

From Static Templates to Dynamic Runtime Graphs: A Survey of Workflow Optimization for LLM Agents

Ling Yue¹, Kushal Raj Bhandari¹, Ching-Yun Ko², Dhaval Patel²,
Shuxin Lin², Nianjun Zhou², Jianxi Gao¹, Pin-Yu Chen^{2*}, Shaowu Pan^{1*}

¹Rensselaer Polytechnic Institute, ²IBM Research

Abstract

Large language model (LLM)-based systems are becoming increasingly popular for solving tasks by constructing executable workflows that interleave LLM calls, information retrieval, tool use, code execution, memory updates, and verification. This survey reviews recent methods for designing and optimizing such workflows, which we treat as agentic computation graphs (ACGs). We classify the literature according to when workflow structure is determined. This lens distinguishes static methods, which fix a reusable workflow scaffold before deployment, from dynamic methods, which select, generate, or revise the workflow for a particular run before or during execution. We further categorize existing work along three dimensions: when structure is determined, which aspect of the workflow is optimized, and which evaluation signals govern optimization (e.g., task metrics, verifier signals, human preferences, or trace-derived feedback). We also distinguish reusable workflow templates, run-specific realized graphs, and execution traces, separating design-time choices from runtime behavior. Finally, we outline a structure-aware evaluation perspective that complements downstream task metrics with graph-level properties, execution cost, robustness, and structural variation across inputs. This work aims to establish a principled vocabulary, a unified framework for situating new methods, a systematic comparison of existing work, and a minimum reporting protocol for LLM agent workflow optimization.

1 Introduction

Large language model (LLM) systems are increasingly being integrated into executable workflows that coordinate multiple actions over time. By workflow, we mean an executable organization of steps such as LLM calls, tool use, information retrieval, code execution, memory updates, and ver-

ification. A coding assistant, for example, may retrieve files, propose edits, run tests, and use verifier feedback to decide whether to revise or stop.

In multi-agent systems (MAS), these actions may be distributed across multiple specialized agents that communicate over a defined communication pattern, which specifies how agents are connected and how messages pass between them. What matters in practice is not only the quality of each model call, but also the overall workflow structure: which components or agents are present, how they depend on one another, and how information flows between them. A workflow can be naturally represented as a directed graph (Qian et al., 2025), where nodes correspond to components or agents and edges encode dependencies and information flow. This graph perspective enables reasoning about topology, communication density, scheduling, verification placement, and cost. These design choices can affect both effectiveness and efficiency (Zhang et al., 2025e; Zhou et al., 2026; Li et al., 2025). A weak scaffold may be partially compensated by stronger prompts, but deeper improvements require structural intervention: adding verifiers such as unit tests or schema checkers, pruning redundant communication, restructuring manager–worker hierarchies, or replacing a fixed pipeline with a context-aware one. However, improvements in agent capability often come with hidden structural costs, such as excessive depth, fragile control flow, and high communication overhead.

In this survey, we use the term *agentic computation graph* (ACG) as a unifying abstraction for executable LLM-centered workflows, including pipelines, orchestration graphs, communication graphs, plans, and code-defined agent systems. Our goal is not to impose new terminology for its own sake, but to make structure itself the primary object of comparison: how much of the workflow is fixed before deployment, how much is determined for a

*Corresponding authors: pin-yu.chen@ibm.com; pans2@rpi.edu.




Figure 1: An overview of workflow optimization for LLM agents

particular run, and how much can be revised during execution.

A growing body of work now treats workflow design as an optimization problem. Some works search for reusable templates offline (Zhang et al., 2025e; Hu et al., 2025a; Zhou et al., 2026); others optimize prompts, demonstrations, or collaboration behavior within a fixed scaffold (Khattab et al., 2024; Yang et al., 2024; Guo et al., 2024; Zehle et al., 2025; Agrawal et al., 2026; Chen et al., 2025); and a third group generates, selects, or edits the workflow used for a particular run before or during execution (Wang et al., 2026a; Ren et al., 2026; Bhandari et al., 2026; Li et al., 2026; Zhang et al., 2025d; Li et al., 2025; Gao et al., 2025a; Wang et al., 2026c,d). Existing surveys provide important foundations on agent workflows, planning, tool use, multi-agent collaboration, and broad agent optimization (Yu et al., 2025; Li et al., 2024a; Huang et al., 2024; Xu et al., 2025b; Chen et al., 2024b; Tran et al., 2025; Zhang et al., 2025f; V et al., 2026; Du et al., 2026; Yue et al., 2026a), but they typically treat workflow structure as a system ingredient rather than as a first-class optimization object. We therefore ask when structure is determined, what part of the workflow changes, and what evidence authorizes that change.

This survey studies workflow optimization for LLM-centered agentic systems, where workflow structure determines how LLM calls, retrieval, tool use, code execution, memory, and verification are composed. Figure 1 illustrates the scope and organization of the topics covered. We focus on *static workflow optimization*, which improves a reusable template offline, and *dynamic workflow optimization*, which determines part of the realized graph at inference time through subgraph selection,

pre-execution generation, or in-execution editing. Our evidence base contains 84 in-scope works: 46 CORE papers that directly optimize reusable templates or executable realized graphs, 7 ADJACENT papers on routing, team selection, and pruning, and 31 BACKGROUND resources that shape workflow optimization or evaluation. We include planning-only, tool-use-only, memory-only, or evaluation-only work only when it directly changes workflow structure or provides workflow-relevant evaluation assets. We also organize 27 workflow-relevant evaluation assets.

This survey makes the following contributions:

- We introduce a workflow-centered view of LLM agent systems as agentic computation graphs (ACGs) and distinguish reusable templates, run-specific realized graphs, and execution traces. This separation clarifies whether a method optimizes reusable design, per-run structure, or realized runtime behavior.
- We propose a taxonomy organized by when workflow structure is determined, refined by graph determination time (GDT) and graph plasticity mode (GPM) to compare fixed-template optimization, pre-execution generation or selection, and in-execution editing under a common lens. This taxonomy helps resolve gray cases such as offline-trained workflow generators and runtime subgraph selection.
- We synthesize the literature along three dimensions—optimization target (node, graph, or joint), evidence source (metric, verifier, preference, or trace), and update mechanism—to explain what part of the workflow each method changes, what evidence authorizes those changes, and how quality–cost trade-offs are managed.

- We organize workflow-relevant evaluation assets and propose a minimum reporting protocol that separates downstream task performance from graph properties, execution cost, robustness, and structural variation across inputs, making workflow evaluation more comparable and reproducible.

The remainder of this paper is organized as follows. Section 2 introduces the conceptual framework and the hierarchical taxonomy shown in Figure 2. Sections 3 and 4 review static and dynamic workflow optimization methods. Section 5 discusses evaluation and design trade-offs. Section 6 concludes.

2 Conceptual Framework and Taxonomy

This section introduces the distinctions that organize the rest of the survey: reusable templates versus realized graphs, node-level versus graph-level optimization, and static versus dynamic structure determination. The six method families organize the literature at a high level; GDT, GPM, optimization target, evidence, and update mechanism are used to compare individual methods. Appendix A gives additional examples and the full quality–cost formulation; Appendix Table 3 provides a compact reference for the main terms and notation.

2.1 Agentic computation graphs as executable workflows

An *agentic computation graph* (ACG) is our unifying abstraction for an executable LLM-centered workflow. Nodes perform atomic actions such as LLM calls, information retrieval, tool use, validation, or message passing. Edges encode control, data, or communication dependencies. Typical nodes cover computation, information access, state management, and coordination; in MAS, agents correspond to LLM-driven nodes with distinct prompts, tools, or models, and their messages appear as graph edges. Beyond nodes and edges, many workflows include a scheduler or router that decides which node executes next, which nodes can run in parallel, when to terminate, and whether replanning is allowed.

Different papers realize this abstraction through code, domain-specific languages, plain-text specifications, or explicit graph intermediate representations. Representation matters because it shapes what can be searched, verified, or edited. An ACG may include loops, conditionals, and dy-

namic communication. We include explicit or trace-reconstructable structure, but exclude emergent coordination with no observable workflow artifact from the comparison tables.

2.2 Template, realized graph, and trace

We distinguish three related but different objects throughout the survey: a reusable workflow template, the realized graph for a particular run, and the execution trace produced by running that graph.

Definition 1 (ACG template). An *ACG template* is a reusable executable specification $\bar{\mathcal{G}} = (\mathcal{V}, \mathcal{E}, \Phi, \Sigma, \mathcal{A})$, where $\mathcal{V}$ is the set of nodes, $\mathcal{E}$ is the set of directed edges, $\Phi = \{\phi_v\}_{v \in \mathcal{V}}$ contains node parameters such as prompts, tool schemas, model choices, or verifier settings, Σ is a scheduling or routing policy, and $\mathcal{A}$ is the set of admissible activation or edit actions.

The template is the reusable design object. It specifies the structural and parametric space available to the system before a concrete input is observed.

Definition 2 (Realized graph). Given an input x and the evolving runtime state, a *realized graph* $\mathcal{G}^{\text{run}}$ is the workflow structure actually used for a particular run. It may coincide with the reusable template, or it may be obtained by selecting a sub-graph, instantiating optional nodes, or applying allowed edits before or during execution.

The realized graph is run-specific. It captures the structure that is actually deployed for one run, rather than the structure merely available in principle. Every executed workflow has a realized graph; a method is dynamic not because such a graph exists, but because some of its structure is chosen or revised at inference time rather than being fully fixed in the reusable template.

Definition 3 (Execution trace). Executing $\mathcal{G}^{\text{run}}$ yields an *execution trace* $\tau = \{(s_t, a_t, o_t, c_t)\}_{t=1}^T$, where s_t is the system or environment state, a_t is the action taken, o_t is the resulting observation, and c_t is execution cost such as token usage, tool calls, latency, or monetary expense.

These definitions are intentionally lightweight. The template captures what a method makes reusable. The realized graph captures what structure is actually deployed for a particular run. The trace records what happened during execution, including tool failures, retries, verifier outputs, and cost accumulation. Many recent papers differ pre-




Figure 2: Title-only taxonomy of workflow-optimization methods for LLM agents. Methods are first divided into static and dynamic optimization, and then grouped into six method families

cisely in which of these three objects is optimized, generated, or reused.

For a coding agent, retrieve–patch–test is a reusable template; issue-specific dependencies form the realized graph; adding a debugger after a failed test is an in-execution edit. A pre-authored retry changes only the trace.

2.3 Optimization targets

Across a wide range of methods, workflow optimization balances task quality against execution cost. The most important distinction for this survey is what part of the workflow changes. In *node-level optimization*, the high-level scaffold is fixed and local parameters in Φ —such as prompts, tools, models, or verifier policies—are improved. In *graph-level optimization*, structural variables such as $\mathcal{E}$, Σ , and $\mathcal{A}$ are updated, changing topology, communication structure, scheduling, branching logic, or the admissible edit space. In *joint optimization*, both are updated together, either simultaneously or in alternating stages.

2.4 When is workflow structure determined?

The static–dynamic distinction.

Our main organizing principle is *when workflow structure is determined*. We distinguish between static structure determination and dynamic structure determination, and we use two lightweight descriptors to clarify gray cases.

Static structure determination. A method is *static* if its deployed structure is a reusable template whose structural degrees of freedom are fixed after

training or search. The template can still contain conditional execution, loops, or stochasticity, but those behaviors are already encoded in the reusable scaffold.

Dynamic structure determination. A method is *dynamic* if some part of the realized graph is constructed, selected, or edited at inference time. This may happen once before execution, or repeatedly during execution in response to observations, failures, or verifier signals. Adding a new verifier, spawning a new agent, or rewiring communication only after observing execution feedback are dynamic changes.

To compare gray cases, we use two lightweight descriptors. *Graph determination time* (GDT) records when the realized structure is decided: offline if the reusable template is optimized before deployment, pre-execution if a run-specific graph is generated once before execution, and in-execution if structure is revised during execution. *Graph plasticity mode* (GPM) records how structure may vary at inference time: none if the structure is fixed, select if the method activates or prunes parts of a fixed super-graph, generate if it constructs a run-specific graph before execution, and edit if it adds, removes, rewires, or rewrites structure during execution.

These descriptors are deliberately lightweight. Their purpose is to clarify common ambiguities rather than to impose a rigid ontology. For example, a workflow generator trained offline but used to emit a new plan for each input is dynamic in our taxonomy, because the realized graph is determined pre-execution at inference time. AutoAgents is a

hybrid: it drafts before execution and revises from observer feedback (Chen et al., 2024a). Corpus counts use one primary deployed mechanism, while the comparison card retains hybrid labels.

