

SIAgent: Spatial Interaction Agent via LLM-powered Eye-Hand Motion Intent Understanding in VR

Zhimin Wang, Chenyu Gu, and Feng Lu, *Senior Member, IEEE*

Abstract—Eye-hand coordinated interaction is becoming a mainstream interaction modality in Virtual Reality (VR) user interfaces. Current paradigms for this multimodal interaction require users to learn predefined gestures and memorize multiple gesture-task associations, which can be summarized as an “Operation-to-Intent” paradigm. This paradigm increases users’ learning costs and has low interaction error tolerance. In this paper, we propose *SIAgent*, a novel “Intent-to-Operation” framework allowing users to express interaction intents through natural eye-hand motions based on common sense and habits. Our system features two main components: (1) intent recognition that translates spatial interaction data into natural language and infers user intent, and (2) agent-based execution that generates an agent to execute corresponding tasks. This eliminates the need for gesture memorization and accommodates individual motion preferences with high error tolerance. We conduct two user studies across over 60 interaction tasks, comparing our method with two “Operation-to-Intent” techniques. Results show our method achieves higher intent recognition accuracy than gaze + pinch interaction (97.2% vs 93.1%) while reducing arm fatigue and improving usability, and user preference. Another study verifies the function of eye gaze and hand motion channels in intent recognition. Our work offers valuable insights into enhancing VR interaction intelligence through intent-driven design. Our source code and LLM prompts will be made available upon publication. Our project page is at <https://zhimin-wang.github.io/SIAgent.html>

Index Terms—Virtual reality, spatial interaction, intent recognition, agent-based execution, large language models

I. INTRODUCTION

Spatial interaction between users and virtual objects is central to creating immersive experiences in Virtual Reality (VR). Current natural spatial interaction inputs for VR include hand gesture, speech, and eye gaze [1]–[3]. While eye gaze offers rapid target selection, addressing the Midas Touch problem remains critical, where unintentional gaze triggers unwanted actions [4]. Hand motion enables complex operations and provides silent input, yet prolonged use may induce physical fatigue [5], [6]. Speech interaction provides a powerful hands-free modality, though its reliance on vocalization raises social acceptance concerns in shared environments [7]. To leverage these complementary strengths, eye-hand coordinated interaction combines the speed of gaze with the precision

of gestures, increasingly serving as a dominant paradigm in VR. For example, Apple Vision Pro has incorporated this multimodal interaction as a fundamental approach [8].

Common interaction paradigms for eye-hand coordination in VR divide the interaction process into three stages [9]–[12], as shown on the left of Fig. 1: 1) primary pointing: users direct eye gaze rays to point at a target; 2) confirmation: users utilize pinch gestures to confirm target selection; 3) manipulation: users employ one-hand or two-hand gestures, including grab, drag, rotate, pinch, and tap, in various combinations to complete tasks such as pouring water, writing, *etc.* However, this paradigm adopts an “Operation-to-Intent” approach that prioritizes low-level operation design over high-level intent. It focuses on mapping predefined gestures to interaction tasks, such as assigning “come/go/fist” gestures to control image zooming functions (zoom in, zoom out, or stop zooming) [2]. This approach presents several limitations as shown in Fig. 2. First, users must memorize multiple gestures and associate them with appropriate tasks, significantly increasing learning costs [9]. Second, users need to coordinate attention between gaze positioning and hand gesture execution, which compounds complexity and mental burden. Third, users must perform precise hand gesture matching, and recognition accuracy issues both increase interaction errors [6], [13].

To address these limitations, we propose an interaction approach based on an “Intent-to-Operation” framework, where users first express their intent, the system translates this intent into concrete operations, and then generates an agent to execute the task accordingly. For tasks such as pouring or writing, users express their intent through natural eye-hand movements informed by their own common sense and habits, rather than learning predefined interaction stages or gestures. This design enables intuitive interaction, removing the need for learning or adaptation to specific protocols. By focusing on high-level intent recognition instead of precise gesture matching, our approach accommodates flexible intent expressions and is highly tolerant of recognition errors, as long as the user’s actions align with their intents.

Building on this foundation, we leverage Large Language Models (LLMs), which have shown strong capabilities in understanding user intent from text and improving interaction efficiency and usability [14], [15]. Trained on extensive text corpora, LLMs effectively encode common-sense knowledge about human activities [16], [17], enabling comprehension of everyday tasks and associated tools. However, applying LLMs to VR spatial interaction remains challenging due to the

Zhimin Wang, Chenyu Gu, and Feng Lu are with the State Key Laboratory of Virtual Reality Technology and Systems, School of Computer Science and Engineering, Beihang University, Beijing 100191, China. e-mail: {zm.wang \ gucy \ lufeng}@buaa.edu.cn.

Feng Lu is the corresponding author.

Manuscript received July 17, 2025.

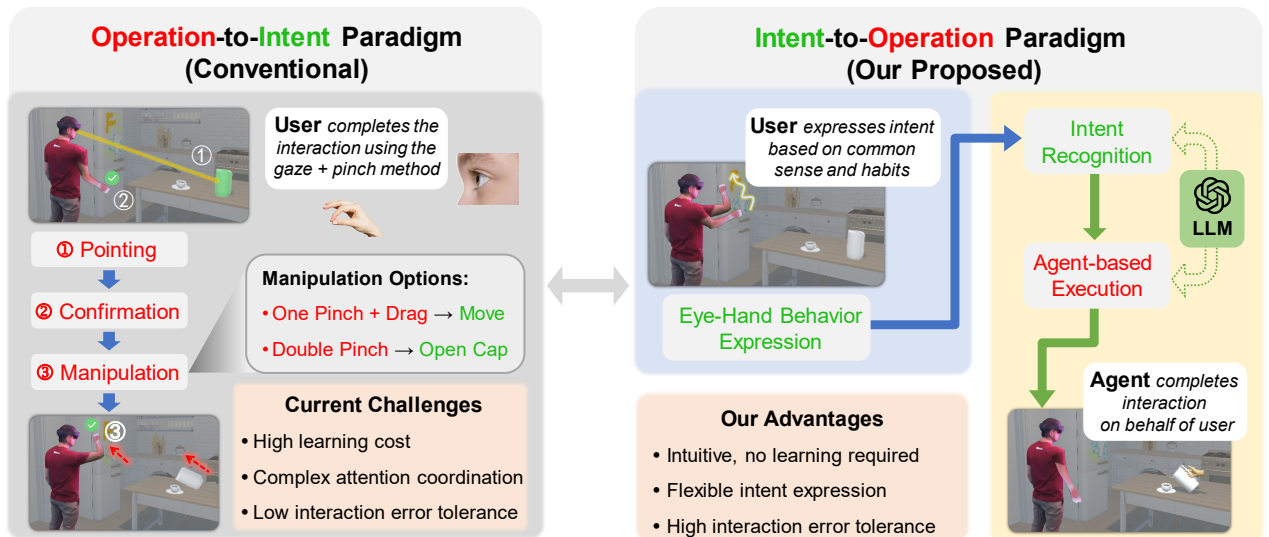


Fig. 1: The conventional Operation-to-Intent paradigm (left) requires users to perform sequential operations such as gaze-based pointing, gesture confirmation, and manipulation, resulting in high learning costs and low error tolerance. In contrast, our proposed Intent-to-Operation paradigm (right) allows users to naturally express intent through eye-hand motion, which is then recognized and executed by an agent powered by LLMs, enabling intuitive, flexible, and robust interaction.

inherent gap between natural language processing and the predominantly spatial, non-linguistic nature of VR interactions.

To bridge this gap, we introduce *SIAgent*, a spatial interaction agent via LLM-powered eye-hand motion intent understanding in VR, as shown on the right of Fig. 1. *SIAgent* comprises intent recognition and agent-based execution. After users demonstrate their intent through actions, *SIAgent* analyzes eye gaze and hand motion data and translates them into natural language descriptions. The system then employs the LLM to infer user intent and present likely options for user selection. Finally, the LLM generates motion or effect trigger data to guide the agent to execute tasks. We validate our approach through two user studies. The first compares our method with two “Operation-to-Intent” techniques (controller-based and gaze + pinch (GP) interaction) across over 60 spatial interaction tasks. Our method achieves higher intent recognition accuracy than GP (97.2% vs 93.1%) and superior performance across arm fatigue, usability, novelty, and user preference metrics. The second study examines the individual contributions of eye gaze and hand motion channels in intent recognition and explores intent expression via eye gaze combined with speech. Our findings provide valuable insights for advancing intelligent VR interactions.

