

Extracting Events Like Code: A Multi-Agent Programming Framework for Zero-Shot Event Extraction

Quanjiang Guo¹, Sijie Wang², Jinchuan Zhang¹, Ben Zhang¹, Zhao Kang¹, Ling Tian¹, Ke Yan^{1*}

¹University of Electronic Science and Technology of China, China

²Nanyang Technological University, Singapore

guochance1999@163.com, wang1679@e.ntu.edu.sg, {jinchuanz, zhangben}@std.uestc.edu.cn, {zkang, lingtian, kyan}@uestc.edu.cn

Abstract

Zero-shot event extraction (ZSEE) remains a significant challenge for large language models (LLMs) due to the need for complex reasoning and domain-specific understanding. Direct prompting often yields incomplete or structurally invalid outputs—such as misclassified triggers, missing arguments, and schema violations. To address these limitations, we present Agent-Event-Coder (AEC), a novel multi-agent framework that treats event extraction like software engineering: as a structured, iterative code-generation process. AEC decomposes ZSEE into specialized subtasks—retrieval, planning, coding, and verification—each handled by a dedicated LLM agent. Event schemas are represented as executable class definitions, enabling deterministic validation and precise feedback via a verification agent. This programming-inspired approach allows for systematic disambiguation and schema enforcement through iterative refinement. By leveraging collaborative agent workflows, AEC enables LLMs to produce precise, complete, and schema-consistent extractions in zero-shot settings. Experiments across five diverse domains and six LLMs demonstrate that AEC consistently outperforms prior zero-shot baselines, showcasing the power of treating event extraction like code generation.

Code — <https://github.com/UESTC-GQJ/AEC>

Introduction

Event extraction (EE) aims to identify event triggers and their associated arguments from unstructured text (Xu et al. 2024), we provide an illustration of the task in Figure 1. As a structured prediction task, it plays a vital role in applications such as knowledge base population, information retrieval, and question answering. Traditional EE methods rely on supervised learning and require labeled examples for each event type. However, the growing diversity of event types and the high cost of annotation make it impractical to collect training data for all possible events.

Zero-shot event extraction (ZSEE) seeks to address the limitations of supervised event extraction by enabling models to identify event types that have not been observed during training, using only the event type’s name or natural language definition. Rather than relying on annotated examples,

*Corresponding author




Figure 1: An illustrative example of the event extraction task. The blue box denotes the event type, while the green boxes represent the argument roles. Underlined words indicate the event trigger or the corresponding event arguments.

the model must leverage these textual definitions or ontological descriptions of event types and roles to guide prediction. Although ZSEE offers significant potential for scalable event extraction, it remains highly challenging (Chen et al. 2024; Cai et al. 2024), and existing approaches frequently fall short due to two key issues. (i) *Contextual ambiguity*: candidate event trigger words are often polysemous, and their correct interpretation depends on subtle contextual cues. For example, the word “strike” may refer to either a labor protest or a physical attack depending on the surrounding context. In the absence of training examples, capturing such nuances is difficult. As illustrated in Figure 2(a), large language models (LLMs) may misinterpret the trigger or fail to exploit contextual clues that indicate the correct event type and its arguments. (ii) *Structural fidelity*: event extraction is a structured prediction task that requires outputs to conform to a predefined schema, such as a JSON object or database entry. While LLMs can be prompted to produce structured outputs, they often fail to strictly follow schema constraints, particularly without fine-tuning (Liu et al. 2024; Guo et al. 2025b). This can result in malformed or incomplete event records that disrupt downstream processing. As shown in Figure 2(b), direct zero-shot prompting of LLMs frequently leads to misidentified triggers or invalid output structures, highlighting the need for a more guided and robust approach to zero-shot event extraction.

To address these limitations, we introduce **Agent-Event-Coder (AEC)**, a novel framework that reconceptualizes zero-shot event extraction as a collaborative and verifiable code-generation process. Instead of

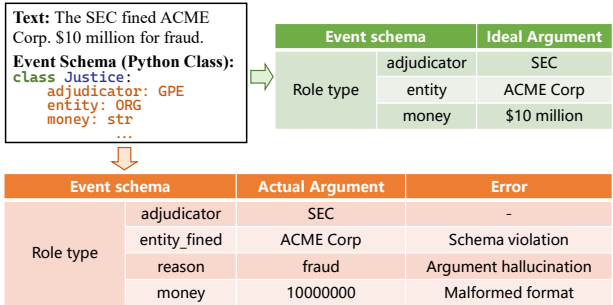



Figure 2: (a) Conceptual illustration of attention failure. The ideal model effectively leverages contextual information to correctly interpret the trigger word “*strike*” as an instance of the *Protest* event type. In contrast, direct zero-shot prompting of LLMs tends to over-rely on the trigger word itself, often leading to misclassification. (b) Illustration of extraction errors caused by insufficient structural fidelity. Outputs generated by direct zero-shot prompting of LLMs may violate the target event schema by: (1) including a non-existent argument role, (2) hallucinating an undefined argument, or (3) producing an argument with an incorrect data type.

prompting a single LLM to directly produce a structured output, AEC decomposes the extraction pipeline into four specialized agents—Retrieval, Planning, Coding, and Verification—each responsible for a distinct subtask. AEC is built upon two core principles. (i) **Multi-agent decomposition**: the overall task is divided into interpretable reasoning stages. For instance, the Planning Agent generates trigger-type hypotheses accompanied by explanatory rationales, while the Coding Agent converts the highest-confidence hypothesis into executable Python code that instantiates a schema-compliant event class. (ii) **Schema-as-code verification**: event schemas are represented as executable Python classes, enabling deterministic structural validation at runtime. A dedicated Verification Agent evaluates the generated code for semantic compatibility, type correctness, and structural validity. When validation fails, a dual-loop refinement procedure is initiated, iteratively patching the code based on compiler-like diagnostic feedback and, if necessary, exploring lower-confidence hypotheses. By combining step-wise reasoning with deterministic schema validation, this programming-inspired architecture allows AEC to sys-

tematically resolve trigger ambiguity and enforce structural fidelity in zero-shot settings, producing precise, complete, and schema-compliant event extractions without requiring annotated training examples.