2.5 Comparison card

Appendix Table 5 summarizes surveyed methods using a unified comparison card. Each method is coded by its *setting* (static or dynamic, with GDT and GPM), *optimized level* (node, graph, or joint), *feedback/evidence* (metric, verifier, supervision, preference, trace, or proxy), and *update mechanism* (search, generator, controller, supervised learning, RL, preference optimization, repair/edit, continuous relaxation, or hybrid); Appendix Table 4 gives the full field definitions and coding rules.

3 Static Optimization of Agent Workflows

Static methods optimize a reusable template or a fixed collaboration scaffold before deployment. Their practical appeal is clear: they are easier to inspect, constrain, ablate, and benchmark under stable budgets. The main limitation is equally clear: once a template is frozen, distribution shift, tool drift, or unanticipated branching can expose brittle structural assumptions. Appendix Table 5 summarizes representative CORE static and dynamic methods using a unified comparison card.

3.1 Offline template search over constrained design spaces

A central line of work treats the workflow as a discrete design object and searches for a reusable template by repeated execution and evaluation. AFlow (Zhang et al., 2025e) searches over typed operator graphs using Monte Carlo Tree Search (MCTS), combining LLM-guided expansion with executable evaluation and explicit dollar cost. Automated Design of Agentic Systems (ADAS) (Hu et al., 2025a) instead searches in code space: a meta-agent proposes runnable agentic systems, evaluates them, archives strong designs, and iteratively improves them. Evolutionary Generation of Multi-Agent Systems (Hu et al., 2026) adopts a more classical population-based view and treats multi-agent system design as an evolvable genotype of roles, topology, and protocol. In a domain-specific setting, VFlow (Wei et al., 2026) combines cooperative evolution and Monte Carlo Tree Search with strong hardware verifiers to discover Verilog-generation workflows under functional and resource objectives.

Archon (Saad-Falcon et al., 2025) searches reusable inference architectures and model assignments over generation, critique, ranking, and fusion modules.

These methods differ in representation, but they share three assumptions. First, there must be an executable search space, whether defined by typed operators, code templates, or structured workflow languages. Second, evaluation must be reliable enough to discriminate candidates. Third, the search space must embody a useful inductive bias: if candidate workflows are mostly invalid or semantically incoherent, black-box search quickly becomes prohibitively expensive. This is precisely why typed operators, code scaffolds, and constrained graph languages are so important in practice.

A notable variation is to learn the operator library itself instead of assuming it is fixed. A²Flow (Zhao et al., 2026) extracts abstract operators from demonstrations, clusters them into reusable patterns, and then refines them for later workflow search. This reduces dependence on hand-designed primitives and makes explicit a design choice that many search papers leave implicit: the quality of the operator vocabulary often matters as much as the search algorithm. SEW (Liu et al., 2025a) sits at the opposite end of the representation spectrum. It encodes candidate workflows directly as text and iteratively self-evolves them using benchmark pass rates in code-generation settings, showing that reusable workflow optimization need not always rely on a typed graph IR if the task feedback is strong enough. More generally, this line of work suggests a recurring theme of the survey: search quality is often limited less by the nominal optimizer than by the representation and evaluator it is allowed to use.

3.2 Node-level optimization inside fixed scaffolds

When full graph search is too expensive or the scaffold is already adequate, a natural fallback is to optimize node-level parameters inside a fixed structure. DSPy (Khatab et al., 2024) is the clearest programmatic example: it treats an LLM pipeline as a composition of typed modules and compiles prompts or demonstrations that optimize a user-specified metric on examples. The underlying graph is fixed, but the local node policies are no longer hand-written strings.

Black-box prompt optimizers such as Large Lan-

language Models as Optimizers (OPRO) (Yang et al., 2024), EvoPrompt (Guo et al., 2024), CAPO (Zehle et al., 2025), and GEPA (Agrawal et al., 2026) fit naturally into this picture when prompts are attached to workflow nodes rather than viewed as isolated artifacts. OPRO (Yang et al., 2024) uses an LLM as the optimizer over a textual history of candidates and scores. EvoPrompt (Guo et al., 2024) turns prompt search into evolutionary optimization. CAPO (Zehle et al., 2025) makes the same family explicitly cost-aware through racing and prompt-length penalties. GEPA (Agrawal et al., 2026) uses trajectory-level reflection and Pareto-style retention of prompt edits, showing that textual feedback can outperform pure scalar-reward optimization when failures are semantically structured. Because these methods optimize workflow-local instructions rather than global topology, we treat them as node-level optimizers; Appendix Table 9 summarizes them in a compact comparison matrix.

The value of node-level optimization is practical: it usually requires fewer assumptions than graph search and can yield large gains quickly. The main limitation is interpretive rather than purely empirical. Better prompts can compensate for weak topology, making a poor scaffold look competitive while leaving the real structural bottleneck unaddressed (Zhou et al., 2026; Wang et al., 2025a). For example, on HotpotQA, Maestro’s configuration-only optimization (prompts, tools, and control knobs) raises an initial scaffold from 38% to 70.33%, within two points of the 72.00% reached once the agent graph is jointly optimized with the configuration; the remaining gap is attributed to structural failure modes that prompt tuning alone cannot fix. Static node tuning is therefore often a necessary baseline, but rarely a sufficient explanation for why a workflow works.

A closely related but more training-oriented line optimizes collaboration behavior within a largely fixed multi-agent scaffold. Optima (Chen et al., 2025) generates, ranks, selects, and trains on multi-agent interaction trajectories to improve the effectiveness–efficiency trade-off of collaboration itself. In our taxonomy, this remains static at deployment time if the collaboration scaffold is fixed, even though the underlying models are updated.

3.3 Joint optimization of structure and local configuration

Structure and local configuration are coupled. Changing the topology affects what information

each node receives. Changing prompts or tool settings changes what the topology effectively computes. The most direct attack on this coupling is staged joint optimization. The MASS framework (Zhou et al., 2026) alternates local prompt optimization, topology optimization, and workflow-level prompt optimization under the revised topology. Its practical significance is not only that joint optimization helps, but also that a staged schedule can make it stable enough to use in practice.

GPTSwarm (Zhuge et al., 2024) jointly optimizes node prompts and graph connectivity. Symbolic learning (Ou et al., 2025) uses language-based loss and symbolic backpropagation to update prompts, tools, and workflow structure.

A complementary route relaxes the discrete communication graph into a continuous parameterization and jointly optimizes graph structure and parameters through gradient-based learning. Hu et al. (2024) demonstrate this approach outside the LLM setting, but the core question transfers directly: communication topology should not be assumed fixed if it materially affects cooperative performance. Maestro (Wang et al., 2025a) offers the clearest recent statement of this joint view within the LLM context: it alternates graph edits with node-level configuration updates and uses both numeric scores and reflective textual feedback to drive the optimization loop. What these papers share is the recognition that prompt tuning alone cannot supply missing structural capabilities such as validation, conditional routing, or intermediate decomposition.

3.4 Verifiability in static workflow optimization

Static optimization is particularly compatible with verification because reusable templates can be checked before deployment. This has motivated growing interest in constrained intermediate representations and verifier-driven search. Mermaid-Flow (Zheng et al., 2025) makes the search space itself more reliable by encoding workflows in a structured, human-readable Mermaid IR, together with static validity checks and safety-constrained evolutionary programming. VFlow (Wei et al., 2026) integrates stronger external verification directly into the search loop, using syntax, functional correctness, synthesizability, and hardware objectives rather than relying on downstream success alone. In both cases, verification is not added after search; it is part of the optimization process itself.

More broadly, static workflow optimization is most effective when the search space is constrained enough and candidate evaluation is trustworthy enough that offline search can meaningfully separate good structure from bad. When those conditions no longer hold, the central question shifts from how to optimize a reusable template to how much structure should remain plastic at inference time. That shift motivates the dynamic methods reviewed next.

4 Dynamic Optimization and Runtime Adaptation

Dynamic methods determine some part of the workflow at inference time. They are motivated by heterogeneity: different queries require different tools, different reasoning depth, different numbers of agents, and different verification regimes. Their appeal is flexibility, but that flexibility also enlarges the action space, complicates credit assignment, and increases the need for explicit budget guards. Within this family, it is useful to distinguish three increasingly flexible settings: runtime selection over a fixed super-graph, pre-execution workflow generation, and in-execution editing. Appendix Table 5 summarizes representative CORE dynamic methods using the same comparison card.

4.1 Selection and pruning as the lightest form of runtime adaptation

The lightest form of dynamism keeps a super-graph fixed and makes run-specific activation or pruning decisions. Adaptive Graph Pruning (Li et al., 2025) learns to prune both edges and agents from a complete communication graph using task and agent embeddings, producing task-adaptive sparse topologies. DAGP (Wang and Tong, 2025) further conditions pruning on estimated realized graph difficulty so that easy examples use leaner communication while hard examples retain richer collaboration. AgentDropout (Wang et al., 2025f) similarly removes redundant agents and communication links across rounds to improve the accuracy–cost trade-off. DyLAN (Liu et al., 2024), MasRouter (Yue et al., 2025), and SkillOrchestra (Wang et al., 2026b) occupy a similar design point from a routing perspective: they select teams, collaboration modes, models, or tools so that less structure is activated unless the task demands it. Appendix Table 10 summarizes this ADJACENT family.

AgentNet (Yang et al., 2025) combines within-

task decentralized routing with across-task edge updates and pruning based on completed trajectories.

This design point has two advantages. First, validity is inherited from the super-graph, which makes optimization easier than unconstrained generation. Second, selective activation often captures a large fraction of the cost savings available from more ambitious dynamic methods. The limitation is expressivity. A pruning policy cannot introduce a missing verifier, a missing tool, or a new decomposition strategy if the super-graph never contained those options. Selection and pruning offer an attractive first layer of dynamism when most of the relevant structure is known in advance and only its activation should vary across inputs.

4.2 Construct-then-execute: pre-execution workflow generation

Difficulty-aware routing and learned graph generators. A stronger form of dynamism generates or selects a run-specific workflow before execution begins. The lightest variant estimates query difficulty and allocates workflow depth, operator choice, and model routing before execution starts, without requiring full graph synthesis from scratch (Su et al., 2025a). A more expressive alternative autoregressively samples roles and edges to produce a query-conditioned DAG rather than modifying a reusable template (Li et al., 2026). G-Designer (Zhang et al., 2025d) learns a graph generator with a variational graph autoencoder, while Dynamic Generation of Multi-LLM Agents Communication Topologies with Graph Diffusion Models (Jiang et al., 2025) synthesizes sparse communication topologies through proxy-guided discrete diffusion and multi-objective reward guidance. Multi-Agent Architecture Search via Agentic Supernet (MaAS) (Zhang et al., 2025b) learns a query-conditioned distribution over architectures through an agentic supernet. These methods differ in parameterization, but they share the same goal: move from one-size-fits-all collaboration patterns to structures that reflect task complexity and cost sensitivity for a particular run.

Preference- and reinforcement-learning-based generators. Preference and reinforcement learning offer two other routes to pre-execution generation. ScoreFlow (Wang et al., 2025e) executes sampled workflows, converts numeric scores into score-aware preferences, and trains a workflow generator with a direct preference optimization (DPO)-like

Table 1: Minimum reporting protocol for workflow-optimization studies

Dimension	Minimum information
Structure	representation; static/dynamic setting; GDT/GPM; admissible edits; routing and stopping rules
Models and tools	base models; decoding; tools; verifier placement; memory policy
Cost	offline budget; online tokens, model/tool calls, latency, dollars, and cost per success
Graph and trace	nodes, depth, communication, retries, edits, fallbacks, termination, and structural variance
Robustness	paraphrases, retrieval noise, tool failures or drift, unseen tools, and budget caps
Protocol	splits, seeds, repeated runs, canonicalization, structural ablations, and representative failures

Table 2: Best-suited conditions and principal limitations of the six method families

Family	Best suited for	Main limitation
Offline template search	stable repeated tasks; strong evaluators	search cost; no instance adaptation
Fixed-scaffold node optimization	adequate scaffold; local errors	cannot repair missing structure
Joint structure/configuration	coupled topology and node policies	large search space; unclear credit
Selection, routing, or pruning	known motifs; varying difficulty or budget	limited to the fixed super-graph
Pre-execution generation	tasks requiring different decompositions	validity risk and generation overhead
In-execution editing	runtime evidence changes needed structure	hardest to constrain and reproduce

objective that preserves the magnitude information in the scores. FlowReasoner (Gao et al., 2025a) trains a query-level meta-agent with reinforcement learning so that a new workflow is generated for each query from an operator library. Workflow-R1 (Kong et al., 2026) reframes workflow construction as a multi-turn decision process and aligns reinforcement learning with grouped think-act subsequences rather than with a single monolithic generation step. These papers collectively show that once workflow generation is treated as a policy problem, score design and action granularity become central design choices.

