

# IntentCUA: Learning Intent-level Representations for Skill Abstraction and Multi-Agent Planning in Computer-Use Agents

Seoyoung Lee\*  
Sookmyung Women’s University  
Seoul, Republic of Korea  
leesy3891@gmail.com

Seobin Yoon\*  
Sookmyung Women’s University  
Seoul, Republic of Korea  
binsong2@sookmyung.ac.kr

Seongbeen Lee†  
Sookmyung Women’s University  
Seoul, Republic of Korea  
seongbeen@sookmyung.ac.kr

Yoojung Chun  
Sookmyung Women’s University  
Seoul, Republic of Korea  
yj.chun@sookmyung.ac.kr

Dayoung Park  
Sookmyung Women’s University  
Seoul, Republic of Korea  
pdysicist@sookmyung.ac.kr

Doyeon Kim  
Sookmyung Women’s University  
Seoul, Republic of Korea  
ehdus@sookmyung.ac.kr

Joo Yong Sim†  
Sookmyung Women’s University  
Seoul, Republic of Korea  
jysim@sookmyung.ac.kr

## ABSTRACT

Computer-use agents operate over long horizons under noisy perception, multi-window contexts, evolving environment states. Existing approaches, from RL-based planners to trajectory retrieval, often drift from user intent and repeatedly solve routine subproblems, leading to error accumulation and inefficiency.

We present IntentCUA, a multi-agent computer-use framework designed to stabilize long-horizon execution through intent-aligned plan memory. A Planner, Plan-Optimizer, and Critic coordinate over shared memory that abstracts raw interaction traces into multi-view intent representations and reusable skills. At runtime, intent prototypes retrieve subgroup-aligned skills and inject them into partial plans, reducing redundant re-planning and mitigating error propagation across desktop applications.

In end-to-end evaluations, IntentCUA achieved a 74.83% task success rate with a Step Efficiency Ratio of 0.91, outperforming RL-based and trajectory-centric baselines. Ablations show that multi-view intent abstraction and shared plan memory jointly improve execution stability, with the cooperative multi-agent loop providing the largest gains on long-horizon tasks. These results highlight that system-level intent abstraction and memory-grounded coordination are key to reliable and efficient desktop automation in large, dynamic environments.

## KEYWORDS

Computer-use agents, Long-horizon automation, Noisy perception, Multi-window context, Multi Agent Planning

\*Both authors contributed equally to this research.

† Corresponding author.



This work is licensed under a Creative Commons Attribution International 4.0 License.

## ACM Reference Format:

Seoyoung Lee, Seobin Yoon, Seongbeen Lee, Yoojung Chun, Dayoung Park, Doyeon Kim, and Joo Yong Sim. 2026. IntentCUA: Learning Intent-level Representations for Skill Abstraction and Multi-Agent Planning in Computer-Use Agents. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/BRAG3288>

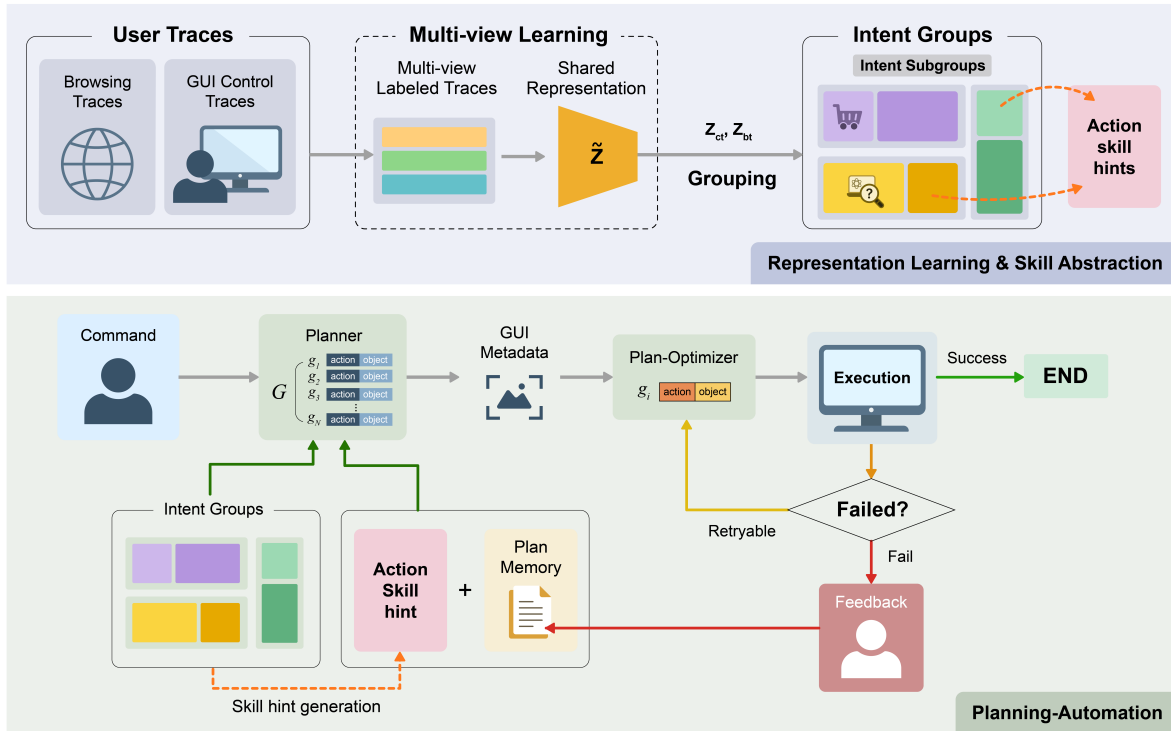
## 1 INTRODUCTION

Rule-based macros and RPA systems enabled early forms of computer-use automation. However, they lack adaptability [11, 33] when compared to recent GUI agents powered by large language models (LLMs) that can interpret screens and generate actions dynamically.

Research on GUI agents has rapidly expanded, spanning web, mobile, and increasingly desktop environments [26, 40]. As highlighted by Tang et al. [32], automation across all desktop environments remains particularly challenging due to multi-window operations, OS-level shortcuts and APIs, and the need to adapt to frequent updates and complex software ecosystems. Within such environments, achieving robust long-horizon planning and managing multi-context workflows emerge as central challenges that current systems have yet to overcome.

Recent multi-modal agents attempt to address these challenges by perceiving screens and generating actions with large models [1, 38]. However, robust long-horizon automation across heterogeneous desktop applications remains unresolved [25, 32]. We identify two recurring failure modes: (i) plans spanning multiple substeps often drift from the original intent and redundantly re-solve already completed routines [23, 24], (ii) local perception errors accumulate and lead to cascading retries [8, 15, 19, 44]. These factors collectively hinder robust long-horizon planning, as agents frequently fall into inefficient and repetitive re-planning cycles. Actions are often retried or nullified when context drifts, leading to prolonged latency and unstable completion rates.

