



HAL
open science

Flows: Building Blocks of Reasoning and Collaborating AI

Martin Josifoski, Lars Klein, Maxime Peyrard, Nicolas Baldwin, Yifei Li, Saibo Geng, Julian Paul Schnitzler, Yuxing Yao, Jiheng Wei, Debjit Paul, et al.

► **To cite this version:**

Martin Josifoski, Lars Klein, Maxime Peyrard, Nicolas Baldwin, Yifei Li, et al.. Flows: Building Blocks of Reasoning and Collaborating AI. 2024. hal-04727425

HAL Id: hal-04727425

<https://hal.science/hal-04727425v1>

Preprint submitted on 9 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Flows: Building Blocks of Reasoning and Collaborating AI

Martin Josifoski^{*1} Lars Klein^{*1} Maxime Peyrard² Nicolas Baldwin¹ Yifei Li^{**1} Saibo Geng^{**1}
Julian Paul Schnitzler¹ Yuxing Yao¹ Jiheng Wei³ Debjit Paul¹ Robert West¹

Abstract

Recent advances in artificial intelligence (AI) have produced highly capable and controllable systems. This creates unprecedented opportunities for structured reasoning as well as collaboration among multiple AI systems and humans. To fully realize this potential, it is essential to develop a principled way of designing and studying such structured interactions. For this purpose, we introduce the conceptual framework *Flows*. Flows are self-contained building blocks of computation, with an isolated state, communicating through a standardized message-based interface. This modular design simplifies the process of creating Flows by allowing them to be recursively composed into arbitrarily nested interactions and is inherently concurrency-friendly. Crucially, any interaction can be implemented using this framework, including prior work on AI–AI and human–AI interactions, prompt engineering schemes, and tool augmentation. We demonstrate the potential of *Flows* on competitive coding, a challenging task on which even GPT-4 struggles. Our results suggest that structured reasoning and collaboration substantially improve generalization, with AI-only Flows adding +21 and human–AI Flows adding +54 absolute points in terms of solve rate. To support rapid and rigorous research, we introduce the aiFlows library embodying *Flows*. The aiFlows library is available at <https://github.com/epfl-dlab/aiflows>. Data and Flows for reproducing our experiments are available at https://github.com/epfl-dlab/cc_flows.

1. Introduction

The success of large language models (LLMs) largely lies in their remarkable emergent ability to adapt to informa-

tion within their context (i.e., prompt) (Brown et al., 2020; Wei et al., 2022; Kojima et al., 2022). By strategically crafting the context, LLMs can be conditioned to perform complex reasoning (Wei et al., 2022; Nye et al., 2021) and effectively utilize external tools (Schick et al., 2023), significantly enhancing their capabilities. Some of the most exciting recent developments involve defining *control flows*, wherein LLMs, with the ability to control a set of tools, are called in an orchestrated fashion to solve increasingly complex tasks. Examples of such control flows include ReAct (Yao et al., 2023b), AutoGPT (Richards, 2023), BabyAGI (Nakajima, 2023), PromptBreeder (Fernando et al., 2023) and FunSearch (Romera-Paredes et al., 2023). Even the ubiquitous ChatGPT (OpenAI, 2023b) application is an instance of a control flow built around the GPT-3.5 and GPT-4 models (Brown et al., 2020; OpenAI, 2023a). However, these represent but a few of the many conceivable control flows, offering only a glimpse into the vast potential of structured LLM interactions. To realize this potential, we need to develop ways to study such interactions systematically.

In software engineering, simple processes can be implemented in an unstructured fashion, perhaps in a single file. However, as the size and complexity of the systems increase, choosing the right abstractions and architecture becomes critical (Garlan & Shaw, 1993). Currently, for structured LLM interactions we want to model, implement, and study, we are at a point where this become unwieldy. Yet, no general efficient abstraction exists for effectively modeling arbitrarily complex structured interactions. Previous work and existing frameworks, such as LangChain (Chase, 2022), Chameleon (Lu et al., 2023), and HuggingGPT (Shen et al., 2023), have converged to an ad-hoc abstraction that models agents as entities that use LLMs to select and execute actions towards specific tasks, where the set of possible actions is pre-defined by the available tools. In this view, tools serve a narrow, well-defined goal and can perform sophisticated tasks (e.g., querying a search engine or executing code). However, their behavior is limited to a single interaction. To highlight the implications of this limitation, consider the following scenario: Alice wants to apply for a job at HappyCorp. If Alice is an agent, she would need to explicitly plan the entire process, including preparing the application, sending it, and evaluating it, which may involve

¹EPFL ²Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG
³PSL University *, ** Equal contribution
Correspondence to: martin.josifoski@epfl.ch, lars.klein@epfl.ch, maxime.peyrard@univ-grenoble-alpes.fr, robert.west@epfl.ch




Figure 1. **Flows framework exemplified.** The first column depicts examples of tools. The second column depicts Atomic Flows constructed from the example tools. The third column depicts examples of Composite Flows defining structured interaction between Atomic or Composite Flows. The fourth column illustrates a specific Composite competitive coding Flow as those used in the experiments. The fifth column outlines the structure of a hypothetical Flow, defining a meta-reasoning process that could support autonomous behavior.²

a background check, organizing interviews, and more. Alice would need the knowledge and the “computational” ability to account for every detail, including unforeseen events that may arise (e.g., the interviewer being on parental leave), and require her to adapt. In reality, most of the complexity is hidden from Alice behind an interface to HappyCorp’s hiring process that might itself be composed of sub-processes involving many other *agents* and *tools*. Therefore, Alice is completely agnostic to the process(es) happening behind the interface and the respective logistics. On the other hand, the hiring process, carefully designed by experts, can be reused by many agents, and its sub-processes can be modified or improved with minimal or no impact on the other components beyond an updated interface. This makes it evident that agents and tools should be able to interact in complex, dynamic or static, ways as parts of nested, modular processes (that run locally or remotely), and the distinction between the two becomes blurred as they both serve as computational units in a complex computational process.

Starting from the observation that all processes are (control) flows defining a potentially complex interaction between many diverse components; we introduce a conceptual framework where Flows are the fundamental building blocks of computation. Flows are independent, self-contained, goal-driven entities able to complete semantically meaningful units of work. To exchange information, Flows communicate via a standardized message-based interface. The framework is depicted in Fig. 1.

The *Flows* abstraction ensures modularity. Alice, a higher-level meta-reasoning Flow that can support autonomous behavior, does not need to know anything beyond how to interface with HappyCorp’s hiring Flow. This substantially reduces complexity (Alice is interacting with a deeply nested, compositional structured interaction through a simple interface) and provides flexibility, allowing sub-Flows to be swapped without consequences as long as they have the same interface. Indeed, HappyCorp’s pre-filtering Flow can be swapped from a rule-based system to an AI model or even a human Flow without affecting the structure of the overall process. The abstraction also enables reusability and the composition of sub-Flows into new Flows for different tasks. Furthermore, the framework shares key design choices with the Actor model, one of the most prominent models of concurrent computation (cf. Sec. 3). Certainly, once Alice submits her application to HappyCorp, she does not need to wait for the response; she can move to her next goal while the other Flows run concurrently.

We showcase the potential of the proposed framework and library by investigating complex collaborative and structured reasoning patterns on the challenging task of competitive coding, a mind sport involving participants trying to solve problems defined by a natural language description.

Contributions. (i) We propose *Flows*, a conceptual framework providing an abstraction that simplifies the design and

²For more details on meta-reasoning Flows see Sec. 7

implementation of arbitrarily nested interactions while enabling concurrency. *Flows* can represent *any* interaction and provides a common framework for reasoning about interaction patterns, specifying hypotheses, and structuring research more broadly. (ii) We open-source the aiFlows library, which embodies *Flows*, together with FlowVerse, which is a repository of *Flows* that can be readily used, extended, and composed into novel, more complex *Flows*. (iii) We leverage *Flows* and the accompanying library to systematically investigate the benefits of complex interactions for solving competitive coding problems and develop AI-only *Flows* adding +21 and human-AI *Flows* adding +54 absolute points in terms of solve rate.

2. Related Work

Existing libraries for modeling structured interactions.

LangChain (Chase, 2022) has become the go-to library for creating applications using large language models. However, most recent works involving structured interaction, such as Cameleon (Lu et al., 2023), Camel (Li et al., 2023), HuggingGPT (Shen et al., 2023), and the concurrent works MetaGPT (Hong et al., 2023) and AutoGen (Wu et al., 2023) all come with their own library. Researchers opt to implement bespoke solutions due to the lack of a general yet efficient abstraction for modeling and designing structured interactions as well as the infrastructure to implement them, that should enable and facilitate open-ended exploration of novel ideas. In this work, we develop such an abstraction, *Flows*, which, in concert with aiFlows, fills this lacuna.

