. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2017.
- [50] Mario Selvaggio, Marco Cognetti, Stefanos Nikolaidis, Serena Ivaldi, and Bruno Siciliano. Autonomy in physical human-robot interaction: A brief survey. *IEEE Robotics and Automation Letters (RA-L)*, 6(4):7989–7996, 2021.
- [51] Dhruv Shah, Ajay Sridhar, Nitish Dashora, Kyle Stachowicz, Kevin Black, Noriaki Hirose, and Sergey Levine. ViNT: A foundation model for visual navigation. In *Conference on Robot Learning (CoRL)*, 2023.
- [52] Lucy Xiaoyang Shi, brian ichter, Michael Robert Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James Tanner, Anna Walling, Haohuan Wang, Niccolo Fusai, Adrian Li-Bell, Danny Driess, Lachy Groom, Sergey Levine, and Chelsea Finn. Hi Robot: Open-ended instruction following with hierarchical vision-language-action models. In *International Conference on Machine Learning (ICML)*, 2025.
- [53] Ut Na Sio and Thomas C. Ormerod. Does incubation enhance problem solving? a meta-analytic review. *Psychological bulletin*, 135 1:94–120, 2009.
- [54] Megha Srivastava, Erdem Biyik, Suvir Mirchandani, Noah Goodman, and Dorsa Sadigh. Assistive teaching of motor control tasks to humans. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [55] Megha Srivastava, Noah Goodman, and Dorsa Sadigh. Generating language corrections for teaching physical control tasks. In *International Conference on Machine Learning (ICML)*, 2023.
- [56] Megha Srivastava, Reihaneh Iranmanesh, Yuchen Cui, Deepak Gopinath, Emily Sarah Sumner, Andrew Silva, Laporsha Dees, Guy Rosman, and Dorsa Sadigh. Shared autonomy for proximal teaching. In *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2025.
- [57] Saiying Steenbergen-Hu and Harris Cooper. A meta-analysis of the effectiveness of intelligent tutoring systems on k–12 students’ mathematical learning. *Journal of Educational Psychology*, 105(4):970–987, 2013.
- [58] Kyle Strabala, Min Kyung Lee, Anca Dragan, Jodi Forlizzi, Siddhartha S Srinivasa, Maya Cakmak, and Vincenzo Micelli. Toward seamless human-robot handovers. *Journal of Human-Robot Interaction (JHRI)*, 2(1):112–132, 2013.
- [59] Merryanna L. Swartz and Masoud Yazdani, editors. *Intelligent Tutoring Systems for Foreign Language Learning: The Bridge to International Communication*, volume 80 of *NATO ASI Series F: Computer and Systems Sciences*. Springer, 1992.
- [60] Ran Tian, Masayoshi Tomizuka, Anca D. Dragan, and Andrea V. Bajcsy. Towards modeling and influencing the dynamics of human learning. *ACM/IEEE International Conference on Human-Robot Interaction(HRI)*, 2023.
- [61] Chao Wang, Stephan Hasler, Daniel Tanneberg, Felix Ocker, Frank Joublin, Antonello Ceravola, Joerg Deigmoeller, and Michael Gienger. Lami: Large language models for multi-modal human-robot interaction. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA)*, 2024.
- [62] Kevin H. Wilson, Yan Karklin, Bojian Han, and Chaitanya Ekanadham. Back to the basics: Bayesian extensions of irt outperform neural networks for proficiency estimation. In *Educational Data Mining*, 2016.
- [63] World Health Organization. Blindness and vision impairment. <https://www.who.int/news-room/fact-sheets/detail/blindness-and-visual-impairment>, 2026. Accessed 2026-05-11.
- [64] Zhanxin Wu, Bo Ai, Tom Silver, and Tapomayukh Bhattacharjee. SAVOR: Skill affordance learning from visuo-haptic perception for robot-assisted bite acquisition. In *Conference on Robot Learning (CoRL)*, 2025.
- [65] Anxing Xiao, Wenzhe Tong, Lizhi Yang, Jun Zeng, Zhongyu Li, and Koushil Sreenath. Robotic guide dog: Leading a human with leash-guided hybrid physical interaction. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [66] Cunjun Yu, Yiqing Xu, Linfeng Li, and David Hsu.

Coach: Cooperative robot teaching. In *Conference on Robot Learning (CoRL)*, 2022.

- [67] G.A. Zijlstra, Judith Ballemans, and Gertrudis Kempen. Orientation and mobility training for adults with low vision: a new standardized approach. *Clinical Rehabilitation*, 27:3–18, 2013.