To address these limitations, we bridge user interaction and multi-agent planning. Rather than simply replaying trajectories or storing textual reflections [28], we transform interaction traces into



**Figure 1: Overview of IntentCUA.** *Offline:* raw user traces are multi-view labeled, embedded into a shared intent space, and clustered into intent groups (IG) and subgroups (SG); SG action patterns are converted into parameterized skill schemas (“skill hints”) and stored with their SG in the IG/SG index, while plan memory stores only user-approved global plans (G). *Online:* the Planner/Plan-Optimizer/Critic query and reuse skills; cache-first reuse and template-based gap filling reduce re-planning on long-horizon desktop tasks.

labeled units, induce generalized skills from sub-intent clusters, and learn multi-view representations across environment, action, keyword, and description.

These skills are organized hierarchically in a plan memory and retrieved via semantic search during planning, which supports cross-application transfer and helps stabilize long roll-outs. At runtime, intent prototypes are projected into a shared embedding space, where centroid-based retrieval augments partial plans with relevant skills.

In end-to-end evaluations, IntentCUA achieves a 74.83% task success rate with a Step Efficiency Ratio (SER) of 0.91, outperforming both RL-based (UI-TARS-1.5 [27]) and trajectory-centric (UFO<sup>2</sup> [41]) baselines in success rate, efficiency, and latency. Ablation studies confirm that multi-view intent abstraction and shared plan memory jointly improve execution stability, with the cooperative multi-agent loop providing the largest gains on long-horizon tasks. These results indicate that system-level intent abstraction and memory-grounded coordination are central to reliable desktop automation.

Our contributions are summarized as follows:

- (1) We propose **IntentCUA**, a multi-agent computer-use framework that stabilizes long-horizon execution through intent-aligned plan memory and coordinated planning.
- (2) We introduce a trace-to-skill abstraction pipeline that learns multi-view intent representations and induces hierarchical, reusable skills from raw user interaction traces.
- (3) We design a planning-time memory mechanism that retrieves subgroup-aligned skills to augment partial plans, reducing intent drift and redundant re-planning in dynamic desktop environments.
- (4) We demonstrate through extensive ablations and end-to-end evaluations that intent abstraction and memory-grounded coordination significantly improve execution stability, efficiency (SER 0.91), and task success (74.83%) on complex desktop workflows.

In summary, IntentCUA shows that intent-level abstraction and memory-grounded multi-agent coordination are key to stabilizing long-horizon desktop automation in large, dynamic environments. This automation is made possible by a robust planning policy that maintains coherence and efficiency across extended sequences.

## 2 RELATED WORK

### 2.1 Desktop and GUI Automation Agents

GUI automation spans web, mobile, and desktop domains. Web agents such as WebArena and WebVoyager operate under structured DOM feedback [7, 45], but real desktop environments lack such schema-level constraints and require cross-application coordination.

Desktop benchmarks like OSWorld [36] tasks are typically long-horizon, requiring stable execution 10–20 sequential steps. This makes execution latency-sensitive and error-prone. As step count increases, local perception errors compound and agents often enter loops of repeated or failed actions.

Recent desktop agents such as UI-TARS [22], UFO [42], ScreenAgent [20] extend vision-language models with planner-critic loops. However, surveys highlight a persistent challenge: determining actions that align with specific user contexts and preferences in dynamic, interruption-prone interfaces [32]. Even with improved GUI grounding [8, 10, 19], intent drift and redundant re-planning remain common in long-horizon workflows.

These findings indicate that stable long-horizon planning, rather than perception alone, remains the key bottleneck for reliable desktop automation.

### 2.2 Agents Leveraging Interaction Traces

One approach to addressing long-horizon instability is to learn directly from large-scale interaction traces.

Macro-mining and process-mining techniques cluster demonstrations into recurrent procedures or labeled schemas [4, 5, 9]. Large-scale corpora such as OS-ATLAS support perception pretraining across millions of GUI elements [35]. Offline reinforcement learning has also been explored for device agents [29], while systems such as AppAgentv2 [14], AgentBank [31], and UI-TARS-1.5 [27] leverage hierarchical feedback or large-scale trajectory tuning to improve control robustness.

These works demonstrate that interaction traces improve policy generalization and low-level stability. However, most approaches operate at the trajectory or action level, emphasizing replay or large-scale tuning rather than structured intent abstraction. As a result, redundancy and error accumulation often persist in long-horizon execution [24], and reliance on controlled environments or explicit reward signals limits applicability to open-ended desktop workflows [26].

### 2.3 Plan Memory, Intent Identification, and Skill Abstraction

A complementary direction enhances robustness through memory retrieval and skill abstraction.

Memory-based methods such as Reflexion [28], Conversational Memory [34], and Contextual Experience Replay [18] retrieve prior trajectories, manuals, or reflections to guide future decisions [2]. Skill-level prompting approaches such as SkillAct [17] show that abstracted routines can improve interactive performance, while UFO<sup>2</sup> [41] manages app-specific demonstrations as reusable references.

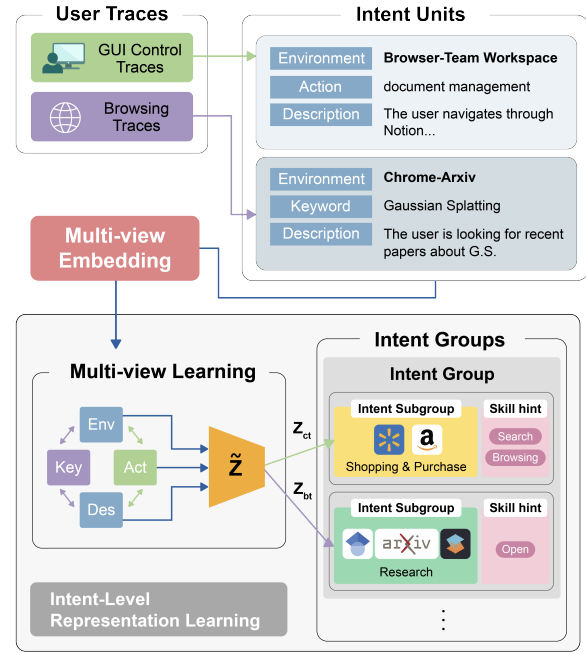


Figure 2: Multi-view intent representation. Control traces use [E,A,D], browsing traces [E,K,D]. A multi-view encoder aligns views into a shared space, inducing environment-centric IG and finer SG. SG centroids enable retrieval, and SG action patterns are converted into parameterized skill schemas (“skill hints”) with verb-argument structure for planning.

Parallel work investigates intent recognition from UI logs [13, 23] and representation learning of screens (e.g., Screen2Vec [12], Aria-UI [37]), while GUI grounding methods reduce perceptual ambiguity [15, 19]. More recent systems explore adaptive planning and dependency modeling from demonstrations [39, 43].

Despite these advances, structured and hierarchical skill abstractions that remain transferable across heterogeneous desktop workflows are still relatively underexplored. As a result, maintaining stable long-horizon execution under dynamic user contexts continues to be an active area of research. Our approach complements these directions by learning multi-view intent representations that integrate environment, action, and description signals. Skills are stored as hierarchical intent prototypes in plan memory and retrieved to augment partial plans, supporting stable long-horizon execution [6, 30].

## 3 INTENT-LEVEL REPRESENTATION LEARNING & SKILL ABSTRACTION

### 3.1 Intent-level Representation Learning

In this section, we describe how raw user traces are transformed into unified intent-level representations that can be clustered and later abstracted into reusable skills.