Data-centric workflow generation. Data-centric workflow generation provides a complementary approach. CapFlow (Wang et al., 2026a) learns reusable workflow-capability bases and sparsely recomposes them for each query, producing a task-specific workflow in one pass while using counterfactual attribution to identify capabilities that drive success. AutoFlow (Li et al., 2024b) uses a lightweight workflow DSL and an interpreter model to generate and score workflows from natural-language requirements. WorkflowLLM (Fan et al., 2025) instead constructs a large workflow corpus from real automation artifacts and fine-tunes a model to emit workflow code directly. RobustFlow (Xu et al., 2025a) adds a robustness objective so that semantically equivalent queries map to more consistent structures. In a more tightly constrained setting, ComfyUI-R1 (Xu et al., 2025c) shows that reasoning-oriented workflow generation can be highly effective when the schema is strict and executability is easy to validate. Although these models are trained offline, they are dynamic in our taxonomy because

they produce new realized graphs at inference time. FlowEvo (Ren et al., 2026) additionally compiles successful workflows into reusable skills for later retrieval or generation.

Hybrid: generate then revise. A hybrid point in this design space drafts both team composition and execution plan before execution, then uses observer agents and execution-time refinement to revise them online, blurring the boundary between construct-then-execute and in-execution editing (Chen et al., 2024a). Taken together, these methods move structural choice to just before execution: it is more expressive than subgraph selection, yet typically easier to validate than fully interleaved runtime editing.

4.3 In-execution editing: interleaving execution with structural change

Designer-executor interleaving and on-demand agent creation. The most flexible dynamic methods treat structural change as a first-class runtime action. DyFlow (Wang et al., 2025c) interleaves a designer with an executor, using intermediate feedback to revise sub-goals and dynamically choose operators as execution unfolds. AgentConductor (Wang et al., 2026c) generates a YAML topology, executes it, then regenerates the topology for the same problem based on validity, code-execution, and cost feedback until success or budget exhaustion. Aime (Shi et al., 2025) uses a dynamic planner, an actor factory, and a progress manager so that specialized agents can be instantiated and coordinated on demand during long-horizon execution. AOrchestra (Ruan et al., 2026) turns sub-agent creation itself into a callable action, making dynamic specialization part of the orchestration policy rather than a preprocessing step.

Training-free in-execution evolution. Other methods focus on training-free evolution during inference. MetaGen (Wang et al., 2026d) initializes a task-adaptive graph and then revises both roles and edges based on lightweight feedback such as contradictions, failures, and cost signals. ProAgent (Ye et al., 2023), framed as agentic process automation, constructs workflows in structured JSON and repairs them incrementally through testing-on-constructing. Flow (Niu et al., 2025) revises an activity-on-vertex graph during execution as new intermediate results arrive. QualityFlow (Hu et al., 2025b) is more conservative: it keeps much of the scaffold fixed but lets a quality checker decide whether to submit, clarify, revert, or continue, thereby making control-flow edits conditional on verifier-like judgments.

Online search and debate. A final cluster uses online search or debate. EvoFlux (Bhandari et al., 2026) treats compact-model tool use as execution-grounded repair of typed workflow graphs, evolving them through structured edits, adaptive search intensity, meta-guided redesign, and diversity pruning. EvoFlow (Zhang et al., 2025a) maintains a diverse population of heterogeneous workflows and evolves them on the fly. DebFlow (Su et al., 2025b) uses multi-agent debate and reflection to improve workflow creation under limited resources. These methods are especially useful when multiple structurally distinct solutions may be reasonable and when diversity helps hedge against search myopia.

Taken together, dynamic methods form a spectrum of structural plasticity, from lightweight subgraph selection to pre-execution generation and full in-execution editing. The relevant design question is not whether more flexibility is intrinsically better, but what minimum plasticity the task actually requires. Moving rightward along this spectrum increases expressivity, but it also raises the burden of validation, stopping, and budget control. Appendix H returns to this trade-off from a design perspective.

5 Evaluation and Design Trade-offs

Workflow structure should be evaluated as an output. Table 1 covers representation and adaptivity, models and tools, offline and online cost, graph and trace statistics, robustness, repeated runs, and failures; Appendix G details suitable metrics and evaluation assets. Evidence should be described as within-study, cross-study, preliminary, or con-

ceptual, since raw gains under different backbones, tools, budgets, and benchmarks are not directly comparable.

The practical rule is to use the least structural plasticity the task requires: reuse a template for stable workloads, select from known motifs when possible, generate for different decompositions, and edit during execution only when observations require structural change. Table 2 compares suitability and limitations. Key open problems are structural credit assignment, verifiability, drift, and when plasticity justifies its cost; Appendix H provides detail.

6 Conclusion

This survey proposes agentic computation graphs (ACGs) as a unifying abstraction for LLM-centered workflows and uses this lens to organize and compare methods for optimizing them. The fundamental taxonomy distinguishes static approaches, which commit to reusable templates prior to deployment, from dynamic methods that instantiate, select, or modify workflow structures conditioned on the current input or runtime feedback. This perspective unifies disparate techniques, such as prompt compilation, topology pruning, workflow generation, and runtime editing, within a coherent quality–cost analytical framework.

We argue that workflow optimization is best understood through three linked objects: the reusable template, the run-specific realized graph, and the execution trace. We further refine the static–dynamic distinction using graph determination time and graph plasticity mode, which helps organize gray cases such as offline-trained workflow generators and runtime subgraph selection. In the literature, metric-based search, verifier-guided optimization, preference learning, and trace-derived textual feedback can all be viewed as ways to learn better local policies, better structures, or both.

Under this view, workflow structure should be elevated from an implementation detail to a first-class design object, recognizing its direct impact on system capability, operational cost, reliability, and scientific comparability. This survey provides the conceptual foundations and analytical apparatus necessary to rigorously analyze, optimize, and evaluate these structures.

Limitations

Literature snapshot. This survey reflects a literature snapshot taken at the time of writing, and the field of LLM-agent workflow optimization is moving rapidly. New preprints, framework releases, and system versions are likely to appear after our cutoff. Our taxonomy and synthesis should therefore be read as a unifying framework for situating current methods, not as an exhaustive or up-to-the-minute index of the literature.

No re-execution of surveyed methods. We did not re-execute any of the surveyed systems. Reported results across papers use different backbones, budgets, benchmarks, verifier configurations, and evaluation protocols, and many papers also provide incomplete details about costs, prompts, or tool access. Raw numbers are therefore not directly comparable, and we avoid strong cross-paper performance claims unless they are directly supported by the original studies.

Acknowledgment

This work received funding support from IBM-Rensselaer Future of Computing Research Collaboration (2026). Shaowu Pan is supported by the Google Research Scholar Program. AI assistants were used for language polishing. All scientific claims, taxonomy decisions, literature screening, and citations were independently verified by the authors.

References

- Lakshya A Agrawal, Shangyin Tan, Dilara Soylu, Noah Ziemis, Rishi Khare, Krista Opsahl-Ong, Arnav Singhvi, Herumb Shandilya, Michael J Ryan, Meng Jiang, Christopher Potts, Koushik Sen, Alex Dimakis, Ion Stoica, Dan Klein, Matei Zaharia, and Omar Khattab. 2026. [GEPA: Reflective prompt evolution can outperform reinforcement learning](#). In *The Fourteenth International Conference on Learning Representations*.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. 2021. [Program synthesis with large language models](#). *arXiv preprint arXiv:2108.07732*.
- Victor Barres, Honghua Dong, Soham Ray, Xujie Si, and Karthik Narasimhan. 2025. [\$\tau^2\$ -Bench: Evaluating Conversational Agents in a Dual-Control Environment](#). *arXiv preprint arXiv:2506.07982*.
- Kushal Raj Bhandari, Ling Yue, Ching-Yun Ko, Dhaval Patel, Shaowu Pan, Pin-Yu Chen, and Jianxi Gao. 2026. [Evoflux: Inference-time evolution of executable tool workflows for compact agents](#). *arXiv preprint arXiv:2606.12674*.
- Guangyao Chen, Siwei Dong, Yu Shu, Ge Zhang, Jaward Sesay, Börje Karlsson, Jie Fu, and Yemin Shi. 2024a. [Autoagents: A framework for automatic agent generation](#). In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 22–30.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#).
- Shuaihang Chen, Yuanxing Liu, Wei Han, Weinan Zhang, and Ting Liu. 2024b. [A survey on LLM-Based multi-agent systems: Recent advances and new frontiers in application](#). *arXiv preprint arXiv:2412.17481*.
- Weize Chen, Jiarui Yuan, Chen Qian, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2025. [Optima: Optimizing effectiveness and efficiency for LLM-Based multi-agent system](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 11534–11557, Vienna, Austria. Association for Computational Linguistics.
- Zehui Chen, Weihua Du, Wenwei Zhang, Kuikun Liu, Jiangning Liu, Miao Zheng, Jingming Zhuo, Songyang Zhang, Dahua Lin, Kai Chen, and Feng Zhao. 2024c. [T-Eval: Evaluating the tool utilization capability of large language models step by step](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9510–9529, Bangkok, Thailand. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *arXiv preprint arXiv:2110.14168*.
- Shimin Di, Xujie Yuan, Hanghui Guo, Chaoqian Ouyang, Yongxu Liu, Ling Yue, Zhangze Chen, Libin Zheng, Jia Zhu, Shaowu Pan, Jian Yin, Yong Rui, and Min-Ling Zhang. 2026. [ToolRosella: Translating code repositories into standardized tools for scientific agents](#). *arXiv preprint arXiv:2603.09290*.
- Shangheng Du, Jiabao Zhao, Jinxin Shi, Zhentao Xie, Xin Jiang, Yanhong Bai, and Liang He. 2026. [A survey on the optimization of large language model-based agents](#). *ACM Computing Surveys*, 58(9):1–37.

- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2368–2378.
- Shengda Fan, Xin Cong, Yuepeng Fu, Zhong Zhang, Shuyan Zhang, Yuanwei Liu, Yesai Wu, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2025. [WorkflowLLM: Enhancing workflow orchestration capability of large language models](#). In *The Thirteenth International Conference on Learning Representations*.
- Hongcheng Gao, Yue Liu, Yufei He, Longxu Dou, Chao Du, Zhijie Deng, Bryan Hooi, Min Lin, and Tianyu Pang. 2025a. [FlowReasoner: Reinforcing query-level meta-agents](#). *arXiv preprint arXiv:2504.15257*.
- Xuanqi Gao, Siyi Xie, Juan Zhai, Shiqing Ma, and Chao Shen. 2025b. [MCP-RADAR: A multi-dimensional benchmark for evaluating tool use capabilities in large language models](#). *arXiv preprint arXiv:2505.16700*.
- Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. 2024. [Connecting large language models with evolutionary algorithms yields powerful prompt optimizers](#). In *The Twelfth International Conference on Learning Representations*.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Shengchao Hu, Li Shen, Ya Zhang, and Dacheng Tao. 2024. [Learning multi-agent communication from graph modeling perspective](#). In *The Twelfth International Conference on Learning Representations*.
- Shengran Hu, Cong Lu, and Jeff Clune. 2025a. [Automated design of agentic systems](#). In *The Thirteenth International Conference on Learning Representations*.
- Yaojie Hu, Qiang Zhou, Qihong Chen, Xiaopeng Li, Linbo Liu, Dejiao Zhang, Amit Kachroo, Talha Oz, and Omer Tripp. 2025b. [QualityFlow: An agentic workflow for program synthesis controlled by LLM quality checks](#). *arXiv preprint arXiv:2501.17167*.
- Yuntong Hu, Yuting Zhang, Matthew Trager, Yi Zhang, Shuo Yang, Wei Xia, and Stefano Soatto. 2026. [EvoMAS: Evolutionary generation of multi-agent systems](#). *arXiv preprint arXiv:2602.06511*.
- Xu Huang, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang, Ruiming Tang, and Enhong Chen. 2024. [Understanding the planning of LLM agents: A survey](#). *arXiv preprint arXiv:2402.02716*.
- Zhenhan Huang, Tejaswini Pedapati, Pin-Yu Chen, Chunheng Jiang, and Jianxi Gao. 2026. [Unveil fundamental graph properties for neural architecture search](#). *Advanced Science*, 13(25):e16574.
- Chunheng Jiang, Zhenhan Huang, Tejaswini Pedapati, Pin-Yu Chen, Yizhou Sun, and Jianxi Gao. 2024. [Network properties determine neural network performance](#). *Nature Communications*, 15(1):5718.
- Eric Hanchen Jiang, Mengting Li, Guancheng Wan, Sophia Yin, Yuchen Wu, Xiao Liang, Xinfeng Li, Yizhou Sun, Wei Wang, Kai-Wei Chang, and Ying Nian Wu. 2025. [Dynamic generation of multi-LLM agents communication topologies with graph diffusion models](#). *arXiv preprint arXiv:2510.07799*.
- Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. 2024. [SWE-bench: Can language models resolve real-world GitHub issues?](#) In *The Twelfth International Conference on Learning Representations*.
- Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan A, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2024. [DSPy: Compiling declarative language model calls into state-of-the-art pipelines](#). In *The Twelfth International Conference on Learning Representations*.
- Mingze Kong, Zikun Qu, Zhongquan Zhou, Pengyu Liang, Xiang Li, Zhiwei Shang, Zhi Hong, Kaiyu Huang, Zhiyong Wang, and Zhongxiang Dai. 2026. [Workflow-R1: Group sub-sequence policy optimization for multi-turn workflow construction](#). *arXiv preprint arXiv:2602.01202*.
- Boyi Li, Zhonghan Zhao, Der-Horng Lee, and Gaoang Wang. 2025. [Adaptive graph pruning for multi-agent communication](#). *arXiv preprint arXiv:2506.02951*.
- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023a. [CAMEL: Communicative agents for "mind" exploration of large language model society](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023b. [API-Bank: A comprehensive benchmark for tool-augmented LLMs](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3102–3116, Singapore. Association for Computational Linguistics.

