

F.A.C.U.L.: Language-Based Interaction with AI Companions in Gaming

Wenya Wei^{1*}, Sipeng Yang^{2*}, Qixian Zhou¹, Ruochen Liu¹, Xuelei Zhang¹,
Yifu Yuan³, Yan Jiang¹, Yongle Luo¹, Hailong Wang¹, Tianzhou Wang¹, Peipei Jin¹,
Wangtong Liu¹, Zhou Zhao², Xiaogang Jin², Elvis S. Liu^{1†}

¹Tencent Games, Shenzhen, China

²Zhejiang University, Hangzhou, China

³Tencent AI Lab, Shenzhen, China

Abstract

In cooperative video games, traditional AI companions are deployed to assist players, who control them using hotkeys or command wheels to issue predefined commands such as “attack”, “defend”, or “retreat”. Despite their simplicity, these methods, which lack target specificity, limit players’ ability to give complex tactical instructions and hinder immersive gameplay experiences. To address this problem, we propose the FPS AI Companion who Understands Language (F.A.C.U.L.), the first real-time AI system that enables players to communicate and collaborate with AI companions using natural language. By integrating natural language processing with a confidence-based framework, F.A.C.U.L. efficiently decomposes complex commands and interprets player intent. It also employs a dynamic entity retrieval method for environmental awareness, aligning human intentions with decision-making. Unlike traditional rule-based systems, our method supports real-time language interactions, enabling players to issue complex commands such as “clear the second floor”, “take cover behind that tree”, or “retreat to the river”. The system provides real-time behavioral responses and vocal feedback, ensuring seamless tactical collaboration. Using the popular FPS game *Arena Breakout: Infinite* as a case study, we present comparisons demonstrating the efficacy of our approach and discuss the advantages and limitations of AI companions based on real-world user feedback.

Homepage — <https://sites.google.com/view/FACUL-demo>

Introduction

Over the past few decades, video games have attracted billions of players, with new techniques continuously introduced to enhance their appeal. In particular, recent breakthroughs in AI have enabled the integration of advanced capabilities in games, such as those seen in Go (Silver et al. 2016) and real-time strategy games like StarCraft II (Vinyals et al. 2019), where AI has reached expert-level performance. However, most AI research in gaming focuses on competition with players, while the exploration of AI companions

*These authors contributed equally.

†Corresponding author: elvissyliu@tencent.com

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

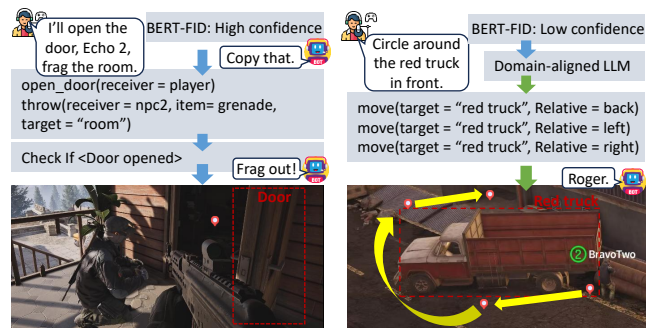


Figure 1: F.A.C.U.L. is the first open-interaction and real-time language-operated AI companion system for commercial video games.

for cooperative gameplay has been largely overlooked (Gao et al. 2023). In First-Person Shooter (FPS) games, players’ enjoyment often depends on effective teamwork, making cooperation crucial for gameplay enhancement and immersion. When it is challenging to find human teammates online, or when players prefer to avoid teaming up with others, the development of AI companions becomes crucial. Such companions must be capable of comprehending player instructions and delivering reliable, real-time assistance.

For AI companions, automated goal execution is relatively straightforward to implement, requiring only predefined objectives, paths, and action sequences. However, the methods for players to communicate with and command AI companions remain limited and simplistic. Existing interaction approaches, such as **hotkeys** and **command wheels** in FPS games, allow players to issue basic commands like “attack” and “retreat”. These commands are overly simplistic and lack target specificity, which prevents players from issuing more complex tactical instructions, such as “clear the second floor” and “take cover behind that tree”, hindering human-AI collaboration and reducing immersion.

Intuitively, language is the most natural and immersive form of interaction, making it highly appealing to collaborate with AI companions in games. Although recent advances in powerful large language models (LLMs) (Yang et al. 2024; Zhang et al. 2024a) have enabled this capa-

bility, integrating them into real-world games for effective use presents several challenges. First, the framework must operate with minimal latency to deliver real-time feedback to players, while ensuring high concurrency and low operational costs on the server side. Second, the language model must accurately understand player commands within the specific game context, translating and decomposing commands into actions, objects, and other relevant elements. Finally, AI companions must precisely identify the objects in the game environment that correspond to the player’s commands to execute them effectively.