The specific contributions of our work are three-fold:

- 1) We propose a novel “Intent-to-Operation” interaction paradigm for VR that allows users to express interaction intent through natural eye-hand motions without learning predefined gestures, fundamentally shifting from operation-centric to intent-centric interaction design.
- 2) We design and implement *SIAgent*, an LLM-powered spatial interaction agent that translates multimodal eye-hand motion data into natural language descriptions for intent recognition and generates executable agents to complete corresponding tasks.
- 3) Through extensive user studies across over 60 spatial interaction tasks, we demonstrate that our method

achieves higher intent recognition accuracy (97.2% vs 93.1%) than gaze + pinch interaction and significantly outperforms traditional “Operation-to-Intent” techniques in arm fatigue reduction, usability, and user preference.

II. RELATED WORKS

In this section, we review single-modal and multimodal interactions in VR, followed by a discussion of AI and LLM applications in VR.

A. Single-Modal Natural Interaction in VR

Single-modal natural interaction has been extensively researched in VR. We examine three primary modalities: hand gestures, eye gaze, and speech.

Hand gestures. Valued for their intuitiveness and ability to simulate real-world actions [5], [18], gestures facilitate actions such as pinching, non-linear interaction [19], and dual-hand manipulation [6]. However, traditional gesture-based systems require users to memorize predefined gesture vocabularies and perform precise matching [20], [21]. Extended use leads to muscle fatigue [22], while tracking errors and occlusion compromise accuracy [23], [24].

Eye gaze. This method is favored for its minimal physical effort and seamless integration [25], [26], enabling rapid object selection. However, challenges include the “Midas touch” problem and reliability issues [27], [28]. Gaze-based systems typically require explicit confirmation mechanisms that interrupt natural interaction flow.

Speech interaction. Speech enables natural expression through LLM integration [29] and silent speech recognition offers solutions to social acceptance concerns [7]. However, speech interactions face recognition latency, accuracy issues, and increased cognitive load for command memorization [30]. Traditional speech systems require users to learn specific command vocabularies and syntax.

Single-modal methods fundamentally require users to learn and execute predefined operations to express intentions, highlighting the need for more intelligent approaches.

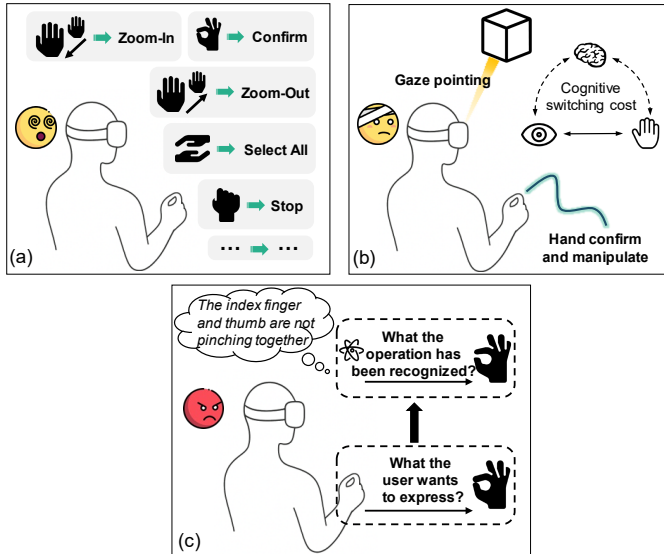


Fig. 2: This Operation-to-Intent paradigm presents several limitations. (a) Users must memorize multiple gesture-task associations, raising learning costs. (2) Users need to coordinate gaze and hand movements, increasing complexity and mental burden. (c) Users must achieve precise gesture matching, with recognition accuracy issues boosting interaction errors.

B. Multimodal Interaction in VR

We discuss two multimodal interactions relevant to our work: hand-eye coordination and gaze-speech interaction.

Hand-eye coordination. Recent VR studies leverage gaze for rapid selection and hand gestures for precise manipulation [12], [31]. However, this approach typically mandates sequential low-level operations, involving gaze-based selection followed by gesture-based confirmation and manipulation. This requirement increases cognitive load by demanding strict coordination adherence [28].

Gaze-speech interaction. This hands-free approach combines gaze-based selection with voice commands for manipulation [32], [33]. However, it may hinder precise object manipulation in complex spatial tasks, increasing cognitive load [9], [34]. Speech ambiguity and lack of real-time feedback further limit its utility [3]. Users must still decompose their intents into gaze selection followed by appropriate voice commands.

Although recent work has begun incorporating visual attention into multimodal LLMs to refine intent understanding [35], the interaction paradigm remains constrained by predefined operation sequences. This dependency imposes high learning costs, highlighting the need for intelligent systems that can interpret user intent from natural multimodal expressions.

C. Applications of AI and LLMs in VR

AI and LLMs are increasingly integrated into VR systems to enhance intelligence and user experience.

AI Applications in VR. AI improves VR by enabling intelligent NPCs [36], automating virtual environment creation [37], enhancing motion prediction for real-time accuracy [38], and adapting content based on eye movement [4]. These applications demonstrate AI’s potential to reduce user workload by automating complex behaviors and adapting to user patterns.

LLM Applications in VR. LLMs represent a paradigm shift in VR, enhancing capabilities ranging from fundamental text entry [14] and spatial interactions [16] to intelligent assistance in creation [39] and social contexts [15]. These applications collectively leverage LLMs’ superior understanding capabilities to interpret natural language descriptions and generate appropriate virtual content.

LLMs possess extensive common-sense knowledge about human activities, enabling understanding of underlying user intents. By leveraging LLMs’ understanding capabilities, VR systems could interpret user intent from natural multimodal expressions, supporting the transition from Operation-to-Intent to Intent-to-Operation paradigms.

III. MOTIVATIONS AND CHALLENGES

To address the limitations of Operation-to-Intent paradigms, we aim to build an Intent-to-Operation system that enables users to interact naturally based on their own habits and intents. This approach presents several key challenges:

Q1: How can we enable users to express interaction intent through natural movements without learning predefined gestures? Current Operation-to-Intent paradigms impose high learning costs by requiring users to memorize specific gesture-to-task mappings [9], [11]. This motivates the design of systems that interpret intent from natural hand movements rooted in daily common sense, thereby eliminating the need for explicit gesture training.

Q2: How can we bridge the gap between spatial VR interactions and language-based AI understanding? While LLMs excel at understanding user intent from text and possess extensive common-sense knowledge about human activities, VR interactions are predominantly spatial and non-linguistic in nature. How can we effectively convert spatial interaction data into natural language descriptions that preserve essential intent information for LLM processing?

Q3: How can we achieve robust intent recognition that accommodates flexible intent expressions? Unlike predefined gesture systems that require precise matching, an Intent-to-Operation approach must handle variations in how different users express the same intent. Users may perform similar tasks with different hand movements, gaze patterns, or interaction styles based on their personal habits and preferences. How can we maintain high tolerance for recognition errors while accurately inferring user intents across diverse motion expressions?

IV. DESIGN OF SIAGENT

Building on the Intent-to-Operation paradigm, we design SIAGENT to include a three-stage pipeline architecture as illustrated in Fig. 3. **Stage 1: Spatial-to-Linguistic Translation** employs LLMs to convert multimodal spatial interaction data (eye gaze and hand motion) into natural language descriptions. This module captures user demonstrations through three complementary channels: gaze target analysis, hand pose analysis, and finger shape analysis, ensuring robustness against individual variations while preserving semantic intent. **Stage 2: Intent Recognition** leverages LLM’s common-sense reasoning to infer user intents from linguistic descriptions and provides ranked intent options for user confirmation. We separate spatial interaction understanding from intent recognition