As shown in Figure 2, each user trace is first labeled across four views: environment ( $E$  with instances  $e_i$ ), action ( $A$  with instances  $a_i$ ), keyword ( $K$  with instances  $k_i$ ), and description ( $D$  with instances  $d_i$ ), where  $i$  indexes the sequential intent units that together compose a user’s interaction trace. Each view  $v \in \{E, A, K, D\}$  is represented as an embedded textual vector, capturing its semantic content. Control traces produce intent units,  $u_i$ , of the form  $[e_i, a_i, d_i]$ , while browsing traces yield  $[e_i, k_i, d_i]$ . Formally, let  $x_i^{(v)}$  denote the feature representation of intent unit  $u_i$  in view  $v \in \{E, A, K, D\}$ . A multi-view encoder  $\phi(x^{(v)})$  maps these view-specific features into a single shared representation  $z_i$ :

$$z_i = \phi\left((x_i^{(v)})_{v \in V}\right) \in \mathbb{R}^d, \quad V \subseteq \{E, A, K, D\}. \quad (1)$$

Building on prior multi-view clustering objectives [16], we train the model to ensure that representations from heterogeneous views are (i) contradistinctively aligned, (ii) cross-view predictive, and (iii) reconstructible. The overall loss is defined as the weighted sum of these three components, as shown in Equation 2:

$$\mathcal{L} = \mathcal{L}_{\text{con}} + \lambda_{\text{pred}} \mathcal{L}_{\text{pred}} + \lambda_{\text{rec}} \mathcal{L}_{\text{rec}} \quad (2)$$

where  $\mathcal{L}_{\text{con}}$ ,  $\mathcal{L}_{\text{pred}}$ ,  $\mathcal{L}_{\text{rec}}$  are cross-view contrastive loss, dual prediction loss, and within-view reconstruction loss, respectively.

$\mathcal{L}_{\text{con}}$ , enforces consistency between embeddings from different views of the same intent unit while separating embeddings from different instances:

$$\mathcal{L}_{\text{con}} = \frac{1}{|P(V)|} \sum_{(p,q) \in P(V)} \left[ -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\langle z_i^{(p)}, z_i^{(q)} \rangle / \tau)}{\sum_{j \neq i} \exp(\langle z_i^{(p)}, z_j^{(q)} \rangle / \tau)} \right] \quad (3)$$

where,  $P(V)$  denotes the set of all ordered positive view pairs  $(p, q)$  within the selected view set  $V$ . The temperature parameter  $\tau$  controls the sharpness of the contrastive distribution, whereas  $N$  denotes the number of intent units sampled in a minibatch.

$\mathcal{L}_{\text{pred}}$  introduces two projection heads  $G_{p \rightarrow q}$  and  $G_{q \rightarrow p}$  that learn to predict the embedding of one view from another. Their averaged mapping  $G = (G_{p \rightarrow q} + G_{q \rightarrow p})/2$  acts as a symmetric predictor encouraging cross-view consistency—ensuring that one view can reconstruct another within the latent space. The scalar coefficient  $\lambda_{\text{pred}}$  balances this term with the others in Equation 2.

$$\mathcal{L}_{\text{pred}} = \frac{1}{|P(V)|} \sum_{(p,q) \in P(V)} \frac{1}{2N} \sum_{i=1}^N \left[ \|G_{p \rightarrow q}(z_i^{(p)}) - z_i^{(q)}\|_2^2 + \|G_{q \rightarrow p}(z_i^{(q)}) - z_i^{(p)}\|_2^2 \right] \quad (4)$$

Finally,  $\mathcal{L}_{\text{rec}}$  ensures that the shared embedding  $z_i$  retains view-specific semantics by reconstructing each original feature  $x_v^{(u)}$  through a decoder  $g_v(\cdot)$ . The weight  $\lambda_{\text{rec}}$  determines the relative strength of this reconstruction constraint within the total loss in Equation 2.

$$\mathcal{L}_{\text{rec}} = \frac{1}{|V|N} \sum_{v \in V} \sum_{i=1}^N \|g_v(z_i^{(v)}) - x_i^{(v)}\|_2^2 \quad (5)$$

Together, these three objectives jointly align, predict, and reconstruct heterogeneous views, yielding a coherent embedding space where intent-level semantics are preserved. The resulting representation  $z_i$  compactly captures user intent across multiple modalities.

We then organize  $\{u_i\}$  hierarchically by  $z_i$ : first into higher-level intent groups ( $IG$ ) driven by environment context, and then into subgroups ( $SG$ ) based on action/keyword and description. These  $SG$  representations provide the foundation for extracting recurrent action patterns and constructing abstract skills, as described in Section 3.2.

### 3.2 Skill abstraction based on Intent Subgroups

Given the per-unit embeddings  $z^{(u)}$  from Section 3.1, we organize intents into a two-level index for planning. First, we run HDBSCAN [3] over  $\{z^{(u)}\}$  to obtain higher-level *intent groups* ( $IG$ ) driven primarily by environment/context. Within each  $IG$ , a second HDBSCAN partitions units into finer *subgroups* ( $SG$ ) using action/keyword and description signals. For every  $SG$ , we compute and store its centroid  $c_{SG}$  in the same embedding space. At retrieval time, we rank subgroups by the cosine similarity between  $c_{SG}$  and a query *intent prototype*. **Retrieval index:** for each  $SG$  we store (i) the centroid  $c_{SG}$  (for cosine-based ranking); (ii) top- $k$  *representative traces* preselected by proximity to  $c_{SG}$  and reranked at query time by similarity to the intent prototype; and (iii) a *support* count for  $SG$  (defined below) used to prefer stable patterns during planning.

We then consolidate the low-level action traces inside each  $SG$  into a reusable skill. Each intent unit  $u_i$  in a subgroup is linked to a low-level action sequence:  $M_i = (a_1, a_2, \dots, a_m)_i$ . To make traces comparable, we map every atomic action  $a_i$  to a pair [*verb predicate*, *typed argument fields*] by applying an alias map  $\Phi$  that collapses surface variants to a fixed predicate and a fixed set of typed fields. For example, "focus URL bar" and "open web site"  $\rightarrow$  verb=press, arg=address\_bar, verb=text\_input, arg=address\_bar, text:"https://example.com". Here, a *verb signature* is the ordered list of canonical predicates in a trace, and an *typed argument field* is a placeholder (e.g., <url>, <query>, <file\_path>) that will be bound at runtime.

We collect candidates  $\text{sg}_{\text{skills},k} = \{M_i \mid u_i \in SG_k\}$  and induce a *skill prototype* as the medoid under a signature-level dissimilarity  $d_{\text{sig}}$  over verb-predicate sequences. The function  $d_{\text{sig}}$  is computed on the canonicalized verb-predicate sequences (after applying  $\Phi$ ), comparing action patterns at the predicate level while deferring literal-argument handling to the parameterization stage. Let  $\mathcal{A}$  denote the verb-predicate alphabet. Thus the subgroup’s *skill prototype* is defined as:

$$\mathcal{S}_{SG} = \arg \min_{a \in \mathcal{A}^*} \sum_{s \in \text{sg}_{\text{skills}}} d_{\text{sig}}(a, s) \quad (6)$$

Next we convert  $\mathcal{S}_{SG}$  into a reusable, *parameterized schema*, *skill hint*: a verb-predicate sequence together with a typed argument structure. This conversion (i) replaces literal values with typed parameters (the runtime-filled fields), (ii) removes incidental or recovery-specific steps that do not affect goal attainment, and (iii) enforces canonical predicate and field names via  $\Phi$ . We refer to this parameterized schema as a *skill hint*. Both the skill hints  $\{\mathcal{S}_{SG}\}$  and the representative traces are stored in plan memory and retrieved at planning time.