In this paper, we propose F.A.C.U.L., the first language-based AI companion framework for FPS video games. F.A.C.U.L. enables natural language interaction with players, supports tactical collaboration, and understands the game environment. The framework comprises three key components: instruction reasoning, environmental perception, and decision execution. To achieve high-throughput language instruction processing, we introduce a confidence-based method that combines a lightweight model, BERT-FID (BERT for Fuzzy Intent Detection), with a larger LLM. BERT-FID quickly decomposes commands into action–target templates and provides both reasoning results and confidence scores. High-confidence (simple) commands are executed directly, while low-confidence (complex) ones are routed to the LLM for deeper interpretation. This design reduces computational overhead for routine commands while preserving accuracy for ambiguous instructions, balancing efficiency and comprehension. Next, a multimodal entity retrieval module matches player commands to real-time game assets, ensuring robust environmental awareness. Finally, commands are translated into atomic actions and merged with environmental data to generate the AI companion’s responses, which are executed using behavior trees with real-time feedback to the player.

Using the popular FPS game *Arena Breakout: Infinite* (Tencent 2024) as a case study, we demonstrate the application of F.A.C.U.L., offering an immersive human-AI collaboration experience and supporting a wide range of open-ended commands within the game. Fig. 1 illustrates the performance of our algorithm in task completion. A full match visualization video is available on our project page. We summarize the contributions of our work as follows:

- F.A.C.U.L., the first FPS AI companion that understands human language, offering natural language interaction and contextual environment recognition for commercial video games.
- BERT-FID, a command decomposition and confidence-based model for fuzzy intent detection, enhancing both the speed and accuracy of command processing.
- Developed and evaluated in a real-world FPS game, showing high reliability and strong player adoption, with a comprehensive analysis of AI companions through player feedback.

Related Work

LLMs in gaming have recently garnered significant attention (Gallotta et al. 2024; Xu et al. 2024; Park et al.

2023), particularly in the development of non-player characters (NPCs) and player assistance. For NPCs, LLMs are leveraged to enhance character realism and interactivity. For instance, Shao et al. (2023) propose the Character-LLM, a framework that trains LLMs to simulate specific characters, thereby enriching role-playing experiences. Similarly, the integration of LLMs in 3D games enables players to collaborate with LLM-powered NPCs to complete tasks (Rao et al. 2024). Another promising direction is the use of LLMs to generate dynamic, interactive game content. A notable application is the text-to-game engine (Zhang et al. 2024b), which transforms text inputs into dynamic RPGs, generating content such as storylines and game mechanics in real-time.

A recent commercial game, *Naraka* (NetEast and Entertainment 2020), adopts a complex AI companion system, which is close to our work. However, it focuses primarily on language communication and behavioral interactions with players, while neglecting the ability to effectively understand real-time game states and identify relevant objects within these environments. This limitation restricts its ability to assist players in task completion and immersion. In contrast, our work addresses this gap by integrating instruction reasoning, scene recognition, as well as strategic execution (FAIR et al. 2022), thereby enabling AI companions to more effectively collaborate with players and adapt to the complexity of in-game scenarios

Instruction reasoning refers to the ability of AI companions to understand players’ language commands and intentions in gaming. Early methods (Huang, Xu, and Yu 2015; Liu and Lane 2016) relied on traditional machine learning, which struggled with complex linguistic phenomena. The advent of pre-trained language models, particularly BERT-based methods (Vaswani et al. 2017; Kenton and Toutanova 2019; Shao et al. 2020; Zhang et al. 2021), has led to significant improvements in handling downstream tasks. More recently, LLMs have demonstrated remarkable generalization capabilities. Models developed by (Brown et al. 2020) and (Patel et al. 2022) demonstrate powerful generative and comprehension abilities, enabling zero-shot learning and significantly reducing the reliance on labeled data. Moreover, recent work (Leviathan, Kalman, and Matias 2023; Lin et al. 2023; Chen et al. 2023, 2024) explores collaborative paradigms between small and large models to accelerate inference and reduce deployment costs for various text-related tasks. Inspired by these advancements, we have enhanced BERT with fuzzy intent detection capabilities, leveraging the generalization strengths of LLMs to more effectively infer player intents.

Image-text retrieval (ITR) is a core cross-modal task that aligns visual and textual information, supporting scene recognition for AI companions. A common approach encodes image and text features separately, then aligns them via similarity computation (Frome et al. 2013; Faghri et al. 2018). Recent advances in large-scale cross-modal pre-training have significantly improved ITR by leveraging web-scale data. For example, CLIP (Radford et al. 2021; Yang et al. 2022) uses contrastive learning on image-text pairs, while ALIGN (Jia et al. 2021) jointly trains encoders for efficient retrieval. In our case, we adopt a similar ITR strat-

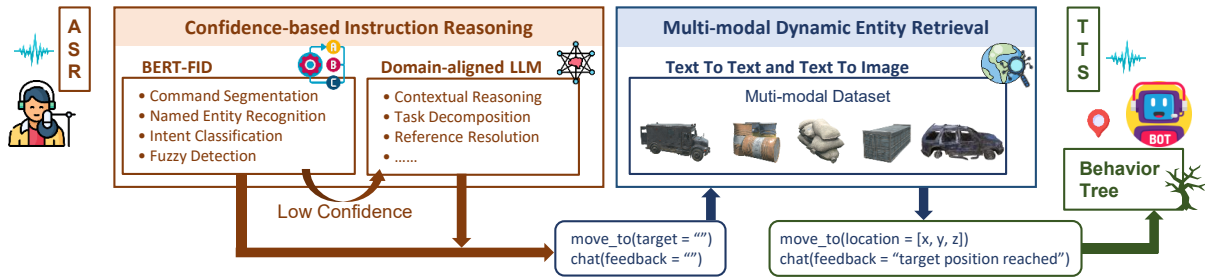


Figure 2: Overview of F.A.C.U.L.. Players can cooperate with the agent in combat and freely speak into the microphone and get behavior response with voice feedback.

egy to match human instructions with game entities, further incorporating real-time attributes such as player-object distance and orientation to improve scene recognition accuracy.