In summary, we make the following contributions:

- We are the **first** to reformulate ZSEE as a multi-agent code generation task, providing a new paradigm that unifies schema constraints, iterative planning, and code-based validation for structured event extraction.
- We introduce a new multi-agent workflow AEC, where specialized agents work together to retrieve knowledge, design extraction plans, and generate structured event representations.
- We design a schema-as-code verification loop, in which a dedicated verification agent applies deterministic programming language rules to constraint outputs and provide precise feedback for iterative refinement.
- Through comprehensive evaluations across five diverse domains and six LLMs, we demonstrate the robustness, generalizability, and effectiveness of AEC as a state-of-the-art ZSEE framework.

Related Work

ZSEE with Prompting

Recent works have explored using LLMs for information extraction tasks by formulating extraction as a prompting or question-answering problem. In event extraction, approaches such as ChatIE (Wei et al. 2023) engage in structured dialogues with ChatGPT to iteratively refine event outputs, while others, such as CODE4STRUCT (Wang, Li, and Ji 2023) and Code4UIE (Guo et al. 2024), represent events and schemas as code or templates to leverage the reasoning capabilities of LLMs. These methods enable models to perform zero-shot or few-shot extraction by providing task descriptions or examples to the model. Additional studies have incorporated event definitions or constraints into prompts to guide the model—for example, by using positive and negative instructions about the event type and trigger (Srivastava, Pati, and Yao 2025). However, purely prompt-based single-agent strategies struggle with complex structured tasks (Guo et al. 2025a). Without explicit decomposition, an LLM may overlook subtle interactions between an event’s trigger and its arguments. Moreover, these methods are highly sensitive to prompt design and the choice of demonstration examples, which can easily mislead the model in zero-shot settings.

Multi-Agent Collaboration for Information Extraction

Multi-agent systems have been proposed to improve information extraction by having specialized agents or models cooperate or debate to reach better results. (Talebirad and Nadiri 2023) explore frameworks in which multiple LLM agents collaborate (either cooperatively or adversarially) through iterative dialogues, showing that such interactions can refine outputs for complex reasoning tasks. In event extraction, (Wang and Huang 2024) introduce a debate-style optimization (DoA) in a few-shot setting, where

