Motius: An AI-Adaptive Behavior Layer for Service Robots

Reference Data, Adaptive Profile Bands, and Cross-Task Humanoid Proof

Motius Robotics

May 2026

Abstract

Robots are advancing quickly in perception, planning, and manipulation, yet human-facing behavior is still commonly implemented through hidden runtime parameters, scene-local scripts, and operator tuning rather than as an adaptive system that improves with data. We propose an AI-adaptive behavior layer for robots: a software surface above the controller that organizes interaction style into profile bands, learns from community-contributed reference data, and is intended to adapt runtime parameters to user and scene characteristics in real time. We instantiate this formulation in Motius, a prototype that already exposes an explicit profile surface, a controller-preserving adapter boundary, and a reference-and-validation loop grounded in human behavior examples and robot proof clips. In the current proof-of-concept, the adaptive logic is demonstrated through a rule-based instantiation rather than a fully learned online predictor. Across three interaction types within a single recorded session on one Unitree G1 humanoid platform, one profile pair (Standard vs. Gentle) produces an 18.0% reduction in configured speed scale, a 110.0% increase in configured pause duration, an 18.7% increase in configured stopping distance, and a 6.7% increase in configured end-effector smoothing. Clip-level proxy checks from the same recording partially corroborate these intended shifts through longer visible exchange and arrival windows and a smaller image-space robot size at the settled stop frame. A clip-based online rating study ($n = 24$) finds that Gentle is rated more polite and more appropriate across all three tasks, while naturalness improves in two of three tasks and clarity effects are mixed. Preference votes favor Gentle most strongly in the approach-and-stop and push-object clips, while handover judgments are more mixed and often marked as no clear difference. The paper therefore makes a narrow systems claim: robot behavior can be surfaced as an inspectable adaptive layer above the controller, grounded in reference data and legible to external observers under controlled comparison, even before a fully learned profile predictor is complete.

1. Introduction

Robots can already move, navigate, and complete tasks. That does not mean they already know how to behave in front of humans, nor does it mean their behavior gets smarter from one deployment to the next.

In many human-facing settings, the central question is not only whether a robot finishes a task, but how that task feels while it is being executed. A robot can stop at the correct place and still feel abrupt. It can hand over the correct object and still feel rushed. It can approach a user without colliding and still feel socially awkward. Prior work in human–robot interaction and socially aware navigation shows that users are highly sensitive to social presence, proxemics, motion style, and perceived appropriateness [1, 2, 3, 4, 5, 6]. Yet in many real deployments, these differences are still implemented as local scripts, hidden runtime parameters, and scene-specific QA adjustments rather than as reusable software objects.

This creates a persistent gap between robot capability and robot interaction quality. The hardware can remain the same, the task family can remain the same, and yet the desired interaction style can vary meaningfully by deployment context and by user. A reception robot, a delivery robot, and a public-facing assistance robot may all share locomotion and manipulation primitives, but users will still care about how directly the robot approaches, how long it pauses, how softly it exits, and whether its behavior feels calm, efficient, reassuring, or abrupt.

We treat this gap as both a software problem and a data problem. Specifically, we propose an *AI-adaptive behavior layer* for robots: a layer above the controller that packages human-facing interaction differences into reusable *interaction profiles*, treats those profiles as learnable from reference data rather than only hand-tuned engineering constants, and ultimately adapts runtime parameters to the current user and scene. A profile is not merely an adjective. It resolves into explicit execution fields that can be surfaced, compared, versioned, and mapped into robot-specific runtimes while preserving low-level control and safety.

Unlike prior work that treats behavior quality as a fixed property or an engineering constant, Motius frames interaction profiles as learnable from data. The long-term goal is not a hand-tuned profile library, but a self-improving system in which every real deployment and every contributed reference clip can help make the next deployment smarter.

The present paper does not claim a complete benchmark, a finished deployment stack, or generalized cross-platform transfer. Instead, it makes a narrower systems contribution. We instantiate the adaptive behavior-layer proposal in a prototype called Motius and study a controlled within-session proof asset on the same Unitree G1 humanoid robot across three interaction types: object handover, approach-and-stop, and push object. Within this scope, the paper makes four contributions:

1. It proposes an AI-adaptive behavior layer in which interaction profiles are treated as learnable from real interaction data rather than only engineering intuition.
2. It instantiates an adapter-based runtime boundary and a user-conditioned profile-band formulation that can support real-time parameter adaptation while preserving controller ownership.
3. It establishes a Reference Network flywheel: contributor uploads, structured reference intake, later AI training, improved adaptive profiles, and follow-on robot behavior improvement.
4. It documents a controlled proof on one Unitree G1 humanoid robot across three interaction types together with a 24-respondent pilot rating study.

The intended contribution is therefore not a full autonomy claim, but an early robotics systems claim: behavior can become legible, parameterized, adaptive, and eventually data-improving software.

2. Related Work

Our work sits at the intersection of social robot navigation, human–robot interaction, expressive motion design, human-to-robot transfer, and preference-aware evaluation for human-facing robot behavior.

2.1. Social Navigation and Perceived Appropriateness

Recent review work makes clear that social robot navigation is still unresolved not only at the algorithmic level, but also at the level of requirements, datasets, and evaluation criteria [1, 2]. What counts as "good" behavior remains fragmented across efficiency, comfort, proxemics, perceived appropriateness, trust, and human interpretation. PARSNiP treats perceived appropriateness as a first-class prediction target and augments social-navigation analysis with emotional and attentional features [3]. Related work on robot gaze during navigation shows that small interaction-level cues can alter social presence even when the underlying navigation task remains unchanged [4]. More recent work argues that social-navigation evaluation should explicitly include human-centered criteria and clearer reporting standards [7], while studies of implicit behavioral cues and perception modeling further underscore that trust, comfort, and judged performance are shaped by subtle behavioral variation rather than task success alone [8, 9].

2.2. Expressive Motion and Legible Interaction Quality

The literature on expressive robot motion argues that movement itself is a major channel for conveying intent, affect, and interaction quality. Foundational work on legibility and predictability established that the same physical task can be interpreted very differently depending on motion style [10]. More recent work studies expressive motion generation for both non-anthropomorphic and humanoid robots, including in-context expressive sequence design, interactive motion-generation pipelines, and LLM-guided gesture design for HRI [12, 13, 11, 14]. These directions align with our central premise: the robot's motion surface has interaction meaning independent of binary task completion.