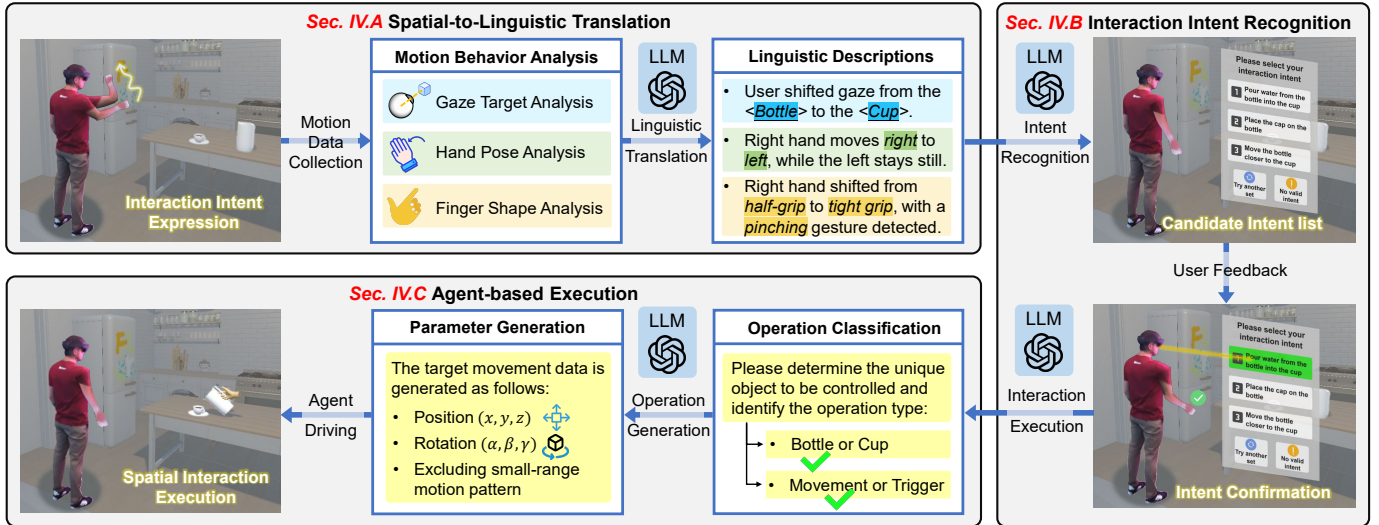


Fig. 3: The pipeline of SIAgent. User eye-hand motions are first captured and translated into natural language descriptions through spatial-to-linguistic translation. The LLM then performs intent recognition to infer possible user intents for selection. Subsequently, the LLM generates executable parameters for agent-driven spatial interaction based on the confirmed intent.

to allow each stage to specialize in distinct cognitive processes, thereby improving overall system accuracy. **Stage 3: Agent-Based Execution** translates recognized intents into spatial operations, distinguishing between movement operations and trigger operations while using a virtual agent to execute realistic interaction sequences. The detailed implementation of spatial-to-linguistic translation (Section IV-A), intent recognition (Section IV-B), and agent-based execution (Section IV-C) are presented in the following subsections.

A. Spatial-to-Linguistic Translation

For a given task, users first express their intent through demonstrative actions guided by their common sense and habits. To infer user intent, we analyze both visual focus and hand motion data. Visual targets are precomputed by integrating the user’s eye gaze with the 3D positions and sizes of scene objects. Hand data analysis encompasses the spatial position and rotation of both hands, as well as gesture configurations, which capture specific hand poses and finger shapes. We utilize LLMs to interpret this multimodal interaction data. The specific analyses are detailed below.

1) Gaze Target Analysis. We identify objects of user interest and attention shifts. Due to the limited accuracy of gaze tracking for VR headsets, our analysis focuses on object-level targeting rather than precise spatial coordinates or specific regions. Based on pilot exploration indicating that users can express their intent within 3 seconds, we set a fixed data acquisition window of 3 seconds. Eye-tracking data is sampled at 30 Hz. However, we observed that providing all frames to the LLM significantly increases processing time and decreases analysis accuracy due to excessive token length. To address this issue, we downsample the data by selecting every fifth frame, resulting in a total of 18 gaze data points as input. Each data point includes the timestamp, whether the user is fixating on an interactive object, and the target object’s name.

The LLM prompt is designed to capture three gaze states: (1) the user continuously gazes at object A; (2) the user shifts

their gaze from object A to object B; and (3) the user shifts their gaze from object A to no object. This translation is formulated as:

$$LLM(\{T_i\}_{i=1}^{18}, P_{\text{gaze}}) \mapsto D_{\text{gaze}}, \quad (1)$$

where T_i denotes the gaze target data, P_{gaze} is the prompt template, and D_{gaze} represents the natural language description of gaze attention. An example of the LLM input, output, and prompt is shown in Fig. 1(a) of the supplemental material.

2) Hand Pose Analysis. To infer user intent during object manipulation tasks, it is essential to analyze hand pose changes relative to the user’s body coordinate system, as these changes convey rich movement semantics. For example, in a water-pouring from bottle task, the right hand typically moves from right to left relative to the body. Capturing such semantics enables more accurate intent recognition. We establish the body coordinate system by taking the head position at the start of the task as the coordinate origin. Hand data are sampled at 30 Hz. To reduce data redundancy and processing cost, we compute hand position and rotation angle changes every five frames, resulting in 18 data points per trial.

The LLM prompt is designed to describe hand movement patterns: (1) the movement direction of each hand relative to the body (*e.g.*, left/right, up/down), (2) the magnitude of hand rotation (*e.g.*, whether the rotation is significant), and (3) whether the hands are moving closer to or farther from each other or the head. We formalize the LLM-based translation as:

$$LLM(\{H_i^L, H_i^R\}_{i=1}^{18}, P_{\text{hand}}) \mapsto D_{\text{hand}}, \quad (2)$$

where $H_i^{L/R}$ denotes poses of each hand at frame i , including palm position and rotation angle change. P_{hand} is the prompt template, and D_{hand} provides a description of hand movement patterns. An example of the LLM input, output, and prompt is shown in Fig.1(b) of the Supplementary Material.

3) Finger Shape Analysis. Rather than relying on predefined gestures, we analyze natural interaction patterns to identify five intuitive hand states without requiring user training




Fig. 4: Based on observations of finger shape patterns during hand-object interaction, five hand types are defined: (a) open, (b) half-grip, (c) tight-grip, (d) tip pinch, and (e) index tap.

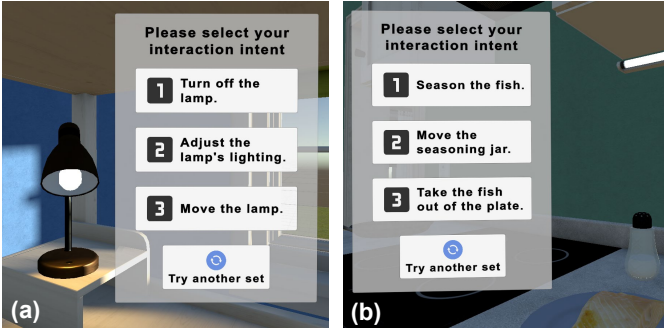


Fig. 5: We demonstrate intent recognition results for two tasks: (a) adjusting the lamp’s lighting, and (b) seasoning the fish.

(Fig. 4). We categorize fundamental grasping into open for reaching or releasing, half-grip for partially closing around large objects, and tight-grip for securing items requiring stability [40]. Furthermore, we distinguish fine manipulation states, specifically tip pinch which employs fingertips to grasp slender items, and index tap which extends the finger for activation. This categorization abstracts natural interactions into interpretable states directly tied to user intent.

Specifically, we track dynamic grasp adjustments by quantifying finger flexion at the metacarpophalangeal joint and curl at the distal joints based on established thresholds [41]. These continuous values are encoded into Boolean vectors representing the bent or extended status of each finger. The LLM then analyzes these vectors against typical configurations to classify fundamental states (*e.g.*, open, half-grip, tight-grip) and identify fine manipulation gestures like tip pinch or index tap. The translation is formalized as:

$$LLM(\{F_i^L, F_i^R\}_{i=1}^{18}, P_{\text{finger}}) \mapsto D_{\text{finger}}, \quad (3)$$

where $F_i^{L/R}$ represents vectors of flexion and curl for each finger at frame i , P_{finger} denotes the prompt template, and D_{finger} provides description of grasp patterns and transitions. An example of LLM input, output, and prompt is shown in Fig.1(c) of the Supplementary Material.

B. Interaction Intent Recognition

Based on the linguistic representations derived from eye-hand motions, we infer the user’s possible interaction intents. We combine the aforementioned descriptions with contextual

object state information to disambiguate user intentions and generate ranked intent predictions. The LLM-based inference is formalized as:

$$LLM(D_{\text{linguistic}}, O_{\text{state}}, P_{\text{intent}}) \mapsto I_{\text{intent}}, \quad (4)$$

where $D_{\text{linguistic}} = D_{\text{gaze}} + D_{\text{hand}} + D_{\text{finger}}$ represents the concatenated descriptions of eye-hand motions, O_{state} denotes the current states of target objects, P_{intent} represents the prompt template for intent recognition, and I_{intent} provides the ranked list of probable interaction intents with confidence scores.

Target state information O_{state} captures objects with distinguishable states (*e.g.*, light on/off, door open/closed) to resolve ambiguities where identical gestures may indicate opposing intents. For objects without clear state variations (*e.g.*, a fruit knife), the state is marked as “no special state.”