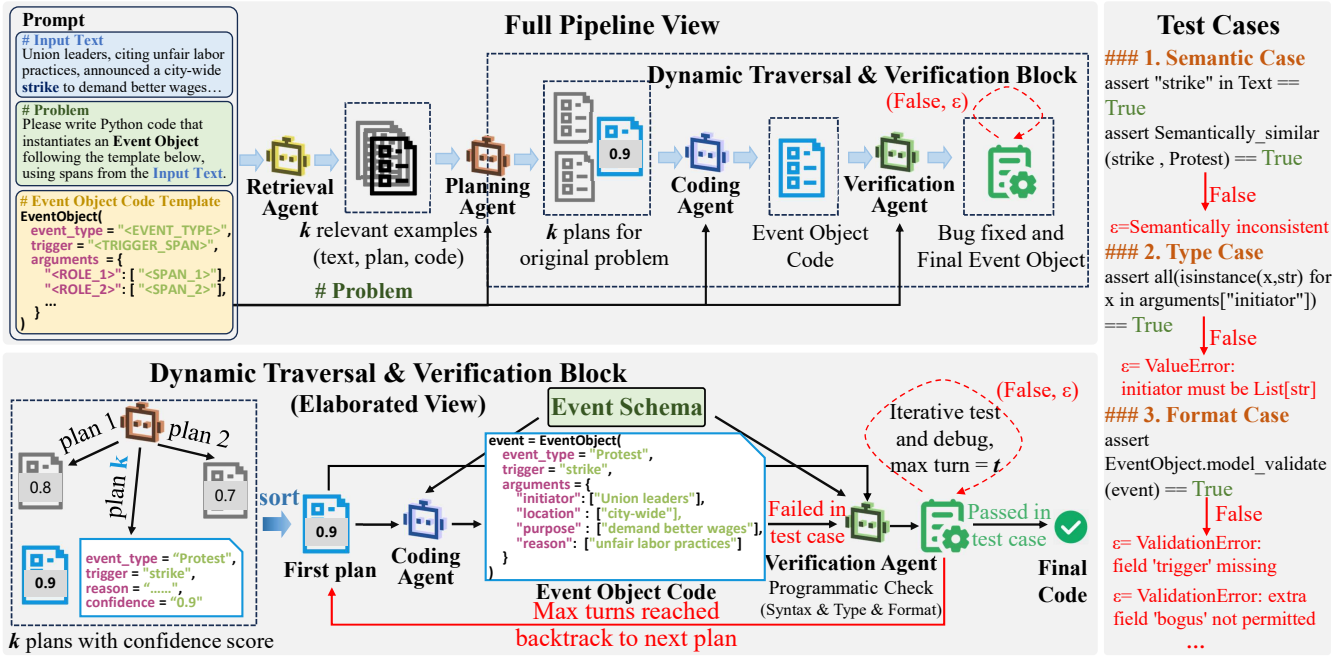


Figure 3: Overview of the proposed AEC framework. The **Full Pipeline View** (top) illustrates four specialized agents—*Retrieval*, *Planning*, *Coding*, and *Verification*—collaborating to generate schema-compliant event objects from unstructured text. The *Retrieval Agent* self-generates relevant exemplars to bridge the gap between schema definitions and textual context. The *Planning Agent* produces k trigger-type hypotheses, each with a confidence score and explanatory rationale. The *Coding Agent* converts the highest-confidence hypothesis into executable Python code that instantiates a predefined event schema. The **Dynamic Traversal and Verification Block** (bottom) depicts the iterative refinement loop. The generated code is evaluated by the *Verification Agent* through three deterministic **test cases**: *semantic*, *type*, and *format* checks (right). If a test fails, the agent patches the code using compiler-like diagnostics. When refinement attempts for the current plan are exhausted, the system backtracks to the next-best hypothesis. This dual-loop architecture ensures that the final output satisfies both semantic correctness and structural fidelity—without requiring any labeled examples.

two agents discuss and revise event predictions to reduce errors. For relation extraction, (Hou et al. 2024) propose a dual-agent approach (EPASS) that jointly models entity-pair extraction and supporting-evidence identification, demonstrating the benefit of agents focusing on different sub-tasks. (Lu et al. 2024) present TriageAgent, a heterogeneous multi-agent system for clinical IE, in which multiple LLM-based agents role-play with turn-taking, confidence scoring, and early stopping criteria to extract medical events more accurately. Most relevant to our work, (Wang et al. 2025) recently applied a cooperative multi-agent system to zero-shot named entity recognition (NER). Their framework (CMAS) uses four agents to handle entity-span detection, type-specific feature extraction, demonstration discrimination, and final prediction, yielding improved zero-shot NER performance by addressing context correlations and filtering prompt examples.

AEC adopts this collaborative paradigm but adapts it to the unique challenges of event extraction, where predictions must account for both triggers and multiple arguments under strict schema constraints. Unlike existing multi-agent IE frameworks, AEC introduces dedicated agents for trigger-type hypothesis generation, event coding,

and schema-level verification, thereby integrating step-wise reasoning with deterministic output validation. To the best of our knowledge, AEC is the first framework to employ a multi-agent, code-generation-based strategy for ZSEE, enabling both contextual disambiguation and structural fidelity through its division of labor and iterative refinement process.

Methodology

Task Definition. ZSEE takes as input an unstructured text span $T = w_1, \dots, w_n$ and an unseen event schema $S_e = \langle e, R_e \rangle$, where e is the event type and $R_e = (r_j, \tau_j)_{j=1}^m$ denotes a set of argument roles r_j with their expected value types τ_j . The objective is to generate a fully specified event instance:

$$y = \langle e, z, A \rangle, \quad (1)$$

where $z \in T$ is the predicted trigger span and $A = (r_j, a_j)_{j=1}^m$ is the set of argument-role pairs with $a_j \subseteq T$ or $a_j = \emptyset$. No labeled examples for e are available.

Overall Architecture

As illustrated in Figure 3, AEC addresses ZSEE through a structured four-agent pipeline, designed with two nested

feedback loops that enable iterative reasoning and verification.

Retrieval Agent. The retrieval agent A_{ret} is responsible for *self-generating* a set of k high-quality exemplar sentences tailored to the given event schema S_e :

$$A_{\text{ret}}(S_e) \rightarrow D_{\text{ex}} = \{s_1, \dots, s_k\}, \quad (2)$$

Drawing inspiration from the *analogical prompting* paradigm proposed by Yasunaga et al. (2024), these exemplars serve as schema–textual “analogies” that align abstract constraints with concrete linguistic realizations. By embedding step-by-step guidance into the demonstration space, they help disambiguate polysemous triggers and ground the model’s reasoning in context, thereby reducing early commitment errors and improving planning agent performance.

Planning Agent. The planning agent A_{plan} examines the input text T in the context of the retrieved exemplars D_{ex} and, by leveraging both lexical and semantic cues, produces a ranked list of *trigger–type* hypotheses accompanied by natural-language rationales:

$$A_{\text{plan}}(T, S_e, D_{\text{ex}}) \rightarrow P = \left\{ \left((z_i, e), \beta_i, \rho_i \right) \right\}_{i=1}^k, \quad (3)$$

where each z_i denotes a potential trigger in T , $\beta_i \in [0, 1]$ represents the model-assigned confidence that z_i evokes event type e , and ρ_i provides a concise natural-language explanation for why the pair (z_i, e) is plausible. These rationales ρ_i are retained for subsequent error analysis and ablation studies, enabling a better understanding of the agent’s decision-making process.

Coding Agent. Following (Srivastava, Pati, and Yao 2025), we compile every new schema $S_e = \langle e, \{(r_1, \tau_1), \dots, (r_m, \tau_m)\} \rangle$ into a Python `BaseModel` whose constructor enforces role types τ_j . Schema compliance therefore reduces to constructing a valid class instance, which can be *deterministically* verified at run time. A_{code} converts the highest-scoring hypothesis $((z^*, e), \beta^*, \rho^*)$ into executable Python that instantiates a Python template, as shown in Figure 4.

```
class EventObject (BaseModel):
    event_type: str
    trigger: str
    arguments: Dict[str, List[str]]
```

Figure 4: Fixed output template used by all agents.