2.3. Human Behavior References and Cross-Embodiment Transfer

Our formulation is also informed by work that uses human demonstrations, human videos, and synthesized hand-object motions as structured inputs for robot behavior and manipulation learning [15, 16, 17, 18, 19, 20]. These papers aim at transfer across embodiment, task domains, or data modalities. We do not attempt full one-to-one imitation of human embodiment. Instead, we focus on extracting transferable interaction primitives—approach, pause, handover rhythm, and exit behavior—that can be mapped onto robot-executable fields. This positions our contribution closer to a systems abstraction and runtime interface than to an end-to-end learning method.

2.4. Preference, Personalization, and Interaction Evaluation

Behavior-layer claims are only meaningful if they can be evaluated. Work on robot behavior personalization, object handover, and service-oriented interaction suggests that human preference is not fixed and that timing, rhythm, and contact strategy strongly affect judged quality [21, 22, 23]. This motivates a behavior representation that can be rated, compared, and revised under the same language used to execute it. In this sense, our work is complementary to recent end-to-end humanoid system papers that integrate perception, reasoning, and embodiment into a unified stack [24]; our emphasis is narrower and sits one layer above low-level autonomy, at the interface where behavior becomes explicit and evaluable.

3. Problem Formulation

The problem is not that robots cannot complete tasks. The problem is that the same robot often has to be behaviorally rebuilt for each deployment setting.

Today, many human-facing behavior differences are handled through:

- one-off scene scripts,
- runtime parameters with weak semantic naming,
- local QA and operator tuning,
- property-specific deployment adjustments.

This leads to a recurring implementation pattern in which the behavior that users actually perceive remains difficult to compare and reuse. The robot may already know how to execute a task, but the deployment team still has no explicit software object for the robot's interaction tone.

We formulate the behavior-layer problem around four requirements:

- **R1: Explicitness.** Behavior differences should be surfaced through named fields rather than hidden local tuning.
- **R2: Runtime compatibility.** The behavior layer should not require replacement of the robot's low-level controller or safety system.
- **R3: Cross-task legibility.** A profile change should remain visible across more than one interaction type.
- **R4: Reference-grounded evaluation.** Behavior quality should be evaluable against structured references and human judgment, not only task success.

The present paper addresses these requirements at prototype level. It does not solve full transfer or benchmarking, but it makes the software boundary and the experimental claim explicit.

4. Behavior-Layer Formulation and System Instantiation

4.1. Interaction Profiles

The core abstraction in Motius is an **interaction profile**. An interaction profile is a reusable behavior object that specifies how the robot should feel during human-facing tasks. Rather than relying only on broad adjectives, a profile resolves into explicit fields that can be mapped into runtime behavior.

Typical fields include:

- approach speed scale,
- stopping distance,
- pause duration,
- hold duration during handover,
- finish softness,
- motion smoothing,
- end-state control preferences.

The profile is therefore not the controller itself. It is a structured software interface above the controller. In the long-term system, a profile is also not intended to be a single frozen vector. Instead, it is a *profile band*: a learned region of acceptable timing, spacing, and finish behavior that can shift for different users or scene types while still remaining recognizably within the same interaction style.

4.2. Adapter Layer and Runtime Boundary

We separate the system into two layers:

1. a **behavior layer**, which defines how the robot should behave in a scene, and
2. a **runtime layer**, which still owns low-level execution, planning, actuation, and safety.

This boundary is a key design decision. It allows behavior to become inspectable and reusable without claiming that high-level profile logic replaces the robot's full motion or autonomy stack.

Different robots and runtimes expose different control surfaces. A shared profile therefore needs an adapter layer that maps robot-agnostic behavior fields into robot-specific runtime values. The adapter handles:

- normalization of units and timing,
- scaling relative to the robot platform,
- mapping into runtime-specific fields,
- safety-aware fallback behavior.

In the current prototype, this mapping is concrete rather than only rhetorical. `speed_scale` acts as a multiplicative factor on nominal locomotion speed, `pause_ms` defines dwell time around arrival or transfer events, `approach_distance_m` sets the terminal stopping threshold for the approach phase, and `ee_smoothing` modulates the sharpness of the end-effector finish trajectory. These mappings do not replace the underlying controller; they parameterize a narrow runtime surface above it.

4.3. Human Reference Network

The system is not intended to rely only on robot clips. We therefore propose a second input surface: **human behavior references**. These references do not attempt to capture full human embodiment as a direct imitation target. Instead, they provide scene-specific examples of transferable interaction primitives such as:

- approach style,
- waiting behavior,
- handover rhythm,
- corridor etiquette,
- exit timing.

These references can be paired with labels describing task type, scene context, user cues, and behavioral tone. Over time, such a reference corpus can help define what robot behavior should feel like in a way that is more structured than local operator intuition. In the intended system design, this corpus is not only a

Table 1: Minimal schema for a human behavior reference item in the current Motius intake surface.

Field	**Example values**	**Role in the reference layer**
task_type	handover; approach-stop; wait-behavior	Binds the clip to a specific interaction primitive rather than a general robot scene
scene_type	hotel; front-desk; corridor; elevator	Encodes context in which the behavior should be interpreted
behavior_tags	warm; direct; patient; soft-finish	Supplies compact descriptors for downstream comparison and filtering
clip_duration	8–30 s	Keeps the contribution at the level of a specific interaction episode
notes	optional free text	Allows annotators to describe why the example is useful or atypical
review_status	validating; quality review; claimable	Tracks whether the contribution has passed intake and evaluation gates

Table 2: Representative human-reference entries compatible with the current Motius intake surface.

task_type	**scene_type**	**behavior_tags**	**duration**	**notes**
handover	hotel	warm; patient; soft-finish	14 s	guest-facing transfer example with delayed retraction after contact
approach-stop	front-desk	direct; calm; spaced	10 s	reception arrival example emphasizing conservative final stand-off
corridor-etiquette	corridor	yield; wait; pass-left	12 s	shared-space hallway example showing human-priority yielding behavior

passive archive; it is the front end of a data flywheel in which contributed references improve later adaptive parameter prediction, and later deployments generate new clips that further refine the reference base.