If multiple predicate sequences are well supported, we keep several schemas ranked by their *support*, where support counts subgroup members whose similarity to  $\mathcal{S}_{SG}$  exceeds a fixed threshold  $\tau$ . *Representative traces* are the top- $k$  members minimizing  $d_{\text{sig}}$  to

*SSG* and serve as concrete exemplars. At planning time, when a retrieved plan is only a partial match, we perform **gap filling**: we instantiate the selected *skill hint SSG* with current-context bindings and insert the resulting steps to complete the missing plan units/steps (Section 4.1).

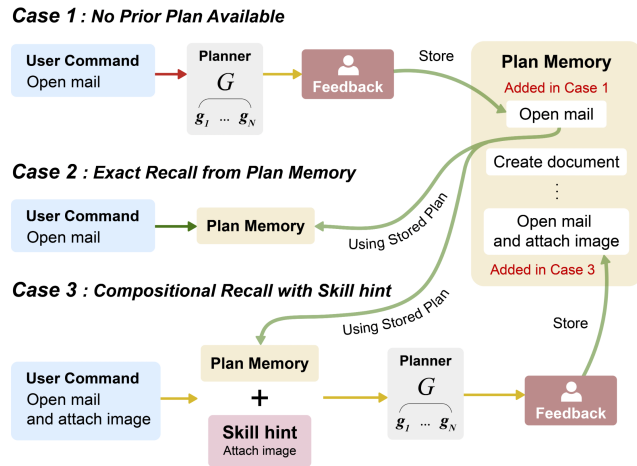
## 4 INTENT-AWARE PLANNING & FEEDBACK MEMORY

Building on the intent-level DB introduced in Section 3, we now focus on how these abstractions are leveraged during planning and execution. This section details the end-to-end workflow in which the Planner, Plan-Optimizer, and Critic cooperate through plan memory to compose, refine, and verify long-horizon automation. The Planning-Automation part in figure 1 represents the overall automation process after the user request.

### 4.1 Planning with Plan Memory and Skill Hints

This section explains how the Planner agent composes a high-level plan  $G$  for a given command; Figure 3 summarizes three pathways: cache miss synthesis (Case 1), direct reuse on exact hit (Case 2), and reuse-with-injection on partial hit (Case 3). At a high level, Planner consults  $IG/SG$  centroids in the shared embedding space, plan-memory entries, and  $M(u)$  sequences; missing spans are completed with normalized  $S_{SG}$  templates.

When no suitable plan exists in plan memory (Figure 3, Case 1), the Planner embeds the command into the shared space, gates by the active  $IG$ , and ranks candidate  $SG$  by centroid similarity. Plan memory stores previously synthesized global plans  $G$  that received *binary* user approval. Plans not approved are discarded. For each plan-unit slot in the intent prototype, it selects the nearest  $SG$ , retrieves the top-2 intent units ( $u_i$ ) from that  $SG$ , and uses



**Figure 3: Cache-first planning with plan memory.** A query intent is gated by  $IG$  and ranked over  $SG$ . Case 1 (miss): synthesize a plan from retrieved skill templates. Case 2 (hit): reuse the stored plan. Case 3 (partial): align to the nearest plan and fill gaps with  $SG$ -derived skill hints, reducing retries.

the GPT-4o model [21] to generate the slot’s steps conditioned on the retrieved  $M_i$  sequences. A *plan unit* is a contiguous block of low-level steps in the global plan  $G$  that achieves one subgoal. A *plan-unit slot* is the placeholder for such a block in the intent prototype that the Planner must populate with executable steps.

Concatenating all slots yields a high-level plan  $G = \{g_1, \dots, g_n\}$ , where each  $g_j$  expands to a contiguous, execution-ordered list of low-level actions. After the user reviews the generated plan and provides optional feedback, we incorporate the edits and then store  $G$  in plan memory for future reuse. Concretely,  $G$  is materialized as *plan units*—intent-prototype-level chunks of the  $G$  derived from the initial user command: Plan Unit 1: [step\_1, step\_2, ...] Plan Unit 2: [step\_k, ..., step\_l] ...

When a high-similarity plan is found in active  $IG$  (Figure 3, Case 2), the stored  $G$  is retrieved and its slots are bound to the current context; because no gaps remain, the planner skips the synthesis and executes the plan as-is.

When only a partial match is found (Figure 3, Case 3), the closest stored plan  $G$  is aligned with the intent prototype. Insert missing plan units or steps by injecting the matched subgroup’s  $S_{SG}$  with current-context bindings (gap filling), after which the finalized  $G$  is executed. This cache-first pipeline reduces re-planning and stabilizes long-horizon execution by combining centroid gating, plan reuse, and hint-based gap filling.

### 4.2 Optimizing steps by memory & feedback loop

Given a finalized plan  $G$ , execution shifts into a cooperative loop between the Plan-Optimizer and the Critic Agent, utilizing the Plan Memory. The Plan-Optimizer refines each plan unit by referencing similar traces stored in the memory, dynamically adapting its substeps to current screen contexts. The Critic, in turn, monitors the execution and provides immediate feedback signals—success, retryable, or blocked—to correct local deviations or trigger partial replanning when necessary.

For each plan unit  $pu$ , we compute its representation  $z^{(pu)} \in \mathbb{R}^d$  (Section 3.1) and compare it with subgroup centroids  $\{c_{SG}\}$  from plan memory (Section 3.2). If a subgroup is relevant, its traces are injected as *hints* into the Plan-Optimizer to guide step execution. After each unit, the Critic evaluates the post-execution state and returns  $q \in \{\text{success, retryable, blocked}\}$  with an observation  $o$ . If  $q$  is retryable, the Plan-Optimizer is re-invoked on the latest state  $s^{\text{after}}$  with observation observation of the current GUI context ( $o$ ) to produce an adjusted subplan  $g'_{\text{new}}$ , which updates  $G$  before re-execution.

As shown in Algorithm 1, the Planner hands over the plan units to the Plan-Optimizer, which integrates hints from prior traces to refine step execution. The Critic then decides whether to proceed, request an adjustment, or terminate. Through this memory-guided collaboration, specialized agents coordinate to minimize redundant re-planning and improve robustness by reusing traces that previously led to success.

GUI Grounding refers to the process of enumerating all actionable GUI components on the current screen, similar to the screen parsing method used in UFO [42]. The resulting state  $s$  includes such component data together with summary metadata, composed

**Table 1: Component-wise ablation on planning. We report task success (%) and averaged plan completion (%).**

Method	Success (%) $\uparrow$	Completion (%) $\uparrow$
$B$	22.73	33.78
$B + T_g$	46.43	57.41
$B + T_{SG} + Z$	54.64	77.56
$B + T_{SG} + S_{SG} + PM$	53.85	81.23
$B + T_{SG} + Z + PM$	62.51	85.00
$B + T_{SG} + Z + S_{SG} + PM$	<b>74.83</b>	<b>91.14</b>

**Algorithm 1** Execution of the Plan utilizing Memory and Feedback Loop

**Require:** Final global plan  $G$  with plan units  $PU = \{pu_1, \dots, pu_M\}$ , where each  $pu$  is an ordered list of steps; subgroup collection  $SG$  with centroids  $c_{SG} \in \mathbb{R}^d$  (representation space from Section 3.1); for every  $pu \in PU$ , its representation  $z^{(pu)} \in \mathbb{R}^d$  (precomputed via the encoder in Section 3.1); action space  $A$ .