Verification Agent. The verification agent A_{verify} evaluates the generated code object C_{obj} by executing a comprehensive three-stage test suite that checks semantic correctness, type validity, and structural integrity. Based on the outcome of these checks, the agent produces a binary verdict and, in case of failure, a diagnostic message ε indicating the first failed test:

$$A_{\text{verify}}(C_{\text{obj}}) \rightarrow (V, \varepsilon), \quad V \in \{\text{True}, \text{False}\}. \quad (4)$$

This design ensures that only code satisfying all verification criteria is accepted, while informative feedback is provided to guide subsequent error correction.

Three-Stage Verification

The verification agent A_{verify} executes a test suite consisting of three checks:

- **Semantic Check \mathcal{T}_1 :** Ensures that the predicted trigger z^* appears in the input text T and is semantically compatible with the event type e , based on lexical matching and contextual similarity.
- **Type Check \mathcal{T}_2 :** Verifies that each argument value conforms to the datatype specified in the schema S_e , while multiplicity constraints are enforced via Pydantic validation.
- **Structural Check \mathcal{T}_3 :** Confirms that the generated code compiles successfully, contains exactly the fields `{event_type, trigger, arguments}`, and produces a serializable event object.

The agent returns $V = \text{True}$ only if all three checks $\mathcal{T}_1 \wedge \mathcal{T}_2 \wedge \mathcal{T}_3$ succeed; otherwise, ε indicates the first failed test.

Dual-Loop Refinement Algorithm

Let k be the number of hypotheses and t the maximum patch attempts per hypothesis, the selection and verification procedure is shown in algorithm 1.

Algorithm 1: Dual-Loop Refinement Algorithm

```
Require: Candidate pool  $P$ 
Ensure: Valid code object  $C_{\text{obj}}$ 
1: Pick  $((z^*, e), \beta^*, \rho^*) = \arg \max_P \beta$  ▷ Selection
2: while  $P \neq \emptyset$  do
3:   for  $j = 1$  to  $t$  do ▷ Inner Loop
4:     Generate code  $C_{\text{obj}}$  with  $A_{\text{code}}$ 
5:     Get  $(V, \varepsilon) = A_{\text{verify}}(C_{\text{obj}})$ 
6:     if  $V$  is true then
7:       return  $C_{\text{obj}}$ 
8:     else
9:       Patch code using  $\varepsilon$ 
10:    end if
11:  end for
12:  Remove current hypothesis from  $P$ 
13:  Pick  $((z^*, e), \beta^*, \rho^*) = \arg \max_P \beta$ 
14: end while
```

The algorithm explores $O(kt)$ candidate paths yet guarantees that the final output is both semantically correct and schema-consistent, thereby realising reliable ZSEE without labelled examples.

Experiments

Experimental Setup

Datasets. We evaluate AEC on five widely used event extraction benchmarks covering diverse domains: FewEvent (General) (Deng et al. 2020), ACE 2005 (News) (Dodington et al. 2004), GENIA (Biomedical), SPEED (Epidemiological), and CASIE (Cybersecurity). For datasets without argument annotations (FewEvent and SPEED), we report only Trigger Identification (TI) and Event Identification

| LLM | Strategy | FewEvent(100) | | ACE(33) | | | | GENIA(9) | | | | SPEED(7) | | CASIE(5) | | | |
|------------|-------------|---------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | TI | EI | TI | EI | AI | AC | TI | EI | AI | AC | TI | EI | TI | EI | AI | AC |
| Llama3-8B | DirectEE | 21.5 | 17.5 | 26.4 | 25.7 | - | - | 27.8 | 27.4 | - | - | 34.3 | 41.5 | 11.8 | 47.6 | - | - |
| | GuidelineEE | 15.8 | 16.4 | 32.4 | 30.5 | 25.4 | 23.7 | 27.1 | 26.7 | 22.4 | 21.8 | 35.0 | 36.9 | 12.5 | 43.1 | 31.7 | 27.8 |
| | DecomposeEE | 20.7 | 20.5 | 30.1 | 35.2 | 26.1 | 24.9 | 28.9 | 28.5 | 22.8 | 21.7 | 31.2 | 38.4 | 10.5 | 50.7 | 27.0 | 25.1 |
| | CEDAR | 25.2 | 18.7 | 36.1 | 30.9 | - | - | 29.8 | 29.4 | - | - | 34.9 | 37.3 | 15.8 | 48.3 | - | - |
| | ChatIE | - | 24.8 | - | 44.2 | 32.4 | 30.8 | - | 23.8 | 21.7 | 20.3 | - | 42.9 | - | 33.3 | 22.2 | 20.8 |
| | AEC | 27.0 | 27.6 | 40.5 | 48.8 | 33.7 | 31.8 | 32.0 | 31.5 | 25.3 | 21.5 | 36.3 | 41.8 | 16.5 | 55.7 | 30.7 | 28.5 |
| Llama3-70B | DirectEE | 32.1 | 30.3 | 50.7 | 46.9 | - | - | 37.3 | 45.3 | - | - | 44.7 | 48.9 | 13.5 | 62.4 | - | - |
| | GuidelineEE | 29.8 | 32.1 | 46.2 | 51.4 | 30.7 | 27.5 | 31.5 | 34.8 | 27.5 | 27.0 | 41.5 | 44.1 | 12.3 | 48.3 | 36.2 | 31.6 |
| | DecomposeEE | 35.3 | 33.5 | 45.8 | 49.3 | 30.4 | 27.9 | 39.2 | 44.8 | 28.3 | 26.9 | 40.0 | 42.8 | 17.5 | 60.7 | 33.5 | 31.9 |
| | CEDAR | 34.5 | 33.9 | 51.5 | 48.7 | - | - | 36.8 | 47.7 | - | - | 45.3 | 49.8 | 16.7 | 54.3 | - | - |
| | ChatIE | - | 40.7 | - | 47.5 | 36.6 | 34.5 | - | 34.2 | 27.9 | 26.1 | - | 50.5 | - | 50.8 | 28.7 | 25.5 |
| | AEC | 42.1 | 40.5 | 57.0 | 54.6 | 38.4 | 34.7 | 39.4 | 48.1 | 31.2 | 30.1 | 43.8 | 52.3 | 18.7 | 65.9 | 36.4 | 33.9 |

Table 1: Main results comparing the ZSEE performance of our proposed AEC with all other baselines for the Llama3-8B-Instruct and Llama3-70B-Instruct LLMs. bold = best performance. (·) = number of distinct event types.