Overview of F.A.C.U.L.

Fig. 2 shows the pipeline of the proposed F.A.C.U.L. system. In a cooperative FPS game, players interact with the F.A.C.U.L. agent via voice input through a microphone. The speech is transcribed by an Automatic Speech Recognition (ASR) module (Malik et al. 2021). Then, a confidence-based instruction reasoning approach follows: BERT-FID processes high-confidence commands, while low-confidence (often complex or ambiguous) ones are routed to a domain-aligned LLM for deeper analysis. After extracting the player’s intended action and target, the multi-modal entity retrieval module identifies relevant 3D objects in the game scene. The parsed command is then sent to the game engine, where the AI companion executes it via behavior trees and delivers real-time feedback through a text-to-speech (TTS) system (Kumar, Koul, and Singh 2023). These interconnected modules form the complete F.A.C.U.L. framework, supporting interpretability and enabling independent optimization of each component. Instruction reasoning and scene recognition are elaborated in the following sections.

Confidence-Based Instruction Reasoning

We present BERT-FID, a fast-processing module for interpreting tactical instructions and generating confidence scores. High-confidence instructions are processed directly, while low-confidence ones are routed to the LLM for deeper reasoning. This design allows F.A.C.U.L. to efficiently handle simple commands while reserving LLM resources for complex cases.

BERT-FID: BERT with Fuzzy Intent Detection

Multi-task Label Structure. In gaming, player commands often need to be decomposed into a sequence of executable actions for the behavior tree, and extracting the target entities for each sub-command’s intent is crucial. Based on this, we design the data label structure shown in Fig. 3. The segment labels in the second row are used for common segmentation tasks, where “B-seg” marks the beginning of a sub-command. The labels for the named entity recognition task identify entity types and syntactic categories. The intent

classification labels encompass both the types of intent and their respective subjects. For example, when a player issues the command, “Hold position at the red door, I will grab the loot”, the BERT model can parse the input, output the intent as “move”, and identify the target location as “the red door”.

Model Architecture. As shown in Fig. 4, we use BERT (Kenton and Toutanova 2019) as the pre-training representation model in the feature extraction layer. For downstream tasks, we introduce command segmentation to determine whether each token starts a new sentence, thereby facilitating the precise identification of executable sub-commands. We subsequently integrate the token-level features of the segmented sub-commands through an attention mechanism to generate new features, which serve as input for subsequent hierarchical intent classification and named entity recognition. The recognition of named entities is treated as a sequence classification task, using Conditional Random Fields (CRF) (Lafferty et al. 2001) to capture contextual dependencies and accurately calculate the conditional probability distribution. In hierarchical intent classification, a regressive embedding mechanism is used across layers to improve feature extraction.

We employ this BERT base architecture for joint training across three interrelated tasks, effectively standardizing high-dimensional categories within the same data distribution. These tasks include text segmentation with a maximum length of 128 tokens, named entity recognition for six different target entities, and a hierarchical intent classification task with a three-layer intent recognition network, covering a total of 25 categories.

Fuzzy Intent Detection in Training. We define the fuzzy intent as natural language commands that are either: 1) Unsupported intents for BERT model; 2) Contextual understanding; 3) Semantic-related reference resolution; 4) Instructions requiring inference. Such intents are challenging for lightweight models, as overconfident predictions often lead to incorrect execution in real-time FPS scenarios.

To enable the model to effectively leverage predicted confidence and detect fuzzy intents, we introduce an entropy optimization objective for the unsupported intent classification head in the hierarchical classifier. Specifically, given a batch of $N = N_c + N_e$ samples, where N_c denotes the supported data, the cross-entropy loss is applied to compute the general loss function. For the remaining N_e unsupported data, entropy maximization is used to increase the uncertainty of

Command	Hold	position	at	the	red	door	I	will	grab	the	loot
Segment	B-seg	I-seg	I-seg	I-seg	I-seg	I-seg	B-seg	I-seg	I-seg	I-seg	I-seg
Entities	O	O	O	O	B-entity	I-entity	B-name	O	O	O	B-entity
Syntactic	B-subject	O	O	O	B-object	I-object	B-subject	O	O	O	B-object
Intent	Tactic-move						Tactic-move				
Receiver Type	Null						Player				

Figure 3: Multi-task label structure designed for command segmentation, intent classification, and named entity recognition.