Impact of *Flows*. Crucially, the framework can implement any algorithm and efficiently covers all prior works on AI-AI, human-AI interactions, as well as prompt engineering (cf. Appendix A.3). These works focusing on specific *Flow* instantiations have demonstrated that structured interactions *can* yield performance gains across tasks and models. However, recent results put the universality of previously published results into question (e.g., Huang et al. (2023)) and highlight the necessity for more systematic research. To support these research efforts, we develop the theoretical and practical infrastructure for modeling, implementation, and systematic study structured interactions of arbitrary complexity. We demonstrate the benefits of the proposed infrastructure by conducting experiments that thoroughly investigate multiple core interaction patterns, including Human-AI collaboration, and their combinations, while accounting for data contamination and variance in the results, both of which are, surprisingly, not currently a standard.

Competitive coding (CC). With the advent of transformers, Li et al. (2022) finetuned an LLM on GitHub code repositories, and a dataset scraped from Codeforces. Recently,

Zelikman et al. (2022) proposed decomposing CC problems into function descriptions and, for each function description, using an LLM to generate the implementation in a modular way. While these methods yield promising results, CC remains a challenging task far from being solved (OpenAI, 2023a). As such it presents itself as an ideal test bed for thoroughly studying the benefits of collaborative and structured reasoning interactions.

3. *Flows*

This section introduces *Flows* as a conceptual framework, describes its benefits, and presents the aiFlows library, which embodies the framework.

3.1. *Flows* as a Conceptual Framework

The framework is centered around *Flows* and *messages*. *Flows* represent the fundamental building block of computation. They are independent, self-contained, goal-driven entities able to complete a semantically meaningful unit of work. To exchange information, *Flows* communicate via a standardized message-based interface. Messages can be of any type the recipient *Flow* can process.

We differentiate between two types of *Flows*: Atomic and Composite.³ Atomic *Flows* complete the work directly by leveraging *tools*. Tools can be as simple as a textual sequence specifying a (simple) *Flow*'s fixed response or as complex as a compiler, a search engine, powerful AI systems like LLaMA (Touvron et al., 2023a;b), Stable Diffusion (Rombach et al., 2021), and GPT-4; or even a human. Notably, in the *Flows* framework, AI systems correspond to tools. An Atomic *Flow* is effectively a minimal wrapper around a tool and achieves two things: (i) it fully specifies the tool (e.g., the most basic Atomic *Flow* around GPT-4 would specify the prompts and the generation parameters); and (ii) it abstracts the complexity of the internal computation by exposing only a standard message-based interface for exchanging information with other *Flows*. Examples of Atomic *Flows* include wrappers around chain-of-thought prompted GPT-4 for solving math reasoning problems, few-shot prompted LLaMA for question answering, an existing chatbot, a search engine API, or an interface with a human.

Composite *Flows* accomplish more challenging, higher-level goals by leveraging and coordinating other *Flows*. Crucially, thanks to their local state and standardized interface, Composite *Flows* can readily invoke Atomic *Flows* or other Composite *Flows* as part of compositional, structured interactions of arbitrary complexity. Enabling research on effective patterns of interaction is one of the main goals

³The concept of a *Flow* is sufficient for modeling any interaction. We introduce this distinction as it improves the exposition and simplifies the implementation.

of our work. General examples of such patterns include (i) factorizing the problem into simpler problems (i.e., divide and conquer); (ii) evaluating (sub-)solutions at inference time (i.e., feedback); and (iii) incorporating external information or a tool. Importantly, *Flows* can readily invoke other, potentially heavily optimized, specialized *Flows* to complete specific (sub-)tasks as part of an interaction, leading to complicated behavior. One example of a Composite Flow is ReAct (Yao et al., 2023b). ReAct is a sequential Flow that structures the problem-solving procedure in two steps: a Flow selects the next action out of a predefined set of actions, and another Flow executes it. The two steps are performed until an answer is obtained. Another prominent example, AutoGPT, extends the ReAct Flow with a Memory Flow and an optional Human Feedback Flow. More generally, our framework provides a unified view of prior work, which we make explicit in Appendix A.3.

Importantly, as illustrated in Fig. 1, Composite Flows can script an arbitrarily complex pattern (i) precisely specifying an interaction (e.g., generate code, execute tests, brainstorm potential reasons for failure, etc.); or (ii) defining a high-level, meta-reasoning process in which a Flow could bring about dynamic unconstrained interactions.

Key properties. The proposed framework is characterized by the following key properties:

- Flows are the compositional building blocks of computation.
- Flows encapsulate a local, isolated state.
- Flows interact only via messages.
- Flows’ behaviour depends only on their internal state and the input message.
- Flows can send messages to other Flows and create new Flows.

Connection to the Actor model. *Flows* is fundamentally a framework modeling the computation underlying interactions. As such, it shares key design principles with the Actor model (Hewitt et al., 1973) — a mathematical model of concurrent computation. Similarly to *Flows*, in the Actor model, an Actor is a concurrent computation entity that can communicate with other Actors exclusively through an asynchronous message-passing interface. By encapsulating the state and the computation within individual Actors, the model provides a high-level abstraction for effectively managing and reasoning about complex concurrent and distributed systems, completely avoiding issues associated with shared states, race conditions, and deadlocks. These benefits are similar in nature to those observed in the domain of interactions. The main distinction between the proposed framework and the Actor model lies in their respective communication protocols. Concretely, while the Actor model

prescribes purely asynchronous communication, *Flows* natively supports synchronous communication, which is essential for the implementation of structured reasoning. Interestingly, a similar deviation from the “pure” Actor model can be identified in the implementation of Erlang, a concurrent programming language based on it (Armstrong, 2003). Overall, the shared design choices still make *Flows* inherently concurrency-friendly, from the practical perspective, and suffice for important results from the five decades of extensive studies of the Actor model, such as the fact that every physically possible computation can be directly implemented using Actors (Hewitt, 2010), to transfer to *Flows*.

3.2. Why *Flows*?

Modularity. *Flows* introduces a higher-level abstraction that isolates the state of individual Flows and specifies message-based communication as the only interface through which Flows can interact. This ensures perfect modularity by design.

Reduction of complexity. The framework ensures the complexity of the computation performed by a Flow is fully abstracted behind the universal message-based interface. This enables an intuitive and simple design of arbitrarily complex interactions from basic building blocks.

Systematicity, flexibility, and reusability. The separation of responsibility allows for modules to be developed and studied systematically in isolation or as part of different interactions. Once the correctness and the benefits of a Flow have been established, it can be readily used in developing novel Flows or as a drop-in replacement for less effective Flows leveraged in completing similar goals.

Concurrency. The proposed framework’s design is consistent with the Actor model, one of the most prominent models of concurrent computation. As a consequence, *Flows* can readily support any setting in which Flows run concurrently.

3.3. The aiFlows Library

Accompanying *Flows*, we release the aiFlows library, which embodies the framework. In addition to the inherent benefits that come with the framework, the library comes with the following add-ons: (i) FlowVerse: a repository (to which anyone can contribute) of Flows that can be readily used, extended, or composed into novel, more complex Flows. *Flows* allows for existing “tools” (as well as “models”, “chains”, “agents”, etc.) to be readily incorporated by wrapping them in an Atomic Flow; (ii) a detailed logging infrastructure enabling transparent debugging, analysis, and research in optimizing (i.e., learning or fine-tuning) Flows.

4. Competitive Coding Flows

This work investigates the potential of structured interactions for solving competitive coding (CC) problems. In CC, given a natural language description and a few input–output examples, the task is to generate code that will produce the expected output for all of the hidden input–output test cases associated with the problem. Fig. 4 provides examples.

We focus the analysis on three canonical dimensions of interactions: (i) problem decomposition as structured reasoning; (ii) human-AI collaboration; and (iii) refinement with various feedback types. By providing a common language for clearly specifying interactions as well as the capability to flexibly compose, exchange, and extend them, the framework makes it possible to study the space of complex interactions in a principled fashion. In the rest of the section, we describe the specific Flows used in the experiments, depicted in Fig. 2.

Problem decomposition. Planning has been an integral intermediate step in recent work (Lu et al., 2023; Shen et al., 2023; Yao et al., 2023b). Similar decomposition is natural in the context of CC as well. In particular, we approach the task in two steps: generating a solution strategy by a Plan Flow and then generating the corresponding code by a Code Flow. This is depicted by panel A in Fig. 2.