(EI). For ACE 2005, GENIA, and CASIE, we additionally evaluate Argument Identification (AI) and Argument Classification (AC).

To mitigate potential distributional biases, we follow the TEXTEE evaluation protocol (Huang et al. 2024; Parekh et al. 2025) and uniformly sample 250 test instances from each dataset to form evaluation splits. For CASIE, due to its smaller size, we sample 50 test instances. Our experiments are conducted under a purely zero-shot setting, i.e., no training data are used.

Baselines. We benchmark AEC against five strong zero-shot event extraction baselines:

- **DirectEE** (Gao et al. 2023) prompts LLMs directly to extract structured events in a single inference step without intermediate reasoning or decomposition.
- **CEDAR** (Li et al. 2023) adopts a multi-stage detection framework explicitly designed for large-ontology event detection, involving hierarchical reasoning over event types and triggers.
- **DecomposeEnrichEE** (Shiri et al. 2024) decomposes event extraction into event detection and argument extraction stages, utilizing dynamic, schema-aware retrieval augmentation to reduce hallucinations.
- **GuidelineEE** (Srivastava, Pati, and Yao 2025) leverages annotation guidelines, converting event extraction into a structured Python code-generation task guided by textual schema descriptions.
- **ChatIE** (Wei et al. 2023) transforms zero-shot IE into a conversational, multi-turn question-answering process, iteratively querying the LLM to progressively refine extraction outputs.

For fair comparison, all baselines were adapted to output structured, schema-conformant event objects. Moreover, we incorporate a unified Verification component into each baseline (if not already present) to ensure robustness and consistency in benchmarking performance.

Base LLMs. We evaluate AEC using a diverse set of instruction-tuned LLMs from three prominent model families: Llama3-8B-Instruct and Llama3-70B-Instruct from the Llama3 family (Dubey et al. 2024); Qwen2.5-14B-Instruct and Qwen2.5-72B-Instruct from the Qwen2.5 family (Hui et al. 2024); and GPT-3.5-turbo and GPT-4o from OpenAI (Achiam et al. 2023).

Evaluation Metrics. Following prior work (Srivastava, Pati, and Yao 2025; Parekh et al. 2025), we adopt four standard event extraction metrics: (1) **Trigger Identification (TI)**, which measures the exact match of predicted trigger spans; (2) **Event Identification (EI)**, which further requires correct classification of event types; (3) **Argument Identification (AI)**, which evaluates the accurate extraction of argument spans linked to the predicted triggers; and (4) **Argument Classification (AC)**, the most comprehensive metric, which additionally requires correct role-type assignment for each argument. We report micro-averaged F1 scores for all metrics on the constructed test splits.

Implementation Details. We use TextEE (Huang et al. 2024) for our benchmarking, datasets. Specifically, AEC is implemented on top of LLM backbones without performing any additional fine-tuning. We primarily use the instruction-tuned LLaMA3-8B and LLaMA3-70B models as the backbone LLMs. These models power the different agents in the AEC framework and operate entirely in a zero-shot prompting setting. Each agent is prompted with natural language instructions, and interagent communication is achieved through structured outputs rather than parameter updates (no task-specific training of the LLMs is performed). For robust evaluation, we report results averaged over three independent runs and set both the number of exemplars and inner-loop iterations to $k = t = 3$. All open-source models are executed locally on NVIDIA RTX A800 machines equipped with 4 GPUs.