We design a structured prompt P_{intent} to guide the LLM’s intent recognition process based on three key principles:

- 1) *Multi-object Priority*: When gaze data involves multiple objects, the system prioritizes intents that involve both objects simultaneously and considers interaction sequences. Single-object intents are analyzed only when multi-object interactions are implausible.
- 2) *State-based Disambiguation*: The LLM initially analyzes intent without state information to avoid bias. State data is incorporated only when the inferred intent exhibits ambiguity (*e.g.*, open/close actions), serving as a disambiguation factor rather than a primary determinant.
- 3) *Ranked Output Generation*: The system generates six ranked intents with probability scores ranging from 0 to 100, with higher scores indicating higher likelihood. The interaction target names in the intents must match those identified in the gaze target analysis.

This design leverages gaze as a natural attention indicator. We maintain semantic understanding as the primary basis while using state information for supplementary disambiguation. The multi-candidate approach accommodates LLM uncertainty and provides alternatives when the primary prediction fails. We highlight high-confidence intents ($\geq 90\%$) for immediate attention. The prompt structure and example outputs are illustrated in Fig.1(d) of the Supplementary Material.

Interactive Intent Confirmation. Users select from LLM’s ranked predictions, providing explicit feedback for intent determination. The top three intents are displayed initially, with additional options accessible through interface controls, as shown in Fig. 5. This interactive process ensures accurate intent capture while generating ground truth data for potential system refinement through learning mechanisms in the future.

C. Agent-based Execution

Once user intent is recognized, we translate it into executable VR operations via a structured LLM-based approach. Our system decomposes spatial interactions into two real-world manipulation patterns: 1) *movement operations* that relocate objects in 3D space, such as placing a pencil in a holder, and 2) *trigger operations* that activate object-specific behaviors like turning on a lamp. For movement, the LLM generates target position and rotation descriptions to drive

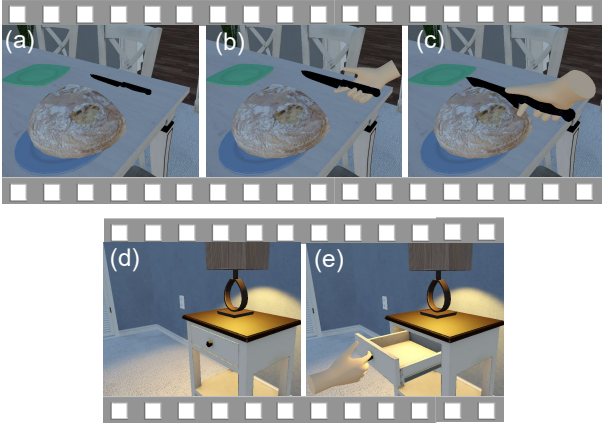


Fig. 6: The execution of spatial interactions is based on interaction intents: a) using a knife to cut bread, and b) opening the drawer.

object relocation. For triggers, since the LLM cannot directly render visual effects, it selects from predefined object-specific effects based on interaction intent. While current implementation relies on these definitions, recent advances in 3D diffusion models [42] offer promising avenues for future intelligent visual generation. The generation process is formalized as:

$$LLM(I_{\text{intent}}, O_{\text{info}}, E_{\text{effects}}, P_{\text{execution}}) \mapsto M_{\text{move}} \vee T_{\text{trigger}} \quad (5)$$

where I_{intent} represents the recognized intent, O_{info} contains target attributes (pose, mobility, size), E_{effects} denotes defined trigger effects for each object, $P_{\text{execution}}$ represents the prompt template for interaction execution, M_{move} specifies target position and rotation angles for movement operations, and T_{trigger} indicates selected visual effect for trigger operations.

We design a structured prompt $P_{\text{execution}}$ to guide the LLM’s execution generation process based on three key principles, with the example of prompt, input data and output result illustrated in Fig.1(e) of the Supplementary Material:

- 1) *Single-Object Focus*: The LLM identifies a single manipulation target from the intent. Complex multi-step tasks are pre-decomposed into sequential single-object operations (e.g., “place apple in refrigerator” becomes: open refrigerator → move apple inside → close refrigerator).
- 2) *Operation Type Classification*: The system analyzes intent against object capabilities to select between movement operations for spatial displacement and trigger operations for object-specific behaviors. If neither condition is met, no operation is output.
- 3) *Parameter Generation*: For movement, the LLM generates target poses and determines additional small-scale motion patterns are needed to simulate more realistic interactions (e.g., writing with a pencil). For triggers, it selects appropriate visual effects from the predefined set via semantic matching.

Virtual Agent. All LLM-inferred operations are executed by a virtual hand agent (Fig. 6). For movement tasks, the agent grasps and relocates the target to the generated pose; for example, in “pouring water,” the hand agent picks up the bottle and, using the target position provided by the LLM, moves




Fig. 7: The scenes used in the user study. (a) Study room, (b) Bedroom, (c) Living room, and (d) Kitchen. Scenes (c) and (d) constitute a single, navigable environment in the experiment.

it near the cup and tilts it at the required angle. For trigger tasks, the agent performs contact gestures (e.g., tapping) while the system activates corresponding visual effects, such as illuminating a lamp upon a tap. This approach mirrors real-world manipulation patterns. While we currently use a hand model, recent diffusion models for human-object interaction [43], [44] suggest potential for sophisticated humanoid agents in future work.

D. System Implementation and Apparatus

Our SIAgent is implemented using a Pico 4 Pro VR headset and remote server, with data transmitted via UDP protocol. The headset is equipped with eye tracking, hand tracking, and speech recognition capabilities. The VR scene was developed using Unity engine (v2022.3.21f1) and rendered at native resolution (2160×2160 pixels per eye) and framerate (90 Hz). The remaining technical steps are executed on the remote server, implemented in Python (v3.10.14). The server calls the LLM API either remotely or locally. The locally deployed LLM runs on an NVIDIA RTX 3090 GPU.

Our prompts employ an instruction-examples structure that accommodates LLMs without relying on version-specific features. This design ensures seamless adaptation to newer LLM versions, such as GPT-4. In the user study, we use ChatGLM-4 (cloud) [45] unless otherwise specified.

V. EVALUATION DESIGN

A. Research Objectives

To comprehensively evaluate our SIAgent system, we conduct user studies where participants perform typical VR interaction tasks across multiple environments. This evaluation involves both objective and subjective comparisons across different interaction modalities and input channels. Our objective is to address the following research questions:

RQ1: How does intent-driven agent execution perform compared to conventional VR interaction modalities? We investigate whether the Intent-to-Operation paradigm can achieve comparable or superior performance to conventional Operation-to-Intent approaches across objective metrics (time, accuracy) and subjective measures (usability, fatigue, user preference). This question will be investigated in Study 1.

TABLE I: Twenty-one ambiguous interaction tasks demonstrating multiple possible intents for identical gaze targets.

Objects	Interaction Tasks
Laptop	Open it ◦ Close it ◦ Power on it ◦ Power off it
Desk lamp	Turn on it ◦ Adjust its brightness ◦ Turn off it
Window	Open it ◦ Close it ◦ Clean it
Bottle	Open its cap ◦ Shake it ◦ Close its cap ◦ Move it
Pen and book	Place the pen on the book ◦ Write on the book with the pen
Guitar	Fetch it from a distance ◦ Play it
Washing machine	Open its lid ◦ Close its lid ◦ Start it

RQ2: What are the individual contributions of eye gaze and hand gesture channels in intent recognition? We investigate the effectiveness of combining gaze information with hand gestures compared to using gaze information alone for intent recognition, particularly in ambiguous interaction scenarios where multiple intents are possible. The contribution of each modality channel to the overall system performance will be investigated in Study 2.

RQ3: How does our LLM-based intent recognition framework perform with alternative input modalities? To investigate the generalizability of our approach beyond eye-hand motion, we explore the effectiveness of eye-speech input as an alternative modality for intent expression in VR environments. This question will also be investigated in Study 2.

B. Participants

We recruited 16 participants (12 males, 4 females) aged 21-29 years ($M = 24.3$, $SD = 2.2$) from our university campus. Pre-study questionnaires using a 5-point Likert scale revealed varying familiarity levels: low with gaze-based interaction ($M = 2.5$), medium with hand gesture interaction ($M = 3.1$) and VR systems ($M = 3.5$), and high with controller-based interaction ($M = 3.9$). Two participants were left-handed. All participants had normal or corrected-to-normal vision and were fluent in English. The studies were approved by the Biology and Medical Ethics Committee of Beihang University under Application No. BM20250179, and performed in line with the Declaration of Helsinki, and informed consent was obtained from all participants. All participants received a gift valued at approximately 15 CNY (or \$2 USD) as compensation for their time.

C. Task Design and Scenarios

To address our research questions, we designed 60 interaction tasks within realistic VR environments depicting daily indoor scenarios, as detailed in Tab. I of the Supplementary Material. These environments provide rich diversity of object types and interaction semantics that users encounter in everyday life, allowing us to evaluate whether our intent-driven approach can handle familiar interaction patterns without requiring users to learn new gestures or protocols. Tasks are categorized into *movement operations* and *trigger operations* as described previously.