Human-AI collaboration. When designing human-AI collaborations, it is essential to take the costs of human interaction into account (Horvitz, 1999; Amershi et al., 2019; Mozannar et al., 2023). By providing immense flexibility, *Flows* can support research in the design of interactions involving humans as computational building blocks in a way that maximizes the utility of the overall computation with a minimal human effort. In the context of CC, we hypothesize that a human can be effectively incorporated at the plan level to provide a short “oracle” plan in natural language. We operationalize this by an (Atomic) Human Flow, illustrated in Panel B of Fig. 2 as the *Oracle Plan Flow*.

Refinement with various feedback types. Iterative refinement is a general problem-solving strategy successfully deployed across various disciplines (Perrakis et al., 1999; Reid & Neubig, 2022; Schick et al., 2022; Saharia et al., 2021). The strategy revolves around the idea that a solution can be gradually improved through a mechanism for analysis, modification, and re-evaluation. The design of this “*feedback*” mechanism is critical for the effectiveness of the problem-solving strategy. The conceptual framework, paired with the accompanying library, provides the infrastructure to support the design, implementation, and principled research of effective refinement strategies and feedback mechanisms. In this work, we consider a canonical iterative refinement setup where a *generator* Flow is tasked with generating the solu-

tion, and a *critic* Flow provides feedback on the proposed solution. We consider two feedback types in the context of both the Plan and the Code Flow: (i) Reflection Flow: the feedback consists of a fixed message encouraging the model to reflect on important aspects of the proposed solution; (ii) Collaboration Flow: the feedback is provided by an AI system that “evaluates” the proposed solution. Furthermore, we explore two more code-specific feedback types: (i) Debug Flow: the feedback message corresponds to the results from executing the code and testing it against the examples provided in the problem description; (ii) Debug–Collab Flow: the feedback is provided by an AI system with access to the code testing results, effectively, grounding the feedback and allowing more systematic reasoning about the potential causes of failure.

We refer to Flows using the following convention: *CodeFlowName* when no plan is generated and *PlanFlowName-CodeFlowName* otherwise.

5. Experimental Setup

Data. We scrape publicly available problems from one of the most popular websites hosting CC contests, Codeforces (Mirzayanov, 2023), and LeetCode (LeetCode, 2023), which cover a broad spectrum of problems ranging from easy interview questions to hard CC problems (see Appendix A.1 for more details). The datasets cover problems from 2020-August-21 to 2023-March-26 for CodeForces, and from 2013-October-25 to 2023-April-09 for LeetCode. Importantly, to study the effect of structured interactions (i.e., different Flows) in a principled manner, it is crucial to account for the possibility of *data contamination*, i.e., that some of the test data has been seen during training (Magar & Schwartz, 2022). Containing problems published over an extended period up to a few months ago (at the time of writing), our datasets allow for reliable identification of the training data cutoff date that can help with addressing this issue. Prior code evaluation datasets like APPS (Hendrycks et al., 2021), HumanEval (Chen et al., 2021), and CodeContests (Li et al., 2022) lack problem release dates, and considering the lack of publicly available information about LLMs’ training data, can likely lead to confounded evaluation of models’ memorization and generalization abilities.

Code testing and solution evaluation. Just like a human participant, the Debug Flow has access only to the input–output example pairs contained in the problem description and, at inference time, uses a local code testing infrastructure to evaluate (intermediate) solution candidates. Crucially, these examples cover only a few simple cases, and generating outputs consistent with them does not imply the code corresponds to a correct solution. A solution is considered correct if it passes all the hidden test cases. To determine correctness, we leverage online evaluators that

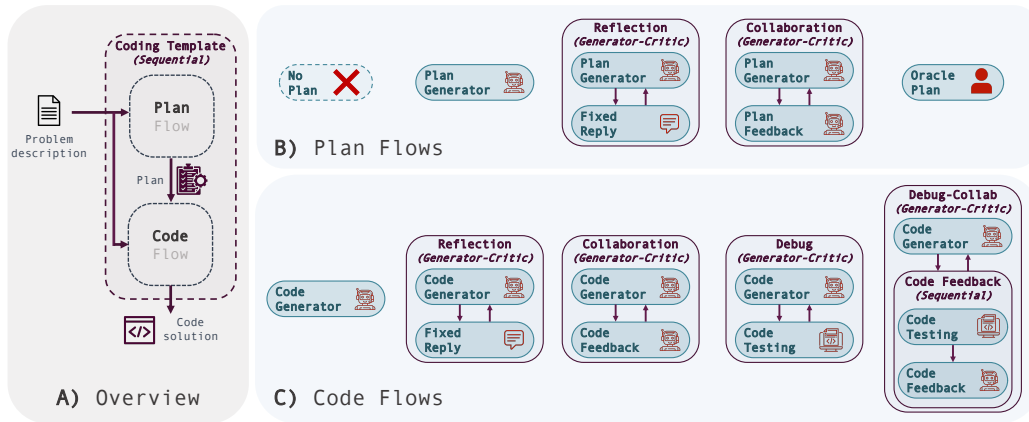


Figure 2. **Competitive coding Flows.** At the highest level, we consider planning as a specific structured reasoning pattern for problem decomposition. In particular, the Plan Flow generates a solution strategy and passes it to the Code Flow, which implements it, as depicted in A). B) and C) depict the different choices of sub-Flows used as Plan and Code Flows in the experiments. Notably, we explore the impact of human-AI collaboration at the plan level and refinement with different types of *feedback*: i) fixed reply encouraging reflection; ii) AI generated feedback; iii) code testing results as feedback; iv) AI generated feedback grounded in code testing results.

submit candidate solutions to the websites’ online judges, ensuring authoritative results. For many of the Codeforces problems, we also support local evaluation based on a comprehensive set of hidden test cases we managed to scrape. For more details, see Appendix A.2.

Models and Flows. We experiment with the competitive coding Flows described in Sec. 4, and GPT-4 (OpenAI, 2023a) as the LLM tool of choice. See Appendix A.4 for the specific prompts. Also, the code to reproduce the experiments in the paper is available in the project’s GitHub repository.

Evaluation metrics. The most common evaluation metric for code generation is $pass@k$, corresponding to the probability that in a set of k sampled candidates, there will be at least one correct solution (Chen et al., 2021). To better align with practical use cases, we focus on $pass@1$, i.e. the solve rate when averaged across the problem set. We report a point estimate and a 95% confidence interval constructed from 1000 bootstrap resamples.

Compute and cost. All the experiments, including the most complex Flows, can be performed on commodity hardware relatively cheaply. For instance, the costs associated with querying the OpenAI API for generating Table 1 amount to \$1000.

6. Experimental Results

We first study the generalization ability of representative Flows and empirically identify GPT-4’s knowledge-cutoff date. Next, we perform a focused analysis along the dimensions described in Sec. 4.

6.1. Performance of Coding Flows on Pre- vs. Post-Knowledge-Cutoff-Date Data




Figure 3. **Temporal analysis.** Performance is averaged over a sliding window of two months. The substantial drop in performance around the reported knowledge cutoff date for GPT-3/4 (the crimson vertical line) reveals limited generalization ability that can be alleviated through structured interactions.

In this experiment, we consider three representative Flows: (i) Code: the simplest Code Generator Flow corresponding to a single GPT-4 API call; (ii) Code_Debug_Collab: the most complex code Flow; (iii) Plan_Oracle-Code_Debug_Collab: the most complex code Flow with human guidance at the plan level. We perform the analysis by running the three Flows on Codeforces problems released from October 2020 to April 2023 and averaging the performance over a sliding window of two months. The results are reported in Fig. 3.

We observe a substantial drop in performance centered around September 2021, consistent with the knowledge

Table 1. **Main Results.** Performance of competitive coding Flows on Codeforces and LeetCode, with direct inference (Code) as baseline.