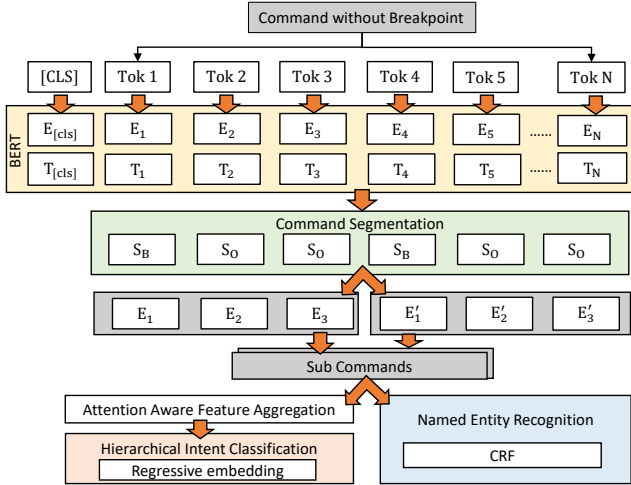


Figure 4: Model Architecture of BERT-FID.

the probability vector. The total loss function is defined as:

$$\mathcal{L} = -\frac{1}{N_c} \sum_{i=1}^{N_c} \sum_{j=1}^C w_j y_{ij} \log p_{ij} - \lambda \frac{1}{N_e} \sum_{i=1}^{N_e} \left(\log C - \sum_{j=1}^C p_{ij} \log p_{ij} \right) \quad (1)$$

where C is the number of categories for one head, y_{ij} is the one-hot encoded label for the j -th category in the i -th sample, and p_{ij} is the probability vector representing the likelihood of the i -th sample belonging to the j -th category. w_j is the weight of the j -th category, which is determined based on the statistical distribution of intent labels $\frac{N_e}{n_i}$, where n_i is the number of samples in the j -th category. This method increases the weight of sparse labels and reduces the weight of dense labels, which is crucial for addressing the unbalanced intent distribution in our task.

This approach allows the model to reduce confidence in any category when encountering data outside the predicted categories. In the hierarchical intent classification setup, the first-level labels determine which prediction heads are activated at the second level. For instance, if the first intent label is “throw”, the model activates item prediction heads for grenade, smoke, food, or water. If the relevant prediction head is not activated, the model reduces confidence in any specific category by maximizing entropy regularization. The pseudocode can be found in the supplementary material.

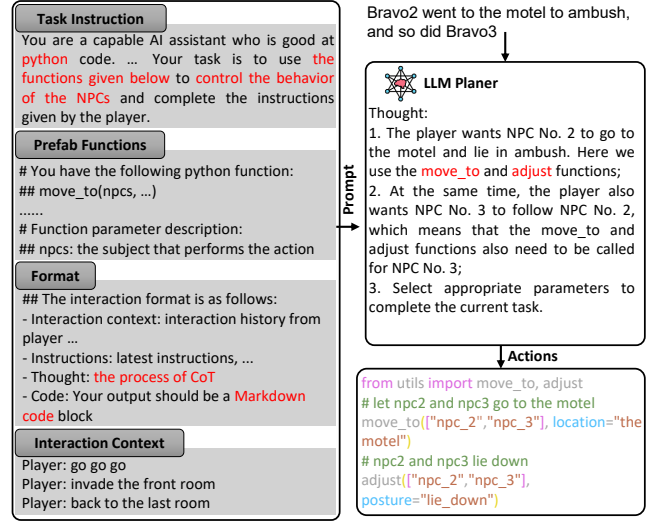


Figure 5: LLM for actions planning. The LLM planner synthesizes Python code to call functions to orchestrate the actions of F.A.C.U.L. agents, reasoning and acting to address complex tactical tasks.

Domain-Aligned LLM

To enable F.A.C.U.L. agents to accurately execute complex or ambiguous commands that BERT-FID classifies with low confidence, we fine-tune a lightweight yet domain-specialized LLM using our in-house dataset for FPS games, aligning it with the preferences of real-life tactical companions. As shown in Fig. 5, we adopt a code generation approach with LLMs to control the actions of F.A.C.U.L. agents. Python is chosen for this task due to its compatibility with the training corpora of most existing LLMs and its flexibility in representing common code constructs, such as conditionals and loops, which JSON format lacks.

We design a series of atomic tasks (e.g., move, attack, open_door) as Python functions, implemented via behavior trees. The LLM Planner is trained to generate Python code that invokes these functions, enabling agents to reason and execute complex tactical actions. Expected actions of FPS game agents, in response to player commands, are generated using the general-purpose LLM, Qwen2-1.5B (Yang et al. 2024), and validated by human annotators. This command-action dataset is then used to supervise the fine-tuning of the lightweight LLM.