The current prototype already instantiates a minimal reference-ingestion surface rather than leaving this layer purely conceptual. In the live contribution flow, each uploaded reference clip is associated with a small structured record including task type, scene type, behavior tags, free-form notes, clip duration, and review status. Intake is currently constrained to short MP4/MOV/WebM clips in an 8–30 s range so that the collected references stay tied to specific interaction moments rather than broad, ambiguous scene footage. This means the present paper still does not contribute a public dataset, but it does define a concrete schema for what a reference item is supposed to contain.

Representative examples from the current intake design are shown in Table 2. These rows are design-time, schema-conforming examples rather than collected study data or a released benchmark corpus.

4.4. Validation Workflow

A behavior layer is only useful if it can be validated. We therefore define a structured validation workflow with three inputs:

1. human reference clips,
2. behavior labels and preference annotations,
3. robot execution proof clips.

Table 3: Parameter surface for the current Standard/Gentle comparison.

Field	**Standard**	**Gentle**	**Relative change**
speed_scale	1.00	0.82	-18.0%
pause_ms	200	420	+110.0%
approach_distance_m	0.80	0.95	+18.7%
ee_smoothing	0.90	0.96	+6.7%

Relative change is computed against the Standard profile.

```
profile: Gentle
task: object_handover
runtime_adapter:
  locomotion.speed_scale: 0.82
  arrival.pause_ms: 420
  proxemics.stop_distance_m: 0.95
  arm.ee_smoothing: 0.96
```

Figure 1: Minimal adapter instantiation for the Gentle profile in the current prototype. The behavior layer does not output motor torques or trajectories directly; it writes a small runtime-facing parameter surface above the controller.

This framework is designed to support not only task-completion checks, but also follow-on preference-aware evaluation: for example, whether one profile is perceived as more natural, more professional, or more appropriate than another in a given scene.

4.5. Current Prototype Surface

The current Motius prototype exposes a small profile family and evaluates one explicit pair: **Standard** and **Gentle**. A third profile, **Attentive**, is present in the prototype surface but is not part of the present comparison. The paper therefore treats Standard and Gentle as the evaluated conditions and the remaining profile family only as evidence that the abstraction is being instantiated as a reusable surface rather than as a one-off clip label.

The current prototype surfaces four explicit execution fields in the Standard/Gentle comparison:

- `speed_scale`,
- `pause_ms`,
- `approach_distance_m`,
- `end_effector_smoothing`.

These values are not presented as a benchmark. They are presented as evidence that the behavior layer is already explicit and parameterized rather than purely descriptive.

At the current stage, the Standard/Gentle proof pair should be read as the *base* configuration of a larger adaptive system rather than the final adaptive system itself. The present comparison uses fixed configured values for experimental readability, but the intended learned system treats those values as centers or defaults inside broader user-conditioned ranges. For example, a Gentle-style arrival may remain recognizably Gentle while slowing further and widening stand-off distance for an elderly user, or tightening toward a faster and more efficient exchange for a business traveler in a time-sensitive service scene. The current paper does not claim that such adaptation is already produced by a trained online model; it claims that the profile surface and reference-network design already support that direction.

Table 4: Illustrative adaptive parameter bands for the Gentle profile. The current paper does not claim that these user-conditioned values are already predicted by a trained model; they describe the target adaptation surface that the Reference Network is intended to support.

Condition	speed_scale	pause_ms	distance_m	Interpretation
Gentle base	0.82	420	0.95	Field-validated base setting used in the current proof pair
Elder-facing arrival	0.65	650	1.10	Slower approach, longer dwell, wider spacing for reassurance
Child-nearby arrival	0.50	800	1.25	More conservative braking and stand-off under higher uncertainty
Business-user arrival	0.90	250	0.90	Faster and more efficient exchange while preserving profile tone

Table 5: Clip-level proxy checks extracted from the current Unitree G1 proof asset. These measurements do not fully recover controller-state variables; they provide partial corroboration that the configured profile differences are reflected in visible timing and spacing cues in the recorded session.

Field	Std.	Gentle	Proxy	Std.	Gentle
speed_scale	1.00	0.82	terminal arrival window (s)	1.25	1.75
pause_ms	200	420	visible exchange window (s)	2.25	3.00
approach_distance_m	0.80	0.95	visor bbox height (px)	161	151
ee_smoothing	0.90	0.96	not recovered from monocular clip	n/a	n/a

The arrival and exchange windows were manually phase-coded from the raw clip. The image-space stand-off proxy uses HSV-based visor bounding-box height at the settled stop frame; a smaller box is consistent with greater final distance to the camera. End-effector smoothing could not be metrically recovered from the current monocular video and therefore remains unvalidated by this proxy method.

5. Prototype Implementation and Experimental Protocol

5.1. Research Questions

The present prototype study is designed to answer three narrow research questions:

- **RQ1:** Can explicit profile fields create visible cross-task behavior differences on the same Unitree G1 humanoid platform?
- **RQ2:** Can those differences be surfaced without replacing the robot's controller or autonomy stack?
- **RQ3:** Does an explicit profile surface provide a credible basis for subsequent human-reference and preference-based evaluation?