	Codeforces		Leetcode					
	Pre-cutoff	Post-cutoff	Easy	Pre-cutoff Medium	Hard	Easy	Post-cutoff Medium	Hard
Code	71.8 ±11.0	26.9 ±11.0	97.8 ±3.1	93.4 ±5.4	66.7 ±10.9	76.3 ±8.6	25.1 ±8.9	8.0 ±5.5
Code_Reflection	+9.3 ±9.7	+0.0 ±10.6	+0.0 ±3.1	+0.0 ±5.4	+1.2 ±10.6	+0.9 ±8.1	+5.4 ±9.4	+3.5 ±6.6
Code_Collaboration	+4.8 ±10.5	+9.6 ±11.8	+0.0 ±3.1	-2.3 ±6.0	-0.1 ±10.9	-3.2 ±8.7	+0.0 ±8.7	+1.2 ±5.9
Code_Debug	+12.7 ±8.6	+7.9 ±11.6	+0.0 ±3.1	+1.1 ±5.0	+6.9 ±10.0	+7.7 ±7.3	+7.7 ±9.6	+2.4 ±6.3
Code_Debug_Collab	+12.6 ±8.9	+20.6 ±12.1	+0.0 ±3.1	+0.0 ±5.4	+5.5 ±10.4	+7.5 ±7.4	+9.8 ±9.7	+1.2 ±6.0
Plan-Code	-1.6 ±11.0	+8.0 ±11.6	-3.1 ±4.5	-2.3 ±5.9	-9.7 ±11.2	+2.3 ±8.3	+3.2 ±9.1	-3.4 ±4.3
Plan_Reflection-Code	-3.3 ±11.6	+4.8 ±11.6	-2.1 ±4.1	-4.5 ±6.6	-3.1 ±10.7	+1.2 ±8.3	-3.3 ±8.5	+0.0 ±5.5
Plan_Collaboration-Code	-4.8 ±11.5	+6.3 ±11.4	-1.1 ±3.7	-2.3 ±6.1	-7.2 ±11.2	-2.0 ±8.6	+0.1 ±9.0	+1.2 ±5.8
Plan_Oracle-Code	+11.0 ±9.4	+47.6 ±10.7	-	-	-	-	-	-
Plan_Oracle-Code_ Debug_Collab	+23.0 ±5.2	+53.9 ±9.5	-	-	-	-	-	-

cutoff date reported by OpenAI, and denote it by a vertical line on the plot. With Codeforces problems appearing in contexts outside of the contest itself (e.g., editorials), it is reasonable to assume the model has been exposed to older problems more frequently during training. This would explain why the drop spans multiple months, from May 2021 to November 2021, depending on when which data was published and crawled.

Notably, there is a stark difference in the performance of the Code Flow on problems published before and after the knowledge cutoff data, with the solve rate decreasing from around 80% to 23%. While still experiencing a substantial performance drop, the Code_Debug_Collab Flow doubles the solve rate on novel problems to around 45%. Provided with human input at the plan level, the same Flow reaches 85%. Overall, this highlights that GPT-4 performs poorly on novel complex reasoning problems, but structured interactions have the potential to enhance its generalization capabilities. As both GPT-4 (i.e., the Code Flow) and the more complex interactions (Flows) exhibit qualitatively different behavior on novel data, to draw accurate conclusions, it is critical that data contamination is taken into serious consideration when designing experiments and interpreting results.

6.2. Comparing Competitive Coding Flows

Table 1 reports the performance of the systematically chosen set of Flows described in Sec. 4. Rows 6–10 correspond to Flows comprising planning and coding, while rows 1–5 perform the coding directly. In line with the findings of the previous section, we separately consider the performance on problems published before and after the knowledge cutoff date of September 2021.

Problem decomposition. The idea behind planning before implementing the solution is to decouple the high-level reasoning from the code implementation. To analyze the effectiveness of this pattern, we compare the Code and the Plan-Code Flow. Looking at the point estimates, in the pre-cutoff problems, introducing the plan Flow leads to decreased performance (-1.6 for Codeforces and -3.1/2.3/-9.7 for LeetCode easy/medium/hard). However, in the post-cutoff problems, incorporating a plan Flow leads to gains for Codeforces (+8) and LeetCode easy and medium (+2.3 and +3.2). While these trends are consistent, considering the confidence intervals, we see that they are not statistically significant. Crucially, these results do not imply that this specific problem decomposition is not valuable as it creates a lot of potential in designing an effective human-AI collaboration.

Human-AI collaboration. After every contest, the Codeforces community publishes an editorial that, in addition to the code implementation, provides a short natural language description of the solution. To simulate a Flow where a human provides high-level guidance at the core of the reasoning process, we scrape the solution descriptions and pass them as human-generated plans. The results are striking: despite being only a few sentences long, human-provided plans lead to a substantial performance increase (from 26.9% to 74.5% and from 47.5% to 80.8% on novel problems, when the code is generated by Code and Code_Debug_Collab Flows, respectively). First and foremost, these results showcase the opportunities created by Flows for designing, implementing, and studying Human-AI collaboration as a key component of structured interactions. Second, specific to the problem of competitive coding, they validate the hypothesis that high-quality plans are important, suggesting that the design of more effective plan Flows

is a promising direction to explore in the future. Last but not least, the results highlight the necessity of more systematic research, as patterns seemingly not valuable in one Flow, such as the simple plan-code structured reasoning problem decomposition, can provide immense value as part of another Flow.

Refinement with various feedback types. We find that Code_Reflection and Code_Collaboration lead to limited improvements among the code Flows. The two exceptions are Codeforces pre-cutoff (+9.3) for the former and Codeforces post-cutoff (+9.6) for the latter pattern. While close, these results are not statistically significant. On the other hand, the Flows providing grounded feedback, Code_Debug and Code_Debug_Collab, lead to consistent and statistically significant improvements, most notable on the novel Codeforces problems where performance increases from 26.9, without feedback, to 47.5, when the refinement is based on AI-generated feedback grounded in tests. On LeetCode, these improvements are smaller in magnitude. We suspect this is a consequence of the examples provided with the problem description being more simple than those in Codeforces, leading to false positives and, thereby, incorrect grounding, affecting the feedback quality. This could be addressed by generating additional tests with a Test_Case_Generator Flow, a direction we leave for future work to explore. Finally, in the plan Flows, where we consider Reflection and Collaboration (without grounding), we find that refinement does not provide statistically significant benefits.

Overall, our findings provide several important insights: (i) the direct benefit of problem decomposition hinges on the quality of the intermediate steps; (ii) involving humans at the core high-level reasoning process yields major improvements as humans can easily provide high-quality, grounded feedback; (iii) strategic problem decomposition is a powerful strategy for creating opportunities for effective Human-AI collaboration; (iv) the effectiveness of refinement patterns is not universal and depends on the quality of the starting solution and the feedback (e.g., the level of grounding), and the model’s ability to incorporate that feedback modulated through the feedback’s specificity and the model’s capabilities. This analysis paints a more complex picture than what is reported by prior work for simple interactions.

7. Discussion

Simplicity and systematicity. Thanks to its key properties, *Flows*, together with aiFlows, provides an infrastructure that greatly simplifies the design and implementation of open-ended interactions, with a capability to flexibly isolate, compose, replace, or modify sub-Flows. The experiments demonstrate that carefully designed interactions can substantially improve generalization. However, they also reveal

that the effectiveness of particular interaction patterns is not universal; instead, there are many factors at play. As researchers, we need to clearly specify the patterns we are studying, clearly communicate our hypotheses, and study them both in isolation and as parts of other interactions across different datasets or/and tasks. Furthermore, it is critical that data contamination is taken into serious consideration when designing experiments and drawing conclusions, and error bars become a standard in the field.

Cost and performance optimization. In our experiments, we used “off-the-shelf” LLMs that have not been specifically optimized for collaboration. Performance (and compute costs) can be substantially improved by fine-tuning models to collaborate more effectively, generally or toward specialized roles (e.g., controller or critic). Learning requires data, and to support research in this direction, aiFlows implements detailed logging mechanisms of Flow runs.

Meta-reasoning Flows and asynchronous execution. Cognitive science research in metacognition and meta-reasoning suggests the existence of meta-level monitoring and control processes underlying cognition (Ackerman & Thompson, 2017). Since *Flows* supports asynchronous execution of sub-Flows, it makes it possible to achieve similar asynchronous meta-cognition for autonomous AI systems moving beyond a single LLM call serving as a controller (Nakajima, 2023; Richards, 2023). For example, distributed and asynchronous execution of Flows such as FunSearch (Romera-Paredes et al., 2023) is naturally supported by *Flows*.