To evaluate the contribution of each input channel, we design 21 ambiguous tasks where the same interaction object can support multiple interaction types. These tasks require

comprehensive analysis of both gaze targets and hand motion to distinguish user intent. These tasks comprise 11 from the previous task set and 10 newly supplemented scenarios, as detailed in Tab. I. The experimental environments are illustrated in Fig. 7 and include the following:

Study Room (19 tasks): This environment provides a study workspace with office supplies (*e.g.*, books, laptop, markers, erasers). This setting allows us to evaluate intent recognition for fine-grained manipulation tasks involving writing, drawing, and document handling where precise hand movements and tool usage patterns are critical.

Bedroom (18 tasks): This environment offers a mix of personal care items and household appliances (*e.g.*, bedside lamp, bottle, washing machine, cups) requiring different interaction strategies. This setting enables evaluation of intent recognition for state-based interactions where the same objects can be opened, closed, turned on, or moved depending on user intent.

Living Room & Kitchen (23 tasks): This combined environment provides the rich variety of home management activities (*e.g.*, TV cabinet, knife, bread, coffee cup). This setting allows evaluation of intent recognition for complex manipulation sequences involving food preparation, appliance operation, and multi-step cooking tasks that require understanding of procedural interactions.

D. Evaluation Metrics

To comprehensively evaluate our intent-driven agent system, we define metrics that assess task completion performance, system accuracy, and user experience across different interaction modalities.

Objective Measures. We define quantitative metrics to evaluate system performance and interaction efficiency.

- 1) **Task Completion Time:** Mean time to complete tasks. For SIAgent, we measure time consumed by each component: user intent expression time (U), LLM inference time (L), intent confirmation time (I), and agent execution time (A). We report three time metrics: Total time $Ag_t = U + L + I + A$; total interaction time $Ag_t^* = U + I + A$; and user operation time $Ag_t^{**} = U + I$.
- 2) **Task Completion Accuracy:** Average accuracy of task completion. For SIAgent, we define three accuracy metrics: Ag_t (overall accuracy), Ag_t^1 (intent recognition accuracy), and Ag_t^2 (agent-based execution accuracy, assuming 100% intent recognition accuracy).
- 3) **LLM Performance Comparison:** We evaluate different LLM models (ChatGLM-4 [45], ChatGPT-3.5-turbo [46], ChatGPT-4o [46]) in terms of inference speed and intent recognition accuracy to identify the optimal model for our framework. These models were selected because they represented the state-of-the-art during the study period (June 2024–January 2025).

Subjective Measures. We assess user experience and acceptance of the Intent-to-Operation paradigm compared to conventional Operation-to-Intent approaches.

- 1) **NASA-TLX [47]:** We measure cognitive workload including mental demand, physical demand, temporal demand, effort, performance, and frustration using a 21-point Likert scale.

- 2) **Arm Fatigue Level:** We assess physical fatigue using the Borg CR10 scale [48] to evaluate whether our approach reduces the physical burden associated with precise gesture execution in traditional paradigms.
- 3) **Ease of Use:** A 10-point Likert scale measuring learning curve and intuitiveness, specifically whether our method eliminates gesture memorization requirements.
- 4) **Novelty:** A 10-point Likert scale assessing users' perception of interaction innovation and naturalness compared to conventional VR interaction paradigms.
- 5) **System Preference:** We collect user preference rankings among the three interaction techniques (SIAgent, Controller, Gaze+Pinch).

VI. STUDY 1: INTENT-DRIVEN AGENT PERFORMANCE EVALUATION

In this study, our objective is to assess the overall effectiveness of Intent-to-Operation paradigm in VR. We compare our SIAgent system with two commonly used VR interaction techniques across multiple interaction scenarios.

A. Comparative Setups

We conduct a within-subjects comparative study evaluating three interaction methods:

Controller-based Interaction (CTRL). We compare controller-based object interaction, extensively explored in object manipulation [49], [50]. This interaction consists of two phases: selection and manipulation. **In the Selection Phase,** users point to the target object using a controller-emitted ray and select it by pressing the trigger button, displaying an operation menu with movement and trigger options. **In the Manipulation Phase,** for movement tasks, users press the “A button” to select the movement option and drag the object to the desired location by moving the controller. For trigger operations, users press the “B button” to display available trigger effects, scroll through options using the joystick, and press the “A button” to confirm selection, triggering the corresponding visual feedback.

Gaze+Pinch Interaction (GP). This interaction incorporates the three stages mentioned previously. **In the Pointing Phase,** users locate the target object by looking at it. **In the Confirmation Phase,** for movement tasks, users pinch their fingers for more than 0.5 seconds to select the object. For trigger operations, users perform two quick successive pinches, displaying available trigger effects. **In the Manipulation Phase,** for movement tasks, users maintain the pinch gesture while the object follows their hand movements. For trigger tasks, users select the desired effect by gazing at the item and using the pinch gesture to trigger effect selection. To reduce learning costs, we present tasks as menus rather than implementing complex gesture-task associations from previous studies [2], [5], though this approach sacrifices the semantic expressiveness of diverse gesture patterns.

SIAgent (Agent): Our method allows users to express interaction intent through natural eye-hand motions. The system recognizes these intents, translates them into concrete operations, and generates agents to execute tasks accordingly.

B. Procedure

Preparation Phase: Participants completed a pre-study questionnaire assessing their familiarity with VR systems and interaction techniques, then signed informed consent. After watching an introductory video explaining the interaction methods and experimental procedure, participants wore the VR headset and completed gaze calibration. To evaluate learning curves, participants practiced two tasks per technique in a selected VR scene for familiarization.

Experimental Phase: Participants performed tasks using all interaction techniques, with technique order counterbalanced via Latin Square design. Scene order was similarly counterbalanced for each technique. Within each scene, participants completed 6 randomly selected tasks. After completing all tasks for a technique, participants took a break and completed an evaluation questionnaire for that technique. Following completion of all techniques, participants ranked their method preferences in a post-study questionnaire. Each participant completed 54 trials total (3 techniques \times 3 scenes \times 6 tasks), with the experiment lasting approximately 70 minutes.

C. Results

We performed repeated-measures ANOVAs ($\alpha = 0.05$) and post hoc pairwise t-tests to assess whether there were statistically significant differences in a given metric across the various techniques. Prior to using ANOVA, Shapiro-Wilk (S-W) tests and homogeneity of variance tests were conducted. S-W tests indicated that some data met normality assumptions, while for data that did not satisfy normality, visual inspection of normality test histograms suggested acceptable approximation to normal distribution. Therefore, parametric tests were employed. For data violating homogeneity of variance assumptions, Welch's ANOVA was used.

Completion Time. The results are presented in Fig. 8 (a). Statistical analysis revealed a significant main effect of the interaction method on completion time ($F(4, 60) = 130.442$, $p < 0.001$, $\eta^2 = 0.901$). Specifically, *Agt* was significantly slower than both *CTRL* and *GP* (both $p < 0.001$). This delay is primarily attributed to the three stages of LLM inference, which currently require substantial time. When LLM inference time is excluded, *Agt** remained slower than *CTRL* (7.7 s vs 5.8 s, $p = 0.001$) but showed no significant difference compared to *GP* (6.7 s, $p = 0.075$). Regarding active user interaction time, *Agt*** was significantly faster than *Agt** (6.4 s vs 7.7 s, $p = 0.031$) and showed no significant difference compared to the other methods. We have included a supplementary study regarding the impact of latency in the Section I of the Supplementary Material.

Completion Accuracy. Results are presented in Fig. 8 (b). Statistical analysis revealed that the effect of different techniques and method stages on accuracy was statistically significant ($F(4, 60) = 22.99$, $p < 0.001$, $\eta^2 = 0.529$). Specifically, *Agt* had significantly lower accuracy compared to *CTRL* and *GP* (both $p < 0.001$). Further analysis showed that *Agt*¹ intent recognition accuracy was not significantly different from *CTRL* (97.2% vs 97.2%), and was significantly higher than *GP* (97.2% vs 93.1%, $p = 0.02$). However, for *Agt*², which considers agent-based execution accuracy, there was a




Fig. 8: The bar charts display the performance of the techniques across different measurements, with error bars representing the standard error. Statistical significance is indicated by ** for $p < 0.05$. Both *Agt* and *Agent* are our methods.




