

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 will greatly expand 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.

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 fundamental challenges. First, these tasks are complex processes



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

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

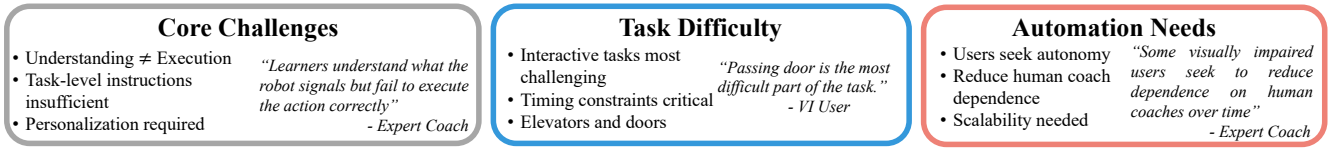


Fig. 2: **Expert coach insights from formative study.** Analysis reveals common learner challenges, effective coaching strategies (timing and content), and design implications for automated coaching systems.

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.

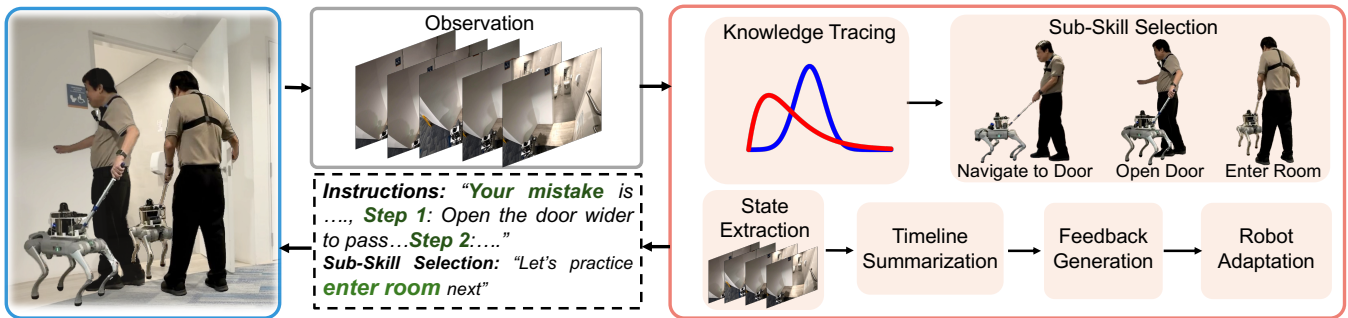


Fig. 3: **Overview of CANINE.** CANINE employs a two-level coaching strategy. The **inter-skill** coaching (up-right) tracks proficiency across sub-skills using knowledge tracing and selects the sub-skill to practice next. The **intra-skill** coaching (down-right) takes in video observations and generates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

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

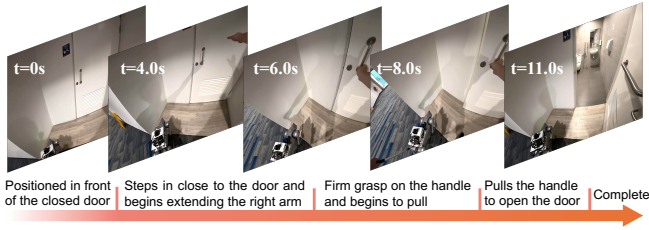


Fig. 4: **Illustration of the summarized timeline.** The timeline aggregates per-frame observations sampled every 0.5 s into a structured episode-level summary.

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

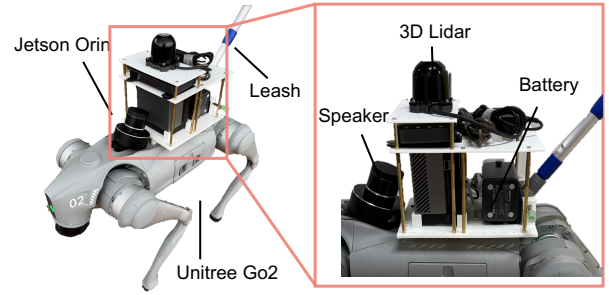


Fig. 5: **Robot guide dog hardware setup.** Unitree Go2 quadruped robot equipped with Jetson AGX Orin for onboard computation, Hesai JT16 LiDAR for navigation, and chest-mounted camera on the user for first-person view analysis.

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

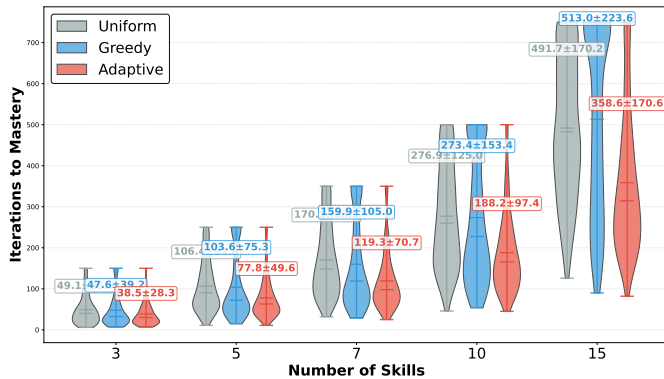


Fig. 6: **Adaptive curriculum sequencing achieves faster learning.** Violin plots show the distribution of teaching actions needed to reach mastery for different numbers of sub-skills. Adaptive curriculum sequencing achieves faster learning.

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. All studies were IRB-approved.

- 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

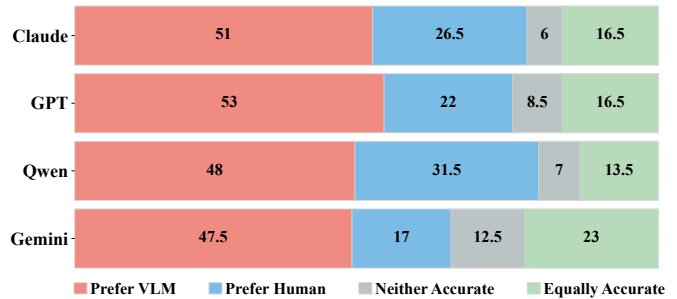


Fig. 7: **Frame analysis accuracy comparison.** Preference rates of VLM-generated descriptions versus human annotations across different language models (GPT-5.1, Claude-4.5-Sonnet, Gemini-3.0-Pro and Qwen3-VL).

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, Factuality, 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). None of the participants had prior experience using the robot guide dog.

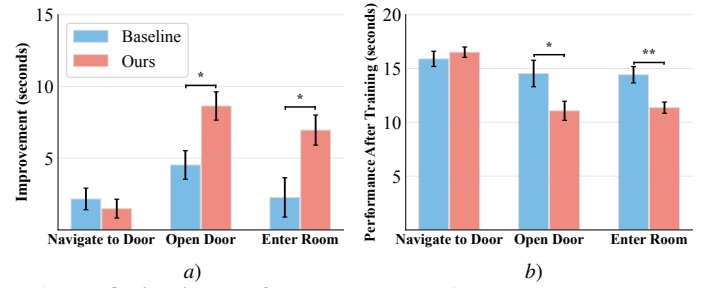


Fig. 9: **Objective performance.** a) Performance improvement across sub-skills. b) Final completion time in evaluation trials.




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

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 I).

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 straightforward 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. $M = 2.27, SD = 1.37$), $t(18) = -2.61, p = .019, g = -1.11$, leading to faster completion times ($M = 11.36, SD = 0.52$ vs. $M = 14.42, SD = 0.76$),




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).

$t(18) = 3.00, p = .006, g = 1.35.$

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 substantial improvement in the “Open Door” task. In contrast, “Enter Room” and “Navigate to Door” showed mi-




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

nor 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 coaching systems must adapt not only feedback content but also its timing based on task dynamics. We treat this handover



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

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 to mitigate camera field-of-view limits, automating skill discovery to reduce reliance on expert-defined sub-skill decomposition, conducting factorial ablations of adaptive curriculum, personalized feedback, and robot-side scaffolding, 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. 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. doi: 10.1177/02783649251333479.
- [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 (UMUAI)*, 4(4):253–278, 1995.
- [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 snoo pie 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] R.K. Hambleton and H. 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)*, June 2025.
- [27] Matthew Jörke, Shardul Sapkota, Lyndsea Warkentien, 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 *Robotics Research: Volume 1*. 2017.
- [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, January 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)*, number 18, 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, Ra-

- jat 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. *Intelligent tutoring systems: Lessons learned*. 1988.
- [45] Peizhu Qian, Filip Bajraktari, Carlos Quintero-Pea, Qingxi Meng, Shannan K Hamlin, Lydia E. Kavrakı, 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. *Intelligent tutoring systems for foreign language learning: The bridge to international communication*, volume 80. 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, 2023.
- [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 Balleman, and Gertrudis Kempen. Orientation and mobility training for adults with low vision: a new standardized approach. *Clinical Rehabilitation*, 27:3–18, 06 2013.