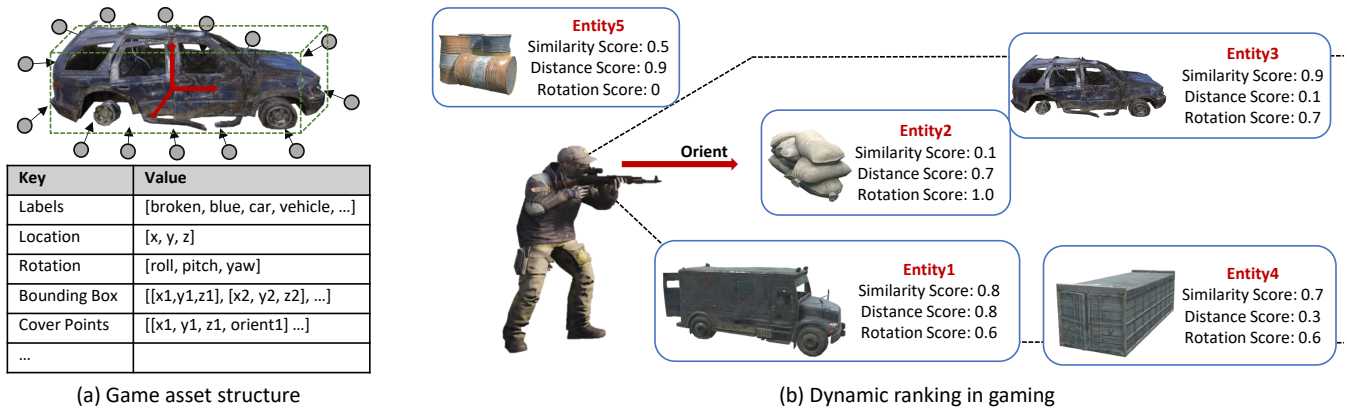


Figure 6: Game asset structure and dynamic ranking. (a) Overview of 3D game assets, including labels, location, orientation, and so on. (b) Dynamic entity ranking, considering player position and command context, to identify the target entity.

Scene Recognition

This section details the implementation of a real-time multi-modal scene recognition system that enables fine-grained understanding of a wide range of complex objects from the player’s view. As shown in Fig. 6 (a), we structure the 3D game assets dataset with tactical information, including coordinate positions, orientations, bounding boxes, and cover points. Coordinate positions are used for F.A.C.U.L. to perform navigation, while orientation data helps F.A.C.U.L. understand spatial relationships between game assets. Bounding boxes enable F.A.C.U.L. to perceive the size and range of physical assets. Cover points, derived from the game’s cover system, assist F.A.C.U.L. in finding optimal positions within the FPS environment.

For entity description input, we first perform a similarity search against the embeddings of all entity labels in the scene. If the retrieved entities are insufficient, we use a fine-tuned CLIP (Radford et al. 2021) model to search for image embeddings, generating a list of candidate entities. Real-time dynamic game data, such as the player’s position and orientation, is then considered for fine-ranking the candidates. As shown in Fig. 6 (b), although entity3 has the highest similarity to the command, it is too far away. Entity5 is closer but exceeds the range specified by the directional term. At last, entity1 is selected as the final result, as it ranks highest across all factors.

It is worth noting that our dynamic entity search method can handle complex hierarchical relationships. For instance, when a player issues the command: “Go behind the red sofa on the first floor of the motel and find a cardboard box”, we first locate the primary target, “the first floor of the motel”, then use its coordinates to search for the secondary target, “red sofa”. Finally, the sofa is used as a reference to find the nearby “cardboard box”.

Furthermore, since we have cover points data in the game, our entity search method can also guide F.A.C.U.L. to seek cover at specific locations. Each object has multiple cover points, and each cover point has its corresponding orientation—only when facing the entity can it act as a cover. In this search task, we need to prioritize the orientation re-

lationship. By assigning different weighted parameters for specific tasks, our entity search system can handle a wide variety of functions.

Experiments

To evaluate the performance of F.A.C.U.L., we collect real player data from the video game *Arena Breakout: Infinite*. A fully deployed version of F.A.C.U.L. is integrated into the game, and players are invited to interact with the system. This section presents two parts: a quantitative analysis comparing our method with baseline models, and a user study capturing player feedback on AI companions.

Quantitative Analysis

Evaluation Dataset and Metrics. We develop a real-time test and evaluation system incorporating live player feedback. During gameplay, players issue verbal commands to AI teammates, which are transcribed into text. A thumbs-up marks successful execution, while a thumbs-down flags failures for review. These failed cases are reviewed by at least two professional annotators (more than 6 hours of FPS gaming per week), who correct the AI response. Each command is then paired with the corrected action to form the test dataset. Each game session yields 10 to 30 data points, resulting in 2,550 samples collected from 100 sessions. In our experiments, the instruction reasoning modules are executed on two NVIDIA GeForce RTX 3080 GPUs, while entity retrieval tasks run on two NVIDIA Tesla T4 GPUs. Detailed experimental settings and reference code are available on our project page. The evaluation metrics include:

- **Accuracy:** Measures the proportion of correct actions executed by the AI in response to player commands.
- **Parameters:** Refers to the total number of parameters in the ACLI model, indicating the model’s complexity and computational demands.
- **Precision, Recall, and F1 Score:** Evaluate the performance of the entity retrieval models.
- **Queries Per Second (QPS):** Measures the system’s ability to process player commands in real-time, indicating

Method	Accuracy	Parameters	Decision count	
			BERT	LLM
BERT	0.714	104M	2550	0
BERT-FID	0.749			
Qwen2-7B	0.807			
Qwen2-1.5B	0.642	1.5B	0	2550
LLM*	0.908			
BERT & LLM*	0.831	1.6B	2321	229
BERT-FID & LLM*	0.886		2148	402

Note: “Decision count” indicates the number of final decisions selected from either BERT or LLM. “LLM*” refers to our domain-aligned LLM (Qwen2-1.5B)

Table 1: The comparison of various methods in about instruction reasoning demonstrates that our approach, “BERT-FID with fine-tuned LLM”, can route more samples to the LLM and effectively achieve higher accuracy than the pure BERT-based method.