- Shiyuan Li, Yixin Liu, Qingsong Wen, Chengqi Zhang, and Shirui Pan. 2026. [Assemble your crew: Automatic multi-agent communication topology design via autoregressive graph generation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 23142–23150.
- Xinyi Li, Sai Wang, Siqi Zeng, Yu Wu, and Yi Yang. 2024a. [A survey on LLM-Based multi-agent systems: Workflow, infrastructure, and challenges](#). *Vicnearth*, 1(1):9.
- Zelong Li, Shuyuan Xu, Kai Mei, Wenyue Hua, Balaji Rama, Om Raheja, Hao Wang, He Zhu, and Yongfeng Zhang. 2024b. [AutoFlow: Automated workflow generation for large language model agents](#). *arXiv preprint arXiv:2407.12821*.
- Siwei Liu, Jinyuan Fang, Han Zhou, Yingxu Wang, and Zaiqiao Meng. 2025a. [SEW: Self-evolving agentic workflows for automated code generation](#). *arXiv preprint arXiv:2505.18646*.
- Zhiwei Liu, Jieli Qiu, Shiyu Wang, Jianguo Zhang, Zuxin Liu, Roshan Ram, Haolin Chen, Weiran Yao, Shelby Heinecke, Silvio Savarese, Huan Wang, and Caiming Xiong. 2025b. [MCPEval: Automatic MCP-based deep evaluation for AI agent models](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 373–402, Suzhou, China. Association for Computational Linguistics.
- Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2024. [A dynamic LLM-powered agent network for task-oriented agent collaboration](#). In *First Conference on Language Modeling*.
- Ziyang Luo, Zhiqi Shen, Wenzhuo Yang, Zirui Zhao, Prathyusha Jwalapuram, Amrita Saha, Doyen Sahoo, Silvio Savarese, Caiming Xiong, and Junnan Li. 2025. [MCP-Universe: Benchmarking large language models with real-world model context protocol servers](#). *arXiv preprint arXiv:2508.14704*.
- Mike A Merrill, Alexander Glenn Shaw, Nicholas Carlini, Boxuan Li, Harsh Raj, Ivan Bercovich, Lin Shi, Jeong Yeon Shin, Thomas Walshe, E. Kelly Buchanan, Junhong Shen, Guanghao Ye, Haowei Lin, Jason Poulos, Maoyu Wang, Jenia Jitsev, Marianna Nezhurina, Di Lu, Orfeas Menis Mastromichalakis, and 65 others. 2026. [Terminal-bench: Benchmarking agents on hard, realistic tasks in command line interfaces](#). In *The Fourteenth International Conference on Learning Representations*.
- Grégoire Mialon, Clémentine Fourier, Thomas Wolf, Yann LeCun, and Thomas Scialom. 2024. [GAIA: A benchmark for general AI assistants](#). In *The Twelfth International Conference on Learning Representations*.
- Guozhao Mo, Wenliang Zhong, Jiawei Chen, Qianhao Yuan, Xuanang Chen, Yaojie Lu, Hongyu Lin, Ben He, Xianpei Han, and Le Sun. 2025. [LiveMCPBench: Can agents navigate an ocean of MCP tools?](#) *arXiv preprint arXiv:2508.01780*.
- Subhrangshu Nandi, Arghya Datta, Rohith Nama, Udit Patel, Nikhil Vichare, Indranil Bhattacharya, Prince Grover, Shivam Asija, Giuseppe Carenini, Wei Zhang, Arushi Gupta, Sreyoshi Bhaduri, Jing Xu, Huzefa Raja, Shayan Ray, Aaron Chan, Esther Xu Fei, Gaoyuan Du, Zuhaib Akhtar, and 5 others. 2025. [SOP-Bench: Complex industrial SOPs for evaluating LLM agents](#). *arXiv preprint arXiv:2506.08119*.
- Boye Niu, Yiliao Song, Kai Lian, Yifan Shen, Yu Yao, Kun Zhang, and Tongliang Liu. 2025. [Flow: Modularized agentic workflow automation](#). In *The Thirteenth International Conference on Learning Representations*.
- Yixin Ou, Wangchunshu Zhou, Shengwei Ding, Long Li, Jialong Wu, Tiannan Wang, Jiamin Chen, Shuai Wang, Xiaohua Xu, Ningyu Zhang, Huajun Chen, and Yuchen Eleanor Jiang. 2025. [Symbolic learning enables self-evolving agents](#). *AI Open*, 6:314–322.
- Chaoqian Ouyang, Ling Yue, Shimin Di, Libin Zheng, Linan Yue, Shaowu Pan, Jian Yin, and Min-Ling Zhang. 2026. [Code2mcp: Transforming code repositories into mcp services](#). In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, KDD '26, page 3835–3846, New York, NY, USA. Association for Computing Machinery.
- Dhaval Patel, Shuxin Lin, James Rayfield, Nianjun Zhou, Roman Vaculin, Natalia Martinez, Fearghal O'Donncha, and Jayant Kalagnanam. 2025. [AssetOpsBench: Benchmarking AI agents for task automation in industrial asset operations and maintenance](#). *arXiv preprint arXiv:2506.03828*.
- Chen Qian, Zihao Xie, YiFei Wang, Wei Liu, Kunlun Zhu, Hanchen Xia, Yufan Dang, Zhuoyun Du, Weize Chen, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2025. [Scaling large language model-based multi-agent collaboration](#). In *The Thirteenth International Conference on Learning Representations*.
- Shuofei Qiao, Runnan Fang, Zhisong Qiu, Xiaobin Wang, Ningyu Zhang, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. 2025. [Benchmarking agentic workflow generation](#). In *The Thirteenth International Conference on Learning Representations*.
- Zeyu Ren, Ling Yue, Ran Li, Yishu Wang, Shengxiang Xu, Hanmo Liu, Shaowu Pan, and Shimin Di. 2026. [FlowEvo: Self-evolving agents through the co-evolution of workflows and executable skills](#). In *Third Conference on Language Modeling*, San Francisco, California, USA.
- Jianhao Ruan, Zhihao Xu, Yiran Peng, Fashen Ren, Zhaoyang Yu, Xinbing Liang, Jinyu Xiang, Bang Liu, Chenglin Wu, Yuyu Luo, and Jiayi Zhang. 2026. [AOrchestra: Automating sub-agent creation for agentic orchestration](#). *arXiv preprint arXiv:2602.03786*.
- Jon Saad-Falcon, Adrian Gamarra Lafuente, Shlok Natarajan, Nahum Maru, Hristo Todorov, Etash Kumar Guha, E. Kelly Buchanan, Mayee F. Chen, Neel

- Guha, Christopher Ré, and Azalia Mirhoseini. 2025. [An architecture search framework for inference-time techniques](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 52475–52507. PMLR.
- Yexuan Shi, Mingyu Wang, Yunxiang Cao, Hongjie Lai, Junjian Lan, Xin Han, Yu Wang, Jie Geng, Zhenan Li, Zihao Xia, Xiang Chen, Chen Li, Jian Xu, Wenbo Duan, and Yuanshuo Zhu. 2025. [Aime: Towards fully-autonomous multi-agent framework](#). *arXiv preprint arXiv:2507.11988*.
- Jinwei Su, Qizhen Lan, Yinghui Xia, Lifan Sun, Weiyou Tian, Tianyu Shi, Xinyuan Song, Lewei He, and Jingsong Yang. 2025a. [Difficulty-aware agentic orchestration for query-specific multi-agent workflows](#). *arXiv preprint arXiv:2509.11079*.
- Jinwei Su, Yinghui Xia, Yiqun Duan, Jun Du, Jianuo Huang, Tianyu Shi, and Lewei He. 2025b. [DebFlow: Automating agent creation via agent debate](#). *arXiv preprint arXiv:2503.23781*.
- Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan, and Hoang D. Nguyen. 2025. [Multi-agent collaboration mechanisms: A survey of LLMs](#). *arXiv preprint arXiv:2501.06322*.
- Harsh Trivedi, Tushar Khot, Mareike Hartmann, Ruskin Manku, Vinty Dong, Edward Li, Shashank Gupta, Ashish Sabharwal, and Niranjan Balasubramanian. 2024. [AppWorld: A controllable world of apps and people for benchmarking interactive coding agents](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16022–16076, Bangkok, Thailand. Association for Computational Linguistics.
- Arunkumar V, Gangadharan G. R., and Rajkumar Buyya. 2026. [Agentic artificial intelligence \(AI\): Architectures, taxonomies, and evaluation of large language model agents](#). *arXiv preprint arXiv:2601.12560*.
- Jialiang Wang, Shengxiang Xu, Hanmo Liu, Jiachuan Wang, Yuyu Luo, Shimin Di, Min-Ling Zhang, and Lei Chen. 2026a. [Learning to compose for cross-domain agentic workflow generation](#). *arXiv preprint arXiv:2602.11114*.
- Jiayu Wang, Yifei Ming, Zixuan Ke, Shafiq Joty, Aws Albarghouthi, and Frederic Sala. 2026b. [SkillOrchestra: Learning to route agents via skill transfer](#). *arXiv preprint arXiv:2602.19672*.
- Siqi Wang and Guangmo Tong. 2025. [DAGP: Difficulty-aware graph pruning for LLM-Based multi-agent systems](#). In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*, pages 5319–5323, Seoul, Republic of Korea. ACM.
- Siyu Wang, Ruotian Lu, Zhihao Yang, Yuchao Wang, Yanzhou Zhang, Lei Xu, Qimin Xu, Guojun Yin, Cailian Chen, and Xinping Guan. 2026c. [Agent-Conductor: Topology evolution for multi-agent competition-level code generation](#). *arXiv preprint arXiv:2602.17100*.
- Wenxiao Wang, Priyatham Kattakinda, and Soheil Feizi. 2025a. [Maestro: Joint graph & config optimization for reliable AI agents](#). *arXiv preprint arXiv:2509.04642*.
- Xingyao Wang, Simon Rosenberg, Juan Michelini, Calvin Smith, Hoang Tran, Engel Nyst, Rohit Malhotra, Xuhui Zhou, Valerie Chen, Robert Brennan, and Graham Neubig. 2025b. [The OpenHands software agent SDK: A composable and extensible foundation for production agents](#). *arXiv preprint arXiv:2511.03690*.
- Yanbo Wang, Zixiang Xu, Yue Huang, Xiangqi Wang, Zirui Song, Lang Gao, Chenxi Wang, Xiangru Tang, Yue Zhao, Arman Cohan, Xiangliang Zhang, and Xiuying Chen. 2025c. [Dyflow: Dynamic workflow framework for agentic reasoning](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Yimeng Wang, Jiaxing Zhao, Hongbin Xie, Hexing Ma, Yuzhen Lei, Shuangxue Liu, Xuan Song, Zichen Zhang, and Haoran Zhang. 2026d. [MetaGen: Self-evolving roles and topologies for multi-agent LLM reasoning](#). *arXiv preprint arXiv:2601.19290*.
- Yingxu Wang, Siwei Liu, Jinyuan Fang, and Zaiqiao Meng. 2025d. [EvoAgentX: An automated framework for evolving agentic workflows](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 643–655, Suzhou, China. Association for Computational Linguistics.
- Yinjie Wang, Ling Yang, Guohao Li, Mengdi Wang, and Bryon Aragam. 2025e. [ScoreFlow: Mastering LLM agent workflows via score-based preference optimization](#). *arXiv preprint arXiv:2502.04306*.
- Zhenting Wang, Qi Chang, Hemani Patel, Shashank Biju, Cheng-En Wu, Quan Liu, Aolin Ding, Alireza Rezazadeh, Ankit Shah, Yujia Bao, and Eugene Siow. 2026e. [MCP-bench: Benchmarking tool-using LLM agents with complex real-world tasks via MCP servers](#). In *The Fourteenth International Conference on Learning Representations*.
- Zhexuan Wang, Yutong Wang, Xuebo Liu, Liang Ding, Miao Zhang, Jie Liu, and Min Zhang. 2025f. [Agent-Dropout: Dynamic agent elimination for token-efficient and high-performance LLM-Based multi-agent collaboration](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 24013–24035, Vienna, Austria. Association for Computational Linguistics.