TABLE II: Parameter distributions for the Simulated Learner. Parameters are randomized per skill or per learner episode to create diverse student profiles.

Parameter	Distribution / Logic	Description
Optimal Performance Time (τ_ω^*)	30%: $\mathcal{U}(1.0, 5.0)$ s (Easy) 30%: $\mathcal{U}(5.0, 12.0)$ s (Medium) 40%: $\mathcal{U}(12.0, 25.0)$ s (Hard)	Minimum time to complete skill
Initial Multiplier (m_ω)	20%: $\mathcal{U}(1.0, 2.0)$ (Near Expert) 30%: $\mathcal{U}(2.0, 5.0)$ (Novice) 50%: $\mathcal{U}(5.0, 15.0)$ (Struggling)	Initial incompetence factor ($T_{init} = m \cdot T_{opt}$)
Variance Factor (v_ω)	30%: $\mathcal{U}(0.05, 0.2)$ (Consistent) 70%: $\mathcal{U}(0.3, 1.2)$ (Noisy)	Performance noise magnitude
Learning Rate (α_ω)	25%: $\mathcal{U}(0.2, 0.5)$ (Fast) 25%: $\mathcal{U}(0.01, 0.08)$ (Slow) 50%: $\mathcal{U}(0.08, 0.25)$ (Average)	Rate of proficiency gain per attempt
Forgetting Rate (ϕ)	$\mathcal{U}(0.02, 0.1)$	Proficiency decay per non-practice step
Fatigue Threshold	$\mathcal{U}_{int}(2, 20)$ attempts	Steps before performance degrades

APPENDIX

A. Study 1: Simulated Learner Model

This section provides the formal mathematical specification of the simulated learner model used to evaluate the adaptive teaching policy. The model integrates key cognitive factors grounded in established learning theories and prior work on simulated learners for intelligent tutoring systems [12, 62, 46, 32].

Model Overview. The simulated learner maintains a latent proficiency state for each sub-skill that evolves over time through practice. At each teaching step t , when the learner practices a sub-skill ω , the model: (1) computes the current expected performance time $\mu_{\omega,t}$ based on accumulated learning progress, forgetting, and fatigue; (2) samples an observed performance time $y_{\omega,t}$ from a Gaussian distribution to simulate behavioral variability; (3) updates the learning progress $p_{\omega,t}$ to reflect skill improvement from practice. This creates realistic learning trajectories that exhibit improvement with practice, performance variability, skill decay without practice, and fatigue from over-practice.

Overview of Cognitive Factors The simulated learner reflects the multifaceted and stochastic nature of human learning through several key factors:

- **Power Law of Practice:** Learners improve proficiency through practice following the Power Law of Practice, where performance time decreases as a power function of the number of practice trials. This captures the well-established phenomenon that early practice yields rapid improvement, with diminishing gains as proficiency approaches mastery.
- **Slips and Guesses:** Following Bayesian Knowledge Tracing, the model accounts for slips (errors on mastered skills due to momentary lapses) and guesses (correct responses on unmastered skills due to chance), reflecting stochastic variability in behavior. This is implemented through Gaussian noise in performance observations.
- **Forgetting:** Performance degrades over time without practice. The forgetting rate λ_{forget} determines how quickly skills decay when not practiced.
- **Fatigue Effects:** Performance degradation from extended practice sessions is modeled following cognitive load theory. After a threshold number of attempts N_{fatigue} , additional practice within the same session yields diminishing returns and may temporarily impair performance.
- **Individual Variability:** Learners exhibit heterogeneous learning rates, initial proficiency levels, and susceptibility to fatigue. We model this by sampling learner parameters from distributions (Table II) to create diverse learner profiles.

Model Initialization For each skill ω , the learner’s initial state is defined by a baseline mean performance ($\hat{\mu}_\omega$) and standard deviation ($\hat{\sigma}_\omega$). These are derived from the optimal performance time τ_ω^* , a performance multiplier $m_\omega > 1$, and a variance factor $v_\omega > 0$.

$$\hat{\mu}_\omega = \tau_\omega^* \cdot m_\omega \quad (2)$$

$$\hat{\sigma}_\omega = \tau_\omega^* \cdot v_\omega \quad (3)$$

Table II provides the distributions used to initialize these parameters for diverse learner profiles. Initially, at $t = 0$, the learning progress $p_{\omega,0}$ and attempt count $N_{\omega,0}$ for all skills are zero.

Performance on an Attempt At any given time step t , the learner’s performance $y_{\omega,t}$ on a task for skill ω is modeled as a random variable drawn from a Gaussian distribution. The distribution is defined by a time-varying mean $\mu_{\omega,t}$ (representing the

learner's true underlying skill level) and standard deviation $\sigma_{\omega,t}$ (representing performance variability due to slips and guesses).

$$y_{\omega,t} \sim \mathcal{N}(\mu_{\omega,t}, \sigma_{\omega,t}^2) \quad (4)$$

The final performance time is clipped to a minimum positive value to ensure physical plausibility, i.e., $\max(0.1, y_{\omega,t})$. The current mean performance, $\mu_{\omega,t}$, represents the learner's expected proficiency. It dynamically updates based on learning progress, forgetting, and fatigue. The equation starts from the baseline novice performance $\hat{\mu}_{\omega}$, subtracts improvements from learning (I_{learn}) and recovery from forgetting (I_{forget}), and adds penalties from fatigue (I_{fatigue}). The result is clipped to never go below the optimal expert time τ_{ω}^* .

$$\mu_{\omega,t} = \max(\tau_{\omega}^*, \hat{\mu}_{\omega} - I_{\text{learn}} - I_{\text{forget}} + I_{\text{fatigue}}) \quad (5)$$

where:

- **Learning Improvement** (I_{learn}): The reduction in performance time due to accumulated knowledge. The learning progress $p_{\omega,t-1}$ ranges from 0 (no learning) to 1 (full mastery) and is updated after each practice attempt (see Learning Update Mechanism below).

$$I_{\text{learn}} = p_{\omega,t-1} \cdot (\hat{\mu}_{\omega} - \tau_{\omega}^*) \quad (6)$$

This establishes a linear relationship between learning progress and expected performance: when $p_{\omega,t-1} = 0$, the learner performs at baseline novice level ($\mu_{\omega,t} = \hat{\mu}_{\omega}$), and when $p_{\omega,t-1} = 1$, the learner achieves optimal expert performance ($\mu_{\omega,t} = \tau_{\omega}^*$).

- **Forgetting Penalty** (I_{forget}): Performance degradation due to elapsed time Δt since the last attempt.

$$I_{\text{forget}} = \lambda_{\text{forget}} \cdot \Delta t \cdot (\hat{\mu}_{\omega} - \tau_{\omega}^*) \quad (7)$$

- **Fatigue Penalty** (I_{fatigue}): Performance degradation from excessive recent practice. This penalty is applied only when the attempt count exceeds a threshold N_{fatigue} .

$$I_{\text{fatigue}} = \lambda_{\text{fatigue}} \cdot \max(0, N_{\omega,t-1} - N_{\text{fatigue}}) \cdot (\hat{\mu}_{\omega} - \tau_{\omega}^*) \quad (8)$$

Current Performance Variance The variance $\sigma_{\omega,t}$ models the consistency of the learner's performance. It decreases as the learner gains more experience, reflecting more stable performance.

$$\sigma_{\omega,t} = \max(0.1 \cdot \tau_{\omega}^*, \hat{\sigma}_{\omega} \cdot (1 - f_{\text{exp}})) \quad (9)$$

where the experience factor f_{exp} is given by:

$$f_{\text{exp}} = \min(0.5, \lambda_{\text{exp}} \cdot N_{\omega,t-1}) \quad (10)$$

Learning Update Mechanism After each attempt, the learner's progress $p_{\omega,t}$ is updated. This updated value is then used in the next attempt to compute I_{learn} (see equation above), which determines how much the learner's expected performance has improved. The model assumes learning always occurs with practice, with the magnitude of the update implementing the Power Law of Practice through diminishing returns.

$$p_{\omega,t} = \min(1, p_{\omega,t-1} + \Delta p_{\omega,t}) \quad (11)$$

The learning increment, $\Delta p_{\omega,t}$, is calculated as:

$$\Delta p_{\omega,t} = \max(0.01, \alpha_{\omega} \cdot (1 - \gamma_{\text{dim}} p_{\omega,t-1})) \quad (12)$$

where:

- α_{ω} is the base learning rate for skill ω , sampled from a distribution to model individual differences.
- γ_{dim} is a coefficient modeling diminishing returns as mastery increases.
- The minimum value of 0.01 ensures some learning always occurs, preventing complete stagnation.