8. Conclusion

In this paper, we propose *Flows*, an abstraction that, in concert with the accompanying library aiFlows, provides the theoretical and practical infrastructure with a modular and concurrency-friendly design, which enables and facilitates the modeling, implementation, and systematic study of arbitrarily complex structured interactions. We thoroughly investigate multiple core interaction patterns, including Human-AI collaboration, and their combinations, while accounting for data contamination and the variance in the results. The investigation shows that the developed AI-only Flows add +21 and human-AI Flows add +54 absolute points in terms of solve rate, and highlights the effect of data contamination, variance, and non-universality of results. Overall, our experiments establish the potential of *Flows*, the necessity of more systematic research, and the value brought by *Flows* and aiFlows in support of these research efforts. On the one hand, *Flows* provides a high-level abstraction enabling the design and implementation of interactions of arbitrary complexity. On the other, it offers a common framework for reasoning about interaction patterns, specifying hypotheses, and structuring research. We hope

the framework will serve as a solid basis for practical and theoretical innovations, paving the way toward ever more useful AI, similar to the Actor model’s role for concurrent and distributed systems.

References

- Ackerman, R. and Thompson, V. A. Meta-reasoning: Monitoring and control of thinking and reasoning. *Trends in Cognitive Sciences*, 21(8):607–617, 2017. ISSN 1364-6613. doi: <https://doi.org/10.1016/j.tics.2017.05.004>. URL <https://www.sciencedirect.com/science/article/pii/S1364661317301055>.
- Amershi, S., Weld, D. S., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S. T., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., and Horvitz, E. Guidelines for human-ai interaction. In Brewster, S. A., Fitzpatrick, G., Cox, A. L., and Kostakos, V. (eds.), *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI 2019, Glasgow, Scotland, UK, May 04-09, 2019*, pp. 3. ACM, 2019. doi: 10.1145/3290605.3300233. URL <https://doi.org/10.1145/3290605.3300233>.
- Armstrong, J. *Making reliable distributed systems in the presence of software errors*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2003. URL <https://nbn-resolving.org/urn:nbn:se:kth:diva-3658>.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf.
- Chase, H. Langchain. <https://github.com/hwchase17/langchain>, 2022.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., Ponde, H., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G., Ray, A., Puri, R., Krueger, G., Petrov, M., Khlaaf, H., Sastry, G., Mishkin, P., Chan, B., Gray, S., Ryder, N., Pavlov, M., Power, A., Kaiser, L., Bavarian, M., Winter, C., Tillet, P., Such, F. P., Cummings, D. W., Plappert, M., Chantzis, F., Barnes, E., Herbert-Voss, A., Guss, W. H., Nichol, A., Babuschkin, I., Balaji, S. A., Jain, S., Carr, A., Leike, J., Achiam, J., Misra, V., Morikawa, E., Radford, A., Knight, M. M., Brundage, M., Murati, M., Mayer, K., Welinder, P., McGrew, B., Amodei, D., McCandlish, S., Sutskever, I., and Zaremba, W. Evaluating large language models trained on code. *ArXiv*, abs/2107.03374, 2021.
- Chen, W., Ma, X., Wang, X., and Cohen, W. W. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *ArXiv*, abs/2211.12588, 2022.
- Chen, X., Lin, M., Schärli, N., and Zhou, D. Teaching large language models to self-debug. *ArXiv*, abs/2304.05128, 2023.
- Fernando, C., Banarse, D., Michalewski, H., Osindero, S., and Rocktäschel, T. Promptbreeder: Self-referential self-improvement via prompt evolution. *CoRR*, abs/2309.16797, 2023. doi: 10.48550/ARXIV.2309.16797. URL <https://doi.org/10.48550/arXiv.2309.16797>.
- Garlan, D. and Shaw, M. An introduction to software architecture. In Ambriola, V. and Tortora, G. (eds.), *Advances in Software Engineering and Knowledge Engineering*, volume 2 of *Series on Software Engineering and Knowledge Engineering*, pp. 1–39. World Scientific, 1993. doi: 10.1142/9789812798039_0001. URL https://doi.org/10.1142/9789812798039_0001.
- Hendrycks, D., Basart, S., Kadavath, S., Mazeika, M., Arora, A., Guo, E., Burns, C., Puranik, S., He, H., Song, D., and Steinhardt, J. Measuring coding challenge competence with apps. *NeurIPS*, 2021.
- Hewitt, C. E. Actor model of computation: Scalable robust information systems. *arXiv: Programming Languages*, 2010.
- Hewitt, C. E., Bishop, P. B., and Steiger, R. A universal modular actor formalism for artificial intelligence. In *International Joint Conference on Artificial Intelligence*, 1973.
- Hong, S., Zheng, X., Chen, J., Cheng, Y., Wang, J., Zhang, C., Wang, Z., Yau, S. K. S., Lin, Z., Zhou, L., Ran, C., Xiao, L., and Wu, C. Metagpt: Meta programming for multi-agent collaborative framework. *CoRR*, abs/2308.00352, 2023. doi: 10.48550/ARXIV.2308.00352. URL <https://doi.org/10.48550/arXiv.2308.00352>.
- Horvitz, E. Principles of mixed-initiative user interfaces. In Williams, M. G. and Altom, M. W. (eds.), *Proceeding of the CHI ’99 Conference on Human Factors in Computing Systems: The CHI is the Limit, Pittsburgh, PA, USA, May 15-20, 1999*, pp. 159–166. ACM, 1999. doi:

- 10.1145/302979.303030. URL <https://doi.org/10.1145/302979.303030>.
- Huang, J., Chen, X., Mishra, S., Zheng, H. S., Yu, A. W., Song, X., and Zhou, D. Large language models cannot self-correct reasoning yet. *CoRR*, abs/2310.01798, 2023. doi: 10.48550/ARXIV.2310.01798. URL <https://doi.org/10.48550/arXiv.2310.01798>.
- Kim, G., Baldi, P., and McAleer, S. Language models can solve computer tasks. *ArXiv*, abs/2303.17491, 2023.
- Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., and Iwasawa, Y. Large language models are zero-shot reasoners. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A. (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 22199–22213. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/8bb0d291acd4ac06ef112099c16f326-Paper-Conference.pdf.
- LeetCode. Leetcode.com, 2023. URL <https://leetcode.com>.
- Li, G., Hammoud, H. A. A. K., Itani, H., Khizbullin, D., and Ghanem, B. Camel: Communicative agents for "mind" exploration of large scale language model society. *arXiv preprint arXiv:2303.17760*, 2023.
- Li, Y., Choi, D., Chung, J., Kushman, N., Schrittwieser, J., Leblond, R., Eccles, T., Keeling, J., Gimeno, F., Dal Lago, A., et al. Competition-level code generation with alpha-code. *Science*, 378(6624):1092–1097, 2022.
- Lu, P., Peng, B., Cheng, H., Galley, M., Chang, K.-W., Wu, Y. N., Zhu, S.-C., and Gao, J. Chameleon: Plug-and-play compositional reasoning with large language models. *ArXiv*, abs/2304.09842, 2023.
- Madaan, A., Tandon, N., Gupta, P., Hallinan, S., Gao, L., Wiegrefe, S., Alon, U., Dziri, N., Prabhunoye, S., Yang, Y., et al. Self-refine: Iterative refinement with self-feedback. *arXiv preprint arXiv:2303.17651*, 2023.
- Magar, I. and Schwartz, R. Data contamination: From memorization to exploitation. In Muresan, S., Nakov, P., and Villavicencio, A. (eds.), *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pp. 157–165. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.acl-short.18. URL <https://doi.org/10.18653/v1/2022.acl-short.18>.
- Mirzayanov, M. Codeforces.com, 2023. URL <https://codeforces.com>.
- Mozannar, H., Bansal, G., Fournay, A., and Horvitz, E. When to show a suggestion? integrating human feedback in ai-assisted programming. *CoRR*, abs/2306.04930, 2023. doi: 10.48550/arXiv.2306.04930. URL <https://doi.org/10.48550/arXiv.2306.04930>.
- Nakajima, Y. Babyagi. <https://github.com/yoheinakajima/babyagi>, 2023.
- Nye, M. I., Andreassen, A. J., Gur-Ari, G., Michalewski, H., Austin, J., Bieber, D., Dohan, D., Lewkowycz, A., Bosma, M., Luan, D., Sutton, C., and Odena, A. Show your work: Scratchpads for intermediate computation with language models. *CoRR*, abs/2112.00114, 2021. URL <https://arxiv.org/abs/2112.00114>.
- OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023a.
- OpenAI. ChatGPT. <https://openai.com/chatgpt>, 2023b. Accessed: 2024-01-28.
- Paul, D., Ismayilzada, M., Peyrard, M., Borges, B., Bosse-lut, A., West, R., and Faltings, B. Refiner: Reasoning feedback on intermediate representations. *arXiv preprint arXiv:2304.01904*, 2023.
- Perrakis, A., Morris, R. J., and Lamzin, V. S. Automated protein model building combined with iterative structure refinement. *Nature Structural Biology*, 6:458–463, 1999. URL <https://api.semanticscholar.org/CorpusID:20292852>.
- Reid, M. and Neubig, G. Learning to model editing processes. In *Conference on Empirical Methods in Natural Language Processing*, 2022. URL <https://api.semanticscholar.org/CorpusID:249062636>.
- Richards, T. B. Autogpt. <https://github.com/Significant-Gravitas/Auto-GPT>, 2023.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. *CoRR*, abs/2112.10752, 2021. URL <https://arxiv.org/abs/2112.10752>.
- Romera-Paredes, B., Barekatin, M., Novikov, A., Balog, M., Kumar, M. P., Dupont, E., Ruiz, F. J., Ellenberg, J. S., Wang, P., Fawzi, O., et al. Mathematical discoveries from program search with large language models. *Nature*, pp. 1–3, 2023.
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D. J., and Norouzi, M. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45:4713–4726, 2021. URL <https://api.semanticscholar.org/CorpusID:233241040>.