- Jiaqi Wei, Yuejin Yang, Xiang Zhang, Yuhan Chen, Xiang Zhuang, Zhangyang Gao, Dongzhan Zhou, Guangshuai Wang, Zhiqiang Gao, Juntao Cao, Zijie Qiu, Ming Hu, Chenglong Ma, Shixiang Tang, Junjun He, Chunfeng Song, Xuming He, Qiang Zhang, Chenyu You, and 8 others. 2025. [From ai for science to agentic science: A survey on autonomous scientific discovery](#). *arXiv preprint arXiv:2508.14111*.
- Yangbo Wei, Zhen Huang, Lei He, Li Huang, Ting-Jung Lin, and Wei W Xing. 2026. [VFlow: Discovering optimal agentic workflows for Verilog generation](#). In *2026 31st Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 355–361. IEEE.
- Qingyun Wu, Gagan Bansal, Jieyu Zhang, Yiran Wu, Beibin Li, Erkang Zhu, Li Jiang, Xiaoyun Zhang, Shaokun Zhang, Jiale Liu, Ahmed Hassan Awadallah, Ryen W. White, Doug Burger, and Chi Wang. 2023. [AutoGen: Enabling next-gen LLM applications via multi-agent conversation](#). *arXiv preprint arXiv:2308.08155*.
- Ruixuan Xiao, Wentao Ma, Ke Wang, Yuchuan Wu, Junbo Zhao, Haobo Wang, Fei Huang, and Yongbin Li. 2024. [FlowBench: Revisiting and benchmarking workflow-guided planning for LLM-Based agents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10883–10900. Association for Computational Linguistics.
- Jiawei Xu, Arief Koesdwiady, Sisong Bei, Yan Han, Baixiang Huang, Dakuo Wang, Yutong Chen, Zheshen Wang, Peihao Wang, Pan Li, and Ying Ding. 2026. [Rethinking the value of multi-agent workflow: A strong single agent baseline](#). *arXiv preprint arXiv:2601.12307*.
- Qiantong Xu, Fenglu Hong, Bo Li, Changran Hu, Zhengyu Chen, and Jian Zhang. 2023. [On the tool manipulation capability of open-source large language models](#). *arXiv preprint arXiv:2305.16504*.
- Shengxiang Xu, Jiayi Zhang, Shimin Di, Yuyu Luo, Liang Yao, Hanmo Liu, Jia Zhu, Fan Liu, and Min-Ling Zhang. 2025a. [RobustFlow: Towards robust agentic workflow generation](#). *arXiv preprint arXiv:2509.21834*.
- Weikai Xu, Chengrui Huang, Shen Gao, and Shuo Shang. 2025b. [LLM-Based Agents for tool learning: A survey](#). *Data Science and Engineering*, 10(4):533–563.
- Zhenran Xu, Yiyu Wang, Xue Yang, Longyue Wang, Weihua Luo, Kaifu Zhang, Baotian Hu, and Min Zhang. 2025c. [ComfyUI-R1: Exploring reasoning models for workflow generation](#). *arXiv preprint arXiv:2506.09790*.
- Yunhe Yan, Shihe Wang, Jiajun Du, Yexuan Yang, Yuxuan Shan, Qichen Qiu, Xianqing Jia, Xinge Wang, Xin Yuan, Xu Han, Mao Qin, Yinxiao Chen, Chen Peng, Shangguang Wang, and Mengwei Xu. 2025. [MCPWorld: A unified benchmarking testbed for API, GUI, and hybrid computer use agents](#). *arXiv preprint arXiv:2506.07672*.
- Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2024. [Large language models as optimizers](#). In *The Twelfth International Conference on Learning Representations*.
- Yingxuan Yang, Huacan Chai, Shuai Shao, Yuanyi Song, Siyuan Qi, Renting Rui, and Weinan Zhang. 2025. [AgentNet: Decentralized evolutionary coordination for LLM-based multi-agent systems](#). In *Advances in Neural Information Processing Systems 38*, pages 107309–107336.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik R Narasimhan. 2025. [\$\tau\$ -Bench: A Benchmark for Tool-Agent-User Interaction in Real-World Domains](#). In *The Thirteenth International Conference on Learning Representations*.
- Yining Ye, Xin Cong, Shizuo Tian, Jiannan Cao, Hao Wang, Yujia Qin, Yaxi Lu, Heyang Yu, Huadong Wang, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2023. [ProAgent: From robotic process automation to agentic process automation](#). *arXiv preprint arXiv:2311.10751*.
- Chaojia Yu, Zihan Cheng, Hanwen Cui, Yishuo Gao, Zexu Luo, Yijin Wang, Hangbin Zheng, and Yong Zhao. 2025. [A survey on agent workflow – status and future](#). In *2025 8th International Conference on Artificial Intelligence and Big Data (ICAIBD)*, pages 770–781.
- Ling Yue, Ching-Yun Ko, Pin-Yu Chen, Shimin Di, and Shaowu Pan. 2026a. [Building mcp-native hierarchical ai scientist ecosystems: A perspective on scaling multi-agent scientific discovery](#). *Frontiers in Artificial Intelligence*, 9:1820375.
- Ling Yue, Nithin Somasekharan, Tingwen Zhang, Yadi Cao, Zhangze Chen, Shimin Di, and Shaowu Pan. 2026b. [Foam-Agent: A large language model-based multi-agent framework for automating computational fluid dynamics workflows](#). *Computer Methods in Applied Mechanics and Engineering*, 461:119271.
- Yanwei Yue, Guibin Zhang, Boyang Liu, Guancheng Wan, Kun Wang, Dawei Cheng, and Yiyang Qi. 2025. [MasRouter: Learning to route LLMs for multi-agent systems](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15549–15572, Vienna, Austria. Association for Computational Linguistics.
- Tom Zehle, Moritz Schlager, Timo Hei, and Matthias Feurer. 2025. [Capo: Cost-aware prompt optimization](#).

- In *Proceedings of the Fourth International Conference on Automated Machine Learning*, volume 293 of *Proceedings of Machine Learning Research*, pages 18/1–45. PMLR.
- Guibin Zhang, Kaijie Chen, Guancheng Wan, Heng Chang, Hong Cheng, Kun Wang, Shuyue Hu, and Lei Bai. 2025a. [EvoFlow: Evolving diverse agentic workflows on the fly](#). *arXiv preprint arXiv:2502.07373*.
- Guibin Zhang, Luyang Niu, Junfeng Fang, Kun Wang, Lei Bai, and Xiang Wang. 2025b. [Multi-agent architecture search via agentic supernet](#). In *Forty-Second International Conference on Machine Learning*.
- Guibin Zhang, Yanwei Yue, Zhixun Li, Sukwon Yun, Guancheng Wan, Kun Wang, Dawei Cheng, Jeffrey Xu Yu, and Tianlong Chen. 2025c. [Cut the crap: An economical communication pipeline for LLM-Based multi-agent systems](#). In *The Thirteenth International Conference on Learning Representations*.
- Guibin Zhang, Yanwei Yue, Xiangguo Sun, Guancheng Wan, Miao Yu, Junfeng Fang, Kun Wang, Tianlong Chen, and Dawei Cheng. 2025d. [G-designer: Architecting multi-agent communication topologies via graph neural networks](#). In *Forty-second International Conference on Machine Learning*.
- Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xiong-Hui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, and Chenglin Wu. 2025e. [AFlow: Automating agentic workflow generation](#). In *The Thirteenth International Conference on Learning Representations*.
- Xueqiang Zhang, Xiaofei Dong, Yiru Wang, Dan Zhang, and Feng Cao. 2025f. [A survey of multi-AI agent collaboration: Theories, technologies and applications](#). In *Proceedings of the 2nd Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence*, pages 1875–1881, Dongguan, China. ACM.
- Mingming Zhao, Xiaokang Wei, Yuanqi Shao, Kaiwen Zhou, Lin Yang, Siwei Rao, Junhui Zhan, and Zhitang Chen. 2026. [A²flow: Automating agentic workflow generation via self-adaptive abstraction operators](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 29930–29938.
- Chengqi Zheng, Jianda Chen, Yueming Lyu, Wen Zheng, Terence Ng, Haopeng Zhang, Yew-Soon Ong, Ivor Tsang, and Haiyan Yin. 2025. [Mermaid-Flow: Redefining agentic workflow generation via safety-constrained evolutionary programming](#). *arXiv preprint arXiv:2505.22967*.
- Han Zhou, Xingchen Wan, Ruoxi Sun, Hamid Palangi, Shariq Iqbal, Ivan Vulić, Anna Korhonen, and Serkan O Arik. 2026. [Multi-agent design: Optimizing agents with better prompts and topologies](#). In *The Fourteenth International Conference on Learning Representations*.
- Yuxuan Zhu, Tengjun Jin, Yada Pruksachatkun, Andy Zhang, Shu Liu, Sasha Cui, Sayash Kapoor, Shayne Longpre, Kevin Meng, Rebecca Weiss, Fazl Barez, Rahul Gupta, Jwala Dhamala, Jacob Merizian, Mario Giulianelli, Harry Crippock, Cozmin Ududec, Antony Kellermann, Jasjeet Sekhon, and 7 others. 2025. [Establishing best practices in building rigorous agentic benchmarks](#). In *Advances in Neural Information Processing Systems*, volume 38, Main Conference. Curran Associates, Inc.
- Mingchen Zhuge, Wenyi Wang, Louis Kirsch, Francesco Faccio, Dmitrii Khizbullin, and Jürgen Schmidhuber. 2024. [GPTSwarm: Language agents as optimizable graphs](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 62743–62767. PMLR.

A Additional Conceptual Details

Two simple examples make the template–realized graph–trace distinction concrete. In a fixed planner–retriever–executor–verifier pipeline, the pipeline definition is the template, while the workflow structure that is actually traversed for one question—often the full pipeline, but sometimes a pre-authored branch-conditioned substructure—is the realized graph; the resulting retrieval calls, code executions, verifier decisions, and retries constitute the trace. By contrast, in a query-conditioned multi-agent workflow generator, the reusable object may specify only the operator vocabulary or graph-generation policy; each query then induces a different realized graph, and the trace records the messages, tool calls, failures, and edits that occur when that generated graph runs.

The quality–cost view used in the main text can be written schematically as

$$\max_{x \sim \mathcal{D}} \mathbb{E}_{x \sim \mathcal{D}} \left[\mathbb{E}_{\mathcal{G}^{\text{run}} | x} \left[\mathbb{E}_{\tau | \mathcal{G}^{\text{run}}, x} [R(\tau; x) - \lambda C(\tau)] \right] \right],$$

where $R(\tau; x)$ is a task-quality score, $C(\tau)$ is execution cost, the inner expectation reflects execution stochasticity conditional on a realized graph, the middle expectation matters when the method generates or selects a run-specific graph, and λ controls the quality–cost trade-off. As concrete instantiations, AFlow (Zhang et al., 2025e) collapses the middle expectation because the realized graph coincides with the template, whereas FlowReasoner (Gao et al., 2025a) makes the middle expectation non-trivial because a new graph is sampled per query. For in-execution editing, the realized graph can be treated as part of the evolving system state, and edit actions appear inside the trace rather than being decided entirely upfront.

B Survey Corpus Profile

Figure 4 summarizes the temporal, venue, and geographic distribution of the 46 CORE and ADJACENT papers included in this survey.

Corpus construction and scope. Each eligible source receives one primary scope tag for counting, even when it spans multiple method families. CORE papers construct or optimize a reusable template or executable realized graph, including its composition rules, topology, or node contents. ADJACENT papers only mask, prune, or route within

a pre-authored scaffold, or select among intact alternatives, without changing composition rules or node contents. BACKGROUND sources provide frameworks, datasets, benchmarks, or conceptual foundations without directly optimizing workflow structure.