This formulation captures the empirically observed pattern that novices improve rapidly with initial practice, while experts require substantially more practice to achieve marginal improvements [62].

TABLE III: Curriculum sequencing performance on simulated learners. Mean±Std shows average teaching actions used, and Success rate (%) for mastering all sub-skills within 200 teaching actions across varying curriculum complexity.

#Sub-skills	Uniform		Greedy		Adaptive	
	Mean±Std	Success	Mean±Std	Success	Mean±Std	Success
3	49.1±35.5	97.6%	47.6±39.2	93.8%	38.5±28.3	98.6%
5	106.4±61.8	96.8%	103.6±75.3	87.4%	77.8±49.6	98.0%
7	170.1±88.0	93.6%	159.9±105.0	86.4%	119.3±70.7	97.0%
10	276.9±125.0	93.0%	273.4±153.4	78.4%	188.2±97.4	98.0%
15	491.7±170.2	90.2%	513.0±223.6	66.0%	358.6±170.6	93.6%

B. Study 1: Simulated User Evaluation Results

This section presents detailed results from the simulated learner evaluation comparing three curriculum sequencing strategies: Uniform, Greedy, and Adaptive. We evaluate each strategy’s ability to help diverse learners master all sub-skills within a fixed budget of 50 teaching actions per skill.

- 1) **Uniform Curriculum** Cycles through all sub-skills in a fixed round-robin order regardless of learner proficiency. This represents a non-adaptive baseline that treats all learners identically and does not respond to individual learning patterns.
- 2) **Greedy Curriculum** Reactively selects the sub-skill with lowest recent performance (highest completion time or failure rate) without maintaining an explicit proficiency model. This myopic strategy may over-practice already-mastered skills due to temporary performance fluctuations (slips) or under-practice skills that appear easy due to lucky guesses. It also uses a fixed exploration rate of 0.1.
- 3) **Adaptive Curriculum (Ours)** Maintains explicit Bayesian proficiency estimates for each sub-skill and selects practice targets to maximize expected learning progress. The curriculum balances exploration (practicing uncertain skills) and exploitation (reinforcing near-mastery skills) using epsilon-greedy with a fixed exploration rate of 0.1 [62].

Experimental Setup For each curriculum strategy and sub-skill count (3, 5, 7, 10, 15), we generated 500 diverse simulated learners by sampling parameters from the distributions in Table II. By adding one additional sub-skill, we increase the allowable number of teaching actions by 50 and evaluate success as the proportion of learners achieving proficiency ≥ 0.8 across all sub-skills.

Results and Analysis Table III shows success rates across curriculum complexity levels. The adaptive curriculum consistently outperforms both baselines across all complexity levels. For 15 sub-skills, adaptive achieves 93.6% success rate compared to 66.0% for greedy and 90.2% for uniform. The performance gap widens as the number of sub-skills increases, demonstrating the importance of maintaining an explicit proficiency model for complex skill hierarchies.

The greedy strategy performs worst, particularly for complex curricula (15 sub-skills: 66.0% success). Analysis of individual learner trajectories reveals two failure modes:

- 1) **Thrashing:** Switching between sub-skills too frequently due to performance noise, preventing sustained practice on any single skill
- 2) **Premature abandonment:** Moving away from near-mastery skills due to temporary slips, missing opportunities to consolidate learning

Uniform Curriculum Robustness The uniform strategy performs surprisingly well (90.2% for 15 sub-skills), nearly matching the adaptive approach. This suggests that for learners with sufficient practice budget, systematic coverage of all skills can be effective. However, uniform wastes resources on already-mastered skills and may frustrate learners with repetitive practice.

Complexity Scaling All strategies show declining success rates as sub-skill count increases, reflecting the fundamental challenge of mastering complex hierarchies within fixed practice budgets. The adaptive curriculum’s advantage grows with complexity.

Statistical Significance We conducted pairwise t-tests comparing teaching actions used for each learner across curriculum strategies. Table IV shows p-values and Cohen’s d effect sizes across all skill counts.

C. Study 1: Frame Analysis Evaluation Protocol

This section describes the detailed methodology for evaluating VLM frame analysis accuracy. Accurate frame analysis is critical for the coaching pipeline, as errors in state extraction propagate through timeline summarization and feedback generation.

Dataset Construction We randomly sample 100 frames from navigation task videos where users are either entering a door or using an elevator with the robot guide dog. Frame selection criteria:

- **Temporal diversity:** Frames sampled from different time points (beginning, middle, end) of episodes to capture varying interaction states
- **Action diversity:** Frames showing different user actions (approaching, positioning, door manipulation, entry)

These frames are randomly assigned to two different human annotators who provide natural-language descriptions of:

- User body position and posture (e.g., “user standing 1.5m from door, facing slightly left”)

TABLE IV: Statistical significance of pairwise comparisons between curriculum strategies. P-values from t-tests and Cohen’s d effect sizes for teaching actions used. Negative d values indicate Adaptive uses fewer actions.

#Skills	Comparison	p-value	Cohen’s d
3	Adaptive vs. Greedy	<0.0001	−0.268
	Adaptive vs. Uniform	<0.0001	−0.332
5	Adaptive vs. Greedy	<0.0001	−0.404
	Adaptive vs. Uniform	<0.0001	−0.509
7	Adaptive vs. Greedy	<0.0001	−0.453
	Adaptive vs. Uniform	<0.0001	−0.636
10	Adaptive vs. Greedy	<0.0001	−0.664
	Adaptive vs. Uniform	<0.0001	−0.792
15	Adaptive vs. Greedy	<0.0001	−0.776
	Adaptive vs. Uniform	<0.0001	−0.781

- Robot position relative to user (e.g., “robot positioned to user’s right, 0.5m ahead”)
- Environmental context (doorway, obstacles, spatial layout)
- User actions or errors visible in the frame (e.g., “user reaching for handle but hand positioned too high”)

Annotation Quality Control To ensure annotation quality, we used 5 randomly selected samples as a quality control check. We manually evaluated these 5 annotations for consistency and accuracy against expert assessments. If the annotations showed inconsistencies or quality issues (e.g., missing critical details, factual errors, vague descriptions), we recruited a new annotator to re-annotate the entire dataset.

Annotator Recruitment and Compensation Annotators were recruited through Prolific with the following requirements:

- Native English speakers
- Approval rate $\geq 95\%$ on previous Prolific tasks
- Passed the quality control checks

Annotators were compensated with 9 GBP per hour. It took approximately 1 hour to complete the 50-frame annotation task (each annotator annotated 50 frames). The UI for annotators can be found in Fig. 15.

VLM Prediction Generation For each frame, we prompt the VLM (GPT-5.1) to generate structured observations following our frame analysis schema (detailed in Appendix O). The VLM receives:

- The frame image
- Task context (“User is navigating with a robot guide dog to open a door”)
- Structured output schema specifying required fields (user state, robot state, environmental context, observed actions/errors)
- No access to human annotations (to ensure independent evaluation)

The prompt instructs the VLM to be factual and specific, avoiding speculation about user intent or internal states that are not directly observable.

