

Who's Who? LLM-assisted Software Traceability with Architecture Entity Recognition

DOMINIK FUCHSS, Karlsruhe Institute of Technology (KIT), Germany

HAOYU LIU, Karlsruhe Institute of Technology (KIT), Germany

SOPHIE CORALLO, Karlsruhe Institute of Technology (KIT), Germany

TOBIAS HEY, Karlsruhe Institute of Technology (KIT), Germany

JAN KEIM, Karlsruhe Institute of Technology (KIT), Germany

JOHANNES VON GEISAU, Karlsruhe Institute of Technology (KIT), Germany

ANNE KOZIOLEK, Karlsruhe Institute of Technology (KIT), Germany

Identifying architecturally relevant entities in textual artifacts is crucial for Traceability Link Recovery (TLR) between Software Architecture Documentation (SAD) and source code. While Software Architecture Models (SAMs) can bridge the semantic gap between these artifacts, their manual creation is time-consuming. Large Language Models (LLMs) offer new capabilities for extracting architectural entities to construct SAMs automatically or establish direct trace links. This paper extends our ICSA 2025 paper [19], which introduced ExArch for LLM-based architecture component name extraction, by contributing the novel ArTEMiS approach, an extended evaluation, and a combined evaluation of both approaches. ExArch extracts component names as simple SAMs from SAD and source code, while ArTEMiS identifies architectural entities in documentation and matches them with SAM entities. Our evaluation compares against state-of-the-art approaches SWATR, TransArC, and ArDoCode. TransArC achieves strong performance (F1: 0.87) but requires manually created SAMs; ExArch achieves comparable results (F1: 0.86) using only SAD and code. ArTEMiS matches SWATR (F1: 0.81) and can replace it when integrated with TransArC. The combination of ArTEMiS and ExArch outperforms ArDoCode, the best baseline without manual SAMs. Our results demonstrate that LLMs can effectively enable automated SAM generation and TLR, making architecture-code traceability more practical and accessible.

CCS Concepts: • **Software and its engineering** → *Software design engineering*; **Documentation**; *Maintaining software*; *Software evolution*; **Software architectures**; • **Computing methodologies** → **Natural language processing**; **Information extraction**.

Additional Key Words and Phrases: Traceability Link Recovery, Large Language Models, Software Architecture, Model Extraction

ACM Reference Format:

Dominik Fuchß, Haoyu Liu, Sophie Corallo, Tobias Hey, Jan Keim, Johannes von Geisau, and Anne Koziolk. 2026. Who's Who? LLM-assisted Software Traceability with Architecture Entity Recognition. *ACM Trans. Autonom. Adapt. Syst.* 1, 1 (January 2026), 30 pages. <https://doi.org/10.1145/3807453>

Authors' Contact Information: **Dominik Fuchß**, dominik.fuchss@kit.edu, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany; **Haoyu Liu**, haoyu.liu@kit.edu, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany; **Sophie Corallo**, sophie.corallo@kit.edu, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany; **Tobias Hey**, hey@kit.edu, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany; **Jan Keim**, jan.keim@kit.edu, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany; **Johannes von Geisau**, johannes.geisau@student.kit.edu, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany; **Anne Koziolk**, koziolk@kit.edu, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany.



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

© 2026 Copyright held by the owner/author(s).

ACM 1556-4703/2026/1-ART

<https://doi.org/10.1145/3807453>

1 Introduction

In software development, numerous artifacts are produced, each representing different levels of abstraction and addressing distinct aspects of the system. Challenges for architects and developers arise because the relationships between these artifacts are often unclear, which prevents their effective use and causes challenges in maintenance. To address this, traceability link recovery (TLR) techniques are used to establish, maintain, and manage explicit trace links between artifacts. Improving software quality is closely tied to the creation and management of trace links [48, 61].