**Ensure:** Execution outcome

```

1: for each plan unit  $pu$  in  $PU$  do
2:    $hint \leftarrow \text{INJECTHINT}(pu, SG, z^{(pu)})$ 
3:   for each step  $g$  in  $pu$  do
4:      $s \leftarrow \text{GUI Grounding of the current screen}$ 
5:      $(a, g', o) \leftarrow \text{PLAN-OPTIMIZER}(s, g, G, pu, o, hint)$ 
6:     Execute action  $a$  in the current GUI context
7:      $s^{\text{after}} \leftarrow \text{GUI Grounding of the screen after finishing } pu$ 
8:      $(q, o) \leftarrow \text{CRITIC}(pu, G, s^{\text{after}})$ 
9:     if  $q == \text{success}$  then
10:      continue to next  $pu$ 
11:     else if  $q == \text{retryable}$  then
12:        $(a, g'_{\text{new}}, o) \leftarrow \text{PLAN-OPTIMIZER}(s^{\text{after}}, g, G, pu, o, hint)$ 
13:       Apply  $g'_{\text{new}}$  to adjust the prior  $g'$ 
14:     else
15:       return  $(G, \text{BLOCKED})$ 
16: return  $(G, \text{SUCCESS})$ 

```

of window title, panel names and component counts captured from the environment. Each step  $g$  denotes an individual operation in the global plan  $G$ , composed of an action  $a$  (e.g., click, text input, open) and its corresponding object targets; thus  $a$  specifies the interaction primitive, whereas  $g$  represents the full executable tuple  $(a, \text{object})$ .

The INJECTHINT function searches the plan memory for previous plan units whose representations  $z^{(pu)}$  are most similar to the current one, and uses their traces as contextual hints guiding the next execution steps.

## 5 ABLATION & CASE STUDIES

### 5.1 Evaluation Setup

We evaluate our design on 286 real-world GUI tasks: 100 in-house, 116 from WebVoyager [7] (643 total), and 70 from ScreenAgent [20] (70 sessions). Tasks span local applications, web platforms, productivity tools, and cross-application workflows.

For task mining, we collect 30 active hours of interaction traces across 18 sessions, yielding 113 trace files. The mined corpus is intentionally distribution-shifted from the test suite: traces skew toward

Local/App, while the evaluation set contains more Web/Crossover tasks. The traces cover 36 domains, whereas the 286-task suite spans 63 domains; only 22 overlap (34.92%), leaving 41 unseen test domains (65.07%). This setup stresses generalization of mining and retrieval rather than memorization.

All agents use the same atomic GUI action interface and identical timeout policies. We report task success (74.83%), average completion ratio (91.14%), Step Efficiency Ratio (successful steps / actual execution steps; higher is better). Differences are reported in percentage points (pp).

### 5.2 End-to-End Execution of Ablated Models

We ablate each component to analyze its impact on long-horizon planning stability and execution depth.

For Table 1, Completion denotes averaged plan completion (executed steps / synthesized plan steps), measuring execution progress beyond binary success.

We denote components as:  $B$  (baseline planner-executor),  $T_g$ ,  $T_{SG}$  (greedy vs.  $IG/SG$ -gated trace retrieval),  $Z$  (intent-level representation),  $S_{SG}$  (subgroup skill hints),  $PM$  (plan memory reuse)

Starting from  $B$  (22.73% success), adding greedy trace retrieval ( $B+T_g$ ) improves success by +23.7, pp showing that user traces substantially reduce cold-start errors but still induce drift in long horizons. Replacing greedy retrieval with  $IG/SG$  gating +  $Z$  brings a further +8.21, pp success gain and +20.15, pp completion gain, showing that organizing traces into representation-learned intent subgroups significantly stabilizes long-horizon execution. Even without  $Z$ , combining  $S_{SG} + PM$  achieves a high task completion rate (81.23%), indicating that skill hints and plan memory alone substantially improve execution depth.

The full system ( $B+T_{SG}+Z+S_{SG}+PM$ ) achieves the best success and completion overall, demonstrating that representation learning, intent-level gating, skill abstraction, and plan reuse act complementarily.

### 5.3 Case Study

To illustrate the framework, we consider a task where the user asks the system to summarize a previously viewed lightweight-ML video and record the result in a personal workspace. This requires retrieval, reasoning, and coordination across multiple applications.

Figure 4 shows the execution process. The Planner retrieves relevant traces from plan memory, including prior interactions with the video platform, AI chatbot, and workspace application, and reconstructs the video source from the user’s history. It decomposes the request into structured intent-level plan units, each grounded into executable GUI actions by the Plan-Optimizer, while the Critic monitors progress and handles local inconsistencies. Through hierarchical reasoning and memory-guided skill retrieval, the system completes the multi-application task while maintaining intent coherence.

A failure case occurs when an unexpected pop-up appears during execution. Because underlying components become occluded, the grounding module fails to detect them, leading to incorrect retries. This highlights a limitation of script-based GUI grounding under transient interface changes.