Human Evaluation Protocol Human judges evaluate VLM prediction accuracy through pairwise comparisons. For each frame, judges are shown:

- The original frame image
- Human annotation (Description A)
- VLM prediction (Description B)

Judges rate which description is more accurate using a multiple choice question:

- A is more accurate
- B is more accurate
- Equally accurate
- Neither is accurate

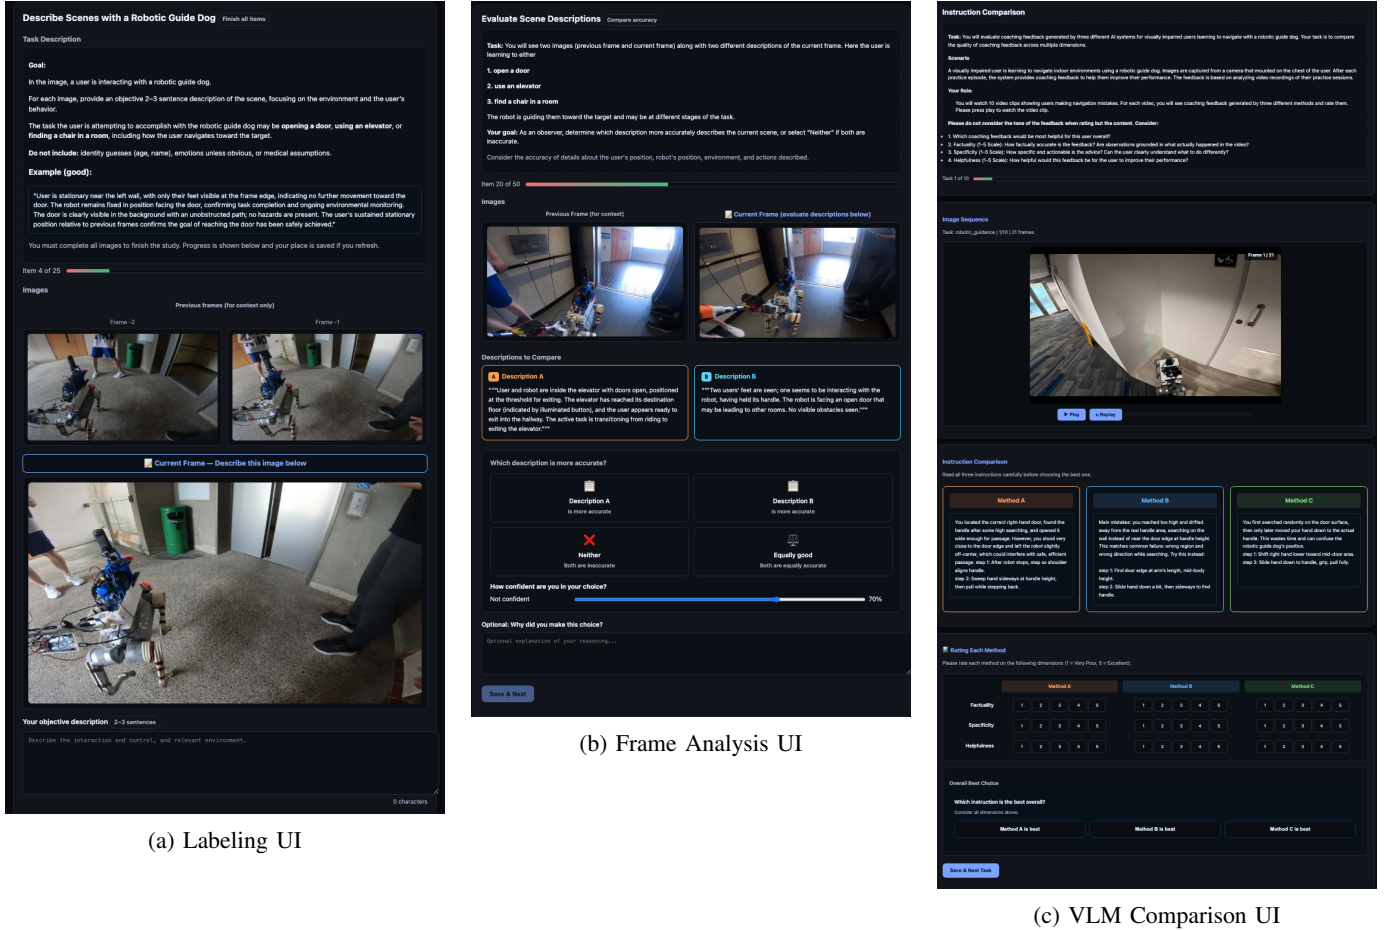
Evaluation Procedure

- 1) The order of A/B is randomized to prevent position bias
- 2) Judges are instructed to evaluate based on the accuracy of the description.
- 3) For each model, we recruited 4 annotators to rate the pairwise comparison
- 4) Each human judge evaluated 50 frames
- 5) Judges were blind to which description came from human vs. VLM

Inter-Rater Agreement Table V shows inter-rater agreement metrics, including Fleiss’s kappa and unanimous agreement rates when judges had clear preferences (excluding “equally accurate” responses).

TABLE V: **Inter-rater agreement for frame analysis evaluation.** Four independent judges rated VLM-generated descriptions versus human annotations. Fleiss’s kappa measures overall agreement. Unanimous agreement shows the percentage of frames. Both are for where judges had a clear preference for the same description (excluding “equally accurate” or “neither accurate” ratings).

Model	Fleiss’s Kappa	Unanimous Agreement (%)
GPT-5.1	0.49	81.7%
Claude-4.5-Sonnet	0.55	84.9%
Gemini-3.0-Pro	0.32	83.9%
Qwen3-VL	0.44	82.5%



(a) Labeling UI

(b) Frame Analysis UI

(c) VLM Comparison UI

Fig. 15: User interfaces used for data collection and evaluation. (a) Interface for human annotators to label ground truth frame descriptions. (b) Interface for judges to rate the faithfulness of VLM outputs against ground truth. (c) Interface for judges to compare coaching feedback from different models (Ours vs. Baselines).

D. Study 1: VLM Comparison Methodology

This section provides detailed methodology for comparing our decomposed four-stage pipeline against end-to-end VLM baselines. The comparison evaluates whether decomposing the coaching task into explicit stages (frame analysis, timeline summarization, feedback generation, robot adaptation) provides benefits over monolithic end-to-end processing.

We compare three approaches:

Baseline 1: VLM (Direct Prompting) The VLM receives the entire video (10-30 seconds, sampled at 2 fps) and a generic prompt requesting coaching feedback. This represents a minimal-effort baseline that relies entirely on the VLM’s general capabilities without task-specific engineering.

Baseline 2: Structured VLM The VLM receives the entire video and structured prompt templates similar to our decomposed pipeline. The prompt includes:

- The same schemas used in our pipeline (user state categories, common failure modes, actionable feedback templates)

TABLE VI: **Pipeline inputs and outputs.** Each stage consumes structured outputs from the preceding stage, reducing open-ended generation and making the worked example in the supplementary material directly traceable.

Stage	Input	Output
f_{frame}	Sampled RGB frames and task context	Per-frame state h_t : user, robot, environment, progress, safety flags
f_{time}	Frame states $h_{1:T}$	Episode summary τ : outcome, metrics, timestamped events
f_{coach}	Summary τ and transcribed user feedback	Coaching action c : diagnosis and terminal verbal instruction
f_{param}	Coaching action c and proficiency belief	Robot adjustment u : conservative scaffolding parameters

- Explicit instructions to: (1) identify what happened, (2) diagnose why errors occurred, (3) provide specific actionable guidance
- Request for structured output matching our feedback format

However, unlike our pipeline, the VLM processes the entire video in a single call without decomposition into stages. This tests whether decomposition provides benefits beyond structured prompting alone.

E. Pipeline I/O Specification

Ours (Decomposed Pipeline) Our four-stage pipeline as described in Section V:

- 1) Frame analysis (f_{frame}): Extract structured observations from individual frames
- 2) Timeline summarization (f_{time}): Aggregate frame observations into episode-level summary
- 3) Coaching generation (f_{coach}): Generate personalized feedback based on timeline and proficiency model
- 4) Robot adaptation (f_{param}): Adjust robot parameters based on diagnosed errors

Table VI summarizes the input and output contract for each stage, making the intermediate representations used by the decomposed pipeline explicit.

Video Selection We selected 10 navigation videos where users made clear errors requiring coaching intervention:

- **Error types:** Wrong door handle search (3 videos), following too closely (3 videos), poor door coordination (3 videos), good performance (1 video)
- **Video length:** 10.5-31.5 seconds (mean 17.45s)

Quality Control To ensure reliable ratings, we used 2 videos as quality control checks:

- 1) **Easy video:** Clear, obvious error with straightforward correction, user followed the robot too closely (expected high agreement)
- 2) **Difficult video:** Subtle, ambiguous error requiring careful observation, user searched for the door handle for a long time on the wall (expected lower agreement but consistent with expert assessment)

Each judge evaluated these control videos, and we manually verified their ratings against expert assessments. If a judge’s ratings on both control videos showed significant discrepancies from the expert baseline (e.g., rated obviously incorrect feedback as helpful), we excluded that judge’s ratings from the final analysis and recruited a replacement judge. Five judges were replaced using this procedure.