Fig. 9: The bar charts display the NASA-TLX questionnaire scores for the three interaction techniques, with error bars representing the standard error. Statistically significant differences are indicated by ** ($p < 0.05$).

significant decrease (86.8%). Analysis revealed that accuracy reduction was primarily caused by system timeouts exceeding the 30-second threshold, which automatically terminated agent execution. Our method requires three sequential LLM calls to complete each interaction, with an average API response time of 9.1s (see Tab. II). During peak usage periods of the GLM-4 cloud service, response times increase substantially, often causing total interaction time to exceed the 30-second limit. We anticipate this limitation will be resolved as LLM computational efficiency continues to improve.

LLM Inference Time and Accuracy. We recruited three participants from the main study to evaluate intent recognition performance across different LLMs. As shown in Tab. II, ChatGPT-3.5-turbo achieved the shortest API call time, followed by ChatGLM-4 (local). ChatGLM-4 (cloud) demonstrated the highest accuracy, with 83.3% success rate for returning the correct intent on the first attempt and 96.3% within the top-6 results. ChatGPT-4o (cloud) also exhibited strong performance. Given that API call times will improve as LLMs continue to evolve, we selected ChatGLM-4 (cloud) as our foundational model based on its superior accuracy.

Task Load. Repeated-measures ANOVA on the NASA-TLX questionnaire revealed significant differences in *mental/physical demand and effort* among techniques. Post-hoc pairwise t-tests are shown in Fig. 9. Results showed that *CTRL* had significantly higher *mental demand* than *GP* and *Agent*, while *GP* had significantly higher *physical demand and effort* than *Agent* (all $p < 0.05$). *CTRL*’s higher mental demand stems from cognitive overhead of navigating menus and remembering button mappings, while *GP*’s elevated physical demand is attributed to sustained pinch gestures and precise hand positioning. Our Intent-to-Operation paradigm reduces both burdens by allowing natural actions without learning.

Subjective Experience Metrics. Fig. 8 (c)-(e) presents

TABLE II: Comparison of different LLMs’ performance, including API Call Time, First Intent Accuracy, Top-3 Intent Accuracy, Top-6 Intent Accuracy.

LLMs	API call time (s)	1st Intent Acc. (%)	Top-3 Acc.(%)	Top-6 Acc. (%)
GLM-4 (cloud)	9.1s	83.3%	90.7%	96.3%
GLM-4 (local)	<u>7.9s</u>	53.7%	75.9%	79.6%
GPT-3.5-turbo (cloud)	4.2s	46.3%	55.6%	61.1%
GPT-4o (cloud)	8.5s	<u>72.2%</u>	87.0%	94.4%

results for arm fatigue, ease of use, and novelty assessments, all showing significant differences (all $p < 0.001$). *Agent* exhibited lower arm fatigue than both baselines by eliminating requirements for sustained precise gestures. Regarding ease of use, *Agent* outperformed both techniques by leveraging users’ existing common-sense knowledge rather than requiring memorization of gesture-operation mappings. *Agent* also achieved higher novelty scores, reflecting positive user perception of the LLM-powered intent recognition approach.

User Preference. Preference ranking questionnaire results indicate that *Agent* is most preferred by users, as illustrated in Fig. 8 (f). Specifically, 81% of participants ranked *Agent* first, 12.5% preferred *CTRL*, and one participant favored *GP*. This overwhelming preference validates the effectiveness of our Intent-to-Operation paradigm in addressing limitations of traditional VR interaction methods.

D. Discussion

Experimental results demonstrate that the Intent-to-Operation paradigm offers superior user experience over traditional Operation-to-Intent approaches. We organize discussion around three advantages that emerged from evaluation:

Intuitive Interaction with No Learning Required. Traditional VR interaction methods impose substantial learning burdens on users. *CTRL* requires memorizing button mappings and menu navigation sequences, while *GP* demands learning specific gesture-operation associations (e.g., double-pinch for trigger operations). Results show that *CTRL* exhibited significantly higher mental demand than *GP* and *Agent*. In contrast, *Agent* leverages users’ existing common-sense knowledge about object interactions. Users naturally demonstrate pouring motions when intending to pour water, drawing on everyday experience without learning system-specific protocols. As P12 commented, “This interaction is intuitive and natural, allowing task completion in your usual way without a fixed operational mode.” This design directly contributed to *Agent*’s superior ease-of-use ratings and reduced mental workload.

Flexible Intent Expression. Our approach accommodates individual differences in motion patterns and interaction pref-



Fig. 10: The illustration of flexible intent expression. Users can express milk-pouring intent through various natural gestures: (a) one-hand holding the bottle, (b) two-hand holding the bottle, or (c) using one hand for the cup and one for the bottle.

ferences. In Fig. 10, users express milk-pouring intent through different eye-hand coordination strategies. Traditional systems enforce uniform gesture vocabularies that ignore personal variations. Our LLM-based recognition adapts to diverse expression styles by understanding intent semantics rather than enforcing fixed templates. This flexibility extends beyond hand gestures to different gaze behaviors, movement speeds, and coordination preferences reflecting users’ natural habits.

High Interaction Error Tolerance. Traditional VR interaction methods suffer from failure modes where minor execution errors lead to complete task breakdown. *GP* fails when pinch gestures are mistimed or hand positions drift, while *CTRL* requires restarting entire sequences if incorrect buttons are pressed. Our Intent-to-Operation paradigm demonstrates resilience through error recovery mechanisms. When primary intent recognition fails, the system provides ranked alternatives, allowing users to select correct intentions without repeating demonstrations. Additionally, approximate hand motions and imprecise eye gaze convey sufficient information for successful intent inference. P15 noted, “With *Agent*, I barely felt any fatigue during interactions,” while P06 and P10 emphasized the system’s tolerance of imperfect execution.

The overwhelming user preference (81% ranked *Agent* first) validates these advantages. P1 remarked, “*Agent* is innovative, understanding my intent through simple gestures, which is very interesting.” Despite current processing speed limitations, users expressed strong confidence in the Intent-to-Operation approach for improving VR interaction experiences.

VII. STUDY 2: MULTIMODAL CHANNEL CONTRIBUTION ANALYSIS AND ALTERNATIVE MODALITY EVALUATION

To evaluate the individual contributions of eye gaze and hand gesture channels (RQ2) and to explore intent recognition with alternative input modalities (RQ3), we conducted two targeted experiments. The first experiment investigates how hand gesture information improves intent recognition accuracy compared to gaze-only input. The second experiment examines the performance of eye-speech interaction for intent recognition and agent execution, and compares its effectiveness with our eye-hand motion approach.

A. Participants and Apparatus

We recruited 3 participants (3 males) aged 24-29 years from Study 1 for both experiments. This sample size focuses on evaluating technical performance of our LLM-based intent recognition system using objective metrics rather than subjective user preferences, where individual variation has minimal

TABLE III: A comparison of the performance of different LLMs in verifying hand motion for intent recognition, measured by First Intent Accuracy, Top-3 Intent Accuracy, and Top-6 Intent Accuracy.

LLMs	1st Intent Acc.	Top-3 Acc.	Top-6 Acc.
Gaze Input Only	30.2%	63.5%	81.0%
Gaze + Hand Motion	58.3%	75.0%	93.3%

impact on measured recognition accuracy. The apparatus was identical to Study 1, using a Pico 4 Pro headset.

B. Task Design and Experimental Procedure

Experiment 1: Hand Motion Contribution Verification.

We designed 21 ambiguous tasks where multiple interaction possibilities exist for the same gaze target (*e.g.*, distinguishing between “open computer,” “close computer,” and “power on computer” when gazing at the same laptop). We compared two conditions: 1) *Gaze Input Only*: LLM receives only eye gaze target sequences without hand motion data; 2) *Gaze + Hand Motion*: LLM receives both gaze targets and complete hand motion information. The experimental procedure followed the same protocol as Study 1.

Experiment 2: Eye-Speech Modality Comparison.

We implemented an eye-speech interaction variant using the same tasks from Study 1. Users gaze at target objects and verbally express their interaction intent in natural language. The system performs speech recognition and treats the recognized speech as interaction intent, proceeding directly to agent-based execution as described in Section IV-C. We compared this modality’s performance with eye-hand motion results from Study 1.

C. Results and Discussion

Hand Motion Information Contribution Analysis.

The results presented in Table III demonstrate the critical importance of hand motion information for intent recognition in ambiguous scenarios. When using gaze target information alone, the system achieved 30.2% first intent accuracy, 63.5% top-3 accuracy, and 81.0% top-6 accuracy. However, incorporating hand motion data (including hand pose and finger shape) significantly improved performance to 58.3% first intent accuracy (+28.1%), 75.0% top-3 accuracy (+11.5%), and 93.3% top-6 accuracy (+12.3%). These findings confirm that hand motion information plays a crucial role in disambiguating user intent, particularly for the most likely intent prediction.