5.2. Platform and Assets

Our current proof is derived from a raw validation recording on the same Unitree G1 humanoid robot. The source video has the following properties:

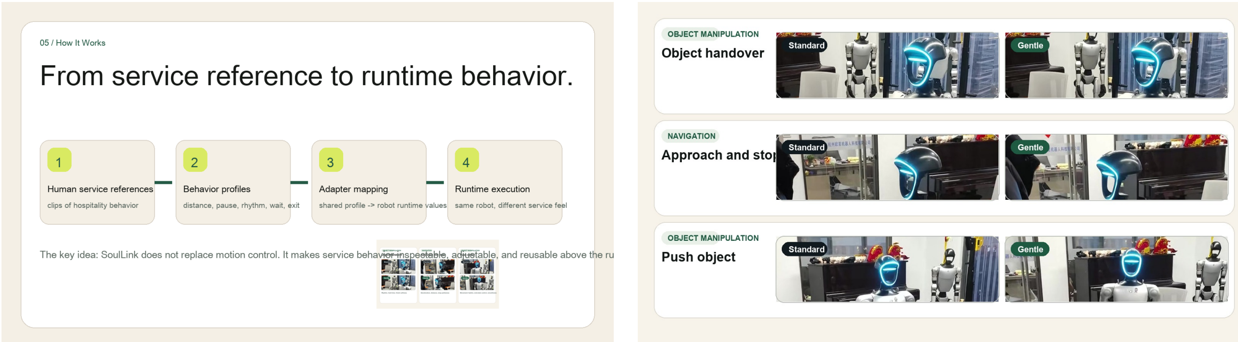

(a) Behavior-layer formulation from reference input to runtime execution.

(b) Controlled A/B proof on the same Unitree G1 humanoid platform.

Figure 2: System overview. The behavior layer does not replace motion control; it surfaces human-facing behavior above the runtime and makes the comparison observable on the same robot and task family.

Table 6: Proof assets used in this paper.

Asset	Duration	Format	Role in this paper
Raw field validation clip	54.1 s	720 × 1280, vertical, 30 fps	Source proof covering the three interaction comparisons on the same Unitree G1 humanoid robot
Edited side-by-side proof clip	43.0 s	1920 × 1080, 30 fps	Comparison artifact used here as a clearer observational layer
Public profile surface	live prototype	web implementation	Provides explicit profile names, scene fit, and parameter fields

- duration: 54.1 seconds,
- frame rate: 30 fps,
- orientation: vertical,
- resolution: 720 × 1280.

The current asset is not a benchmark dataset. It is a controlled prototype recording used to test whether one profile change remains legible across multiple tasks within a single recorded session. To make the comparison more legible, the raw footage was also converted into a 43.0-second side-by-side proof clip in which Standard and Gentle are shown with explicit comparison framing.

5.3. Controlled Comparison Design

The clip contains three interaction comparisons:

1. **Object handover**: the robot approaches a user and hands over an object.
2. **Approach-and-stop**: the robot walks toward the user and stops in front of them.
3. **Push object**: the robot interacts with a wheeled object or platform.

These interactions span both navigation-related and object-interaction-related behavior. This matters because it reduces the chance that a profile difference is only a one-task artifact within the recorded session.

The current comparison is controlled in a limited but meaningful sense. Across the Standard/Gentle comparison, the following are held fixed:

- robot body,
- filming environment,

Table 7: Task-level coverage and phase statistics derived from the edited and raw proof assets. Frame counts assume 30 fps.

Interaction	Duration	Frames	Share of edited clip	Raw-phase coverage	Primary visible shifts
Object handover	14.0 s	420	32.6%	0.0–14.0 s	closing pace, hold duration, exit softness
Approach-and-stop	9.0 s	270	20.9%	15.0–25.0 s	braking timing, stop distance, settle calmness
Push object	9.0 s	270	20.9%	25.0–54.1 s	contact smoothness, movement pace, finish control
Total direct comparison	32.0 s	960	74.4%	54.1 s raw clip	comparison coverage across all coded tasks

- task family,
- interaction ordering within each comparison window.

What changes is the active profile.

5.4. Phase Coding and Coverage Statistics

To make the proof asset more analyzable, we manually phase-coded the edited clip at 0.25 s granularity and the raw validation recording at 0.5 s granularity using timestamped keyframe sheets generated from the source media. The purpose of this coding was not to infer semantics beyond what the clip can support, but to make segment boundaries and evidence density explicit. The resulting task-level coverage statistics are summarized in Table 7.

5.5. Measurement Strategy

The present paper reports six classes of evidence:

1. **Parameter-level evidence**: explicit numerical differences in surfaced profile fields.
2. **Clip-level proxy evidence**: partial measurements from the raw video used to corroborate intended timing and spacing shifts.
3. **Task-coverage evidence**: visibility of the profile change across more than one interaction type.
4. **Phase-coded temporal evidence**: manually coded segment boundaries, durations, and frame counts in the edited and raw proof assets.
5. **Behavioral-reading evidence**: structured visual interpretation of the side-by-side comparison under fixed platform and task conditions.
6. **Human-rating evidence**: paired Likert judgments and profile preferences from external respondents over the same task clips.

The clip is organized around an A/B comparison between two profiles:

- **Standard**: neutral, direct, efficient.
- **Gentle**: warm, patient, soft.

At the execution level, the intended difference is reflected in fields such as shorter versus longer pause, tighter versus wider spacing, quicker versus softer completion, and more direct versus more gradual approach.