- Schick, T., Dwivedi-Yu, J., Jiang, Z., Petroni, F., Lewis, P., Izacard, G., You, Q., Nalmpantis, C., Grave, E., and Riedel, S. Peer: A collaborative language model. *ArXiv*, abs/2208.11663, 2022. URL <https://api.semanticscholar.org/CorpusID:251765117>.
- Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., and Scialom, T. Toolformer: Language models can teach themselves to use tools. *ArXiv*, abs/2302.04761, 2023.
- Shen, Y., Song, K., Tan, X., Li, D. S., Lu, W., and Zhuang, Y. T. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *ArXiv*, abs/2303.17580, 2023.
- Shinn, N., Cassano, F., Labash, B., Gopinath, A., Narasimhan, K., and Yao, S. Reflexion: Language agents with verbal reinforcement learning. *arXiv preprint arXiv:2303.11366*, 2023.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. Llama: Open and efficient foundation language models. *ArXiv*, abs/2302.13971, 2023a.
- Touvron, H., Martin, L., Stone, K. R., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D. M., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A. S., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I. M., Korenev, A. V., Koura, P. S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T. Llama 2: Open foundation and fine-tuned chat models. *arXiv*, 2023b.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35: 24824–24837, 2022.
- Welleck, S., Lu, X., West, P., Brahman, F., Shen, T., Khashabi, D., and Choi, Y. Generating sequences by learning to self-correct. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=hH36JeQZDa0>.
- Wu, Q., Bansal, G., Zhang, J., Wu, Y., Zhang, S., Zhu, E., Li, B., Jiang, L., Zhang, X., and Wang, C. Autogen: Enabling next-gen LLM applications via multi-agent conversation framework. *CoRR*, abs/2308.08155, 2023. doi: 10.48550/ARXIV.2308.08155. URL <https://doi.org/10.48550/arXiv.2308.08155>.
- Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., and Narasimhan, K. Tree of thoughts: Deliberate problem solving with large language models. *ArXiv*, abs/2305.10601, 2023a.
- Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K. R., and Cao, Y. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*, 2023b. URL https://openreview.net/forum?id=WE_vluYUL-X.
- Yoran, O., Wolfson, T., Bogin, B., Katz, U., Deutch, D., and Berant, J. Answering questions by meta-reasoning over multiple chains of thought. *ArXiv*, abs/2304.13007, 2023.
- Zelikman, E., Huang, Q., Poesia, G., Goodman, N. D., and Haber, N. Parsel: A (de-)compositional framework for algorithmic reasoning with language models, 2022. URL <https://arxiv.org/abs/2212.10561>.
- Zhang, Z., Zhang, A., Li, M., Zhao, H., Karypis, G., and Smola, A. J. Multimodal chain-of-thought reasoning in language models. *ArXiv*, abs/2302.00923, 2023.

A. Appendix

A.1. Data

Example Codeforces and LeetCode problems are provided in Fig. 4.

In the first experiment, the temporal analysis, we use 239 Codeforces problems ranging from October 2020 to April 2023. In the second experiment, we have 136 problems for Codeforces (some problems are dropped in order to keep the pre-cutoff and post-cutoff buckets equal to 68) and 558 problems for LeetCode (93 for each of the six buckets). Additionally, to support research in the area, we set up an AI competitive coding challenge based on a dataset of Codeforces problems of various difficulties published after the knowledge cutoff date. More details about the CC competition are available in Appendix A.5.

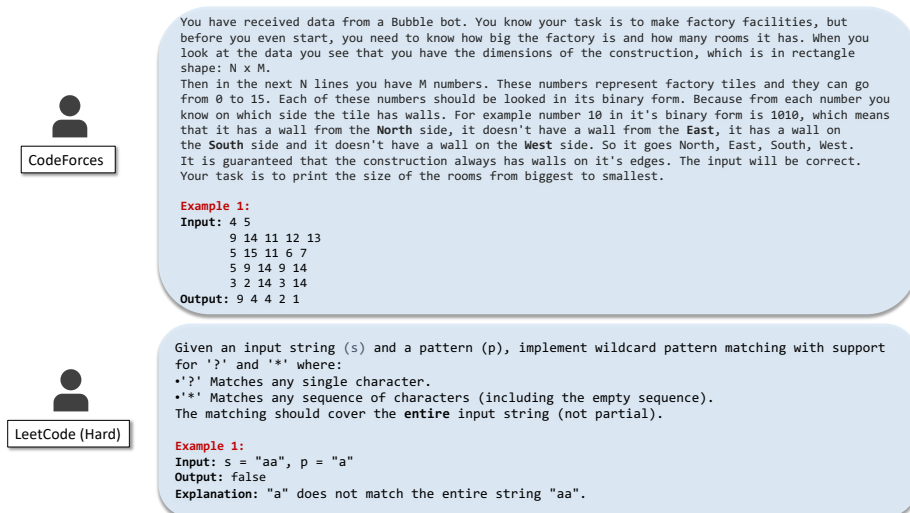


Figure 4. Examples of competitive coding problems from Codeforces and LeetCode.

A.2. Code Testing and Solution Evaluation

The solution evaluation requires a set of input–output pairs, hidden from the user, that comprehensively test the behavior of the program. To compute the final results, we have implemented an online evaluation infrastructure that submits the candidate solutions to the websites' online judges and automatically scrapes the judgment. This mechanism ensures authoritative results.

For many of the Codeforces problems, we managed to scrape (sometimes a subset) of the hidden tests, allowing us to use a faster, local infrastructure for evaluating candidate solutions. On the other hand, LeetCode does not expose any of the hidden tests publicly.

For code testing at inference time, just like a human would, we rely on tests constructed from the (public) input–output example pairs contained in the problem description.

A.3. Concurrent and Previous Works as Specific Instances of Flows

The introduction of LLMs such as BARD, GPT-3, ChatGPT, and its latest version, GPT-4, has led to a breakthrough in AI. This has enabled many exciting developments like CoT, HuggingGPT, AutoGPT, AgentGPT, and BabyAGI. In this section, we demonstrate how *Flows* provides a unified view encompassing concurrent and previous work as specific Flow instances. The details are provided in Figure 5 and Table. 2.

1. **Few shot Prompting (FS)** (Brown et al., 2020) consists in providing a few input-output examples within the prompt, acting as demonstrations to enable the LLM to perform a specific task. This technique relies on the LLM's emergent in-context learning ability to extrapolate from these limited examples and infer how to solve the task in general.




Figure 5. Previous works are specific Flows. We depict a selected subset of previous works incorporating structured reasoning and/or interactions between AI agents, tools, and humans, through the lens of the Flows framework. This demonstrates that Flows is a powerful language for describing, conceptualizing, and disseminating structured interaction patterns.