Evaluation Metrics

- 1) **Preference Rate** Eight human judges picked their preferred generated feedback from the three approaches for each video. Judges were shown:
 - The original video
 - Three feedback texts (order randomized, approach labels hidden)
 - Instructions to select the feedback that would be most helpful for learning
Preference rate is the percentage of videos where an approach received the most votes (plurality winner).
- 2) **Quality Ratings** In addition to preference rankings, judges rate each approach’s feedback independently on four dimensions using 5-point Likert scales (1=poor, 5=excellent):
 - **Faithfulness:** Accuracy of observations about what actually happened in the video (e.g., does the feedback correctly identify user actions and errors?)
 - **Specificity:** Level of actionable detail provided (e.g., “move hand forward 20cm” vs. “adjust hand position”)
 - **Helpfulness:** Overall utility for learning and skill improvement (e.g., would this feedback help the user improve on the next attempt?)
- 3) **Latency** Time from video input to feedback generation, measured on identical hardware calling the same API endpoints. We report average time across all 10 videos.

Latency Breakdown. Direct video prompting required 44.36 s on average, structured full-video prompting required 42.15 s, and our decomposed pipeline required 20.62 s. For our pipeline, the dominant cost is the parallel frame-analysis stage; timeline summarization, coaching generation, and robot-parameter adaptation operate on compact structured text outputs. Because all stages run after an episode completes, this latency supports terminal feedback between attempts but not concurrent feedback during safety-critical navigation.

Judges rated all three approaches for each video, allowing within-video comparisons to control for video difficulty.

a) *Inter-Rater Agreement:* For preference rankings, we report agreement metrics appropriate for ranked data:

- **Unanimous agreement:** All judges ranked the same method first (10.0% of videos)
- **Majority agreement:** $\geq 50\%$ of judges agreed on top method (80.0% of videos)
- **Average modal share:** Mean proportion of judges agreeing on the most-preferred method per video (58.8%)

F. Study 2: Detailed Experimental Protocol

This section provides the full timeline and procedural details for Study 2 (human subject evaluation with blindfolded sighted participants).

Total session time is approximately 40 minutes.

1) Introduction (5 min):

- Consenting and briefing.
- Rigid leash fitting and adjustment; blindfold fitting; camera and microphone adjustment.

2) Practice Phase (7 min, 2 trials):

- Two different routes that lead to different stopping points, distances, and facing directions from the door.
- Purpose: Collect pretest performance to establish baseline.

3) Training Phase (11 min):

- Routes: Varying complexity, including different starting and ending positions, ending facing directions, and directional changes
- *Baseline group:* 3 full end-to-end trials with task-level verbal instructions only (e.g., “Now open the door”) without sub-skill decomposition, personalized error diagnosis, or adaptive curriculum sequencing
- *CANINE group:* Up to 6 sub-skill episodes selected by the adaptive curriculum, with the session ending early if the system determined all sub-skills were mastered. Terminal feedback provided after each episode with personalized error diagnosis and adaptive curriculum sequencing. In our experiments, 4 participants mastered all sub-skills in less than 6 episodes.
- *Safety timeout:* For each sub-skill training, a 25-second timeout was set for each sub-skill episode other than navigate to door to trigger human intervention if the participant was unable to complete the sub-skill within this time window.

4) Evaluation Phase (7 min, 2 trials):

- Different routes with similar complexity to training routes
- No feedback provided to assess skill learning

5) Survey (10 min): NASA-TLX, SUS, Customized Questionnaires for Perceived Usefulness and Perceived Ease of Use.

Route Length Compensation for Navigate to Door When calculating skill improvement for the Navigate to Door sub-skill across different phases (practice, evaluation, retention), we compensate for differences in route length to ensure fair comparison. Since different routes have varying distances from the starting position to the door, raw completion times are not directly comparable.

We normalize performance using expert baseline times for each route configuration:

- **Practice phase routes:** Expert baseline = 14.25s
- **Evaluation phase routes:** Expert baseline = 18.51s
- **Retention phase routes:** Expert baseline = 18.91s

For each participant’s Navigate to Door completion time $t_{\text{participant}}$ on a given route, we compute the normalized performance as:

$$t_{\text{normalized}} = t_{\text{participant}} - (t_{\text{expert}} - 14.25) \quad (13)$$

where t_{expert} is the expert baseline for that specific route. This adjustment removes the effect of route length differences, allowing us to measure genuine skill improvement rather than artifacts of route difficulty. All Navigate to Door performance comparisons reported in the main paper use these normalized times.

Optimal Performance For the knowledge tracing model, we define optimal performance times for each sub-skill based on expert demonstrations:

- **Navigate to Door:** 15s (optimal expert time)

- **Open Door:** 5s (optimal expert time)
- **Enter Room:** 10s (optimal expert time)

For the novice level, we set all sub-skill performance estimates to 40s.

G. Study 2: Baseline Instruction

This section documents the task-level verbal instructions provided to the baseline group during Study 2. Unlike the CANINE group, which received personalized error diagnosis, adaptive curriculum sequencing, and sub-skill decomposition, the baseline group received only simple task-level instructions without detailed feedback.

Initial Instruction (Before Practice Phase)

Participants in both groups received the following initial briefing:

“I am a robot guide dog. I will help you navigate to the bathroom, pull the door, and go into the bathroom.”

Phase-Level Instructions During Training

During the training phase, the baseline group received brief task-level prompts at the start of each phase, without personalized error diagnosis or corrective feedback:

a) *Navigate to Door Phase:*

“I’m leading you to the door.”

b) *Open Door Phase:*

“You are arriving at the door. The handle is about 1.5 meters in front of you, on your right side. Please pull it open.”

c) *Enter Room Phase:*

“The door is opened successfully. I’m leading you to the bathroom.”

H. Study 2: Coaching Feature Preference Survey

This section provides the questionnaire used to assess which coaching features participants found most valuable in Study 2.

Survey Instrument After completing all trials, participants were asked to rank the following coaching features using a 7-point Likert scale from least helpful (1) to most helpful (7):

- **Personalized Feedback:** The system provided feedback tailored to my specific mistakes and performance.
- **Sub-skill Practice:** The ability to practice individual components.
- **Adaptive Curriculum:** The system automatically selected which navigation sub-skills to practice based on my proficiency.
- **Progressive Difficulty:** The system adjusted the navigation difficulty (e.g., distance to the door) based on my performance.
- **Terminal Feedback Timing:** Receiving feedback after each navigation trial rather than during navigation.
- **Feedback System:** I can provide feedback to the system.

Participants were also asked open-ended questions:

- 1) Please rank your top 3 most valuable features from the list above
- 2) What aspects of the navigation task did you find most challenging?
- 3) What aspects of the teaching method were most helpful for your learning?
- 4) Did you develop any specific strategies for successful navigation? If so, please describe them.
- 5) Were there any moments during navigation where you felt particularly unsafe or uncomfortable? If so, please describe them.
- 6) Do you have any additional comments or suggestions about the robot guide dog system or the teaching method?

Rankings were analyzed using mean rank scores (higher is better) and frequency of top-3 selections. The top-3 selections are 1) Personalized Feedback, 2) Adaptive Curriculum, and 3) Sub-skill Practice.

I. Study 2: Perceived Usefulness and Ease of Use Questionnaire

During the pilot study, we observed that participants had difficulty distinguishing between the coaching system component and the robot guide dog system itself. Since both the navigation guidance and the coaching feedback were delivered verbally, participants often conflated the two systems when providing feedback. For example, when asked about the usefulness of the system, participants would sometimes refer to the guide dog’s navigation capabilities rather than the coaching feedback they received. To address this issue, we designed a customized questionnaire that explicitly asks participants to distinguish between two distinct components:

Questionnaire Items Table VII presents the questionnaire items used to assess Perceived Usefulness and Perceived Ease of Use of the coaching system in Study 2, along with their respective Cronbach’s alpha reliability coefficients. Participants rated each item on a 7-point Likert scale (1 = Strongly Disagree, 7 = Strongly Agree).

TABLE VII: Perceived Usefulness and Perceived Ease of Use Questionnaire Items

Perceived Usefulness (Cronbach's $\alpha = 0.836$)
1. Using the robot guide dog training system would improve my performance during robot guide dog handling tasks.
2. Using the robot guide dog training system would make robot guide dog training tasks easier to perform.
Perceived Ease of Use (Cronbach's $\alpha = 0.859$)
1. Learning to operate the robot guide dog training system would be easy for me.
2. My interaction with the robot guide dog training system would be clear and understandable.
3. It would be easy for me to become skillful at using the robot guide dog training system.

J. Study 2: Qualitative Feedback

This section summarizes additional qualitative feedback and system improvement suggestions from Study 2 participants. We conducted thematic analysis of open-ended survey responses and identified recurring themes.