| LLM | Prompt Style | FewEvent(100) | | ACE(33) | | | | GENIA(9) | | | | SPEED(7) | | CASIE(5) | | | |
|--------------|--------------|---------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | | TI | EI | TI | EI | AI | AC | TI | EI | AI | AC | TI | EI | TI | EI | AI | AC |
| Qwen2.5-14B | GuidelineEE | 23.5 | 20.2 | 35.1 | 33.5 | 22.4 | 21.0 | 27.0 | 25.8 | 20.1 | 19.2 | 32.5 | 36.4 | 15.2 | 47.5 | 27.3 | 26.8 |
| | DecomposeEE | 25.7 | 23.9 | 37.6 | 38.8 | 25.1 | 23.2 | 29.3 | 28.5 | 23.2 | 22.0 | 34.7 | 38.9 | 16.4 | 49.8 | 28.8 | 27.9 |
| | AEC | 30.4 | 28.1 | 42.5 | 45.3 | 28.5 | 22.7 | 32.1 | 30.7 | 25.6 | 24.8 | 38.6 | 42.4 | 17.8 | 53.2 | 31.4 | 29.5 |
| Qwen2.5-72B | GuidelineEE | 34.8 | 32.2 | 47.6 | 50.2 | 30.8 | 29.4 | 35.5 | 34.6 | 28.0 | 27.2 | 40.6 | 45.7 | 17.6 | 54.3 | 33.7 | 32.4 |
| | DecomposeEE | 37.1 | 35.8 | 49.9 | 53.7 | 31.6 | 32.7 | 38.4 | 36.9 | 29.7 | 28.6 | 43.1 | 47.9 | 18.2 | 57.5 | 34.8 | 33.7 |
| | AEC | 39.8 | 38.4 | 54.3 | 58.0 | 36.4 | 34.9 | 40.7 | 39.4 | 31.9 | 30.1 | 47.6 | 52.5 | 20.5 | 60.8 | 37.5 | 35.4 |
| GPT3.5-turbo | GuidelineEE | 23.2 | 20.9 | 35.1 | 37.9 | 20.6 | 17.4 | 26.7 | 25.5 | 19.1 | 18.3 | 31.3 | 36.5 | 18.3 | 56.7 | 31.5 | 30.8 |
| | DecomposeEE | 30.5 | 28.3 | 42.6 | 45.8 | 23.6 | 22.9 | 29.0 | 27.4 | 20.8 | 19.4 | 33.7 | 39.2 | 16.5 | 55.4 | 30.9 | 27.6 |
| | AEC | 32.9 | 30.2 | 46.2 | 50.1 | 26.7 | 25.2 | 31.5 | 29.8 | 22.6 | 18.8 | 37.9 | 43.1 | 17.9 | 58.6 | 28.8 | 26.7 |
| GPT4o | GuidelineEE | 40.7 | 38.5 | 53.4 | 55.9 | 32.2 | 33.9 | 38.9 | 37.8 | 30.2 | 29.5 | 44.2 | 50.1 | 19.7 | 59.7 | 36.8 | 35.3 |
| | DecomposeEE | 42.9 | 40.9 | 55.7 | 58.4 | 34.4 | 35.3 | 41.2 | 39.9 | 32.5 | 31.0 | 47.3 | 52.6 | 21.0 | 62.1 | 37.9 | 36.5 |
| | AEC | 44.6 | 42.8 | 58.3 | 61.8 | 38.2 | 36.8 | 43.7 | 41.9 | 34.2 | 32.5 | 50.8 | 56.7 | 22.5 | 65.1 | 39.7 | 37.8 |

Table 2: Generalization results for ZSEE performance comparing AEC with two major baselines for four other LLMs. **bold** = best performance. (·) = number of distinct event types.

| Llama3-70B | | | | | | | |
|-------------------------|-------------|-------------|-------------|-------------|-------------|-------------|--|
| Ablation Setting | FewEvent | | ACE | | | | |
| | TI | EI | TI | EI | AI | AC | |
| AEC (full model) | 42.1 | 40.5 | 57.0 | 54.6 | 38.4 | 34.7 | |
| w/o Retrieval Agent | 36.5 | 34.2 | 49.8 | 47.2 | 33.1 | 30.8 | |
| w/o Planning Rationales | 38.2 | 36.0 | 52.6 | 50.7 | 35.6 | 32.8 | |
| w/o Verification Loop | 35.0 | 32.5 | 47.1 | 44.7 | 30.7 | 28.5 | |
| w/o Structural Check | 39.8 | 37.6 | 54.9 | 52.5 | 37.2 | 33.6 | |
| GPT4o | | | | | | | |
| Ablation Setting | FewEvent | | ACE | | | | |
| | TI | EI | TI | EI | AI | AC | |
| AEC (full model) | 44.6 | 42.8 | 58.3 | 61.8 | 38.2 | 36.8 | |
| w/o Retrieval Agent | 39.0 | 36.5 | 51.6 | 54.2 | 32.6 | 32.0 | |
| w/o Planning Rationales | 41.3 | 39.1 | 54.3 | 57.6 | 35.0 | 33.9 | |
| w/o Verification Loop | 37.2 | 34.8 | 49.0 | 51.3 | 31.1 | 29.5 | |
| w/o Structural Check | 42.5 | 40.6 | 56.7 | 59.9 | 36.4 | 35.2 | |

Table 3: Ablation studies on two different LLMs, evaluated on FewEvent and ACE datasets.

Results and Analysis

Main Results Table 1 summarises the primary results comparing AEC with all baseline methods using two variants of Llama3. Across all benchmarks, AEC consistently achieves the best overall performance, substantially outperforming competitive baselines. On ACE 2005 with Llama3-8B, AEC yields gains of +7.8% and +6.0% in TI and EI, respectively, over ChatIE, while also achieving superior results on argument extraction metrics. With Llama3-70B as the backbone, these improvements become even more pro-

nounced, demonstrating AEC’s strong generalisation ability.

Although simpler baselines such as DirectEE and GuidelineEE perform reasonably well on smaller or less complex datasets, more structured methods show advantages on datasets with richer schemas. Nevertheless, AEC consistently achieves the best results, particularly on datasets with complex event schemas and diverse event types. These findings confirm that AEC’s collaborative multi-agent design, coupled with rigorous schema-driven verification, substantially enhances zero-shot event extraction performance.

Generalization across LLMs Table 2 demonstrates the generalization capability of AEC across four additional LLMs. AEC consistently achieves top performance over all baselines, with notable average improvements of approximately +3–5% TI, +4–6% EI, and +2–4% in argument metrics compared to the strongest baseline, DecomposeEE. We also observe clear parameter scaling effects: GPT4o achieves the highest overall performance, followed closely by Qwen2.5-72B, underscoring the enhanced reasoning capabilities afforded by larger model sizes under the AEC framework.