2. **Chain of Thoughts (CoT)** (Wei et al., 2022) is a prompting method (atomic Flow) that allows LLMs to generate a series of intermediate natural language reasoning steps that lead to the final output.
3. **Tree of Thoughts (ToT)** (Yao et al., 2023a) is a framework that enables (*orchestration*) exploration over coherent units of text (thoughts) that serve as intermediate steps toward problem-solving. ToT allows LLMs to perform deliberate decision-making by considering multiple different reasoning paths and self-evaluating choices to decide the next course of action, as well as looking ahead or backtracking when necessary to make global choices.
4. **Program of Thoughts (PoT)** (Chen et al., 2022) is a prompting method that allows language models (mainly Codex) to express the reasoning process as a program. The computation is relegated to an external program, which executes the generated programs to derive the answer.
5. **Multimodal CoT (M-CoT)** (Zhang et al., 2023) is a method that incorporates language (text) and vision (images) modalities into a two-stage framework that separates rationale generation and answer inference. To facilitate the interaction between modalities in M-CoT, smaller language models (LMs) are fine-tuned by fusing multimodal features.
6. **ToolFormer** (Schick et al., 2023) is a model that is trained to decide which APIs to call, when to call them, what arguments to pass, and how to incorporate the results into future tokens prediction.
7. **ReAct** (Yao et al., 2023b) is a framework that uses LLMs to generate reasoning traces and task-specific actions sequentially. The framework allows for greater synergy between the two: reasoning traces help the model induce, track, and update action plans and handle exceptions, while actions allow it to interface with external sources, such as knowledge bases or environments, to gather additional information.
8. **Parsel** (Zelikman et al., 2022) is a framework that enables the automatic implementation and validation of complex algorithms with code LLMs. The framework first synthesizes an intermediate representation based on the Parsel language and can then apply a variety of postprocessing tools. Code is generated in a next step.
9. **REFINER** (Paul et al., 2023) is a framework for LMs to explicitly generate intermediate reasoning steps while interacting with a critic model that provides automated feedback on the reasoning.

10. **Self-Refine** (Madaan et al., 2023) is a framework for LLMs to generate coherent outputs. The main idea is that an LLM will initially generate an output while the same LLM provides feedback for its output and uses it to refine itself iteratively.
11. **Recursively Criticize and Improve** (RCI) (Kim et al., 2023) showed that a pre-trained large language model (LLM) agent could execute computer tasks guided by natural language using a simple prompting scheme where the agent Recursively Criticizes and Improves its output (RCI). Unlike Self-refine, this method uses two separate LLMs (ChatGPT), one for performing the task and another for criticizing.
12. **Self-Correct** (Welleck et al., 2023) is a framework that decouples a flawed base generator (an LLM) from a separate corrector that learns to iteratively correct imperfect generations. The imperfect base generator can be an off-the-self LLM or a supervised model, and the corrector model is trained.
13. **Self-Debug** (Chen et al., 2023) is a framework that relies on external tools (SQL application or Python interpreter) to help large language models revise and debug SQL commands or Python code with bugs.
14. **Reflexion** (Shinn et al., 2023) is a framework that provides a free-form reflection on whether a step was executed by LLM correctly or not and potential improvements. Unlike self-refine and self-debug, Reflexion builds a persisting memory of self-reflective experiences, which enables an agent to identify its own errors and self-suggest lessons to learn from its mistakes over time.
15. **Meta-Reasoner** (Yoran et al., 2023) is an approach which prompts large language models to meta-reason over multiple chains of thought rather than aggregating their answers. This approach included two steps: (i) ask LLM to generate multiple reasoning chains, (ii) ask another LLM (meta-reasoner) to reason over the multiple reasoning chains to arrive at the correct answer.
16. **HuggingGPT** (Shen et al., 2023) is a framework that leverages LLMs (e.g., ChatGPT) to connect various AI models in machine learning communities (e.g., Hugging Face) to solve numerous sophisticated AI tasks in different modalities (such as language, vision, speech) and domains.
17. **Camel** (Li et al., 2023) is a communicative agent framework involving inception prompting to guide chat agents toward task completion while maintaining consistency with human intentions.
18. **Chameleon** (Lu et al., 2023) is a plug-and-play compositional reasoning framework that augments external tools with LLMs in a plug-and-play manner. The core idea is that an LLM-based planner assembles a sequence of tools to execute to generate the final response. The assumption is that this will be less error-prone, easily expandable to new modules, and user-friendly.
19. **AutoGPT** (Richards, 2023) is an experimental open-source application that leverages the capabilities of large language models (LLMs) and Chatbots such as OpenAI’s GPT-4 and Chat-GPT to create fully autonomous and customizable AI agents. It has internet access, long-term and short-term memory management.
20. **BabyAGI** (Nakajima, 2023) is an intelligent agent capable of generating and attempting to execute tasks based on a given objective. BabyAGI operates based on three LLM flows: Task creation flow, Task prioritization flow, and Execution flow.

A.4. Prompting

We provide the prompts used to obtain the results in Section 6. Our evaluation is made possible thanks to the modular and compositional nature of *Flows*. Some of the experimental setups are deeply nested, and in cases where Flows build on each other, we avoid repetition. Note that the project’s GitHub repository provides the code and data to reproduce all of the experiments in the paper.

Direct prompting for a solution is shown in Listing 1. To add reflection, we use a Generator-Critic Flow to combine the code generation with a fixed reply, as shown in Listing 2. In the collaboration setting, we use Listing 3 as the generator and Listing 4 as the critic.

Debugging is incorporated via a testing Flow that adds formatting to the output of a code executor. The formatting templates are shown in Listing 6. To respond to the debug output, we rely on an adjusted coding Flow 5. Adding collaboration in the

Flows	Flow Type	Interactions				Reasoning Patterns		Feedback	Learning
		Self	Multi-Ag.	Human	Tools	Struct.	Plan		
FS (Brown et al., 2020)	Atomic	✗	✗	✗	✗	✗	✗	✗	✗
CoT (Wei et al., 2022)	Atomic	✗	✗	✗	✗	✓	✗	✗	✗
ToT (Yao et al., 2023a)	Circular	✓	✗	✗	✓	✓	✗	✗	✗
PoT (Chen et al., 2022)	Seq.	✗	✗	✗	✓	✓	✗	✗	✗
M-CoT (Zhang et al., 2023)	Seq.	✗	✗	✗	✗	✓	✗	✗	✓
ToolFormer (Wei et al., 2022)	Seq.	✗	✗	✗	✓	✓	✗	✗	✓
ReAct (Yao et al., 2023b)	Circular	✗	✗	✗	✓	✓	✗	✗	✗
Parsel (Zelikman et al., 2022)	Seq.	✗	✓	✗	✓	✓	✓	✗	✗
REFINER (Paul et al., 2023)	Gen-Crit	✗	✓	✓	✗	✓	✗	✓	✓
Self-Refine (Madaan et al., 2023)	Gen-Crit	✓	✗	✗	✗	✓	✗	✓	✗
RCI (Kim et al., 2023)	Gen-Crit	✓	✗	✗	✓	✓	✗	✓	✗
Self-Correct (Welleck et al., 2023)	Gen-Crit	✓	✗	✗	✓	✓	✗	✓	✗
Self-Debug (Chen et al., 2023)	Gen-Crit	✓	✗	✗	✓	✓	✗	✓	✗
Reflexion (Shinn et al., 2023)	Gen-Crit	✓	✗	✗	✓	✗	✗	✓	✗
Meta-Reasoner (Yoran et al., 2023)	Seq.	✓	✓	✗	✗	✓	✗	✗	✗
HuggingGPT (Shen et al., 2023)	Seq.	✗	✓	✗	✓	✓	✓	✗	✗
Camel (Li et al., 2023)	Circular	✗	✓	✓	✗	✓	✗	✓	✗
Chameleon (Lu et al., 2023)	Seq.	✗	✓	✗	✓	✓	✓	✗	✗
AutoGPT (Richards, 2023)	Circular	✓	✓	✗	✓	✓	✓	✓	✗
BabyAGI (Nakajima, 2023)	Circular	✗	✓	✗	✓	✓	✓	✗	✗

Table 2. **Previous work.** We compare previous work across relevant dimensions.

debugging setting is done by introducing a critic that provides feedback grounded in the test results. This Flow is detailed in Listing 3.

The scenarios explained above also support the addition of a planning Flow. An example of plan generation is shown in Listing 8.