Positive Feedback

- 1) **Personalized Error Diagnosis** Participants appreciated specific, actionable feedback tailored to their mistakes:
 - “The system answered my questions, cleared my confusion. It also suggested tips for how to improve in next trial.” (P7, CANINE group)
 - “The personalized feedback is accurate” (P8, CANINE group)
- 2) **Sub-Skill Decomposition** Breaking down the complex task into manageable components reduced cognitive load:
 - “Practicing where the door was with respect to the robot is helpful” (P10, CANINE group)

However, some participants wanted concurrent feedback for certain sub-skills (see improvement suggestions below).

Improvement Suggestions

Participants suggested several system enhancements:

- 1) **Concurrent Feedback for Specific Sub-Skills:** Some participants (4/10) wanted real-time feedback during execution for certain sub-skills, particularly “Open Door”:
 - “For entering the room, terminal feedback was fine. But for finding the door handle, I want real-time guidance.” (P8, CANINE group)
- 2) **Tone of Feedback:** Some participants suggested the tone of feedback should be more positive:
 - “The feedback could be less harsh” (P4, CANINE group)
 - “The coach sounds condescending” (P10, CANINE group)

K. Study 3: Retention Analysis Protocol

Five participants from Study 2 (N=5 from CANINE) agreed to return 2 weeks after their initial training session.

a) Session Structure (30 min):

- 1) **Re-familiarization (5 min):**
 - Brief reminder of rigid leash and robot guidance system
 - No new instruction or coaching provided
 - 2) **Practiced Routes (10 min, 2 trials):**
 - Same routes from the evaluation phase of Study 2
 - Purpose: Assess retention of learned skills
 - No feedback provided
 - 3) **Novel Routes (10 min, 2 trials):**
 - New routes with similar complexity but different layouts
 - Combine elements from practiced routes (turns and doorways) in new configurations
 - Purpose: Assess transfer and generalization
 - No feedback provided
 - 4) **Brief Interview (5 min):**
 - Qualitative feedback on what they remembered from training
 - What felt easier or harder compared to the initial session
 - Confidence levels navigating with the robot
- b) *Evaluation Metrics:*
- **Performance Maintenance:** Comparison of completion time and navigation efficiency between final training trial (from Study 2) and retention session

All participants demonstrated confident navigation with the robot guide dog. When asked about the skills they learned, they noted the importance of paying attention to the dog's orientation and allowing sufficient space for the dog to adjust.

L. Study 4: Visually Impaired User Case Study Protocol

This exploratory case study evaluates CANINE on both the doorway navigation task and robotic handover task with a totally blind participant (Male, 70, with no residual vision). The participant currently uses a white cane for daily navigation and has prior experience with robot guide dogs. Total session time is approximately 60 minutes for navigation followed by 30 minutes for handover.

a) Navigation Task Session (60 min):

1) Introduction and Consent (10 min):

- Informed consent and compensation explanation
- Demographic questionnaire
- System explanation and safety protocols
- Rigid leash fitting and familiarization with the robot guide dog

2) Practice with Adaptive Coaching (35 min):

- Full CANINE adaptive coaching system active
- Routes of varying complexity (same doorway navigation task as Study 2)
- System adapts curriculum across three sub-skills: Navigate to Door, Open Door, and Enter Room
- Terminal feedback provided after each sub-skill episode
- Evaluation trials to measure improvement
- The trial structure matched Study 2: two initial practice trials, up to six coached sub-skill episodes, and two evaluation trials.

3) Semi-Structured Interview (15 min):

- Coaching feature preference ranking (same instrument as Study 2)
- System usability and learning experience
- VI-specific design considerations (multimodal feedback, system control)
- Audio recorded with participant consent

Results In the initial trial, the participant failed to locate the door handle and complete the task. After coaching, the participant successfully completed the task in the evaluation trial. For the “Open Door” sub-skill, completion time improved from 22.4 s before coaching to 9.1 s after coaching. The attached video illustrates this progression.

b) Interview Topics: The semi-structured interview covered the following key topics:

- **Coaching Feature Preference:** Ranking of coaching features (sub-skill decomposition, personalized feedback, adaptive curriculum, error-specific diagnosis, actionable guidance) - same instrument as Study 2
- **Learning Experience:** Perceived effectiveness of coaching, most/least helpful feedback types, comparison to previous mobility training experiences
- **VI-Specific Design Needs:**
 - Multimodal feedback preferences (audio, haptic, other modalities)
 - Need for system control (ability to pause, replay, skip feedback)
- **Feedback Timing:** Preference for terminal (post-episode) vs. concurrent (real-time) feedback during navigation
- **Most Challenging Sub-skill:** Identification of which sub-skill (Navigate to Door, Open Door, Enter Room) was most difficult and why

c) Robotic Handover Task Session (30 min): Following the navigation session, the participant engaged with a robotic handover task to evaluate system extensibility (detailed setup in Appendix M):

1) Task Introduction and Familiarization (5 min):

- Explanation of robotic handover task and safety protocols
- Haptic familiarization with robot gripper and workspace boundaries
- Demonstration of object transfer process

2) Practice with Adaptive Coaching (20 min):

- Receive objects of varying types from robotic arm (light, heavy, and large objects)
- CANINE coaching system active
- System adapts curriculum across three handover sub-skills based on object category
- Participant requested concurrent (real-time) feedback during execution to adjust hand position
- Metrics: Success rate, completion time, learning progression



Fig. 16: Objects used for handover experiment.

3) Cross-Task Comparison Interview (5 min):

- Comparison of navigation vs. handover coaching experience
- Feedback timing preferences: terminal feedback for navigation vs. concurrent feedback for handover
- Task-specific coaching requirements and design considerations

M. Study 4: Robotic Handover Task Setup

This section details the experimental setup for the robotic handover task used in Study 4 to evaluate the extensibility of CANINE to a different robotic assistance domain.

Robotic System

- **Manipulator:** Franka Research 3 robot with parallel gripper.
- **Workspace:** Robot positioned at table height (75 cm) to facilitate natural handover posture for seated users.
- **Safety Features:** Built-in collision detection, configurable force limiting (10N threshold), emergency stop button within participant reach.
- **Perception:** RGB-D camera (Intel RealSense L515) mounted overhead (1.5m height) providing top-down view of handover workspace for video observation and hand pose tracking.

a) *Handover Objects:* Three object categories representing different manipulation challenges as shown in Fig. 16:

- **Light objects:** Potato chips, plastic spoon, card - requires basic hand positioning and timing coordination.
- **Heavy objects:** A row of yogurt, water bottle, heavy box - requires sufficient grip strength and stability during transfer.
- **Irregular-shaped objects:** Sandwich box, large bag, kettle - requires proper hand configuration to accommodate object shape and size.

Task Structure and Sub-skill Decomposition. In each trial, the robot moves from a home position to the handover location, transfers the object to the user's hand, and returns home. We decompose this interaction into three sub-skills based on object properties:

- 1) **Handpose Alignment:** Put the hand in the correct position relative to the object.
 - Focus: Hand positioning accuracy.
 - Success criterion: Hand positioned appropriately before robot arrival.
 - Common errors: Hand positioned too high/low/far/close.
- 2) **Light Object Handover:** Receiving light objects (100-200g).
 - Focus: Hand positioning accuracy and timing coordination.
 - Success criterion: Object transferred without dropping; hand positioned appropriately before robot arrival.
 - Common errors: Hand positioned too high/low, premature or delayed grasp, insufficient contact area.
- 3) **Heavy Object Handover:** Receiving heavy objects (400-600g)
 - Focus: Grip strength and weight preparation
 - Success criterion: Object transferred without dropping; stable grip maintained during and after transfer
 - Common errors: Insufficient grip force leading to drop, unprepared for object weight, grip released too early
- 4) **Irregular-shaped Object Handover:** Receiving irregular-shaped objects.
 - Focus: Hand configuration and spatial awareness
 - Success criterion: Object transferred without dropping; hand configuration accommodates object dimensions
 - Common errors: Hand configuration too narrow for object, fingers not spread appropriately, unstable grasp due to poor hand shape

Each sub-skill requires mastering coordination without visual feedback, making personalized coaching essential for error diagnosis and skill improvement.

Observation and State Representation

In addition to the chest-mounted camera, the system also uses the overhead RGB-D camera as input to get the user’s hand state: position relative to the handover location, hand orientation (palm facing direction), and hand configuration (open/closed, finger spread).