Ablation Study Table 3 presents the ablation results on LLaMA3-70B and GPT-4o across the FewEvent and ACE datasets. Removing the *Retrieval Agent* leads to substantial declines in trigger identification performance, highlighting the importance of exemplar generation for contextual disambiguation. Excluding *Planning Rationales* also results in performance degradation, confirming that intermediate reasoning plays a crucial role in guiding accurate type and argument decisions. Disabling the *Verification Loop* or the *Structural Check* consistently reduces all evaluation metrics—particularly for argument classification—demonstrating that code-level validation is critical for enforcing schema fidelity. Overall, these results underscore the necessity of both multi-agent reasoning and schema-as-

| Sentence | Best Baseline Prediction | Planning Agent Prediction | Coding Agent Prediction | Verification Agent Prediction |
|---|---|---------------------------------|--|--|
| The company acquired a startup specializing in AI technology. | [[" Transaction ", " startup "], [" Acquisition ", " technology "]] | [["Acquisition", "acquired"]] | [["Acquisition", "acquired"], ["Transaction", "startup"]] | [["Acquisition", "acquired"]] |
| A massive earthquake struck the city on Monday morning. | [[" Disaster ", " struck "]] | [["Earthquake", "earthquake"]] | [["Earthquake", "earthquake"], ["Location", "city"]] | [["Earthquake", "earthquake"], ["Location", "city"]] |
| The president announced new sanctions against the country after the attack. | [[" Attack ", " announced "]] | [["Announcement", "announced"]] | [["Announcement", "announced"], ["Sanction", "sanctions"]] | [["Announcement", "announced"], ["Sanction", "sanctions"]] |

Table 4: Qualitative examples comparing AEC components with the best baseline. Gold triggers and incorrect predictions are in bold.

| k | t | FewEvent | | ACE | | | |
|-----|-----|----------|------|------|------|------|------|
| | | TI | EI | TI | EI | AI | AC |
| 1 | 1 | 40.5 | 39.2 | 52.4 | 52.4 | 31.7 | 29.1 |
| 1 | 3 | 42.7 | 41.4 | 54.6 | 56.8 | 36.8 | 33.7 |
| 3 | 1 | 44.1 | 42.0 | 56.2 | 58.7 | 36.3 | 33.4 |
| 3 | 3 | 44.6 | 42.8 | 58.3 | 61.8 | 38.2 | 36.8 |
| 5 | 3 | 45.2 | 42.3 | 58.4 | 61.5 | 38.9 | 35.8 |
| 5 | 5 | 45.1 | 42.7 | 58.3 | 62.3 | 39.1 | 36.2 |

Table 5: Impact of the number of hypotheses k and patch attempts t on GPT4o. Performance improves with larger k and t but saturates beyond $k = 3$ and $t = 3$.

code verification for achieving robust zero-shot structured event extraction.

Qualitative Study Table 4 presents qualitative examples comparing AEC components with the best baseline. We observe three key trends. (i) The Planning Agent produces plausible trigger-type hypotheses but may omit essential arguments. (ii) The Coding Agent incorporates more structured arguments guided by the event schema, reducing role confusion. (iii) The Verification Agent further corrects type errors and removes inconsistent arguments via schema-level checks. In contrast, the baseline often misclassifies triggers or fails to capture critical arguments. These cases illustrate that AEC’s multi-agent reasoning and schema-as-code verification jointly improve contextual disambiguation and structural fidelity in zero-shot event extraction.

Impact of k and t We study how the number of hypotheses k and the maximum patch attempts t affect AEC’s performance. Table 5 shows results on FewEvent and ACE using GPT-4o as the backbone. Increasing k provides a larger hypothesis pool, while increasing t allows more opportunities for iterative error correction. Both factors lead to performance improvements up to a certain point, beyond which gains saturate or slightly decline due to the introduction of noisy low-confidence hypotheses and redundant refinement steps.

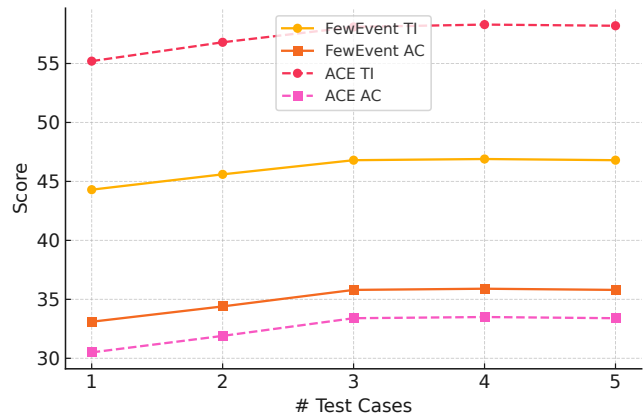


Figure 5: Impact of the number of test cases in verification.

Impact of Number of Test Cases We further examine how the number of test cases in the verification stage influences AEC. Figure 5 reports results on GPT-4o. Using more test cases improves both trigger identification and argument classification, as additional checks reduce structural and semantic errors. However, improvements plateau beyond three cases, indicating that further tests add little benefit while increasing computational overhead.