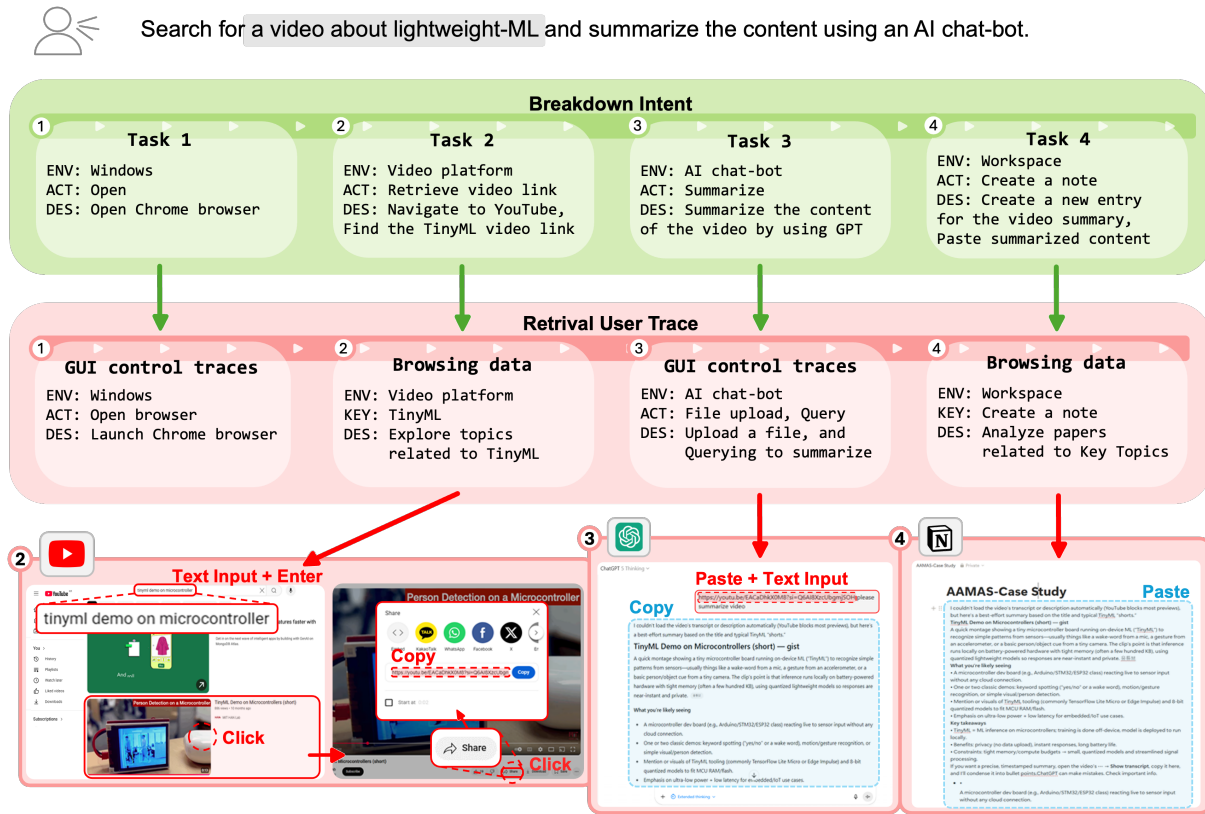


Figure 4: IntentCUA in action: the system recalls intent units from memory and decomposes a multi-application command into intent-level plan units, each executed through learned skills and recomposed into an end-to-end automation plan.

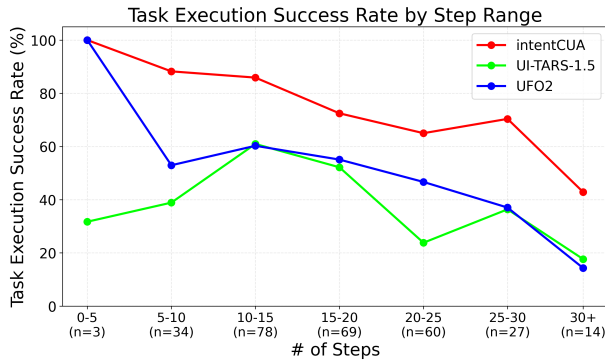


Figure 5: Success rate by step length (bin size = 5 steps). The x-axis shows step-length bins and the y-axis shows task success rate (%).

## 6 EXPERIMENTS

We compare **IntentCUA** against two representative desktop GUI agents chosen for methodological diversity: UI-TARS-1.5 [27], an RL-based visual planner-executor with self-evolving policies and screen grounding, and UFO<sup>2</sup> [41], a trajectory-centric Windows

automation agent that organizes demonstrations as executable sequences. Together these baselines span reinforcement learning-driven automation versus demonstration-driven planning, and both operate at the level of atomic GUI actions, ensuring comparability with our interface. We evaluate 286 tasks (the same evaluation suite described in Section 5.1) and report task success rate, Step Efficiency Ratio(SER), and Latency, further analyzing robustness by step-length bins, each step defined as a atomic action performed by the agent. SER is defined as the ratio of successful steps to total steps, ranging from 0 to 1. Latency is measured as the execution time per task, reflecting not only the number of steps but also the overhead of perception and planning.

Table 2: End-to-end success rate comparison across datasets (%). Columns show WebVoyager, ScreenAgent, our in-house suite, and overall average.

Method	WebVoyager	ScreenAgent	Ours	Total(%)
UI-TARS-1.5 [27]	35.9	42.9	46.0	38.8
UFO <sup>2</sup> [41]	69.0	41.4	38.0	51.2
<b>IntentCUA (ours)</b>	<b>71.6</b>	<b>77.1</b>	<b>78.0</b>	<b>74.8</b>

### 6.1 Robust Long-Horizon Planning Efficiency

We evaluate how each agent sustains task completion as sequence length increases, focusing on the robustness of long-horizon planning. Table 2 and Figure 5 summarize overall and step-wise success trends across 286 evaluation tasks. IntentCUA achieves the highest overall success rate of **74.8%**, compared to 51.2% for UFO<sup>2</sup> and 38.8% for UI-TARS-1.5, yielding relative improvements of about +23.6 and +36 percentage points, respectively.


Notably, IntentCUA performs consistently well across all datasets, achieving 71.6% on the web-based WebVoyager, 77.1% on the cross-application ScreenAgent, and 78.0% on our in-house local suite, demonstrating that its advantage is not confined to a specific benchmark. While agents like UFO<sup>2</sup> specialize in narrow domains such as web navigation, IntentCUA generalizes effectively to heterogeneous desktop environments that include both online and offline contexts, confirming its versatility and domain robustness.

As shown in Figure 5, IntentCUA maintains stable performance even as task length grows: 85.9% at 10–15 steps, 72.5% at 15–20, and 65.0% at 20–25, while still retaining 42.9% beyond 30 steps. Both baselines, in contrast, decline sharply after 20 steps, dropping below 20%. This gradual degradation indicates that IntentCUA’s planning remains consistent and resistant to drift even in extended workflows spanning multiple windows and applications.

The stability across longer horizons can be attributed to its *intent-aware retrieval* and *plan memory reuse*, which enable the planner to recall previously successful subplans aligned with the current intent embedding rather than regenerating them from scratch. Together, these results confirm that IntentCUA achieves robust and generalizable long-horizon planning efficiency, effectively preserving goal coherence and minimizing redundant re-planning under complex, real-world desktop environments.

### 6.2 Stable & Scalable Planning Efficiency and Latency

We examine efficiency using two complementary metrics: the Step Efficiency Ratio (SER; Left) and the average latency per task (Right), as shown in Figure 6. IntentCUA achieves the highest SER of **0.91**, exceeding UI-TARS (0.85) and UFO<sup>2</sup> (0.82). While SER in IntentCUA decreases moderately from 0.93 at 10–15 steps to 0.88 at 20–25, it



**Figure 6: Performance by task length (bin size = 5 steps). Left: Step Efficiency Ratio (SER). Right: Average latency per task (minutes)**

remains consistently above 0.85 even for the longest tasks, indicating that most actions continue to contribute effectively to progress.

In contrast, both baselines show sharper declines across similar ranges, suggesting increased redundancy or re-planning.

Latency patterns further highlight scalability. IntentCUA’s average execution time is 1.46 minutes, approximately 4.5× lower than the baselines (UFO<sup>2</sup>: 6.63 min, UI-TARS: 9.82 min). Its latency increases smoothly with task length—for instance, from 0.95 min at 10–15 steps to 2.01 min at 20–25—showing near-linear growth. By comparison, UI-TARS exhibits irregular delays that expand sharply with step count, and UFO<sup>2</sup> shows unstable spikes on shorter tasks due to looped retries.

These results demonstrate that IntentCUA sustains high planning efficiency and low, predictable latency as task complexity increases. Its memory-guided retrieval and feedback design minimize redundant computation, yielding a scalable and robust planning policy suitable for real desktop automation.