Module	Method	QPS	Latency (ms)
Instruction reasoning	BERT-based method	1672	412
	Qwen2-7B	46	423
	LLM*	245	392
	BERT-FID with LLM*	813	470
Dynamic entity retrieval	CLIP-based method	194	355
	Text similarity	2128	364
	Text similarity with CLIP	1057	363

Table 2: Performance testing shows our method achieves about 4× higher inference concurrency than LLM-based methods and five-fold higher than CLIP-based methods.

how well it can handle high-frequency interactions.

- **Latency:** Tracks the delay between a player’s input and the AI’s response.

Comparison and Results. The F.A.C.U.L. method includes two main components: instruction reasoning and entity retrieval. We first evaluate each component separately against corresponding baselines, followed by an assessment of the integrated system.

For the *instruction reasoning* component, textual commands are fed into the BERT-FID model to generate an action sequence. If the confidence score falls below a pre-defined threshold, the command is routed to a domain-aligned LLM to re-infer the action. We compare the performance of using BERT alone, LLM alone, and combinations of both. The evaluated models include the original BERT, BERT-FID, Qwen2-7B, Qwen2-1.5B, and our fine-tuned domain-aligned LLM* (based on Qwen2-1.5B). We also assess the performance of combining the original BERT with the domain-aligned LLM and BERT-FID with the domain-aligned LLM.

Table 1 presents the accuracy of different instruction reasoning methods. Our enhanced BERT-FID outperforms the original BERT (0.749 vs. 0.714). Among LLMs,

Range	Method	Precision	Recall	F1
Top 5	General CLIP	0.459	0.400	0.429
	Fine-tuned CLIP	0.817	0.763	0.790
	Text similarity	0.904	0.873	0.888
	Text similarity with fine-tuned CLIP	0.926	0.890	0.908
Top 10	General CLIP	0.578	0.521	0.549
	Fine-tuned CLIP	0.875	0.842	0.859
	Text similarity	0.930	0.913	0.921
	Text similarity with fine-tuned CLIP	2*0.953	0.932	0.943
Top 20	General CLIP	0.688	0.636	0.662
	Fine-tuned CLIP	0.923	0.904	0.913
	Text similarity	0.946	0.936	0.941
	Text similarity with fine-tuned CLIP	0.966	0.954	0.960

Table 3: Results on various retrieval ranges and various methods in environment recognition. By combining text matching with the fine-tuned CLIP model, our method can improve the accuracy across several retrieval ranges.

our domain-aligned model significantly improves over the Qwen2-1.5B baseline by 26.6%, and even surpasses the larger Qwen2-7B by 9.9%. Table 2 shows the runtime performance. While BERT-based models achieve the highest speed (1,672 QPS), Qwen2-7B and our domain-aligned LLM reach only 46 and 245 QPS, respectively. To balance speed and accuracy, we combine BERT-FID with the domain-aligned LLM, achieving 813 QPS with only a 2.2% accuracy drop compared to the LLM-only setup. In contrast, combining the original BERT with LLM leads to a 7.7% accuracy drop, further validating the effectiveness of our proposed instruction reasoning strategy.

For evaluating the *entity retrieval* component, we compare our method with several baselines: general CLIP, fine-tuned CLIP, a text similarity method, and our integrated approach combining text similarity with fine-tuned CLIP.

Table 3 displays the results of these methods across retrieval ranges of top 5, 10, and 20¹. Across all ranges, CLIP fine-tuned with our in-house asset label dataset significantly outperforms the general CLIP. The text similarity method also achieves good retrieval results. Furthermore, our combined approach consistently delivers the best outcomes, with F1 scores exceeding 0.9. Although this method results in a lower QPS compared to using only the text similarity method (as shown in Table 1), the operational speed remains within an acceptable range. As enhancing retrieval accuracy is our primary goal, this method effectively achieves it.

Finally, we evaluate the *complete F.A.C.U.L. model*. For comparison, we also test a fast version using BERT-FID with

¹We use the entity retrieval method to give a list of potential entities and then select the most suitable game asset from this list based on real-time in-game data such as the player’s orientation. Therefore, this method necessitates $n \geq 5$, ensuring a more comprehensive selection process than merely opting for the top 1 result.




Figure 7: User study results on player preferences for four game modes across six key aspects.

Method	Accuracy	QPS	Latency
BERT-FID and text similarity	73.9%	2050	599
LLM* and fine-tuned CLIP	89.2%	210	614
F.A.C.U.L. (Our full method)	87.2 %	916	613

Table 4: Evaluation of intent and environment recognition.

text similarity, and a slower version using LLM with CLIP.