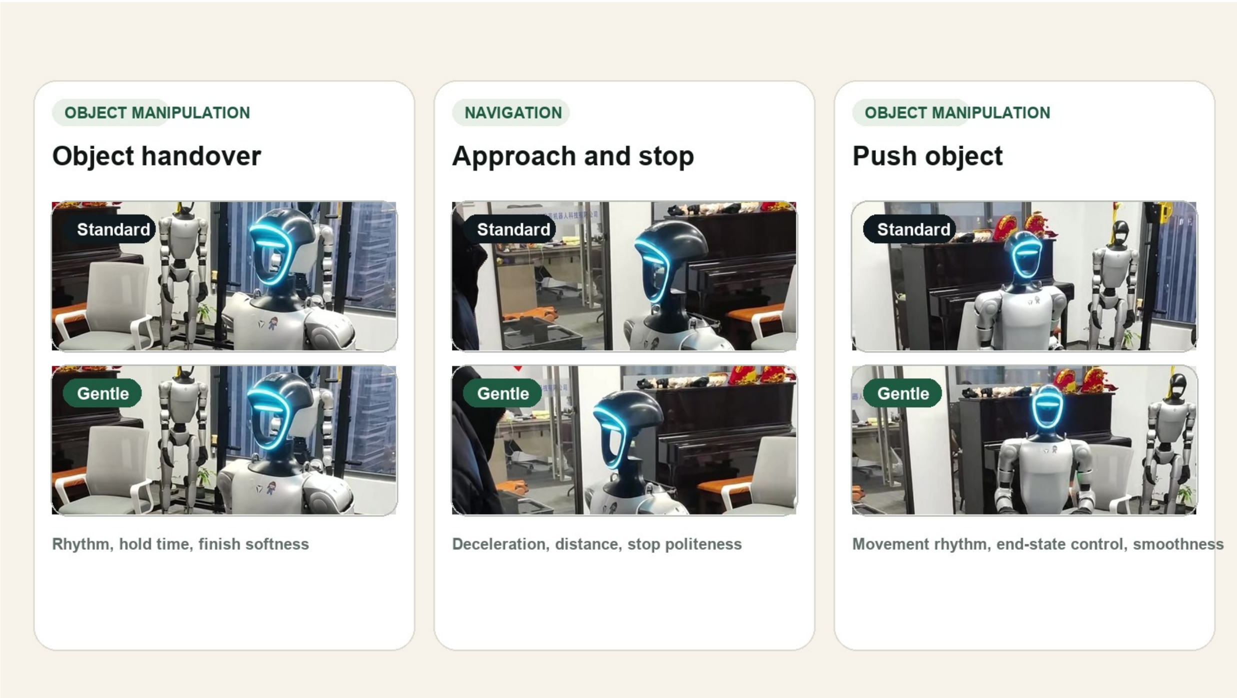

Figure 3: The current proof spans three within-session interactions on the same Unitree G1 humanoid robot: object handover, approach-and-stop, and push object.

Table 8: Interpretive summary of the two interaction profiles used in the current proof.

Profile	Interaction tone	Expected execution traits
Standard	Neutral, direct, efficient	Shorter pause, tighter spacing, quicker completion, more operational finish
Gentle	Warm, patient, soft	Longer pause, wider stop distance, softer finish, calmer arrival and exit

5.6. Lightweight Human Rating Study

To test whether the profile difference is not only visible to the authors but also legible to external observers, we ran a lightweight clip-based rating study over the same three interaction types. Following recent work that frames social navigation and interaction quality in terms of appropriateness, trust, comfort, and judged performance [3, 7, 8, 9], respondents viewed paired Standard/Gentle clips under fixed task framing and scored:

- perceived naturalness,
- politeness,
- clarity,
- appropriateness for the scene.

Each item used a 1–7 Likert scale. Respondents also answered a preference question for each task ('Standard', 'Gentle', or 'No clear difference'). The study was administered as a Google Forms survey. A total of $n = 24$ consented respondents completed the rating task. Recruitment was conducted through

Table 9: Participant profile for the human rating study ($n = 24$).

Dimension	Category	n	Category	n
English comfort	Very comfortable	1	Comfortable	8
English comfort	Neutral	7	Slightly uncomfortable	8
Robotics familiarity	Very familiar	2	Somewhat familiar	2
Robotics familiarity	Slightly familiar	8	Not familiar	12
Background	Student	6	Robotics researcher	5
Background	Product manager	4	Hospitality	3
Background	Operations	2	Engineer	2
Background	Designer / Marketing	2	–	–

a combination of social and messaging-channel circulation and direct offline outreach. Participation was voluntary, and no identifying information beyond self-reported familiarity and background category was collected. Self-reported familiarity with robotics was mixed: 2 respondents reported being very familiar, 2 somewhat familiar, 8 slightly familiar, and 12 not familiar. Backgrounds included students, robotics researchers, product managers, hospitality workers, operations roles, engineers, a designer, and a marketing respondent.

The study is intentionally lightweight and clip-based rather than deployment-based. Its purpose is not to establish broad population-level claims, but to test whether the behavior-layer abstraction aligns with the kinds of human judgment that the literature already treats as central. The resulting ratings should therefore be interpreted as pilot evidence about perceptual legibility under controlled presentation, not as representative estimates of population-level preference. The current manuscript likewise reports an anonymous online media-rating exercise rather than a formal in-person HRI trial. All respondents provided informed consent before proceeding, and the analysis is reported as a minimal-risk exploratory study rather than as a formally reviewed institutional human-subjects protocol. Importantly, the same fields used to execute the behavior—timing, spacing, smoothing, and finish quality—also define the language of evaluation. This closes a gap that often appears in robot deployment work, where runtime changes are measured in one vocabulary and user-facing quality is discussed in another.

6. Observed Cross-Task Findings

The present findings support a narrow claim: **one profile change is visibly legible across three within-session interactions on the same robot.**

6.1. Result 1: The behavior difference is not merely narrative

The Standard/Gentle comparison is backed by explicit surfaced fields rather than only broad descriptive labels. Table 3 shows an 18.0% decrease in speed scale, a 110.0% increase in pause duration, an 18.7% increase in stopping distance, and a 6.7% increase in end-effector smoothing. Table 5 adds clip-level corroboration from the same recording: the Gentle condition shows a longer terminal arrival window, a longer visible exchange window, and a smaller image-space robot size at the settled stop frame, all consistent with the intended slower, softer, and more conservative interaction surface. Even at this early stage, the behavior layer is therefore parameterized in a way that can be inspected, compared, and partially cross-checked against visible execution.

6.2. Result 2: The profile difference is visible across multiple within-session interaction types

The current proof is not limited to a single "hero" moment. It allocates 32.0 seconds of direct comparison coverage, corresponding to 960 comparison frames or 74.4% of the edited proof clip, across three interactions spanning both navigation-related and object-interaction-related behavior (Table 7).

In **object handover**, the visible difference is not only whether the robot reaches the correct position, but how the exchange unfolds. The Gentle variant appears to hold longer and exit more softly, while the

Table 10: Task-level human rating results ($n = 24$ respondents). Values are mean $\pm$ population standard deviation on a 1–7 scale. p-values are from paired two-sided Wilcoxon signed-rank tests comparing Gentle against Standard within each task and metric.