Search procedure. The search snapshot was updated through August 2026. Initial discovery used arXiv, OpenReview, ACL Anthology, and publisher or proceedings indexes, with combinations of *LLM*, *language agent*, or *multi-agent*; *workflow*, *computation graph*, *topology*, or *orchestration*; and *optimization*, *search*, *generation*, *adaptation*, or *evolution*. We merged preprint and venue versions by title and author list, screened titles and abstracts using the scope rules above, and inspected the full text when the structural intervention was unclear. Planning-only, tool-use-only, memory-only, evaluation-only, and general framework papers were excluded from CORE unless they changed or directly evaluated workflow structure. We then followed references and citing papers for GPTSwarm, AFlow, ADAS, and G-Designer, selected as seeds spanning joint graph optimization, reusable template search, code-defined system search, and run-specific graph generation. This check added GPTSwarm, Symbolic Learning, Archon, and AgentNet as archival CORE papers.

Temporal distribution. Panel (a) shows cumulative publication counts over time. The earliest paper in this collection dates to late 2023 (Ye et al., 2023), and growth accelerates sharply in early 2025. Dynamic methods dominate the corpus: 30 of the 46 CORE and ADJACENT papers address runtime selection, generation, or editing, compared with 16 static methods.

Venue distribution. Panel (b) breaks down the papers by venue. ICLR is the largest peer-reviewed venue with 11 papers, followed by ACL (3) and ICML (2), with the remainder spread across machine-learning, NLP, web, and systems venues. ArXiv preprints account for 20 of the 46 papers, consistent with the rapid release cycle in this area.

Geographic distribution. Panel (c) reports author-affiliations rather than papers: each author is counted once per paper and assigned to the country of their affiliated institution. Across 379 author-affiliations, Asia accounts for 64.7%, led by China (211), Singapore (20), and the United Arab Emirates (9). North America accounts for 29.0%,




Figure 3: Full citation-level taxonomy of workflow-optimization methods for LLM agents

led by the United States (103) and Canada (7), with Europe and Oceania making up the remainder.

C Glossary and Coding Guide

This section defines the terms and coding fields used throughout the appendix tables. The goal is to make the comparison card reproducible: each method is coded by when the graph is determined, what part of the workflow is optimized, what evidence authorizes updates, and how costs are handled.

D Comparison of Surveyed Core Methods

Tables 5 and 6 are the main appendix comparison tables. They use the coding guide from Section C to place static template optimizers, pre-execution generators or selectors, and in-execution editors on the same set of axes.

E Survey Positioning

Table 7 clarifies how this survey differs from adjacent surveys on agent workflows, planning, tool use, collaboration, architecture, and broad agent optimization.

F Feedback Signals and Update Mechanisms

The same structural search space can behave very differently depending on the feedback/evidence used to accept, reject, rank, or revise candidate structures. In practice, feedback signal and update mechanism are tightly coupled, so it is useful to discuss them together. A given paper can also use multiple signals at once; throughout this section, categories refer to the primary feedback/evidence used for the main structural decision rather than to an exhaustive list of every signal appearing in the method.

Metric- and score-driven optimization. The most direct signal is a scalar metric such as success, accuracy, F1, pass@k, or an application-specific reward. This signal drives black-box search in AFlow (Zhang et al., 2025e), ADAS (Hu et al., 2025a), Evolutionary Generation of Multi-Agent Systems (Hu et al., 2026), SEW (Liu et al., 2025a), and VFlow (Wei et al., 2026). It also underlies query-conditioned generators such as Difficulty-Aware Agentic Orchestration (Su et al., 2025a), FlowReasoner (Gao et al., 2025a), and MaAS (Zhang et al., 2025b), where performance

and cost jointly define the reward or routing objective used to train a controller. Scalar metrics are appealing because they align directly with the deployment objective, but they are often sparse, noisy, and expensive. Search efficiency, therefore, depends heavily on the quality of the evaluator and on the granularity of the action space.

Verifier-driven optimization. Verifier signals occupy a special place in workflow optimization because they can be used as hard constraints, intermediate checkpoints, or dense reward components. In program synthesis and software tasks, unit tests and executability checks provide strong supervision for workflow repair (Hu et al., 2025b; Ye et al., 2023). In structured workflow synthesis, format validity and schema compliance can be checked before expensive downstream execution (Zheng et al., 2025; Xu et al., 2025c). AgentConductor (Wang et al., 2026c) combines validity, execution, and cost feedback during topology revision. VFlow (Wei et al., 2026) illustrates multi-level verification particularly well: it measures syntax, functional correctness, synthesizability, and hardware objectives separately and feeds each back into search.

Verifiers do more than improve quality. They also shape the feasible edit space. A workflow language with strong static checks enables more aggressive generation or mutation, since invalid candidates can be rejected cheaply. The main limitation is that verifiers are imperfect or expensive themselves. A weak verifier can reward brittle hacks, while an overly costly verifier can dominate the optimization budget.

Preference and ranking signals. Preference signals compare workflows or traces instead of assigning each one an independent scalar reward. ScoreFlow (Wang et al., 2025e) constructs score-aware preference pairs from executed workflows and uses them to stabilize gradient-based optimization of a generator. RobustFlow (Xu et al., 2025a) builds preferences from semantic clusters and self-consistency, preferring structures that remain stable across paraphrases. Optima (Chen et al., 2025) similarly uses ranking and selection over generated multi-agent trajectories before fine-tuning. Preference learning is attractive when absolute rewards are noisy, but relative comparisons are reliable, or when multiple aspects, such as accuracy and cost, must be combined in a ranking rule.




Figure 4: Distribution of the 46 CORE and ADJACENT papers surveyed. (a) Cumulative publication count over time (by paper count), split by static and dynamic methods. (b) Venue distribution (by paper count). (c) Geographic distribution at the author-affiliation level, where each author is counted once per paper and assigned to the country of their affiliated institution; the numbers indicate author-affiliation counts (e.g., 211 means 211 author-affiliations are in China across the 46 papers)

Table 3: Glossary of key terms and notation used in this survey

Term	Definition
Agentic computation graph (ACG)	our unifying abstraction for an executable LLM-centered workflow. Nodes perform atomic actions such as LLM calls, information retrieval, tool use, validation, or message passing. Edges encode control, data, or communication dependencies.
ACG template ($\bar{\mathcal{G}}$)	reusable executable specification ($\mathcal{V}, \mathcal{E}, \Phi, \Sigma, \mathcal{A}$).
Realized graph ($\mathcal{G}^{\text{run}}$)	workflow structure actually used for a particular run.
Execution trace (τ)	sequence of states, actions, observations, and costs produced by executing $\mathcal{G}^{\text{run}}$.
Workflow	an executable organization of multiple steps for solving a task.
Workflow structure	which components or agents are present, how they depend on one another, and how information flows between them.
Topology	the communication pattern among agents or modules.
Graph determination time (GDT)	when structure is decided: <i>offline</i> , <i>pre-execution</i> , or <i>in-execution</i> .
Graph plasticity mode (GPM)	how much structure can change at inference time: <i>none</i> , <i>select</i> , <i>generate</i> , or <i>edit</i> .
Node parameters (Φ)	prompts, role descriptions, tool schemas, model choices, decoding settings, verifier settings, and memory policies.
Structure variables	topology, scheduling or routing policy, activation choices, and admissible edit actions.
Scope tags	CORE: constructs or optimizes a reusable template or executable realized graph, including its composition rules, topology, or node contents. ADJACENT: only masks, prunes, or routes within a pre-authored scaffold, or selects among intact alternatives, without changing composition rules or node contents. BACKGROUND: provides frameworks, datasets, benchmarks, or conceptual foundations without directly optimizing workflow structure. Each paper receives one primary tag for corpus counts.

Trace-derived textual feedback. Textual feedback extracted from traces is increasingly used as a higher-bandwidth signal than scalar reward alone. GEPA (Agrawal et al., 2026) is the clearest

prompt-level example: it reflects on successful and failed trajectories in language and uses those reflections to propose prompt edits. MetaGen (Wang et al., 2026d) updates role descriptions and graph

Table 4: Field definitions and coding rules for the compact classification card used in the comparison card of Tables 5 and 6

Field	Meaning in this survey
Setting	static or dynamic, together with GDT (offline / pre-execution / in-execution) and GPM (none / select / generate / edit)
Optimized level	node / graph / joint
Representation	code / text / DSL / explicit graph / typed operator graph / constrained IR / orchestration state
Feedback / evidence	dominant workflow-relevant evidence used to accept, rank, or revise structures; encode as a coarse tag plus a short instantiation, e.g., metric, verifier, supervision, preference, trace, or proxy
Update mechanism	dominant structural update mechanism; encode as a coarse tag plus a short instantiation, e.g., search, generator, controller, supervised learning, RL, preference optimization, repair/edit, continuous relaxation, or hybrid
Cost handling	whether cost enters the method and how: none / eval-only / soft objective / hard constraint, optionally annotated with a resource type such as tokens, tool calls, latency, dollars, or budget
Reporting rule	when a paper uses multiple evidence types or update mechanisms, the tables report the dominant workflow-relevant entry in each field; a secondary tag may be added only when it is essential to avoid distortion, with finer details discussed in the surrounding text

Table 5: Representative CORE workflow-optimization methods, Part I: static offline template optimization. *Setting* reports graph determination time (GDT) and graph plasticity mode (GPM); *Level* is node, graph, or joint. When a paper uses multiple evidence types or update mechanisms, columns report the dominant workflow-relevant entry

Method	Setting (GDT/GPM)	Level	Feedback / evidence	Update mechanism
AFlow (Zhang et al., 2025e)	offline / none	graph	metric: task score	search: MCTS
ADAS (Hu et al., 2025a)	offline / none	joint	metric: task score	search: archive meta-search
Archon (Saad-Falcon et al., 2025)	offline / none	joint	metric: task quality	search: architecture + model assignment
A ² Flow (Zhao et al., 2026)	offline / none	graph	supervision: demos + execution	hybrid: operator learning + search
Multi-Agent Design (Zhou et al., 2026)	offline / none	joint	metric: task score	hybrid: staged alternation
GPTSwarm (Zhuge et al., 2024)	offline / none	joint	metric: task score	hybrid: node/edge optimization
Symbolic Learning (Ou et al., 2025)	offline / none	joint	trace: language loss + gradients	hybrid: symbolic backpropagation
Learning Multi-Agent Communication from a Graph Modeling Perspective (Hu et al., 2024)	offline / none	joint	proxy: training objective	continuous relaxation
Optima (Chen et al., 2025)	offline / none	node	metric: quality–efficiency reward	hybrid: generate-rank-select-train
SEW (Liu et al., 2025a)	offline / none	graph	metric: pass rate	search: self-evolution
VFlow (Wei et al., 2026)	offline / none	graph	verifier: multi-level checks	hybrid: MCTS + cooperative evolution
MermaidFlow (Zheng et al., 2025)	offline / none	graph	verifier: validity + execution	search: constrained evolution
Maestro (Wang et al., 2025a)	offline / none	joint	trace: reflective text + score	hybrid: alternating graph/config updates
Evolutionary Generation of Multi-Agent Systems (Hu et al., 2026)	offline / none	graph	metric: execution fitness	search: evolutionary

edits using contradiction- and failure-oriented feedback from logs. DebFlow (Su et al., 2025b) uses debate and reflection to refine workflows. Maestro (Wang et al., 2025a) explicitly combines numeric scores with reflective text so that graph edits can respond to semantically meaningful failure de-

scriptions rather than only to outcome deltas. SkillOrchestra (Wang et al., 2026b) occupies a related boundary case: it extracts fine-grained skill abstractions from traces and uses them to guide future routing decisions.