Eye-Speech Modality Performance.

Our eye-speech interaction achieved 94.4% task completion accuracy with an average completion time of 8.4 seconds, demonstrating speech input can accurately express interaction intent. While eye-speech represents a viable alternative, several key advantages favor the hand modality: 1) *Spatial Expression intuitiveness*: Describing complex trajectories (*e.g.*, lift the cup 20 cm while tilting it 45 degrees) imposes high cognitive and verbal effort. Hand motions provide a direct mapping where the input motion mirrors the object’s trajectory, avoiding the burden of describing spatial expression; 2) *Fine-Grained Continuous Control*: While speech is well-suited for discrete commands (*e.g.*, “turn on the lamp”), it is inefficient for smooth, real-time

adjustments (e.g., “make it a bit brighter... a bit more... stop!”). Gestures enable precise, millimeter-level modulation that is cumbersome to achieve through iterative verbal instructions.

VIII. LIMITATIONS, FUTURE WORK AND CONCLUSION

A. Limitations and Future Work

Current LLM-based processing introduces significant latency that limits real-time interaction performance. To address this, future work should explore intent prediction, where LLM inference is initiated based on early gaze and hand cues. Predicting intents in advance allows the system to filter options against ongoing actions, effectively masking latency. Additionally, while our evaluation focused on discrete, single-intent interactions, future research should extend to complex, multi-object sequences. Furthermore, acknowledging potential LLM hallucinations, we propose integrating traditional gaze-and-hand techniques as a robust fallback for precise control or error recovery. Finally, our Intent-to-Operation framework naturally aligns with emerging intelligent concepts in future human-AI collaboration, such as embodied AI paradigms [51]. In particular, AI systems deeply understand human operational intent and execute corresponding actions in real-world environments through robotic platforms, enabling smoother human-robot collaboration in complex physical scenes.

B. Conclusion

In this work, we introduced SIAgent, an LLM-powered spatial interaction agent that transforms VR interaction from the traditional “Operation-to-Intent” paradigm to an “Intent-to-Operation” approach. Our system enables users to express interaction intents through natural demonstrative actions based on their daily experiences and habits, eliminating the need for gesture memorization. The three-stage pipeline effectively translates multimodal spatial data into linguistic descriptions, leverages LLMs for intent recognition, and generates executable interaction agents. Through comprehensive user studies across over 60 interaction tasks, we demonstrated superior intent recognition accuracy compared to conventional gaze-pinch interaction methods, while achieving strong user preference and significantly reduced physical demand. Our multimodal analysis revealed that hand motion information contributes critically to intent disambiguation, substantially improving first intent accuracy. This work advances VR system intelligence through intent-driven design and establishes a foundation for extending such paradigms to human-robot collaboration and embodied AI applications.

ACKNOWLEDGMENTS

This work was supported by the Beijing High Innovation Plan (No. 202504841055), the Natural Science Foundation of China (62502020, 62372019), the Postdoctoral Fellowship Program of CPSF (GZC20252741) and the China Postdoctoral Science Foundation (2025M774238).

REFERENCES

- [1] Z.-M. Wang, M.-H. Rao, S.-H. Ye, W.-T. Song, and F. Lu, “Towards spatial computing: recent advances in multimodal natural interaction for extended reality headsets,” *Frontiers of Computer Science*, vol. 19, no. 12, p. 1912708, 2025.
- [2] L. Cao, H. Zhang, C. Peng, and J. T. Hansberger, “Real-time multimodal interaction in virtual reality—a case study with a large virtual interface,” *Multimedia Tools and Applications*, vol. 82, no. 16, pp. 25 427–25 448, 2023.
- [3] J. Lee, J. Wang, E. Brown, L. Chu, S. S. Rodriguez, and J. E. Froehlich, “Gazepointar: A context-aware multimodal voice assistant for pronoun disambiguation in wearable augmented reality,” in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, New York, NY, USA, 2024.
- [4] Z. Wang and F. Lu, “Tasks reflected in the eyes: Egocentric gaze-aware visual task type recognition in virtual reality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 11, pp. 7277–7287, 2024.
- [5] N. Chaconas and T. Höllerer, “An evaluation of bimanual gestures on the microsoft hololens,” in *Proc. IEEE Conf. Virtual Real. 3D User Interfaces*, Reutlingen, Germany, Mar. 2018, pp. 33–40.
- [6] S. Pei, A. Chen, J. Lee, and Y. Zhang, “Hand interfaces: Using hands to imitate objects in AR/VR for expressive interactions,” in *CHI '22: CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA, 29 April 2022 - 5 May 2022*. ACM, 2022, pp. 429:1–429:16.
- [7] Z. Cai, Y. Ma, and F. Lu, “Robust dual-modal speech keyword spotting for XR headsets,” *IEEE Trans. Vis. Comput. Graph.*, vol. 30, no. 5, pp. 2507–2516, 2024.
- [8] Apple, “Introducing apple vision pro: Apple’s first spatial computer,” 2025.06, <https://www.apple.com/newsroom/2023/06/introducing-apple-vision-pro/>.
- [9] Z. Wang, H. Wang, H. Yu, and F. Lu, “Interaction with gaze, gesture, and speech in a flexibly configurable augmented reality system,” *IEEE Transactions on Human-Machine Systems*, vol. 51, no. 5, pp. 524–534, 2021.
- [10] X. Chen, D. Guo, L. Feng, B. Chen, and W. Liu, “Compass+ring: A multimodal menu to improve interaction performance and comfortability in one-handed scenarios,” in *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, Sydney, Australia, October 16–20, 2023, pp. 473–482.
- [11] Y. Bao, J. Wang, Z. Wang, and F. Lu, “Exploring 3d interaction with gaze guidance in augmented reality,” in *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, 2023, pp. 22–32.
- [12] D. Yu, X. Lu, R. Shi, H. Liang, T. Dingler, E. Velloso, and J. Gonçalves, “Gaze-supported 3d object manipulation in virtual reality,” in *CHI Conference on Human Factors in Computing Systems*. ACM, 2021, pp. 734:1–734:13.
- [13] Z. Song, J. J. Dudley, and P. O. Kristensson, “Hotgestures: Complementing command selection and use with delimiter-free gesture-based shortcuts in virtual reality,” *IEEE Trans. Vis. Comput. Graph.*, vol. 29, no. 11, pp. 4600–4610, 2023.
- [14] L. Chen, Y. Cai, R. Wang, S. Ding, Y. Tang, P. Hansen, and L. Sun, “Supporting text entry in virtual reality with large language models,” in *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2024, pp. 524–534.
- [15] B. Yang, Y. Guo, L. Xu, Z. Yan, H. Chen, G. Xing, and X. Jiang, “Socialmind: Llm-based proactive ar social assistive system with human-like perception for in-situ live interactions,” *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 9, no. 1, Mar. 2025. [Online]. Available: <https://doi.org/10.1145/3712286>
- [16] Z. Li, H. Zhang, C. Peng, and R. Peiris, “Exploring large language model-driven agents for environment-aware spatial interactions and conversations in virtual reality role-play scenarios,” in *2025 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, 2025, pp. 1–11.
- [17] J. Chen, Z. Liu, X. Huang, C. Wu, Q. Liu, G. Jiang, Y. Pu, Y. Lei, X. Chen, X. Wang, D. Lian, and E. Chen, “When large language models meet personalization: Perspectives of challenges and opportunities,” 2023. [Online]. Available: <https://arxiv.org/abs/2307.16376>
- [18] Z. Wang, J. Sun, M. Hu, M. Rao, W. Song, and F. Lu, “Gazing: Enhancing hand-eye coordination with pressure ring in augmented reality,” in *2024 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2024, pp. 534–543.
- [19] I. Poupyrev, M. Billinghurst, S. Weghorst, and T. Ichikawa, “The go-go interaction technique: non-linear mapping for direct manipulation in vr,” in *Proceedings of the 9th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '96, New York, NY, USA, 1996, p. 79–80. [Online]. Available: <https://doi.org/10.1145/237091.237102>
- [20] M. Krüger, T. Gerrits, T. Römer, T. Kuhlen, and T. Weissker, “Intenselect+: Enhancing score-based selection in virtual reality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 5, pp. 2829–2838, 2024.