Listing 1. Prompts for Code Flow (Codeforces)

```
"prompt_templates":
"system_message": |-
    Your goal is to provide executable Python code that solves a competitive
    programming problem. The code should correctly handle all corner cases in
    order to pass the hidden test cases, which are used to evaluate the
    correctness of the solution.

    The user will specify the problem by providing you with:
    - the problem statement
    - input description
    - output description
    - example test cases
    - (optional) explanation of the test cases

    The user will provide you with a task and an output format that you will
    strictly follow.
"query_message": |-
    # Problem statement
    {{problem_description}}

    # Input description
    {{input_description}}

    # Output description
    {{output_description}}
```



```
{{io_examples_and_explanation}}
```

```
The input should be read from the standard input and the output should be
    passed to the standard output.
Return Python code that solves the problem. Reply in the following format:
```python
{{code_placeholder}}
```

"human_message": |-
    {{query}}
```

Listing 2. Prompts for Fixed-Reply Flow

```
"prompt templates":
  "fixed_reply": |-
    Consider the problem statement and the last proposed solution. Are you sure
      that the solution is provided in the requested format, and crucially,
      solves the problem?
    If that is not the case, provide the corrected version of the code in the
      following format:
    ```python
 {{python_code}}
    ```

    otherwise, reply:
    "Final answer."
```

Listing 3. Prompts for Code-Collab Flow (Codeforces)

```
"prompt templates":
  "system_message": |-
    Your goal is to provide executable Python code that solves a competitive
      programming problem. The code should correctly handle all corner cases in
      order to pass the hidden test cases, which are used to evaluate the
      correctness of the solution.

    The user will specify the problem by providing you with:
      - the problem statement
      - input description
      - output description
      - example test cases
      - (optional) explanation of the test cases

    The user will provide you with a task and an output format that you will
      strictly follow.
  "query_message": |-
    # Problem statement
    {{problem_description}}

    # Input description
    {{input_description}}

    # Output description
    {{output_description}}
```

```
{{io_examples_and_explanation}}
```

The input should be read from the standard input and the output should be passed to the standard output.

Return Python code that solves the problem. Reply in the following format:

```
```python
{{code_placeholder}}
```
```

```
"human_message": |-
# Feedback on the last proposed solution
{{code_feedback}}
```

Consider the original problem statement, the last proposed solution and the provided feedback. Does the solution need to be updated? If so, provide the corrected version of the code in the following format:

```
```python
{{code_placeholder}}
```
```

```
otherwise, reply:
"Final answer."
```

Listing 4. Prompts for Code-Collab-Critic Flow (Codeforces)

```
"prompt templates":
"system_message": |-
Your goal is to identify potential issues with a competitive programming
solution attempt.

The user will specify the problem by providing you with:
- the problem statement
- input description
- output description
- example test cases
- (optional) explanation of the test cases
- a Python solution attempt

Crucially, your goal is to correctly identify potential issues with the
solution attempt, and not to provide the code implementation yourself.
The user will provide you with a task and an output format that you will
strictly follow.
"query_message": |-
# Problem statement
{{problem_description}}

# Input description
{{input_description}}

# Output description
{{output_description}}

{{io_examples_and_explanation}}
```

```
# Python solution attempt:
```python
{{code}}
```

Consider the problem statement and the solution attempt. Are there any issues
with the proposed solution or it is correct? Explain your reasoning very
concisely , and do not provide code.
"human_message": |-
{{query}}
```

Listing 5. Prompts for Code-Debug Flow (Codeforces)

```
"prompt templates":
"system_message": |-
Your goal is to provide executable Python code that solves a competitive
programming problem. The code should correctly handle all corner cases in
order to pass the hidden test cases , which are used to evaluate the
correctness of the solution .

The user will specify the problem by providing you with:
- the problem statement
- input description
- output description
- example test cases
- (optional) explanation of the test cases

The user will provide you with a task and an output format that you will
strictly follow.
"query_message": |-
# Problem statement
{{problem_description}}

# Input description
{{input_description}}

# Output description
{{output_description}}

{{io_examples_and_explanation}}
```

The input should be read from the standard input and the output should be
passed to the standard output.
Return Python code that solves the problem. Reply in the following format:

```
```python
{{code_placeholder}}
```

"human_message": |-
{{testing_results_summary}}
```

Consider the problem statement, the last proposed solution, and its issue.

Provide a corrected version of the code that solves the original problem and resolves the issue, without any explanation, in the following format:

```
```python
{{code_placeholder}}
```
```

Listing 6. Formatting templates for Code-Testing Flow (Codeforces)

```
"formatting templates":
"no error template": |-
    ${.issue_title}
    All of the executed tests passed.
"all tests header": |-
    ${.issue_title}
    The Python code does not solve the problem in the problem description due to
    logical errors. It fails on the following tests.
"compilation error template": |-
    ${.issue_title}
    The execution resulted in a compilation error.
    ## Compilation error message:
    {{error_message}}
"timeout error template": |-
    ${.issue_title}
    The execution timed out, the solution is not efficient enough.
"runtime error template": |-
    ${.issue_title}
    The execution resulted in a runtime error on the following test.
    ## [Failed test] Input
    ```

 {{test_input}}
    ```

    ## [Failed test] Runtime error message
    {{error_message}}
"single test error": |-
    ${.issue_title}
    The Python code does not solve the problem in the problem description due to
    logical errors. It fails the following test:
    ## [Failed test] Input
    ```

 {{test_input}}
    ```

    ## [Failed test] Expected output
    ```

 {{expected_output}}
    ```

    ## [Failed test] Generated output
    ```

 {{generated_output}}
    ```
"test error": |-
    ## [Failed test] {{idx}}
    ### [Failed test] {{idx}} Input
    ```
```

```

{{test_input}}
```
### [Failed test {{idx}}] Expected output
```
{{expected_output}}
```
### [Failed test {{idx}}] Generated output
```
{{generated_output}}
```

```

Listing 7. Prompts for Code-Debug-Collab Flow (Codeforces)

```

"prompt templates":
"system_message": |-
  Your goal is to identify the issues with an incorrect competitive programming
  solution attempt.

  The user will specify the problem by providing you with:
  - the problem statement
  - input description
  - output description
  - example test cases
  - (optional) explanation of the test cases
  - an incorrect Python solution attempt and a description of its issue

  Crucially, your goal is to consider all aspects of the problem and pinpoint
  the issues with the solution attempt, and not to provide the code
  implementation yourself.

  Some aspects to consider: Is the input correctly parsed? Is the output
  correctly formatted? Are the corner cases correctly handled? Is there a
  logical mistake with the algorithm itself?

  Use the code execution results provided in the issue description to guide
  your reasoning/debugging.
"query_message": |-
  # Problem statement
  {{problem_description}}

  # Input description
  {{input_description}}

  # Output description
  {{output_description}}

  {{io_examples_and_explanation}}

  # Solution attempt to be fixed
  ```python
 {{code}}
  ```

  {{testing_results_summary}}

```

```

Consider the problem statement , the solution attempt and the issue . Why is
the solution attempt incorrect ? How should it be fixed ? Explain your
reasoning very concisely , and do not provide code .
"human_message": |-
  {{query}}

```

Listing 8. Prompts for Plan Flow (Codeforces)

```

"prompt templates":
"system_message": |-
  Your goal is to provide a high-level conceptual solution that , if implemented
  , will solve a given competitive programming problem .

  The user will specify the problem by providing you with :
  - the problem statement
  - input description
  - output description
  - example test cases
  - (optional) explanation of the test cases

  The proposed algorithm should be computationally efficient , logically correct
  and handle all corner cases .

  The user will provide you with a task and an output format that you will
  strictly follow .
"query_message": |-
  # Problem statement
  {{problem_description}}

  # Input description
  {{input_description}}

  # Output description
  {{output_description}}

  {{io_examples_and_explanation}}

  Return a high-level conceptual solution that would solve the problem . Be very
  concise , and do not provide code .
  Reply in the following format :
  # Conceptual solution
  {{plan_placeholder}}
"human_message": |-
  {{query}}

```

A.5. The CC-Flows-competition: a new form of competitive coding

Solving competitive coding challenges is an eminently hard problem. The solve rate of only 27% by directly attempting the problem and 47% by the best-performing code Flow, paired with a reliable automatic evaluation metric, make competitive programming an ideal benchmark for AI systems. Motivated by this, we propose a competition where instead of people, proposed Flows solve competitive programming problems.

The competition will leverage the comprehensive dataset of publicly available Codeforces problems and the open-source infrastructure for inference and testing used in the experiments, available at https://github.com/epfl-dlab/cc_flows.

Submission and Formatting Instructions for ICML 2024

The competition will only include problems published after the knowledge-cutoff date of GPT-4. Furthermore, not to overload the Codeforces online evaluation infrastructure, we further filter this dataset to problems for which public and private tests are available, and the output format is compatible with our local code testing infrastructure. Codeforces ranks the difficulty of each problem from 800 to 2100. At the time of publishing, we have the following number of problems per difficulty (total of 416):

- difficulty 800: 149
- difficulty 900 to 1500 (inclusive): 185
- difficulty 1600 to 2100 (inclusive): 82

We will curate a leaderboard of best-performing Flows that will be publicly available on FlowVerse and provide the predictions that reproduce the reported scores using the provided infrastructure.

The data will be released and should be used in accordance with Codeforces' Terms and Conditions. Concretely, Codeforces prohibits the material from being sold, sublicensed, or commercialized. For more details, take a look at the project's GitHub page.