One challenge in linking artifacts is the semantic gap between different types of artifacts. Bridging this gap is difficult, and automated approaches often misinterpret the underlying semantics. To address this, some methods suggest using intermediate artifacts to reduce the semantic gap, making it easier to link related artifacts [4, 32, 43, 45]. For example, the descriptions in design documentation are semantically closer to design artifacts like Software Architecture Models (SAMs) than to code, and SAMs, in turn, are closer to code. Consequently, specialized approaches can more easily establish links between Software Architecture Documentations (SADs) and SAMs, or between SAMs and code [26, 33]. Based on these insights, transitive approaches like Transitive links for Architecture and Code (TransArC) [32] have been developed. TransArC recovers trace links between SADs and source code based on manually created SAMs as intermediate artifacts. However, these methods are not always applicable, as intermediate artifacts are often unavailable.

The trace links across the SAD-SAM-Code chain are bi-directional and differ in multiplicity. Bi-directionality means that links can be traversed in both directions: forward (SAD \rightarrow SAM \rightarrow Code) to identify which code units implement a given architectural description, and backward (Code \rightarrow SAM \rightarrow SAD) to identify which architectural descriptions and components are responsible for a given code unit. Accordingly, developers can, for example, trace the root cause of a failure back to the relevant passages in the SAD. The two sub-link types also differ in their multiplicity. SAD-SAM links are many-to-many. This means, on one hand, a single sentence in the SAD can mention multiple components. For example the sentence “The MediaManagement requests the user token from the UserManagement.” links to both components simultaneously. On the other hand, a single component of the SAM may be described across multiple sentences throughout the SAD. SAM-Code links are one-to-many in the forward direction: one architectural component (e.g., MediaAccess or UserManagement) can map to multiple implementing source files. In the backward direction, a source file usually belongs only to one component, making these links many-to-one from code to SAM.

We propose two complementary Large Language Model (LLM)-based approaches that together enable end-to-end TLR from SAD to source code without the need for manually created SAMs. Both approaches follow a common design philosophy: rather than delegating the entire TLR task to an LLM, they decompose the problem and employ LLMs only for the subtasks where they excel. Specifically, they focus on identifying and extracting architectural entities. For the actual trace link recovery, these new approaches are then providing necessary information to the existing heuristic-based approach TransArC [32]. Consequently, we can formulate our overarching research question:

RQ0 *How well can LLMs extract architectural elements to enable and assist in TLR between SAD and source code?*

First, we introduce *Extracting Architecture (ExArch)*, an approach that uses LLMs to extract component names from SAD and/or source code. Since, in practice, SAMs are regularly unavailable, we aim to recover trace links between SAD and code without the need for manually created SAMs.

To achieve this, we leverage the strength of LLMs in understanding natural language and code. We design an approach that uses LLMs to recover SAMs in the form of component names from SAD and/or code. Component names contribute the required information to apply TransArC without manually creating SAMs. In doing so, we bridge the semantic gap between SAD and code and empower TLR between these artifacts. We emphasize that the component names define a simple SAM that provides exactly the information for TLR. In contrast to the research field of architecture recovery that aims to recover more detailed architecture models, we focus on the information needed to enable state-of-the-art SAD-to-code TLR approaches such as TransArC. Consequently, we have the following research questions regarding this approach:

- RQ1.1** Is the performance of architecture TLR with LLM-extracted component names as intermediate artifacts (ExArch) comparable to using manually created SAMs?
- RQ1.2** Does TransArC with ExArch perform better than state-of-the-art TLR between SAD and code without SAMs as intermediates?
- RQ1.3** Is the performance of TransArC with ExArch with current open-source LLMs comparable to the performance of closed-source ones?
- RQ1.4** How does the performance of TransArC with ExArch differ when using different artifacts to generate the simple SAMs?