Table 4 indicates that by integrating four model inferences, F.A.C.U.L. achieves an accuracy of 87.2% and a QPS of 916, with an average response time of approximately 613 ms, providing a high-precision, cost-effective solution for commercial games. All three approaches have similar latency (~600 ms). The BERT-FID and text similarity version reaches the highest speed (2050 QPS) but has the lowest accuracy—about 15% lower than the other two. In contrast, our F.A.C.U.L. model reduces accuracy by only 2% compared to the full-large-model version, while achieving a 4.5× gain in QPS. These results demonstrate that our approach effectively balances accuracy and speed, making it practical for real-world deployment.

User Studies

Study Design. To assess the real-world effectiveness and user reception of our AI companions, we conduct a user study within *Arena Breakout: Infinite*, comparing four collaboration modes: 1) our proposed F.A.C.U.L. system, 2) random human teammate matching, 3) wheel-based teammate control, and 4) solo play. Participants are selected based on frequent gameplay and prior experience with F.A.C.U.L., ensuring familiarity with all modes. A total of 24 regular players participated in this preliminary study. The study evaluates player preferences across six key aspects: 1) most enjoyable mode, 2) mode contributing most to winning, 3) best for teammate interaction, 4) most accurate in command execution, 5) most efficient in collaboration, and 6) most preferred for casual play.

Study Results. Fig. 7 illustrates user preferences across the six key aspects. The results show that F.A.C.U.L. and random human teammate matching are the most preferred modes for overall enjoyment, with F.A.C.U.L. particularly favored for its superior performance in teammate interaction, command execution accuracy, and collaboration efficiency. Participants report that F.A.C.U.L.’s integration of natural language and environmental context provides a more immersive and intuitive experience. In contrast, the command wheel mode consistently ranks lowest in most as-

pects, including enjoyment, interaction, and collaboration efficiency, indicating it is less appealing to players seeking a dynamic and interactive experience. Regarding perceived assistance in winning, both F.A.C.U.L. and random human teammate matching are seen as the most helpful, offering better synergy and coordination. Solo play and the command wheel mode are regarded as less effective in this regard. For casual play, participants overwhelmingly favor F.A.C.U.L., highlighting its suitability for more relaxed and enjoyable gaming experiences.

In summary, the results demonstrate that F.A.C.U.L. excels in effective interaction, command execution accuracy, and collaboration efficiency, making it a preferred mode among players. These advantages contribute to its popularity, positioning F.A.C.U.L. as a highly practical option that enhances player satisfaction in both competitive and casual gaming experiences.

Conclusion and Future Work

This paper presents F.A.C.U.L., a cost-effective and high-precision solution for AI companions with advanced natural language understanding and visual perception. It demonstrates how AI companions can assist players in achieving their goals and enhance the gaming experience, offering a practical solution for commercial cooperative games. We propose BERT-FID, which accounts for the uncertainty of complex and ambiguous inputs, enabling effective collaborative inference with LLMs in terms of both accuracy and speed. For scene recognition, we introduce a multi-modal dynamic entity retrieval system that aligns human intentions with decision-making elements in 3D game environments. A real-world study shows that F.A.C.U.L. understands 87.2% of natural language commands and identifies dynamic entities, resulting in a 77% reduction in inference deployment costs compared to baselines.

F.A.C.U.L. provides an efficient framework for AI teammates in FPS games by leveraging the built-in game engine, with a clear applicable scope tied to its latency: it excels in real-time collaborative FPS scenarios (100–800ms latency acceptable). It also has potential in RPGs and MOBAs with minimal adaptations (e.g., adjusted intent templates, game asset datasets), requiring only BERT datasets, LLM prompts, and game asset datasets from developers. Future work could focus on integrating audio signals to enhance voice processing and reduce reasoning latency. Furthermore, we plan to incorporate distinct personality traits for AI companions, allowing for more engaging player interactions.

Acknowledgments

Xiaogang Jin was supported by the National Natural Science Foundation of China (Grant No. 62472373).

Ethical Statement

Informed consent was obtained from all participants prior to their involvement in the user study. Participants were fully informed about the study’s purpose, procedures involved, and any potential risks and benefits. All data collected were anonymized to maintain confidentiality, with access restricted to the research team. The behavior of NPCs within games is dictated by developers’ intentions and undergoes a thorough review to ensure compliance with social and ethical standards.

We are committed to ensuring our research benefits users while minimizing potential harm. AI companions are designed to enhance gaming experiences and engage players positively. Extensive testing, ongoing player feedback through surveys, community discussions, and game analyses are conducted to ensure that AI behaviors within games do not promote real-world violence or aggressive behaviors. Additionally, any potential psychological impacts on players are closely monitored and addressed.