The strength of textual feedback is semantic rich-

Table 6: Representative CORE workflow-optimization methods, Part II: dynamic routing, graph generation, and in-execution adaptation

Method	Setting (GDT/GPM)	Level	Feedback / evidence	Update mechanism
<i>Pre-execution routing and graph generation</i>				
Difficulty-Aware Agentic Orchestration (Su et al., 2025a)	pre-exec / select	joint	proxy: difficulty estimate	controller: router + allocator
Assemble Your Crew (Li et al., 2026)	pre-exec / generate	graph	supervision: task labels	generator: autoregressive DAG
G-Designer (Zhang et al., 2025d)	pre-exec / generate	graph	metric: quality score	generator: VGAE
Dynamic Generation of Multi-LLM Agents Communication Topologies (Jiang et al., 2025)	pre-exec / generate	graph	proxy: reward proxy	generator: graph diffusion
Multi-Agent Architecture Search via Agentic Supernet (Zhang et al., 2025b)	pre-exec / generate	graph	metric: quality score	generator: architecture distribution
ScoreFlow (Wang et al., 2025e)	pre-exec / generate	graph	preference: score-aware pairs	preference optimization
<i>Pre-execution learning, reasoning, and verification</i>				
CapFlow (Wang et al., 2026a)	pre-exec / generate	graph	metric: success + counterfactual attribution	generator: capability recomposition
FlowReasoner (Gao et al., 2025a)	pre-exec / generate	graph	metric: reward	RL: meta-controller
Workflow-R1 (Kong et al., 2026)	pre-exec / generate	graph	metric: grouped rewards	RL: think-act subsequences
AutoFlow (Li et al., 2024b)	pre-exec / generate	graph	verifier: interpreter feedback	hybrid: generation + execution feedback
WorkflowLLM (Fan et al., 2025)	pre-exec / generate	graph	supervision: workflow labels	supervised learning: finetuning
RobustFlow (Xu et al., 2025a)	pre-exec / generate	graph	preference: clustered preferences	preference optimization
ComfyUI-R1 (Xu et al., 2025c)	pre-exec / generate	graph	verifier: validity + reward	hybrid: reasoning training
<i>Hybrid drafting and in-execution adaptation</i>				
AutoAgents (Chen et al., 2024a)	pre-exec → in-exec / edit	joint	trace: runtime feedback	hybrid: draft + observer refinement
DyFlow (Wang et al., 2025c)	in-exec / edit	joint	trace: intermediate feedback	hybrid: online planning + execution
AgentConductor (Wang et al., 2026c)	in-exec / edit	graph	verifier: validity + execution	RL: topology revision
AgentNet (Yang et al., 2025)	in-exec / select	joint	trace: local outcomes + expertise	routing + across-task pruning
Aime (Shi et al., 2025)	in-exec / edit	graph	trace: runtime outcomes	controller: actor instantiation
AOrchestra (Ruan et al., 2026)	in-exec / edit	joint	metric: outcome	controller: sub-agent creation
MetaGen (Wang et al., 2026d)	in-exec / edit	joint	trace: running feedback	repair/edit: training-free evolution
<i>In-execution repair, evolution, and control</i>				
Evoflux (Bhandari et al., 2026)	in-exec / edit	graph	verifier: execution + schema checks	search: evolutionary graph repair
ProAgent (Ye et al., 2023)	in-exec / edit	graph	verifier: tests	repair/edit: incremental repair
Flow (Niu et al., 2025)	in-exec / edit	graph	trace: intermediate execution	repair/edit: graph refinement
EvoFlow (Zhang et al., 2025a)	in-exec / edit	graph	metric: online fitness	search: online evolution
DebFlow (Su et al., 2025b)	in-exec / edit	graph	trace: debate + reflection	repair/edit: debate refinement
QualityFlow (Hu et al., 2025b)	in-exec / edit	graph	verifier: quality checks	controller: submit/clarify/revert

Table 7: Positioning this survey in the broader survey landscape. Most prior surveys center on adjacent topics such as planning, tool learning, collaboration, architecture, or broad agent optimization, whereas this survey centers workflow structure itself

Survey family	Central lens	How workflow appears	Relation to this survey
Workflow / infrastructure (Yu et al., 2025; Li et al., 2024a)	workflow systems, orchestration abstractions, and MAS organization	workflow as system scaffold, lifecycle, or engineering abstraction	adjacent but not centered on structure search, per-run graph construction, or workflow-specific evaluation
Planning (Huang et al., 2024)	decomposition, plan selection, reflection, memory, and external planning modules	workflow as one internal phase or capability of an agent	covers a crucial ingredient of agents, but not workflow optimization as a broader graph-design problem
Tool learning (Xu et al., 2025b)	tool retrieval, tool planning, tool invocation, and tool-use evaluation	workflow as a tool-use pipeline and environment interface	centers tool competence rather than topology, routing, or graph generation/editing
MAS collaboration / applications (Chen et al., 2024b; Tran et al., 2025; Zhang et al., 2025f)	collaboration mechanisms, communication protocols, and application domains	workflow as a communication pattern or application-specific coordination scheme	studies collaborative behavior and applications, not reusable-template or realized-graph design as the main object
Broad agent optimization (Du et al., 2026)	parameter-driven and parameter-free optimization of agents overall	workflow as one optimization lever among prompts, training, retrieval, memory, and other factors	does not isolate workflow structure as the primary object of comparison
Agent architecture / evaluation (V et al., 2026)	architectures, taxonomies, and evaluation of agentic AI systems	workflow as one architectural manifestation within a larger agent taxonomy	provides a broad architecture-and-evaluation view, but not a workflow-optimization-specific synthesis
Workflow Optimization (ours)	workflow optimization for LLM agents	workflow structure as the primary optimization object	unifies static template optimization, per-run graph generation/selection/editing, template-realized graph-trace distinctions, and structure-aware evaluation/reporting

ness. It can suggest why a workflow failed and what structural change might help. Its weakness is that textual critiques can drift, hallucinate, or overfit to anecdotal traces. In the strongest systems, textual feedback acts as a proposal mechanism, while metrics or verifiers decide whether the proposal should survive.

Matching signals to algorithms. Across the surveyed literature, certain signal–algorithm pairings recur for understandable methodological reasons. Search works best when evaluators are trusted and the action space is discrete but structured, as in typed operator graphs, code templates, or constrained graph IRs. Reinforcement learning is natural when workflow generation is genuinely sequential, as in FlowReasoner (Gao et al., 2025a) or Workflow-R1 (Kong et al., 2026), but it places great pressure on reward design and exploration. Supervised learning is attractive when workflow artifacts exist at scale, as in WorkflowLLM (Fan et al., 2025)

or Assemble Your Crew (Li et al., 2026). Training-free adaptation becomes plausible when the edit space is small, the validator is strong, and trace feedback is informative, which explains the recent interest in MetaGen (Wang et al., 2026d), QualityFlow (Hu et al., 2025b), DyFlow (Wang et al., 2025c), and Aime (Shi et al., 2025)-style systems.

A useful cross-paper pattern is that the signal often determines the safe action granularity. Strong verifiers support more aggressive graph mutation because invalid candidates can be rejected cheaply. Preference data stabilizes generation when scalar rewards are noisy but rankings are reliable. Textual feedback carries richer causal hints but is safest when coupled to an external validator. RL is most compelling when workflow construction is inherently sequential rather than merely formatted as a sequence. In other words, the literature is not merely choosing algorithms; it is choosing what evidence is trusted enough to justify structural change.

This is also why evaluation must report not just final outcomes, but the signals and guards that authorized those changes.

G Evaluation and Reporting

Across the surveyed literature, workflow structure is treated as a first-class design object during method development much more often than it is treated as a first-class output during evaluation. For workflow research to mature, the graph itself must become part of what is reported and compared. To that end, Table 8 organizes representative evaluation assets by their role in this research area.

Treating the workflow as a first-class output.

There are two distinct evaluation goals. The first is *structure-aware evaluation*: did the method generate or select a good workflow or graph? The second is *downstream task validation*: did the resulting system solve the task? Both matter, but they answer different questions. Structure-aware evaluation is most informative when the method explicitly outputs workflows, as in WorFBench (Qiao et al., 2025) and, to a lesser extent, FlowBench (Xiao et al., 2024), or in constrained graphical domains such as ComfyUI-R1 (Xu et al., 2025c). Downstream validation is indispensable when real execution in tools or environments is the true target, as in GAIA (Mialon et al., 2024), SWE-bench (Jimenez et al., 2024), Terminal-Bench (Merrill et al., 2026), or the broader Model Context Protocol (MCP) benchmark family (Luo et al., 2025; Wang et al., 2026e).

A good evaluation protocol should therefore separate the two whenever possible. A workflow generator can output a graph that looks plausible but executes poorly. Conversely, a downstream task can be solved by excessive brute-force compute or by a hidden scaffold that is never analyzed structurally. Reporting only final task success makes it impossible to know which of these actually happened.

Effectiveness, efficiency, robustness, and graph-level properties. Effectiveness metrics remain task-dependent: accuracy, exact match, F1, success rate, pass@k, and domain-specific scores all appear across the literature. Efficiency metrics should be reported alongside them, not after the fact. At a minimum, papers should report prompt and completion tokens, number of LLM calls, number of tool calls, latency or wall-clock time, and monetary

cost when priced APIs are used. Cost-per-success is often more informative than raw cost because it normalizes for utility.

Workflow papers should additionally report graph-level properties. For static templates, size, depth, width, critical path length, and communication volume help disentangle “better structure” from “more compute.” For dynamic methods, edit count, fraction of steps spent editing, and structural variance across repeated sampling or paraphrases are especially informative. RobustFlow’s (Xu et al., 2025a) node- and graph-level similarity metrics, as well as WorFBench’s (Qiao et al., 2025) subsequence and subgraph matching, show how this style of evaluation can be made concrete.

Robustness should be evaluated under perturbation, not only under nominal execution. The most immediate tests are paraphrase invariance, tool failure injection, information retrieval noise, API or schema drift, and generalization to unseen tools or larger registries. Dynamic methods should also be tested under strict budget caps because the ability to replan can easily degenerate into uncontrolled cost growth. LiveMCPBench (Mo et al., 2025) is a useful step toward drift-sensitive evaluation, but controlled drift protocols remain rare.

A minimum reporting protocol. Benchmark design itself deserves scrutiny. Rigorous agentic benchmarks need clear task specifications, reliable evaluation setups, careful budget control, and explicit documentation of what tools and side channels are available (Zhu et al., 2025). Workflow-specific benchmarks additionally need canonicalization rules and structure-comparison procedures; otherwise superficially different but semantically identical workflows will be scored inconsistently. The point is not to force a single canonical graph, but to make evaluation robust enough that structural claims are scientifically meaningful.

The central reporting claim of this survey is simple: workflow papers should not report only *what answer was produced*. They should also report *what workflow was used, how that workflow varied across instances, and what it cost to obtain it*. Without those details, apparent gains are hard to attribute: a method may look stronger because it found a better workflow, because it spent more compute, or because hidden retries and fallback scaffolds were doing the real work. Table 1 distills these considerations into a minimum reporting protocol.

Table 8: Representative evaluation assets, grouped by the role they play in workflow optimization research

Role in workflow research	Representative resources	Main artifact or evidence	Structure-aware?	Best used for
Workflow-generation benchmarks	WorFBench (Qiao et al., 2025), FlowBench (Xiao et al., 2024), constrained graphical domains such as ComfyUI-R1 (Xu et al., 2025c)	reference workflows, schema-valid graphs, workflow plans	✓ / $\triangle$	evaluating structural fidelity, validity, and workflow-level robustness
Training corpora and executable validators	WorkflowBench (Fan et al., 2025), HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021), GSM8K (Cobbe et al., 2021), MATH (Hendrycks et al., 2021), DROP (Dua et al., 2019), HotpotQA (Yang et al., 2018)	workflow code pairs, hidden tests, labeled reasoning tasks	$\triangle$ / $\times$	supervised workflow generation and controlled validation of local or global design choices
Interactive tool-use benchmarks	GAIA (Mialon et al., 2024), τ -Bench (Yao et al., 2025), τ^2 -Bench (Barres et al., 2025), APiBank (Li et al., 2023b), T-Eval (Chen et al., 2024c), ToolBench (Xu et al., 2023), AppWorld (Trivedi et al., 2024), MCP-Universe (Luo et al., 2025), MCP-Bench (Wang et al., 2026e), MCP-RADAR (Gao et al., 2025b), MCP Eval (Liu et al., 2025b), LiveMCPBench (Mo et al., 2025), MCPWorld (Yan et al., 2025)	long-horizon execution traces under realistic tools	$\times$	end-to-end orchestration, schema compliance, recovery, and tool-selection quality
Software and operations environments	SWE-bench (Jimenez et al., 2024), Terminal-Bench (Merrill et al., 2026), SOP-Bench (Nandi et al., 2025), AssetOpsBench (Patel et al., 2025)	repository edits, terminal traces, procedural execution logs	$\times$	planning, verification, and recovery under hard constraints

H Discussion: Trade-offs and Open Problems

The surveyed papers support several broad conclusions, but these should be read as cross-paper empirical patterns rather than as universal laws.

When static is enough. The literature suggests that static optimization is often sufficient when three conditions hold at once: the operator space is constrained enough to search, the evaluator is trustworthy enough to discriminate candidates, and the deployment workload is repetitive enough that a reusable template is worth optimizing. In practice, these conditions often coincide when APIs are stable and verification is strong. In such settings, a well-searched template can dominate ad hoc runtime generation because it is cheaper, easier to debug, and easier to benchmark. In code generation with unit tests and hardware generation with strong external toolchains, disciplined search over a constrained design space is the real bottleneck, not adaptation (Zhang et al., 2025e; Wei et al., 2026).

When any of these conditions breaks—for example, under tool drift, environment shift, or strongly

heterogeneous task structure—the main weakness of static methods becomes clearer. They can optimize a fixed design very effectively, but they cannot determine that a different design is needed for the current input.