Second, we introduce *Architecture Traceability with Entity Matching via Semantic inference* (ArTEMiS), an approach to recover trace links between SAD and SAM using LLMs. The approach uses LLMs for a special form of Named Entity Recognition for software architecture. It is a modern alternative to traditional heuristic-based TLR approaches such as the approach by Keim et al. [34]. We compare our new ArTEMiS approach with existing TLR approaches for SAD-SAM TLR. Finally, we combine ArTEMiS with TransArC and ExArch to evaluate the influence of LLM-based SAD-SAM TLR on the performance of TransArC. Consequently, we have the following research questions regarding this approach:

- RQ2.1** Is the LLM-based TLR from SAD to SAM (ArTEMiS) better than state-of-the-art heuristic-based TLR?
- RQ2.2** Does ArTEMiS improve the performance of TransArC with and without ExArch?

This paper extends our ICSA 2025 paper [19], which initially presented ExArch and presented initial evaluation results. ExArch extracts architecture component names for transitive TLR between SAD and code, removing the need for manually created SAMs. ArTEMiS is a novel contribution of this extension: It extracts architecture entities from SAD similar to Named Entity Recognition, links them to SAMs, and provides a modern alternative to classical heuristic-based approaches. We additionally provide an extended evaluation of ExArch with a broader range of LLMs, a revised benchmark, and the combined evaluation of ArTEMiS and ExArch with TransArC. We provide code, baselines, evaluation data, and results in our replication packages [20, 21].

We structure the remainder of the paper as follows: Related work is examined in [Section 2](#). ExArch [19] is presented in [Section 3](#), including a description of how we incorporated TransArC [32] in [Section 3.3](#). Moreover, in [Section 4](#), we describe the ArTEMiS approach to extract architecture entities from SAD and how to link them to SAM. In [Section 5](#), we describe our experimental design. Here, we present our initial experiment on the performance of ExArch. Moreover, we show the results of our extended evaluation of ExArch, ArTEMiS, and their combination. Lastly, we conclude this paper in [Section 6](#).

2 Related Work & Foundations

In this section, we first discuss the ideas and benefits of transitive links, comparing previous works with ours. Afterward, we focus on the application of LLMs for TLR. Here, we concentrate on the

information extraction ability of LLMs that also motivates our use of LLMs. We also discuss the use of LLMs for Named Entity Recognition (NER) that relates to our ArTEMiS approach. Finally, we discuss architecture recovery and its differences from our component name recovery.

2.1 Transitive Trace Links

Automated TLR approaches primarily involve comparing different terms across textual artifacts to find terms referring to the same concept. Consequently, researchers have leveraged semantic similarity techniques developed within the Natural Language Processing (NLP) field to facilitate this process.

Among those techniques, methods like Vector Space Model (VSM) and Latent Semantic Indexing (LSI) are commonly used [3]. However, software artifacts exist at various levels of abstraction. This complexity challenges NLP models, as they often struggle to handle cross-level artifacts effectively, limiting their accuracy in recovering trace links. Thus, researchers have explored different techniques to address this challenge. Approaches include incorporating fine-grained information [28, 30], considering dependencies [36, 47], and using enriched vocabularies [22]. More recently, transitive linking through intermediate artifacts has been shown to be promising in bridging those semantic gaps [32, 49].

The underlying idea of transitive links is that intermediate artifacts can help to find implicit tracing relationships [44]. An early work by Nishikawa et al. [45] showed the importance of choosing suitable intermediate artifacts for transitive link recovery. Their approach focuses on establishing transitive links between two artifacts using a third artifact [45]. It uses VSM to generate initial trace links between various pairs of artifacts. These links include pairs between use cases, interaction diagrams, code, and test cases. The experiment was conducted under various settings, including scenarios with no intermediate artifacts and with different artifacts used as intermediates. Their findings highlight that suitable intermediate artifacts can significantly influence TLR performance. For example, interaction diagrams are better intermediates than use cases when recovering trace links between code and test cases. Instead of only using VSM to generate initial links for later transitive linking, various NLP techniques can be combined [44, 49]. COMET [44] uses a Bayesian inference framework to treat the recovery process from a probabilistic view. It uses multiple similarity scores to estimate the model's parameters. In contrast to COMET, Rodriguez et al. [49] first combine multiple scores from different sources into a single score, and then evaluate the link based on the final score.