Task	**Metric**	**Standard**	**Gentle**	Δ	p
Handover	Naturalness	3.67 ± 1.37	4.92 ± 1.41	+1.25	0.0035
	Politeness	3.50 ± 1.38	4.88 ± 1.09	+1.38	0.0024
	Clarity	4.29 ± 1.31	3.96 ± 1.17	-0.33	0.35
	Appropriateness	3.50 ± 1.44	5.12 ± 1.30	+1.62	0.0015
Approach-stop	Naturalness	3.62 ± 1.75	4.96 ± 1.40	+1.33	0.011
	Politeness	2.88 ± 1.17	5.46 ± 1.38	+2.58	0.000071
	Clarity	3.50 ± 1.26	4.42 ± 1.32	+0.92	0.026
	Appropriateness	3.17 ± 1.34	5.00 ± 1.29	+1.83	0.0015
Push	Naturalness	4.17 ± 1.31	4.04 ± 1.62	-0.12	0.67
	Politeness	3.12 ± 1.39	5.04 ± 1.24	+1.92	0.0010
	Clarity	4.17 ± 1.60	4.12 ± 1.42	-0.04	0.91
	Appropriateness	3.17 ± 1.57	4.50 ± 1.53	+1.33	0.0056

Table 11: Pooled human rating summary across all task clips.

Metric	**Standard**	**Gentle**	Δ	p
Naturalness	3.82 ± 1.51	4.64 ± 1.54	+0.82	0.0015
Politeness	3.17 ± 1.34	5.12 ± 1.27	+1.96	1.6×10^{-9}
Clarity	3.99 ± 1.44	4.17 ± 1.32	+0.18	0.41
Appropriateness	3.28 ± 1.46	4.88 ± 1.40	+1.60	1.5×10^{-7}

Standard variant appears more direct and operational.

In **approach-and-stop**, the legible difference appears before any object exchange. The Gentle variant shows a calmer arrival with more space, while the Standard variant closes more directly.

In **push object**, the difference is expressed through movement rhythm, contact smoothness, and finish behavior, which suggests that the profile abstraction is not limited to one single interaction type.

This matters methodologically as well as visually. The comparison is not built from one short, hand-picked excerpt. Instead, the edited clip devotes a substantial majority of its runtime to direct A/B comparison, and the raw validation recording shows that these three task families are present as sustained phases rather than isolated single frames.

6.3. Result 3: Human ratings favor Gentle on politeness and appropriateness, with mixed clarity effects

Human ratings partially followed the intended profile contrast. Across all three tasks, the Gentle condition was rated higher than the Standard condition on politeness and scene appropriateness, with statistically significant within-task differences in every case (Table 10). Naturalness also favored Gentle in handover and approach-and-stop, but not in push-object. Clarity effects were weaker: Gentle was significantly clearer only in the approach-and-stop clip, while handover and push-object clarity showed no reliable advantage.

Preference choices were more mixed than in the earlier draft, and that is analytically useful. Gentle received more votes than Standard in all three tasks, but 'No clear difference' remained common in handover and approach-and-stop (Table 12). When restricted to decisive responses, the Gentle preference was strongest in approach-and-stop and push-object, while handover did not clear the same threshold under a two-sided binomial test.

At the pooled level across all task clips, Gentle exceeded Standard by 0.82 points on naturalness, 1.96 on politeness, and 1.60 on appropriateness, all with significant two-sided Wilcoxon results. Clarity, by contrast, showed no reliable pooled advantage for either profile. These results do not establish universal preference,

Table 12: Per-task profile preference choices from the human rating study. Binomial p-values are computed on decisive votes only (Gentle vs. Standard), excluding 'No clear difference'.

Task	**Gentle**	**Standard**	**No clear diff.**	**Gentle share (all)**	p
Object handover	11	3	10	45.8%	0.057
Approach-and-stop	15	1	8	62.5%	0.00052
Push object	17	3	4	70.8%	0.0026

but they do show that the profile surface is aligned with systematic human judgment on some dimensions more than others.

One result deserves separate interpretation: in object handover, Gentle increased politeness and appropriateness while slightly reducing clarity relative to Standard. A plausible reading is that the same longer dwell and softer finish that make the exchange feel warmer can also make the transfer moment less decisive when judged from video alone. In other words, softness and explicitness need not move together. This is useful rather than problematic for the behavior-layer thesis because it suggests that the surfaced profile fields expose real interaction trade-offs rather than merely producing one uniformly "better" style.

6.4. Result 4: The profile surface already extends beyond the proof pair

The current prototype exposes a small profile family—including Standard, Gentle, and Attentive—even though only Standard and Gentle are used in the present A/B study. This matters because it indicates that the underlying abstraction is being developed as a reusable profile surface rather than as a one-off comparison artifact. More importantly for the new system framing, these profiles are being treated as adaptive bands rather than endpoint presets. The present paper reports fixed proof settings for experimental clarity, but the same profile family is intended to support different parameter selections for different user types and scenes once the learned predictor matures.

6.5. Result 5: The software boundary remains narrow but credible

The current prototype does not replace motion control, safety, or planning. Instead, it sits above them. This boundary is important: it keeps the claim narrow enough to be credible while still making a useful research point. The evidence does not imply general robot intelligence or complete deployment readiness. It supports a more specific interpretation: the interaction layer itself can be surfaced as reusable adaptive software, even while the current adaptive instantiation remains rule-based rather than fully learned.

7. Discussion

7.1. Why an AI-adaptive behavior layer is a distinct systems contribution

Existing robotics systems often integrate perception, planning, control, and embodiment, but the human-facing behavior surface is still commonly hidden inside runtime-local decisions. Our contribution differs in emphasis. We treat behavior as an explicit software object that can be named, compared, mapped above the controller, and ultimately improved through data rather than only through local engineering retuning.

This is relevant to both socially aware navigation and expressive motion design. Work on proxemics, social presence, expressive motion, and appropriateness shows that users care deeply about movement quality [3, 4, 5, 11]. A behavior layer provides one way to operationalize those concerns inside an inspectable software surface. The current human rating results are still small-scale, but they now suggest a more precise claim: the surfaced abstraction is legible to outside observers on politeness and appropriateness more reliably than on every possible interaction dimension.

7.2. Why human references matter

Human-facing robot behavior should not be defined only by what is technically executable. It should also be informed by what people already recognize as appropriate, calm, direct, or professional. This is why the reference network matters. Its goal is not full human imitation, but the collection of transferable interaction primitives that can later be used for labeling, preference evaluation, and standard-setting. In the intended Motius system, the Reference Network is also the front end of a data flywheel: contributors upload clips, clips become structured references, references improve later adaptive prediction, and improved deployments produce more useful clips.