When select beats generate, and when edit is unavoidable. Across the dynamic literature, a useful design rule emerges: different forms of runtime adaptation are justified by different kinds of task heterogeneity. If most instances fall within a known motif library and differ mainly in difficulty, communication budget, or required model strength, then runtime selection or pruning is often the right amount of dynamism. Adaptive Graph Pruning (Li et al., 2025), DAGP (Wang and Tong, 2025), Agent-Dropout (Wang et al., 2025f), and Difficulty-Aware Agentic Orchestration (Su et al., 2025a) all suggest that substantial savings can come from activating smaller subgraphs or leaner operator allocations on easier instances, without paying the full complexity cost of graph synthesis. In such cases, selection is attractive not only because it is cheaper, but also because it inherits validity from a fixed super-graph

and is therefore easier to constrain and debug.

Pre-execution generation becomes preferable when task heterogeneity is larger, and the right structure varies in ways that local prompt tuning or subgraph activation cannot absorb. Methods such as Assemble Your Crew (Li et al., 2026), G-Designer (Zhang et al., 2025d), FlowReasoner (Gao et al., 2025a), and Workflow-R1 (Kong et al., 2026) are strongest precisely when different queries call for genuinely different decompositions, communication patterns, or operator sequences rather than merely larger or smaller versions of the same scaffold. The advantage is greater structural expressivity; the cost is a larger search or generation space, greater dependence on workflow validity, and harder structural credit assignment.

In-execution editing becomes necessary when the environment is interactive enough that partial execution reveals information a one-shot plan could not know in advance. DyFlow (Wang et al., 2025c), AgentConductor (Wang et al., 2026c), Aime (Shi et al., 2025), AOrchestra (Ruan et al., 2026), and MetaGen (Wang et al., 2026d) are most compelling in settings where plans must survive tool failures, unexpected intermediate outcomes, or evolving task state, and where structural revision is therefore part of solving the task rather than an optional refinement. Moving from selection to generation to editing increases flexibility but also puts greater pressure on validity checks, stopping policies, and budget controls.

A practical implication is that dynamic methods should be selected based on the minimum structural plasticity required by the workload. Selection is often the right first step when tasks vary mostly in difficulty or communication budget. Generation becomes worthwhile when different tasks require genuinely different workflow structures. Editing is best reserved for environments where runtime observations materially change what the workflow ought to do next. The cost of moving rightward along this spectrum is not just additional computation, but also additional system-level overhead: explicit guards on tokens, tool calls, and wall-clock time; stopping rules tied to confidence or verifier acceptance; fallback behavior for tool failures; and loop-avoidance mechanisms such as repeated-state detection or edit caps. These are not minor engineering details. They are part of the workflow policy itself and should be evaluated and reported as such (Wang et al., 2025c, 2026c; Shi et al., 2025; Ye et al., 2023; Hu et al., 2025b).

When graph optimization matters more than prompt tuning.

Prompt optimization is valuable, but in the surveyed papers, it does not replace missing structure. If a failure is caused by missing verification, redundant communication, insufficient decomposition, or incorrect control flow, graph-level optimization is usually the higher-leverage intervention. Multi-Agent Design (Zhou et al., 2026) and Maestro (Wang et al., 2025a) make this point directly by showing that prompt and topology updates should inform each other rather than compete for credit. A useful practical heuristic is to inspect traces before optimizing prompts: if the error arises because the wrong node executed, the right node never existed, or the information path was wrong, a better prompt is unlikely to be the real fix.

Where verifiers pay off most. Across the surveyed papers, verifiers deliver the largest returns when they are both cheap and semantically meaningful. Unit tests, schema checks, executability checks, and synthesizability checks are especially valuable because they provide dense, actionable feedback on whether the workflow is heading in the right direction (Hu et al., 2025b; Ye et al., 2023; Zheng et al., 2025; Wei et al., 2026). Their value is smaller when they are weak proxies for downstream success or so expensive that they dominate the optimization budget. For example, in a code-generation workflow, running unit tests after a candidate program is produced but before expensive self-debugging or tool-intensive downstream steps can efficiently prune unproductive branches. By contrast, inserting an expensive semantic verifier after every minor retrieval or formatting step may add substantial latency without materially improving final task success. The design question is therefore not merely whether to use a verifier, but where to place it and how often to invoke it.

Baselines, hybrids, and a practical recipe.

Stronger structure should be compared against stronger baselines. Framework papers matter here not only as infrastructure, but also because they make reusable structural abstractions explicit. AutoGen (Wu et al., 2023) provides a general conversation-programming scaffold for multi-agent coordination and tool use, while CAMEL (Li et al., 2023a) offers an early role-conditioned interaction template for autonomous cooperation. OpenHands (Wang et al., 2025b) anchors software-agent workflows in a stable execution loop with realistic tools; related examples include Foam-Agent,

Code2MCP, and ToolRosella (Yue et al., 2026b; Ouyang et al., 2026; Di et al., 2026). A concrete bridge from scaffold design to workflow optimization is given by MacNet (Qian et al., 2025), which elevates DAG-based collaboration to a first-class design object; later topology-optimization work such as AgentConductor (Wang et al., 2026c) can be read as optimizing this design space further by dynamically revising layered topologies under validity, execution, and budget feedback. EvoAgentX (Wang et al., 2025d) provides a complementary example in which the interface between framework and optimizer becomes explicit, integrating workflow generation, execution, and evolutionary refinement within a single open-source platform. At the same time, OneFlow (Xu et al., 2026) offers a useful counterpoint: it argues that some gains attributed to multi-agent workflows can be reproduced by a strong single-agent simulator when roles share the same backbone and can reuse context efficiently. Table 11 summarizes this BACK-GROUND landscape of frameworks and strong baselines.

A practical hybrid recipe follows from the survey:

- Start with a constrained static scaffold or a small operator library, and use node-level compilation or prompt optimization to establish a competent baseline.
- Escalate to graph-level search only when trace analysis reveals structural failure modes rather than purely local instruction errors.
- Under heterogeneous deployment, prefer runtime selection or pruning before full workflow generation.
- Reserve in-execution editing for settings with substantial environmental uncertainty, where partial execution reveals information that cannot be anticipated in a one-shot plan.
- Once a capable design has been found, compress or prune communication through static sparsification or runtime pruning (Zhang et al., 2025c; Li et al., 2025; Wang et al., 2025f).

This recipe is not universal, but it matches a recurring empirical pattern across the surveyed systems.

Structural credit assignment. The hardest methodological problem remains credit assignment for structural decisions. It is still difficult to determine whether a gain came from a new edge, a new verifier, a changed role prompt, or sim-

ply more compute. Better counterfactual replay, ablation-efficient estimators, and critics that operate on traces and graphs together are needed (see the staged ablations and trace-level critics in Zhou et al., 2026; Wang et al., 2025a; Agrawal et al., 2026).

Concrete progress would look like (i) ablation-efficient estimators that isolate the contribution of a single edge or verifier node without rerunning the full workflow; (ii) counterfactual replay benchmarks in which ground-truth structural attributions are known and candidate critics can be scored against them; and (iii) trace-and-graph critics that predict whether a proposed structural edit will help before the edit is executed (Huang et al., 2026; Jiang et al., 2024).

Expressivity versus verifiability. Expressive workflows with loops, dynamic agent creation, and rich conditionals are attractive (Niu et al., 2025; Shi et al., 2025), but they are difficult to compare and hard to validate statically. Constrained IRs improve executability and reproducibility (Zheng et al., 2025; Xu et al., 2025c), but they may exclude the very solutions that make dynamic workflows powerful. Designing representations that balance expressivity and verifiability remains one of the clearest research tensions in this area.

Continual adaptation under tool and environment drift. APIs, websites, and tool registries change. Static templates need repair. Dynamic generators need drift-aware policies. Yet very few papers report adaptation efficiency: how much extra cost is needed to recover performance once the environment shifts. Live registries are beginning to expose this problem, but controlled drift benchmarks would make it much easier to study rigorously (Mo et al., 2025).

Concrete progress would look like (i) drift benchmarks that version tool registries and release scheduled API changes on a known schedule, so that adaptation cost becomes measurable; (ii) metrics that report *recovery budget*—the extra tokens, calls, or rollouts needed to match pre-drift performance—rather than only post-drift accuracy; (iii) policies that separate template repair from dynamic regeneration so the cost of each can be attributed.

Data and benchmark quality for workflow research. Workflow optimization depends unusually heavily on the quality of its evaluators and artifacts. If benchmarks leak answers, if reference

workflows are inconsistent, or if evaluators reward brittle hacks, optimization will exploit those weaknesses. Structure-aware evaluation amplifies the problem because canonicalization, semantic equivalence, and execution validity all interact. Better curation protocols and benchmark diagnostics are therefore not peripheral concerns; they are central to the validity of the field (Zhu et al., 2025).

Toward theory for workflow optimization. The field currently borrows tools from search, program synthesis, reinforcement learning, and bi-level optimization, but lacks a clear theory of when dynamic generation is necessary, when static templates are sufficient, and how sample complexity should scale with structural plasticity (Wei et al., 2025). Even partial answers would be valuable. They could clarify when runtime edits are worth their additional complexity and when simpler fixed scaffolds should be expected to win.

I Supporting Catalog Tables

These catalog tables supplement the core-method comparison by separating three supporting families: node-level optimizers that often operate inside a larger workflow, adjacent methods that change routing or communication without always emitting a full graph artifact, and background frameworks or baselines that anchor workflow experiments.

Table 9: Node-level prompt optimization methods commonly used inside larger workflows

Paper	What is optimized	Feedback signal	Method	Why it matters for workflows
DSPy (Khattab et al., 2024)	module prompts and demonstrations	user-defined task metric	compiler / teleprompting	direct local-node optimization in a fixed pipeline
Large Language Models as Optimizers (OPRO) (Yang et al., 2024)	prompts or instructions	external scalar score	LLM-as-optimizer search	black-box optimizer for workflow-local instructions
EvoPrompt (Guo et al., 2024)	prompt population	development metrics	evolutionary search	node-level instruction search without graph change
CAPO (Zehle et al., 2025)	instructions and few-shot content	quality + token-aware cost	cost-aware evolution	makes node tuning sensitive to deployment cost
GEPA (Agrawal et al., 2026)	prompts inside a larger system	task outcomes + reflection	reflective evolution	example of trace-derived text updating workflow-local prompts

Table 10: ADJACENT methods for selection, pruning, sparsification, and routing

Paper	Setting	What changes at inference time	Signal / method	Why it is adjacent to workflow optimization
Cut the Crap (Zhang et al., 2025c)	offline compression	communication edges and message passing	importance proxies + task quality	post-hoc communication sparsification of a discovered scaffold
Adaptive Graph Pruning (Li et al., 2025)	runtime select	agents and communication edges	joint node/edge pruning	learns sparse task-adaptive subgraphs
AgentDropout (Wang et al., 2025f)	runtime select	redundant agents and links	utility versus token cost	dynamically reduces collaboration overhead
DAGP (Wang and Tong, 2025)	runtime select	difficulty-conditioned edges	difficulty estimation + task reward	allocates denser communication only when needed
DyLAN (Liu et al., 2024)	runtime select	agent team and collaboration pattern	preliminary trials + agent importance	team selection before or during collaboration
MasRouter (Yue et al., 2025)	runtime route	collaboration mode, role allocation, model routing	RL under quality–cost objective	routes across multi-agent designs rather than emitting one graph artifact
SkillOrchestra (Wang et al., 2026b)	runtime route	agent/model/tool choice	trace-derived skill transfer	makes workflow routing more interpretable and reusable

Table 11: BACKGROUND frameworks, scaffolds, and strong baselines that frequently anchor workflow experiments

Paper	Workflow abstraction	Main contribution	Why it matters in this survey
CAMEL (Li et al., 2023a)	role-conditioned multi-agent dialogue	early scalable role-playing template for autonomous cooperation	widely used baseline interaction scaffold
AutoGen (Wu et al., 2023)	conversation programming over agents and tools	generic infrastructure for multi-agent conversations and tools	reusable scaffold later optimized by search or runtime control
EvoAgentX (Wang et al., 2025d)	layered framework for workflow generation, execution, and evolution	open-source platform integrating generation and evolutionary refinement	makes the workflow life cycle explicit in one system
Scaling Large-Language-Model-based Multi-Agent Collaboration (Qian et al., 2025)	DAG-based collaboration network (MacNet)	studies how topology and scale affect performance	motivates topology as a meaningful design variable
OpenHands Software Agent SDK (Wang et al., 2025b)	software-agent execution loop with terminal and repository tools	open platform for SWE-style agents	important scaffold for workflow experiments in software settings
OneFlow (Xu et al., 2026)	single-agent simulator over a role-conditioned graph	shows that many homogeneous multi-agent workflows can be simulated by one LLM	strong baseline against overclaiming multi-agent gains