So far, previous works mainly focused on traceability in requirements [3, 22, 28, 36, 44, 49] or test cases [44, 45, 47], leaving architectural traceability [32, 33] less explored. TransArC [32] has so far achieved the best performance in the architectural traceability benchmark. It leverages SAMs as intermediate artifacts to recover trace links between SAD and code. Although it achieves significant improvement, its reliance on the existence of well-maintained SAMs limits its application to a wider range of projects.

Transitive approaches have shown promising results, but cannot work when the intermediate artifacts are missing. Therefore, it remains a challenge to leverage transitive links when the intermediate artifacts are not completely present. To enable transitive trace link recovery of SAD and source code in more settings, we explore using LLMs to recover the needed information of a SAM for TLR.

2.2 LLMs for Traceability

With the rich general knowledge obtained from pretraining, LLMs have been effective at knowledge-intensive software engineering tasks [11, 42, 70]. LLMs have been used for various architecture and modeling tasks [31, 53], including supporting design decision making [13, 14], modeling tasks [9, 10],

and software architecture analysis and generation [1, 15]. In the following, we discuss the role of LLMs in advancing traceability research. T-BERT is an early adoption of LLMs for TLR in an issue-commit setting [39]. The BERT language model was tested using three variants of neural architectures: Twin, Siamese, and Single. They demonstrated that T-BERT outperforms previous VSM and recurrent neural network approaches.

With the rise of decoder-only LLMs like GPT, prompt phrasing greatly influences the language model's output [50]. They explored the performance of those LLMs on recovering links with different types of prompts: classification (two artifacts are linked or not), ranking (rank all related artifacts), and Chain-of-thought (recovery step-by-step while giving reasons). Besides, they showed that LLMs can understand domain-specific terms from the general knowledge obtained by pre-training. Motivated by this understanding abilities, we explore LLMs to extract intermediate structures from architecture documents and source code to help TLR.

Hassine [24] explored LLMs' zero-shot abilities to trace security-related requirements to goal models. The prompts are tailored to the Goal-oriented Requirements Language's peculiarities. Although achieving positive results, the approach's high dependence on the task and data limits its further application to our architectural TLR task. North et al. [46] considered that using requirements-to-code links helps LLMs to generate code by iteratively reformatting prompts. They use the gradients between generated code and requirements sentences to identify which part of the requirement is overlooked. This overlooking link is later used to reformat the prompt.

Fuchß et al. [18] created the LiSSA framework for generic TLR. They use retrieval-augmented generation and LLM-based classification to trace artifacts like SAD, code, and requirements. LiSSA has also been applied to inter-requirements traceability [29]. However, for TLR related to architecture models, they did not outperform the state of the art.

Existing LLM approaches for TLR show that LLMs can have the software knowledge needed for our task. However, our task differs from those already tackled, or the results of the existing approaches did not reach state-of-the-art, so these approaches can't be directly applied.

2.3 Named Entity Recognition (NER)

NER is a subtask of NLP that focuses on identifying and classifying key entities in text into predefined categories, such as names of persons, organizations, locations, and more. In software engineering, NER is often used to extract relevant entities from Q&A discussions and bug reports, cybersecurity-related texts, and other software artifacts. Thus, extracting software entities, like frameworks, libraries, and tools from StackOverflow posts is a common task [58–60, 64, 68]. Zhou et al. [71] recognizes bug-specific entities, like GUI elements, network entities, drivers, hardware, etc., in bug reports. In cybersecurity, NER is used to extract cybersecurity-related terms, such as vulnerabilities or software entities from security websites, news, or cyber threat intelligence texts [5, 56, 62, 66]. More rarely, NER is used to extract software entities from user stories [27], requirements [12, 41]. In contrast to the mentioned approaches, our approach focuses on extracting component names from architecture documents. Thus, the entities are not only domain-specific but more closely tied to the project context, which makes the task even more challenging.