Performance Metrics

- **Success Rate:** Proportion of trials where object was successfully transferred without dropping (primary outcome)
- **Completion Time:** Duration from robot departure from home position to successful stable grasp by user
- **Sub-skill Proficiency:** Success rate tracked separately for each object category (light, heavy, large) to inform adaptive curriculum
- **Error Distribution:** Frequency of different error types (positioning, timing, grip strength, hand configuration) to characterize learning challenges

Results In the first three trials, the user was only able to receive the heavy object passively. After training with CANINE, the user learned to actively open the palm and grasp the object from the bottom. In the final four evaluation trials, the user successfully grasped all objects, including a challenging case that required grasping a kettle. The attached video illustrates this progression.

Comparison to Navigation Task

The handover task provides a complementary test case to evaluate CANINE’s generalizability:

- **Interaction modality:** Fine-motor proprioceptive control (hand positioning, grip) vs. gross-motor spatial coordination (body movement, leash following)
- **Error manifestation:** Immediate binary failure (object drops) vs. accumulated inefficiency (longer completion time, suboptimal path)
- **Sensory demands:** Tactile and force feedback vs. auditory cues and spatial reasoning

Despite these fundamental differences, both tasks share the core coaching challenge: helping users master coordinated interaction with robotic systems without visual feedback. This allows us to evaluate whether CANINE’s architecture (adaptive curriculum, sub-skill decomposition, personalized feedback) generalizes beyond a single task domain.

N. System Failure Modes

We document observed system failures:

- 1) **Failure Mode 1: Visibility Failure** When the user’s body occludes the camera or the user changes the orientation entirely, the VLM cannot extract accurate state information.
- 2) **Failure Mode 2: Lighting Conditions** Poor lighting (too dark or strong backlighting) degrades VLM performance.
- 3) **Failure Mode 3: Motion Blur** Rapid user movements create motion blur, making frame analysis unreliable.
- 4) **Failure Mode 4: Contradictory Feedback** LLM generates feedback that contradicts previous feedback.
- 5) **Failure Mode 5: Overly Complex Language** Feedback uses technical jargon or complex sentence structures difficult for users to understand.
- 6) **Failure Mode 6: Wrong-format Feedback** The LLM failed to return a parsable JSON object.
- 7) **Failure Mode 7: Communication Latency** Network latency causes delays in feedback delivery.

Lessons for Deployment Key insights from failure mode analysis:

- 1) **History is critical:** Multiple historical frames are necessary for accurate state estimation
- 2) **Conservative defaults:** When uncertain, provide generic feedback rather than potentially incorrect specific feedback
- 3) **Human oversight:** Maintain human supervision for safety-critical applications
- 4) **Graceful degradation:** System should continue functioning (with reduced quality) rather than failing completely
- 5) **User agency:** Allow users to override system decisions and request clarification

O. Prompt Templates

a) VLM Baseline:

(1) VLM Analysis

Task: [Task Description according to the task]

You are analyzing a sequence of images from a robotic task. Please examine all the images provided and generate clear, actionable instructions to help complete the task. Please format your response as clear, concise instructions that can be followed.

Please put your response in the following JSON format:

```
{  
  "terminal_instruction_to_user": "string - terminal instruction to the user to help them improve  
  their navigation technique. First, point out mistakes if the performance is bad. Then, provide  
  corresponding actionable suggestions in relative term. If possible, step by step, start with "
```

```
step 1', 'step 2', maximum two steps. For each step, be each point very concise, less than 10 words, no need to be a full sentence. In the end, answer the human user feedback in a very concise way. if there is no feedback, there is no need to answer the human user feedback.",
}
```

b) Structured VLM:

(1) Structured VLM Analysis

Prompt: You are analyzing a camera image mounted on the human chest for robot guide dog navigation training.

Inputs Provided:

Camera Image: Single current frame from the camera mounted on the human chest showing:

- User's body position and walking path
- Distance and positioning relative to robot
- Environment and navigation obstacles
- User's following behavior and spatial awareness

Analysis Approach:

- Always pay attention to Human's hands behavior, e.g. if the human current pose can reach the door handle or not
- Distance between human and the door and door handle
- User difficulties (hesitation, confusion, searching, getting stuck)
- User mistakes that may leads to task failure or inefficient task completion

Be objective and specific. Base your analysis solely on visual observation of the camera image.

Task Information: [Task Description according to the sub-skill]

Common Failure Modes to Watch For:

```
{
  "1": "User search the wrong region for handles and door",
  "2": "User stands too far away from the door",
  "3": "User search the wrong region for handles and door in the wrong direction",
  "4": "User does not keep the door fully open for the robot to pass through",
  "5": "User is too close to the robot when passing through the door so when the robot adjust the pose,
        user may interfere the robot's adjustment",
}
```

Instruction Generation:

Based on ALL the images you see, please provide:

- 1) A summary of what is happening in the image sequence (chronological analysis)
- 2) Specific step-by-step instructions for what actions should be taken
- 3) Any important observations or warnings about the current state
- 4) Identification of any failure modes or potential issues observed
- 5) Maximum 2 steps and be very concise.

Please format your response as clear, concise instructions that can be followed. Focus on providing actionable guidance based on the complete sequence of images. Please put your response in the following JSON format:

```
{
  "technique_assessment": {
    "following_technique": "enum - excellent|good|needs_improvement|poor",
    "distance_maintenance": "enum - excellent|good|needs_improvement|poor",
    "cue_response": "enum - excellent|good|needs_improvement|poor",
    "trust_level": "enum - excellent|good|needs_improvement|poor"
  },
  "specific_feedback": {
    "strengths": ["list of things done well"],
    "areas_for_improvement": ["list of navigation technique improvements"],
    "safety_notes": ["list of safety observations during navigation"],
    "communication_feedback": ["specific feedback about robot communication"]
  },
  "coaching_recommendations": {
    "immediate_practice": ["navigation skills to practice in next session"],
  }
}
```

```

    "technique_tips": ["specific navigation technique improvements"],
    "communication_tips": ["how to better notice and respond to robot signals"],
    "safety_tips": ["safety practices for navigation"]
  },
  "result_summary": "Brief navigation coaching summary",
  "terminal_instruction_to_user": "string - terminal instruction to the user to help them improve their navigation technique. First, point out mistakes if the performance is bad. Then, provide corresponding actionable suggestions in relative term. If possible, step by step, start with 'step 1', 'step 2', maximum two steps. For each step, be each point very concise, less than 10 words, no need to be a full sentence. In the end, answer the human user feedback in a very concise way. if there is no feedback, there is no need to answer the human user feedback.",
  "user_actionables": ["actionable navigation suggestion 1", "actionable navigation suggestion 2", "actionable navigation suggestion 3"]
}

```

c) CANINE:

(2) Frame Analysis

You are analyzing a camera image mounted on the human chest for robot guide dog navigation training.

Inputs Provided: 1. Camera Image: Single current frame from the camera mounted on the human chest showing:

- User's body position and walking path
- Distance and positioning relative to robot
- Environment and navigation obstacles
- User's following behavior and spatial awareness

Analysis Approach:

- Always pay attention to Human's hands behavior, e.g. if the human current pose can reach the door handle or not
- Distance between human and the door and door handle
- User difficulties (hesitation, confusion, searching, getting stuck)
- User mistakes that may leads to task failure or inefficient task completion

Be objective and specific. Base your analysis solely on visual observation of the camera image.

Task:

[Task Description according to the sub-skill]

Analysis Inputs:

- Image: Current camera view from human chest-mounted camera showing the navigation scene

Analysis Task:

- 1) What does the image show about user position and navigation?
- 2) Is the user correctly positioned relative to the robot (if visible)?
- 3) What is the navigation context and progress?
- 4) What specific observations can be made about the navigation scenario?
- 5) Is there adequate clearance and space for safe passage?

Required per-frame schema:

```

{
  "timestamp_sec": "<float>",
  "user_state": "enum - unknown|approaching|positioned|interacting|completed|stuck|retreating|hesitant|confused|searching",
  "robot_state": "enum - unknown|guiding|waiting|signaling|assisting|monitoring|stopped|adjusting",
  "connection_state": "enum - connected|loose|disconnected|unknown - status of leash/harness connection between user and robot",
  "environment_state": "enum - unknown|clear|cluttered|hazardous|changing",
  "task_progress": "enum - not_started|in_progress|nearly_complete|complete|failed|interrupted",
  "safety_concerns": "boolean - are there any safety issues present",
  "user_difficulties": "boolean - is the user experiencing any difficulties or challenges",
  "frame_description": "<string> - a brief description of the frame, noting any challenges, safety issues, or user difficulties"
}

```

(2) Summarization Analysis

You are analyzing a robot guide dog navigation episode to determine if the user successfully completed the navigation task.