- [21] J. Ma, J. Qian, T. Zhou, and J. Huang, "Focalpoint: Adaptive direct manipulation for selecting small 3d virtual objects," vol. 7, no. 1, Mar. 2023. [Online]. Available: <https://doi.org/10.1145/3580856>
- [22] J. D. Hincapié-Ramos, X. Guo, P. Moghadasian, and P. Irani, "Consumed endurance: a metric to quantify arm fatigue of mid-air interactions," in *Proc. Conf. Human Factors Comput. Syst.*, Ontario, Canada, Apr. 2014, pp. 1063–1072.
- [23] D. Wang, L. Qiu, B. Li, Q. Liu, X. Wei, J. Chen, Z. Wang, and M. Fan, "Focalselect: Improving occluded objects acquisition with heuristic selection and disambiguation in virtual reality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 31, no. 5, pp. 3285–3295, 2025.
- [24] D. Yu, Q. Zhou, J. Newn, T. Dingler, E. Velloso, and J. Goncalves, "Fully-occluded target selection in virtual reality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 26, no. 12, pp. 3402–3413, 2020.
- [25] L. Sidenmark, M. Parent, C. Wu, J. Chan, M. Glueck, D. Wigdor, T. Grossman, and M. Giordano, "Weighted pointer: Error-aware gaze-based interaction through fallback modalities," *IEEE Trans. Vis. Comput. Graph.*, vol. 28, no. 11, pp. 3585–3595, 2022. [Online]. Available: <https://doi.org/10.1109/TVCG.2022.3203096>
- [26] Z. Wang, Y. Zhao, and F. Lu, "Gaze-vergence-controlled see-through vision in augmented reality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 11, pp. 3843–3853, 2022.
- [27] P. Mohan, W. B. Goh, C. Fu, and S. Yeung, "DualGaze: Addressing the midas touch problem in gaze mediated vr interaction," in *Proc. IEEE Int. Symp. Mix. Augmented Real.*, Munich, Germany, Oct. 2018, pp. 79–84.
- [28] Y. Bao, J. Wang, Z. Wang, and F. Lu, "Exploring 3d interaction with gaze guidance in augmented reality," in *2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, 2023, pp. 22–32.
- [29] D. Giunchi, N. Numan, E. Gatti, and A. Steed, "Dreamcodevr: Towards democratizing behavior design in virtual reality with speech-driven programming," in *IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, Orlando, FL, USA, March 16-21, 2024, pp. 579–589.
- [30] J. Saint-Aubert, F. Argelaguet, M. J. Macé, C. Pacchierotti, A. Amedi, and A. Lécuyer, "Persuasive vibrations: Effects of speech-based vibrations on persuasion, leadership, and co-presence during verbal communication in VR," in *IEEE Conference Virtual Reality and 3D User Interfaces (VR)*, Shanghai, China, March 25-29, 2023, pp. 552–560.
- [31] M. Kytö, B. Ens, T. Piumsomborn, G. A. Lee, and M. Billinghurst, "Pinpointing: Precise head-and-eye-based target selection for augmented reality," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–14.
- [32] E. Castellina, F. Corno, and P. Pellegrino, "Integrated speech and gaze control for realistic desktop environments," in *Proceedings of the Eye Tracking Research & Application Symposium, ETRA 2008, Savannah, Georgia, USA, March 26-28, 2008*. ACM, 2008, pp. 79–82.
- [33] A. Jing, G. A. Lee, and M. Billinghurst, "Using speech to visualise shared gaze cues in MR remote collaboration," in *IEEE Conference on Virtual Reality and 3D User Interfaces, VR 2022, Christchurch, New Zealand, March 12-16, 2022*. IEEE, pp. 250–259.
- [34] S. Jadon, M. Faridan, E. Mah, R. Vaish, W. Willett, and R. Suzuki, "Augmented conversation with embedded speech-driven on-the-fly referencing in ar," 2024. [Online]. Available: <https://arxiv.org/abs/2405.18537>
- [35] J. Rekimoto, "Gazellm: Multimodal llms incorporating human visual attention," 2025. [Online]. Available: <https://arxiv.org/abs/2504.00221>
- [36] Y. Karaca, D. Derias, and G. Sarsar, "Ai-powered procedural content generation: Enhancing npc behaviour for an immersive gaming experience," *Available at SSRN 4663382*, 2023.
- [37] J. J. Yang, L. Qiu, E. A. Corona-Moreno, L. Shi, H. Bui, M. S. Lam, and J. A. Landay, "Amma: Adaptive multimodal assistants through automated state tracking and user model-directed guidance planning," in *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2024, pp. 892–902.
- [38] N. M. Gamage, D. Ishtaweera, M. Weigel, and A. Withana, "So predictable! continuous 3d hand trajectory prediction in virtual reality," in *The 34th Annual ACM Symposium on User Interface Software and Technology*, 2021, pp. 332–343.
- [39] C. Zhu, S.-K. Hsia, X. Hu, Z. Liu, J. Shi, and K. Ramani, "agentar: Creating augmented reality applications with tool-augmented llm-based autonomous agents," in *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '25. New York, NY, USA: Association for Computing Machinery, 2025. [Online]. Available: <https://doi.org/10.1145/3746059.3747676>
- [40] F. Lan, S. Wang, Y. Zhang, H. Xu, O. Oseni, Z. Zhang, Y. Gao, and T. Zhang, "Dexcatch: Learning to catch arbitrary objects with dexterous hands," 2024. [Online]. Available: <https://arxiv.org/abs/2310.08809>
- [41] Pico, "Pxr hand pose generator script," 2024.09, <https://developer.picoxr.com/document/unity/about-the-pxr-hand-pose-generator-script/>.
- [42] C. Jiang, Y. Zeng, T. Hu, S. Xu, W. Zhang, H. Xu, and D.-Y. Yeung, "Jointdreamer: Ensuring geometry consistency and text congruence in text-to-3d generation via joint score distillation," in *Computer Vision – ECCV 2024*. Cham: Springer Nature Switzerland, 2025, pp. 439–456.
- [43] Z. Wu, J. Li, P. Xu, and C. K. Liu, "Human-object interaction from human-level instructions," *arXiv preprint arXiv:2406.17840*, 2024.
- [44] N. Jiang, Z. He, Z. Wang, H. Li, Y. Chen, S. Huang, and Y. Zhu, "Autonomous character-scene interaction synthesis from text instruction," in *SIGGRAPH Asia 2024 Conference Papers*, 2024, pp. 1–11.
- [45] T. GLM, "Chatglm: A family of large language models from glm-130b to glm-4 all tools," 2024. [Online]. Available: <https://arxiv.org/abs/2406.12793>
- [46] OpenAI, "Openai api," 2024.09, <https://platform.openai.com/docs/models>.
- [47] S. G. Hart, "Nasa-task load index (nasa-tlx); 20 years later," in *Proceedings of the human factors and ergonomics society annual meeting*, vol. 50, no. 9, 2006, pp. 904–908. [Online]. Available: <https://doi.org/10.1177/154193120605000909>
- [48] G. Borg, "Borg's perceived exertion and pain scales," *Human Kinetics*, vol. 104, 1998.
- [49] W.-j. Hou and X.-l. Chen, "Comparison of eye-based and controller-based selection in virtual reality," *International Journal of Human-Computer Interaction*, vol. 37, no. 5, pp. 484–495, 2021.
- [50] J. Song, R. Shi, Y. Li, B. Gao, and H.-N. Liang, "Exploring controller-based techniques for precise and rapid text selection in virtual reality," in *2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR)*. IEEE, 2024, pp. 244–253.
- [51] F. Lu and Q. Zhao, "Towards embodied/symbolic ai: concept and eight scientific and technical problems," *Science China Information Sciences*, vol. 69, p. 116101, 2026.

IX. BIOGRAPHY SECTION

Zhimin Wang is a postdoctoral researcher with the State Key Laboratory of Virtual Reality Technology and Systems, School of Computer Science and Engineering, Beihang University, China. He received his Ph.D. from Beihang University in 2024 and his B.S. from Chang'an University in 2019. His research focuses on VR/AR, human-computer interaction, and eye-tracking technologies. He serves as a program committee member for AAAI and as a regular reviewer for leading international venues, including IEEE VR, ISMAR, TVCG.

Chenyu Gu is a Ph.D. student with the State Key Laboratory of Virtual Reality Technology and Systems, Beihang University. He received his bachelor's degree in Computer Science and Technology from Beihang University, China, in 2023. His research focuses on virtual reality, human-computer interaction, and the application of AI in education.

Feng Lu (Senior Member, IEEE) received the BS and MS degrees in Automation from Tsinghua University, China in 2007 and 2010, respectively, and the PhD degree in Information Science and Technology from The University of Tokyo, Japan in 2013. He is currently a full Professor with the State Key Laboratory of Virtual Reality Technology and Systems, School of Computer Science and Engineering, Beihang University. His research interests include computer vision, natural interaction and VR/AR. He is a distinguished member of CCF and CSIG. He is serving/has served as an Area Chair for prestigious international conferences such as CVPR, ICCV, ECCV, NeurIPS and ACM MM.