## 7 CONCLUSION

We presented **IntentCUA**, a framework that transforms raw interaction traces into multi-view intent representations, abstracts them into reusable skills, and integrates these with plan memory to support stable long-horizon desktop automation. The system combines representation learning, hierarchical skill induction, and memory-guided planning to reduce re-planning and improve stability across complex workflows.

In experiments, IntentCUA achieved a 74.8% task success rate with a step efficiency ratio of 0.91, outperforming both UI-TARS-1.5 (RL-based) and UFO<sup>2</sup> (trajectory-centric) by 4.5× times reduced latency. It also maintained over 40% success on long-horizon tasks exceeding 30 steps. Ablation studies show that each component contributes to robustness and efficiency, with the full design providing the greatest improvements on longer tasks. Additional implementation details, extended ablations, and analysis are available in the **extended technical report**. While IntentCUA shows consistent reasoning and cross-application generalization, several aspects remain open for refinement. Retrieval efficiency may fluctuate as the plan memory grows, though this mainly affects latency rather than accuracy. Graph-based retrieval and lightweight vision cues could further enhance robustness, allowing the system to adapt more smoothly to dynamic and visually changing interfaces.

## ACKNOWLEDGMENTS

This work was supported by the National Research Foundation of Korea(NRF) grant (No. RS-2025-16070288, No. RS-2025-02216282) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant (No.RS-2022-II220025) funded by the Korea government(MSIT) and Ministry of Trade, Industry and Energy of Korea (MOTIE RS 2023 00258591).

## REFERENCES

- [1] Anthropic. 2024. Introducing computer use, a new Claude 3.5 Sonnet, and Claude 3.5 Haiku. <https://www.anthropic.com/news/3-5-models-and-computer-use>. Published Oct 22, 2024.
- [2] Yuzhe Cai, Shaoguang Mao, Wenshan Wu, Zehua Wang, Yaobo Liang, Tao Ge, Chenfei Wu, WangYou WangYou, Ting Song, Yan Xia, Nan Duan, and Furu Wei. 2024. Low-code LLM: Graphical User Interface over Large Language Models. In *Proc. 2024 Conf of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Vol 3: System Demonstrations)*. Association for Computational Linguistics, Mexico City, Mexico, 12–25. <https://doi.org/10.18653/v1/2024.naacl-demo.2>