7.3. Why the current scope is still useful

The current scope is intentionally bounded, but it is still useful because it supports a concrete systems claim:

- behavior can be abstracted into reusable profile bands,
- those profiles can resolve into explicit fields,
- one profile change is visible across three within-session interactions on the same Unitree G1 humanoid platform,
- a reference-and-validation loop can be defined around this abstraction.

This does not replace the need for fuller quantitative studies, but it establishes an architectural and experimental starting point that is specific enough to be tested, criticized, and extended.

7.4. What the present study does not establish

The present results do not establish that the Gentle profile is unconditionally superior, nor do they establish that an adaptive behavior-layer system has already demonstrated absolute value over an unprofiled controller. A default-controller or no-profile baseline was not included in the present study. The observed Standard/Gentle differences should therefore be interpreted as relative effects between two active profile conditions rather than as proof that any adaptive behavior layer necessarily outperforms unparameterized robot control. This boundary matters because the current paper is about making behavior explicit, inspectable, comparable, and eventually learnable; it is not yet a full baseline-comparison study of all possible control regimes.

8. Threats to Validity

Several threats to validity should be made explicit.

First, the present evidence is based on a single robot and a single proof asset from one recorded session. The results may therefore reflect characteristics of this specific hardware, runtime, and session rather than a fully general behavior-layer property.

Second, the human rating study is small-scale and clip-based. Although all 24 respondents rated the same underlying tasks, they were judging edited comparison media rather than interacting with the robot in person. The results therefore support perceptual legibility under a controlled media setting, not full deployment-level acceptance.

Third, the current online rating task did not explicitly test randomization or counterbalancing effects across clip order, profile order, or task order. The judgments are therefore informative about the presented comparison, but they should not be interpreted as isolating all possible presentation-order confounds.

Fourth, the edited side-by-side clip improves legibility but may also introduce presentation effects. In particular, side-by-side A/B comparison can favor the slower or more deliberate variant even when the deployment setting might reward a clearer or more decisive response. The current paper partially mitigates this by (i) grounding the study in the underlying raw validation asset, (ii) allowing respondents to choose 'No clear difference' rather than forcing a binary preference, and (iii) reporting mixed outcomes where they occur. Additional controlled experiments are still needed.

Fifth, the current adapter boundary has not yet been evaluated across multiple robot platforms. The portability claim is therefore architectural rather than empirical at this stage.

Sixth, the current manuscript still does not document a formal compensation policy or institutional ethics review beyond informed consent at survey entry. That omission does not invalidate the present pilot media-rating evidence, but it does limit how strongly the study can be interpreted and should be corrected in follow-on evaluation work.

9. Limitations and Future Work

This work has several clear limitations. The present human rating study is small, clip-based, and limited to one proof asset from a single recorded session. We do not yet compare against a no-profile or default-controller baseline. We do not yet evaluate the adapter abstraction across multiple robot platforms. We do not yet publish a completed reference dataset or a formal benchmark for scene-specific acceptance thresholds. Most importantly for the adaptive framing, the current system represents an early proof-of-concept: the adaptive parameter logic described in this paper is demonstrated through a rule-based instantiation over surfaced profile fields, while the full learned profile predictor remains under development using accumulated Reference Network data.

Future work will focus on five areas:

1. a larger and more diverse human preference study comparing Standard, Gentle, and future profile pairs across the current task set;
2. structured collection of staged human behavior references for approach, waiting, handover, and exit behavior;
3. learned profile prediction that maps user and scene signals into runtime parameter bands rather than fixed hand-set values;
4. multi-robot adapter experiments that test whether the same profile language remains usable across different runtimes;
5. profile versioning and deployment workflows that make behavior changes easier to audit and roll back;
6. scene-specific behavior metrics that connect surfaced fields to downstream user preference and deployment acceptance.

One concrete near-term target is a follow-on rating study that moves beyond edited clips into more tightly controlled replay conditions or live interaction settings. Another is a baseline comparison against an unprofiled controller condition so that the behavior layer can be evaluated not only internally, but also against a realistic no-layer reference. A third is the publication of an annotation schema for service-oriented interaction primitives so that reference collection and robot validation can be studied under the same language.

10. Conclusion

This paper framed robot behavior as a data-and-software-layer problem rather than only a controller-tuning problem. We instantiated that framing in Motius through interaction profiles that resolve into explicit execution fields, an adapter boundary above the runtime, and a reference-and-validation loop built around human examples, labels, and robot proof clips.

The resulting evidence supports a narrow but concrete systems claim: on the same Unitree G1 humanoid robot, a single profile change can create visibly different interaction qualities across handover, approach-and-stop, and push-object tasks within one recorded session while leaving the low-level runtime in place. The accompanying human rating study further shows that some of these differences are not only visible in engineering terms, but also legible to outside raters, especially on politeness and scene appropriateness. That claim is smaller than full deployment autonomy or a finished learned adaptive stack, but it is specific enough to matter. If behavior can be surfaced in this way, and if the reference loop continues to accumulate real interaction data, then behavior can become something that is inspected, rated, adapted, and improved more systematically than today's deployment-local tuning workflows allow.

References

[1] Jarosław Karwowski, Wojciech Szynkiewicz, and Ewa Niewiadomska-Szynkiewicz. Bridging Requirements, Planning, and Evaluation: A Review of Social Robot Navigation. *Sensors*, 24(9):2794, 2024.

[2] Rashid Alyassi, Cesar Cadena, Robert Riener, and Diego Paez-Granados. Social Robot Navigation: A Review and Benchmarking of Learning-Based Methods. *Frontiers in Robotics and AI*, 12:1658643, 2025.

[3] Yunzhong Zhou, Jered Vroon, Zoltán Rusák, and Gerd Kortuem. PARSNiP: A Novel Dataset for Better Perceived Appropriateness Detection in Robot Social Navigation with Emotional and Attentional Features. *International Journal of Social Robotics*, 17(10):2245–2257, 2025.