Success Criteria:

- 1) User followed the robot guide dog's lead and cues
- 2) Proper distance maintained from the robot (about 1-2 feet)
- 3) User responded appropriately to robot's directional commands
- 4) Safe navigation without incidents or obstacles
- 5) Task objective was achieved completely and safely (e.g., reached door, opened door, entered room)
- 6) All actions were completed with adequate clearance and space for safe passage

Failure Indicators:

- 1) User pulled ahead of robot instead of following
- 2) Incorrect distance maintained (too close or too far)
- 3) User failed to respond to robot's directional cues
- 4) User lost contact with robot's harness or guiding mechanism
- 5) User lacked awareness of surroundings
- 6) Task objective not achieved or completed incompletely
- 7) Insufficient clearance or space created, potentially causing safety issues or blocking passage

Task Information:

[Task Description according to the sub-skill]

Timeline:

[timeline_json]

Determine if what is the main mistake the user made across the timeline and write in the timeline consistently, e.g. too far from the door, wrong search region, dangerous behavior, etc.

Required decision schema:

```
{
  "success": "boolean - did user successfully complete the navigation task",
  "navigation_quality": "enum - excellent|good|fair|poor|very_poor",
  "time_efficiency": "float - actual_time / target_time ratio",
  "safety_rating": "enum - excellent|good|fair|poor|very_poor",
  "following_technique": "enum - excellent|good|fair|poor|very_poor",
  "communication_effectiveness": "boolean - was user-robot communication clear",
  "trust_level": "enum - excellent|good|fair|poor|very_poor",
  "issues_encountered": ["list of any problems"],
  "result_summary": "Brief outcome description",
  "timeline_summary": "string - detailed chronological summary of what the user did at key timestamps,
    explicitly stating the timing and actions (e.g., 'At 0.5s: user established connection with
    guide dog. At 2.1s: user followed left turn cue. At 5.3s: user maintained proper distance. At
    8.7s: user arrived at door.'). This summary will be used as input for coaching instruction
    generation."
}
```

(3) Coaching Generation

You are a mobility instructor providing coaching feedback for a visually impaired person learning to use a robot guide dog. Please be very critical and strict. Please pay attention to every detail.

Decision: [summarization_json]

Task Info: [task_info_json]

Provide clear, easy-to-understand coaching for robot guide dog navigation skills. Write as if speaking directly to the user in a friendly, supportive way. Use simple words, short sentences, and focus on what they can feel or hear.

When providing guidance, consider the importance of completing actions with adequate clearance and space. Help users understand how to feel for full range of motion in their interactions, ensuring smooth and safe navigation.

Common Failure Modes to Watch For:

```
{
  "1": "User search the wrong region for handles and door",
  "2": "User stands too far away from the door",
}
```

```

"3": "User search the wrong region for handles and door in the wrong direction",
"4": "User does not keep the door fully open for the robot to pass through",
"5": "User is too close to the robot when passing through the door so when the robot adjust the pose,
      user may interfere the robot's adjustment",
}

```

Required coaching schema:

```

{
  "technique_assessment": {
    "following_technique": "enum - excellent|good|needs_improvement|poor",
    "distance_maintenance": "enum - excellent|good|needs_improvement|poor",
    "cue_response": "enum - excellent|good|needs_improvement|poor",
    "trust_level": "enum - excellent|good|needs_improvement|poor"
  },
  "specific_feedback": {
    "strengths": ["list of things done well"],
    "areas_for_improvement": ["list of navigation technique improvements"],
    "safety_notes": ["list of safety observations during navigation"],
    "communication_feedback": ["specific feedback about robot communication"]
  },
  "coaching_recommendations": {
    "immediate_practice": ["navigation skills to practice in next session"],
    "technique_tips": ["specific navigation technique improvements"],
    "communication_tips": ["how to better notice and respond to robot signals"],
    "safety_tips": ["safety practices for navigation"]
  },
  "result_summary": "Brief navigation coaching summary",
  "terminal_instruction_to_user": "string - terminal instruction to the user to help them improve
  their navigation technique. First, point out mistakes if the performance is bad. Then, provide
  corresponding actionable suggestions in relative term. If possible, step by step, start with '
  step 1', 'step 2', maximum two steps. For each step, be each point very concise, less than 10
  words, no need to be a full sentence. In the end, answer the human user feedback in a very
  concise way. if there is no feedback, there is no need to answer the human user feedback.",
  "user_actionables": ["actionable navigation suggestion 1", "actionable navigation suggestion 2", "
  actionable navigation suggestion 3"]
}

```

(4) Parameters Tuning

You are a robot parameter tuning engineer specializing in robot guide dog navigation systems.

For guide dog navigation tasks, focus on parameters that affect the safety and effectiveness of navigation guidance:

- Distance to door: How close the robot maintains to the target (door)
- Teaching method: Whether to use scaffolding (more support) or challenging (less support) approach

Coaching Analysis Results:

```
{coaching_json}
```

User Feedback (Natural Language):

```
{user_feedback_json}
```

Navigation-Specific Analysis: For robot guide dog navigation tasks, prioritize safety and effective following during navigation.

Step 1 - Extract Navigation Feedback: Parse the user's natural language feedback for navigation-specific concerns:

- Distance comfort: "too close", "too far", "comfortable distance"
- Guidance clarity: "confusing", "clear", "helpful"

Step 2 - Analyze Navigation Technique:

- User's following behavior and positioning
- Safety considerations during navigation

Step 3 - Recommend Navigation Parameter Adjustments: Based on the navigation analysis, recommend adjustments to:

```

"distance_to_door": closer/farther/maintain (distance robot maintains to door)
"pointing_direction": more_left/more_right/maintain (how the robot's pointing direction shifts
  horizontally: left, center, or right relative to the user)

```

Required navigation parameter tuning schema:

```
{
  "distance_to_door_adjustment": "enum - closer|farther|maintain - adjust distance robot maintains to
  door",
  "pointing_direction_adjustment": "enum - more_left|more_right|maintain - adjust horizontal pointing
  direction (left/center/right)",
  "reasoning": "string - explanation for navigation parameter adjustments based on user safety and
  comfort"
}
```

P. Formative Study Semi-Structured Interview Guide

We used the following 25-question semi-structured interview guide in the formative study. The interviewer asked open-ended questions, allowed participants to elaborate freely, and used follow-up prompts only to clarify participant-raised points. The guide was designed to elicit participants' own experiences rather than ask them to choose from predefined answers.

Opening Script. "We are interested in your experience learning to use a robot guide dog and in what kinds of support would make that learning process easier or safer. There are no right or wrong answers. Please describe your experience in your own words, and feel free to skip any question."

1) Background and Prior Training:

- 1) What navigation aids or assistive technologies do you currently use in daily life?
- 2) How did you learn to use your current navigation aid or assistive technology?
- 3) What parts of that learning process were most helpful?
- 4) What parts of that learning process were difficult or insufficient?

2) Robot Guide Dog Experience:

- 5) Please describe your experience navigating with the robot guide dog. What are the most difficult tasks?
- 6) What aspects of the robot guide dog were easiest to understand?
- 7) What aspects of the robot guide dog were confusing, difficult, or unexpected?
- 8) How did you interpret the robot's motion, leash cues, or other signals?
- 9) Did your behavior or understanding change over the practice session? If so, what changed?

3) Coaching and Feedback:

- 10) Please describe the coaching or feedback you received during practice.
- 11) Which feedback, if any, helped you change what you did on the next attempt?
- 12) Were there moments when feedback would have been useful but was missing?
- 13) When you made a mistake or felt uncertain, what information would have helped you recover?
- 14) How should feedback be timed for this kind of task?

4) Practice Structure and Personalization:

- 15) How would you prefer to practice a complex navigation task with a robot guide dog?
- 16) Are there parts of the task that should be practiced separately? If so, which parts and why?
- 17) How should practice difficulty change as a learner improves?
- 18) What should a coaching system remember about an individual learner's progress?

5) Safety, Trust, and Sensing:

- 19) Please describe any moments when you felt especially safe, unsafe, confident, or uncertain.
- 20) What would increase your trust in the robot guide dog during training?
- 21) Were there moments when the coach or robot seemed to lack information about what you were doing?
- 22) What kinds of cues or interaction controls would help you manage cognitive load during training?

6) Future Automated Coaching:

- 23) What would make an automated coaching system useful for learning to use robotic assistants?
- 24) What should an automated coaching system avoid doing?
- 25) Is there anything else about learning to use a robot guide dog or robotic assistant that we have not asked about?