References

- Brown, T. B.; Mann, B.; Ryder, N.; et al. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, 1877–1901.
- Chen, C.; Borgeaud, S.; Irving, G.; et al. 2023. Accelerating large language model decoding with speculative sampling. *arXiv preprint arXiv:2302.01318*.
- Chen, D.; Zhuang, Y.; Zhang, S.; et al. 2024. Data Shunt: Collaboration of Small and Large Models for Lower Costs and Better Performance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10, 11249–11257.
- Faghri, F.; Fleet, D. J.; Kiros, J. R.; and Fidler, S. 2018. VSE++: Improving Visual-Semantic Embeddings with Hard Negatives. 1–13.
- FAIR; Bakhtin, A.; Brown, N.; et al. 2022. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science*, 378(6624): 1067–1074.
- Frome, A.; Corrado, G. S.; Shlens, J.; et al. 2013. DeViSE: A Deep Visual-Semantic Embedding Model. In *Advances in Neural Information Processing Systems*, 2121–2129.
- Gallotta, R.; Todd, G.; Zammit, M.; et al. 2024. Large language models and games: A survey and roadmap. *arXiv preprint arXiv:2402.18659*.
- Gao, Y.; Liu, F.; Wang, L.; et al. 2023. Towards effective and interpretable human-agent collaboration in moba games: A communication perspective. *arXiv preprint arXiv:2304.11632*.
- Huang, Z.; Xu, W.; and Yu, K. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991*.
- Jia, C.; Yang, Y.; Xia, Y.; et al. 2021. Scaling up visual and vision-language representation with noisy text supervision. In *Proceedings of the International Conference on Machine Learning*, 4904–4916. PMLR.
- Kenton, J. D. M.-W. C.; and Toutanova, L. K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*, 4171–4186.
- Kumar, Y.; Koul, A.; and Singh, C. 2023. A deep learning approaches in text-to-speech system: a systematic review and recent research perspective. *Multimedia Tools and Applications*, 82(10): 15171–15197.
- Lafferty, J.; McCallum, A.; Pereira, F.; et al. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the International Conference on Machine Learning*, 282–289. Morgan Kaufmann.
- Leviathan, Y.; Kalman, M.; and Matias, Y. 2023. Fast Inference from Transformers via Speculative Decoding. In *Proceedings of the International Conference on Machine Learning*, 19274–19286. PMLR.
- Lin, B. Y.; Fu, Y.; Yang, K.; et al. 2023. SwiftSage: A Generative Agent with Fast and Slow Thinking for Complex Interactive Tasks. In *Advances in Neural Information Processing Systems*.
- Liu, B.; and Lane, I. R. 2016. Attention-Based Recurrent Neural Network Models for Joint Intent Detection and Slot Filling. In *Annual Conference of the International Speech Communication Association*, 685–689. ISCA.
- Malik, M.; Malik, M. K.; Mehmood, K.; and Makhdoom, I. 2021. Automatic speech recognition: a survey. *Multimedia Tools and Applications*, 80(6): 9411–9457.
- NetEast; and Entertainment, . 2020. Naraka: BladePoin. Melle-Focused Battle Royale.
- Park, J. S.; O’Brien, J.; Cai, C. J.; et al. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology*, 1–22.
- Patel, A.; Li, B.; Rasooli, M. S.; et al. 2022. Bidirectional language models are also few-shot learners. *arXiv preprint arXiv:2209.14500*.
- Radford, A.; Kim, J.; Hallacy, C.; et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the International Conference on Machine Learning*, 8748–8763. PMLR.
- Rao, S.; Xu, W.; Xu, M.; et al. 2024. Collaborative Quest Completion with LLM-driven Non-Player Characters in Minecraft. *arXiv preprint arXiv:2407.03460*.
- Shao, Y.; Li, L.; Dai, J.; and Qiu, X. 2023. Character-LLM: A Trainable Agent for Role-Playing. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 13153–13187.
- Shao, Y.; Mao, J.; Liu, Y.; et al. 2020. BERT-PLI: Modeling paragraph-level interactions for legal case retrieval. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 3501–3507.

Silver, D.; Huang, A.; Maddison, C. J.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587): 484–489.

Tencent. 2024. Arena Breakout: Infinite. Next-Gen Immersive Tactical FPS on Mobile.

Vaswani, A.; Shazeer, N.; Parmar, N.; et al. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, 5998–6008.

Vinyals, O.; Babuschkin, I.; Czarnecki, W. M.; et al. 2019. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782): 350–354.

Xu, X.; Wang, Y.; Xu, C.; et al. 2024. A survey on game playing agents and large models: Methods, applications, and challenges. *arXiv preprint arXiv:2403.10249*.

Yang, A.; Pan, J.; Lin, J.; et al. 2022. Chinese CLIP: Contrastive Vision-Language Pretraining in Chinese. *arXiv preprint arXiv:2211.01335*.

Yang, A.; Yang, B.; Hui, B.; et al. 2024. Qwen2 Technical Report. *arXiv preprint arXiv:2407.10671*.

Zhang, C.; Yang, K.; Hu, S.; Wang, Z.; Li, G.; Sun, Y.; Zhang, C.; Zhang, Z.; Liu, A.; Zhu, S.-C.; et al. 2024a. Proagent: building proactive cooperative agents with large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 17591–17599.

Zhang, L.; Peng, X.; Yang, S.; and Wang, F. 2024b. A Text-to-Game Engine for UGC-Based Role-Playing Games. *arXiv preprint arXiv:2407.08195*.

Zhang, Q.; Li, J.; Jia, Q.; et al. 2021. UNBERT: User-News Matching BERT for News Recommendation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 3356–3362.