[4] Kerry He, Wesley P. Chan, Akansel Cosgun, Albin Joy, and Elizabeth A. Croft. Robot Gaze During Autonomous Navigation and Its Effect on Social Presence. *International Journal of Social Robotics*, 16(5):879–897, 2024.

[5] Huseyin Bilen, Ahmet Burak Can, and Bulent Yanikoglu. Social Robot Navigation With Adaptive Proxemics Based on Emotions. *arXiv preprint arXiv:2401.10196*, 2024.

[6] Margot M. E. Neggers, Simon Belgers, Raymond H. Cuijpers, Peter A. M. Ruijten, and Wijnand A. IJsselsteijn. Comfortable Crossing Strategies for Robots. *International Journal of Social Robotics*, 16:1541–1560, 2024.

[7] Anthony Francis, Claudia Pérez-D'Arpino, Chengshu Li, Fei Xia, Alexandre Alahi, Rachid Alami, Aniket Bera, Abhijat Biswas, Joydeep Biswas, Rohan Chandra, Hao-Tien Lewis Chiang, Michael Everett, Sehoon Ha, Justin Hart, Jonathan P. How, Haresh Karnan, Tsang-Wei Edward Lee, Luis J. Manso, Reuth Mirsky, Sören Pirk, Phani Teja Singamaneni, Peter Stone, Ada V. Taylor, Peter Trautman, Nathan Tsoi, Marynel Vázquez, Xuesu Xiao, Peng Xu, Naoki Yokoyama, Alexander Toshev, and Roberto Martín-Martín. Principles and Guidelines for Evaluating Social Robot Navigation Algorithms. *ACM Transactions on Human-Robot Interaction*, 14(2), 2025.

[8] Yi Lian, Joanne Taery Kim, and Sehoon Ha. Implicit Behavioral Cues for Enhancing Trust and Comfort in Robot Social Navigation. *Advanced Robotics*, 39(11), 2025.

[9] Qiping Zhang, Nathan Tsoi, Mohammad Nagib, Byungjun Choi, Jenny Tan, Hao-Tien Lewis Chiang, and Marynel Vázquez. Predicting Human Perceptions of Robot Performance During Navigation Tasks. *ACM Transactions on Human-Robot Interaction*, 14(3):1–27, 2025.

[10] Anca D. Dragan, Kenton C. Lee, and Siddhartha S. Srinivasa. Legibility and Predictability of Robot Motion. In *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*, 2013.

[11] Till Hielscher, Andreas Bulling, and Kai O. Arras. Interactive Expressive Motion Generation Using Dynamic Movement Primitives. *arXiv preprint arXiv:2504.06735*, 2025.

[12] Yuhan Hu, Peide Huang, Mouli Sivapurapu, and Jian Zhang. ELEGNT: Expressive and Functional Movement Design for Non-Anthropomorphic Robot. *arXiv preprint arXiv:2501.12493*, 2025.

[13] Peide Huang, Yuhan Hu, Nataliya Nechyporenko, Daehwa Kim, Walter Talbott, and Jian Zhang. EMOTION: Expressive Motion Sequence Generation for Humanoid Robots with In-Context Learning. *arXiv preprint arXiv:2410.23234*, 2024.

[14] Liam Roy, Elizabeth A. Croft, Alex Ramirez, and Dana Kulić. GPT-Driven Gestures: Leveraging Large Language Models to Generate Expressive Robot Motion for Enhanced Human-Robot Interaction. *IEEE Robotics and Automation Letters*, 10(5):4172–4179, 2025.

[15] Nataliya Nechyporenko, Ryan Hoque, Christopher Webb, Mouli Sivapurapu, and Jian Zhang. ARMADA: Augmented Reality for Robot Manipulation and Robot-Free Data Acquisition. *arXiv preprint arXiv:2412.10631*, 2024.

[16] Ri-Zhao Qiu, Shiqi Yang, Xuxin Cheng, Chaitanya Chawla, Jialong Li, Tairan He, Ge Yan, Lars Paulsen, Ge Yang, Sha Yi, Guanya Shi, and Xiaolong Wang. Humanoid Policy $\sim$ Human Policy: Learning Humanoid Behaviors from Human Videos. *arXiv preprint arXiv:2503.13441*, 2025.

[17] Zifan Wang, Ziqing Chen, Junyu Chen, Jilong Wang, Yuxin Yang, Yunze Liu, Xueyi Liu, He Wang, and Li Yi. MobileH2R: Learning Generalizable Human to Mobile Robot Handover Exclusively from Scalable and Diverse Synthetic Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.

[18] Manlin Zhao, et al. DexH2R: Task-Oriented Dexterous Manipulation from Human to Robots. *arXiv preprint arXiv:2410.21903*, 2024.

[19] Simon Christen, et al. SynH2R: Synthesizing Hand-Object Motions for Learning Human-to-Robot Handovers. *arXiv preprint arXiv:2309.05688*, 2023.

[20] Yinhuai Wang, Qihan Zhao, Yuen Fui Lau, Runyi Yu, Hok Wai Tsui, Qifeng Chen, Jingbo Wang, Jiangmiao Pang, and Ping Tan. HumanX: Toward Agile and Generalizable Humanoid Interaction Skills from Human Videos. *arXiv preprint arXiv:2602.02473*, 2026.

[21] Chanwoo Kim, Joonhyeok Lee, Eunwoo Kim, and Kyungjae Lee. Time-Varying Preference Bandits for Robot Behavior Personalization. *Applied Sciences*, 14(23):11002, 2024.

[22] Zixi Wang, Zeyi Liu, Nicolas Ouporov, and Shuran Song. ContactHandover: Contact-Guided Robot-to-Human Object Handover. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.

[23] Haonan Duan, Yifan Yang, Daheng Li, and Peng Wang. Human–Robot Object Handover: Recent Progress and Future Direction. *Biomimetic Intelligence and Robotics*, 4(1):100145, 2024.

[24] Thanh Nguyen Canh, Thang Tran Viet, Thanh Tuan Tran, and Ben Wei Lim. SafeGuard ASF: SR Agentic Humanoid Robot System for Autonomous Industrial Safety. *arXiv preprint arXiv:2603.25353*, 2026.