Conclusion

In this work, we introduced Agent-Event-Coder (AEC), a multi-agent framework that reframes ZSEE as a structured, iterative code-generation task. By decomposing extraction into Retrieval, Planning, Coding, and Verification agents and representing event schemas as executable classes, AEC enables systematic disambiguation, schema enforcement, and error correction. Experiments on five benchmarks and six LLM backbones show that AEC consistently outperforms strong zero-shot baselines, especially on complex schemas. Our results demonstrate that combining multi-agent reasoning with schema-as-code verification provides a robust paradigm for zero-shot structured prediction with LLMs.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Cai, Z.; Kung, P.; Suvarna, A.; Ma, M. D.; Bansal, H.; Chang, B.; Brantingham, P. J.; Wang, W.; and Peng, N. 2024. Improving Event Definition Following For Zero-Shot Event Detection. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, 2842–2863.
- Chen, R.; Qin, C.; Jiang, W.; and Choi, D. 2024. Is a large language model a good annotator for event extraction? In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 17772–17780.
- Deng, S.; Zhang, N.; Kang, J.; Zhang, Y.; Zhang, W.; and Chen, H. 2020. Meta-learning with dynamic-memory-based prototypical network for few-shot event detection. In *Proceedings of the 13th international conference on web search and data mining*, 151–159.
- Doddington, G. R.; Mitchell, A.; Przybocki, M.; Ramshaw, L.; Strassel, S.; and Weischedel, R. 2004. The Automatic Content Extraction (ACE) Program—Tasks, Data, and Evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation*.
- Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Yang, A.; Fan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arxiv:2407.21783*.
- Gao, J.; Zhao, H.; Yu, C.; and Xu, R. 2023. Exploring the feasibility of chatgpt for event extraction. *arXiv preprint arXiv:2303.03836*.
- Guo, Q.; Dong, Y.; Tian, L.; Kang, Z.; Zhang, Y.; and Wang, S. 2025a. BANER: Boundary-aware LLMs for few-shot named entity recognition. In *Proceedings of the 31st International Conference on Computational Linguistics*, 10375–10389.
- Guo, Q.; Zhang, J.; Wang, S.; Tian, L.; Kang, Z.; Yan, B.; and Xiao, W. 2025b. Bridging Generative and Discriminative Learning: Few-Shot Relation Extraction via Two-Stage Knowledge-Guided Pre-training. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 8068–8076.
- Guo, Y.; Li, Z.; Jin, X.; Liu, Y.; Zeng, Y.; Liu, W.; Li, X.; Yang, P.; Bai, L.; Guo, J.; et al. 2024. Retrieval-Augmented Code Generation for Universal Information Extraction. In *CCF International Conference on Natural Language Processing and Chinese Computing*, 30–42.
- Hou, W.; Jia, N.; Liu, X.; Zhao, W.; and Wang, Z. 2024. A multiagent-based document-level relation extraction system with entity pair awareness and sentence significance. *IEEE Systems Journal*, 18(4): 1905–1916.
- Huang, K.-H.; Hsu, I.-H.; Parekh, T.; Xie, Z.; Zhang, Z.; Natarajan, P.; Chang, K.-W.; Peng, N.; and Ji, H. 2024. TextEE: Benchmark, Reevaluation, Reflections, and Future Challenges in Event Extraction. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, 12804–12825.
- Hui, B.; Yang, J.; Cui, Z.; Yang, J.; Liu, D.; Zhang, L.; Liu, T.; Zhang, J.; Yu, B.; Lu, K.; et al. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*.
- Li, S.; Zhan, Q.; Conger, K.; Palmer, M.; Ji, H.; and Han, J. 2023. GLEN: General-Purpose Event Detection for Thousands of Types. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2823–2838.
- Liu, S.; Li, Y.; Li, J.; Yang, S.; and Lan, Y. 2024. Unleashing the Power of Large Language Models in Zero-shot Relation Extraction via Self-Prompting. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 13147–13161.
- Lu, M.; Ho, B.; Ren, D.; and Wang, X. 2024. Triageagent: Towards better multi-agents collaborations for large language model-based clinical triage. In *Findings of the Association for Computational Linguistics: EMNLP*, 5747–5764.
- Parekh, T.; Mehta, K.; Mehrabi, N.; Chang, K.-W.; and Peng, N. 2025. DiCoRe: Enhancing Zero-shot Event Detection via Divergent-Convergent LLM Reasoning. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 20571–20593.
- Shiri, F.; Moghimifar, F.; Haffari, R.; Li, Y.-F.; Nguyen, V.; and Yoo, J. 2024. Decompose, Enrich, and Extract! Schema-aware Event Extraction using LLMs. In *2024 27th International Conference on Information Fusion*, 1–8.
- Srivastava, S.; Pati, S.; and Yao, Z. 2025. Instruction-tuning llms for event extraction with annotation guidelines. *arXiv preprint arXiv:2502.16377*.
- Talebirad, Y.; and Nadiri, A. 2023. Multi-agent collaboration: Harnessing the power of intelligent llm agents. *arXiv preprint arXiv:2306.03314*.
- Wang, S.; and Huang, L. 2024. Debate as Optimization: Adaptive Conformal Prediction and Diverse Retrieval for Event Extraction. In *Findings of the Association for Computational Linguistics: EMNLP*, 16422–16435.
- Wang, X.; Li, S.; and Ji, H. 2023. Code4Struct: Code Generation for Few-Shot Event Structure Prediction. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, 3640–3663.
- Wang, Z.; Zhao, Z.; Lyu, Y.; Chen, Z.; de Rijke, M.; and Ren, Z. 2025. A cooperative multi-agent framework for zero-shot named entity recognition. In *Proceedings of the ACM on Web Conference 2025*, 4183–4195.
- Wei, X.; Cui, X.; Cheng, N.; Wang, X.; Zhang, X.; Huang, S.; Xie, P.; Xu, J.; Chen, Y.; Zhang, M.; et al. 2023. Chatie: Zero-shot information extraction via chatting with chatgpt. *arXiv preprint arXiv:2302.10205*.
- Xu, D.; Chen, W.; Peng, W.; Zhang, C.; Xu, T.; Zhao, X.; Wu, X.; Zheng, Y.; Wang, Y.; and Chen, E. 2024. Large language models for generative information extraction: A survey. *Frontiers of Computer Science*, 18(6): 186357.
- Yasunaga, M.; Chen, X.; Li, Y.; Pasupat, P.; Leskovec, J.; Liang, P.; Chi, E. H.; and Zhou, D. 2024. Large Language Models as Analogical Reasoners. In *The Twelfth International Conference on Learning Representations*.