- [3] Ricardo J. G. B. Campello, Davoud Moulavi, Arthur Zimek, and Jörg Sander. 2015. Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Trans. Knowl. Discov. Data* 10, 1 (2015), 5.
- [4] Daehyoun Choi, Hind R'bigui, and Chiwoon Cho. 2022. Enabling the gab between RPA and process mining: User interface interactions recorder. *IEEE Access* 10 (2022), 39604–39612.
- [5] Mohammadreza Fani Sani, Michal Sroka, and Andrea Burattin. 2023. LLMs and Process Mining: Challenges in RPA: Task Grouping, Labelling and Connector Recommendation. In *Int'l Conf on Process Mining*. Springer, 379–391.
- [6] Difei Gao, Lei Ji, Zechen Bai, Mingyu Ouyang, Peiran Li, Donzong Mao, Qinchun Wu, Weichen Zhang, Peiyi Wang, Xiangwu Guo, Hengxu Wang, Luwei Zhou, and Mike Zheng Shou. 2024. AssistGUI: Task-Oriented PC Graphical User Interface Automation. In *2024 IEEE/CVF Conf on Computer Vision and Pattern Recognition (CVPR)*. IEEE Computer Society, Los Alamitos, CA, USA, 13289–13298. <https://doi.org/10.1109/CVPR52733.2024.01262>
- [7] Hongliang He, Wenlin Yao, Kaixin Ma, Wenhao Yu, Yong Dai, Hongming Zhang, Zhenzhong Lan, and Dong Yu. 2024. WebVoyager: Building an End-to-End Web Agent with Large Multimodal Models. In *Proc. 62nd Annual Meeting of the Association for Computational Linguistics (Vol 1: Long Papers)*. Association for Computational Linguistics, Bangkok, Thailand. <https://doi.org/10.18653/v1/2024.acl-long.371>
- [8] Wenyi Hong, Weihang Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxiao Dong, Ming Ding, et al. 2024. Cogagent: A visual language model for GUI agents. In *Proc. CVPR '24*. 14281–14290.
- [9] Forrest Huang, Gang Li, Tao Li, and Yang Li. 2024. Automatic macro mining from interaction traces at scale. In *Proc. CHI '24 Conf on Human Factors in Computing Systems*. 1–16.
- [10] Yue Jiang, Eldon Schoop, Amanda Swearngin, and Jeffrey Nichols. 2025. Iluvui: Instruction-tuned language-vision modeling of UIs from machine conversations. In *Proc. 30th Int'l Conf on Intelligent User Interfaces*. 861–877.
- [11] Rebecca Krosnick and Steve Oney. 2022. ParamMacros: Creating UI Automation Leveraging End-User Natural Language Parameterization. In *2022 IEEE Symp on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, Liverpool, United Kingdom, 1–10. <https://doi.org/10.1109/VL/HCC53370.2022.9833005>
- [12] Toby Jia-Jun Li, Lindsay Popowski, Tom Mitchell, and Brad A Myers. 2021. Screen2vec: Semantic embedding of gui screens and gui components. In *Proc. CHI '21 Conf on Human Factors in Computing Systems*. 1–15.
- [13] Yang Li, Jiacong He, Xin Zhou, Yuan Zhang, and Jason Baldridge. 2020. Mapping Natural Language Instructions to Mobile UI Action Sequences. In *Proc. 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 8198–8210. <https://doi.org/10.18653/v1/2020.acl-main.729>
- [14] Yanda Li, Chi Zhang, Wenjia Jiang, Wanqi Yang, Bin Fu, Pei Cheng, Xin Chen, Ling Chen, and Yunchao Wei. 2025. AppAgent v2: Advanced Agent for Flexible Mobile Interactions. *arXiv:2408.11824 [cs.HC]* <https://arxiv.org/abs/2408.11824>
- [15] Zhangheng Li, Keen You, Haotian Zhang, Di Feng, Harsh Agrawal, Xiujun Li, Mohana Prasad Sathya Moorthy, Jeff Nichols, Yinfei Yang, and Zhe Gan. 2024. Ferret-UI 2: Mastering universal user interface understanding across platforms. *arXiv:2410.18967 (2024)*.
- [16] Yijie Lin, Yuanbiao Gou, Zitao Liu, Boyun Li, Jiancheng Lv, and Xi Peng. 2021. COMPLETE: Incomplete Multi-view Clustering via Contrastive Prediction. In *2021 IEEE/CVF Conf on Computer Vision and Pattern Recognition (CVPR)*. 11169–11178. <https://doi.org/10.1109/CVPR46437.2021.01102>
- [17] Anthony Z. Liu, Jongwook Choi, Sungrull Sohn, Yao Fu, Jaekyeom Kim, Dongki Kim, Xinhe Wang, Jaewon Yoo, and Honglak Lee. 2024. SkillAct: Using Skill Abstractions Improves LLM Agents. In *Proc. of the 41st Int'l Conf on Machine Learning*.
- [18] Yitao Liu, Chenglei Si, Karthik R Narasimhan, and Shunyu Yao. 2025. Contextual Experience Replay for Self-Improvement of Language Agents. In *Proc. 63rd Annual Meeting of the Association for Computational Linguistics (Vol 1: Long Papers)*. Association for Computational Linguistics, Vienna, Austria, 14179–14198. <https://doi.org/10.18653/v1/2025.acl-long.694>
- [19] Yadong Lu, Jianwei Yang, Yelong Shen, and Ahmed Awadallah. 2024. Omniparser for pure vision based gui agent. *arXiv:2408.00203 (2024)*.
- [20] Runliang Niu, Jindong Li, Shiqi Wang, Yali Fu, Xiyu Hu, Xueyuan Leng, He Kong, Yi Chang, and Qi Wang. 2024. Screenagent: A vision language model-driven computer control agent. *arXiv:2402.07945 (2024)*.
- [21] OpenAI. 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>.
- [22] Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao Li, Yunxin Li, Shijue Huang, et al. 2025. UI-TARS: Pioneering Automated GUI Interaction with Native Agents. *arXiv:2501.12326 (2025)*.
- [23] Adrian Rebmann and Han van der Aa. 2024. Recognizing task-level events from user interaction data. *Information Systems* 124 (2024), 102404.
- [24] Andrei Cosmin Redis, Mohammadreza Fani Sani, Bahram Zarrin, and Andrea Burattin. 2024. Skill Learning Using Process Mining for Large Language Model Plan Generation. In *Int'l Conf on Process Mining*. Springer, 650–662.
- [25] Pascal J. Sager, Benjamin Meyer, Peng Yan, Rebekka von Wartburg-Kottler, Layan Etaawi, Aref Enayati, Gabriel Nobel, Ahmed Abdulkadir, Benjamin F. Grewe, and Thilo Stadelmann. 2025. AI Agents for Computer Use: A Review of Instruction-based Computer Control, GUI Automation, and Operator Assistants. *arXiv:2501.16150 [cs.AI]* <https://arxiv.org/abs/2501.16150>
- [26] Pascal J. Sager, Benjamin Meyer, Peng Yan, Rebekka von Wartburg-Kottler, Layan Etaawi, Aref Enayati, Gabriel Nobel, Ahmed Abdulkadir, Benjamin F. Grewe, and Thilo Stadelmann. 2025. A Comprehensive Survey of Agents for Computer Use: Foundations, Challenges, and Future Directions. *arXiv:2501.16150 [cs.AI]* <https://arxiv.org/abs/2501.16150>
- [27] ByteDance Seed. 2025. UI-TARS-1.5. <https://seed-tars.com/1.5>.
- [28] Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: language agents with verbal reinforcement learning. In *Proc. NeurIPS '23*. Curran Associates Inc., New Orleans, LA, USA, Article 377.
- [29] Yunpeng Song, Yiheng Bian, Yongtao Tang, and Zhongmin Cai. 2023. Navigating interfaces with AI for enhanced user interaction. *arXiv:2312.11190 (2023)*.
- [30] Yunpeng Song, Yiheng Bian, Yongtao Tang, Guiyu Ma, and Zhongmin Cai. 2024. Visiontasker: Mobile task automation using vision based ui understanding and llm task planning. In *Proc. ACM UIST '24*. 1–17.
- [31] Yifan Song, Weimin Xiong, Xiutian Zhao, Dawei Zhu, Wenhao Wu, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. AgentBank: Towards Generalized LLM Agents via Fine-Tuning on 50000+ Interaction Trajectories. *arXiv:2410.07706 (2024)*.
- [32] Fei Tang, Haolei Xu, Hang Zhang, Siqi Chen, Xingyu Wu, Yongliang Shen, Wenqi Zhang, Guiyang Hou, Zeqi Tan, Yuchen Yan, Kaitao Song, Jian Shao, Weiming Lu, Jun Xiao, and Yueting Zhuang. 2025. A Survey on (M)LLM-Based GUI Agents. *arXiv:2504.13865 [cs.HC]* <https://arxiv.org/abs/2504.13865>
- [33] Alok Mani Tripathi. 2018. *Learning Robotic Process Automation: Create Software robots and automate business processes with the leading RPA tool—UiPath*. Packt Publishing Ltd.
- [34] Bryan Wang and Yang Li. 2023. Enabling Conversational Interaction on Mobile with LLMs. <https://research.google/blog/enabling-conversational-interaction-on-mobile-with-llms/>. Google Research Blog.
- [35] Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Yian Wang, Qiushi Sun, Chengyou Jia, Kanzhi Cheng, Zichen Ding, Liheng Chen, Paul Pu Liang, and Yu Qiao. 2025. OS-ATLAS: Foundation Action Model for Generalist GUI Agents. In *Proc. ICLR '25*. <https://openreview.net/forum?id=n9PDaFNi8t>
- [36] Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. 2024. OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments. *arXiv:2404.07972 [cs.AI]* <https://arxiv.org/abs/2404.07972>
- [37] Yuhao Yang, Yue Wang, Dongxu Li, Ziyang Luo, Bei Chen, Chao Huang, and Junnan Li. 2024. Aria-ui: Visual grounding for gui instructions. *arXiv:2412.16256 (2024)*.
- [38] Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023. MM-REACT: Prompting ChatGPT for Multimodal Reasoning and Action. *arXiv:2303.11381 [cs.CV]* <https://arxiv.org/abs/2303.11381>
- [39] Yiwen Yin et al. 2025. Automatic Modeling Cognitive Dependencies from User Demonstrations for GUI Task Automation. In *Proc. ACM/SIGCHI (or appropriate venue)*.
- [40] Chaoyun Zhang, Shilin He, Jiaxu Qian, Bowen Li, Liqun Li, Si Qin, Yu Kang, Minghua Ma, Guyue Liu, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Zhang. 2025. Large Language Model-Brained GUI Agents: A Survey. *arXiv:2411.18279 [cs.AI]* <https://arxiv.org/abs/2411.18279>
- [41] Chaoyun Zhang, He Huang, Chiming Ni, Jian Mu, Si Qin, Shilin He, Lu Wang, Fangkai Yang, Pu Zhao, Chao Du, et al. 2025. Ufo2: The desktop agentos. *arXiv:2504.14603 (2025)*.
- [42] Chaoyun Zhang, Liqun Li, Shilin He, Xu Zhang, Bo Qiao, Si Qin, Minghua Ma, Yu Kang, Qingwei Lin, Saravan Rajmohan, Dongmei Zhang, and Qi Zhang. 2024. UFO: A UI-Focused Agent for Windows OS Interaction. *arXiv:2402.07939 [cs.HC]* <https://arxiv.org/abs/2402.07939>
- [43] S. Zhang et al. 2024. Dynamic Planning for LLM-based Graphical User Interface Automation. In *Findings of EMNLP (2024)*.
- [44] Shaoqing Zhang, Zhuosheng Zhang, Kehai Chen, Xinbei Ma, Muyun Yang, Tiejun Zhao, and Min Zhang. 2024. Dynamic planning for llm-based graphical user interface automation. *arXiv:2410.00467 (2024)*.
- [45] Shuyan Zhou, Frank F Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, et al. 2023. Webarena: A realistic web environment for building autonomous agents. *arXiv:2307.13854 (2023)*.