2.4 Architecture Recovery

Our approach's generation of component names is conceptually related to software architecture recovery research. Therefore, we give a brief overview of architecture recovery approaches and describe the differences. In the end, we discuss how our approach differs from previous ones.

Structural information, including system dependencies [40] and folder structure, is an important source for architecture recovery. An early approach, ACCD [63], uses folder and dependency information gathered through static analysis. Building upon ACCD, dynamic dependency information

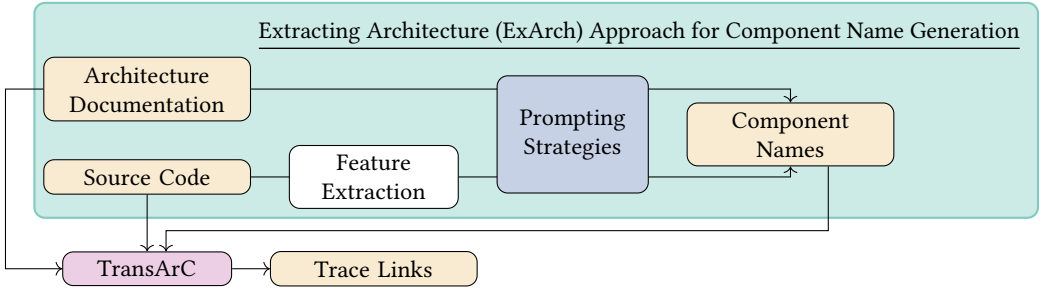


Fig. 1. Overview of the Extracting Architecture (ExArch) Approach for TLR Between SAD and Code. ExArch (green box) uses feature extraction and LLM-based prompting to derive component names from SAD, source code, or both, producing a simple SAM (Component Names). TransArC then uses the SAD, source code, and the generated component names to recover trace links. Artifacts are shown in orange, prompting in blue, feature extraction in white, and TransArC in purple.

has been shown to help recovery in some cases [67]. Lutellier et al. [40] examined the impact of code dependency in detail, finding that symbol dependencies yield better recovery accuracy than include dependencies. In addition to structural information, textual information can also help. The primary idea is that artifacts belonging to the same component may have similar variable names. ARC [23] recovers architecture using system concerns, which are application-specific features extracted from software corpora. Corazza et al. [8] studied the impact of six different types of identifier names. They found it would be better to treat identifiers separately by their types and assign suitable weights than to treat them equally in a large vocabulary. Structural and semantic information can also be combined. They can be integrated to build weighted call graphs [6], create recovery patterns [2], and be combined using information fusion models [69]. Rukmono et al. [51] showed architecture could also be recovered deductively. They start from a reference architecture and iteratively use LLMs to refine it with implementation details. Although we share similar ideas with architecture recovery, the purpose and final result are different. Our approaches are designed to retrieve the required information to reduce the gap for TLR between architecture documentation and code, i.e., component names. The recovered model thus does not serve as a complete, functional artifact. In contrast, architecture recovery focuses on recovering detailed software architecture for more general purposes, like understanding the system’s functionality or performance prediction.

3 ExArch: An Approach to Generate Simple SAMs

This section describes Extracting Architecture (ExArch), our approach to generate a simple SAM from SAD and source code to enable transitive TLR. The approach was introduced by Fuchß et al. [19] at the International Conference on Software Architecture (ICSA) 2025.

Figure 1 provides an overview of ExArch. ExArch consists of two main parts: generating the component names for a SAM and recovering trace links. The input for ExArch is the SAD and the software project’s source code. Since the whole source code is typically too large to be used as input for an LLM, we extract features from it. We then use prompting strategies to extract and generate component names for the intermediate SAM. Using these SAMs, we apply the TransArC [32] approach to recover trace links between the SAD and the source code. In the following, we describe the individual steps in detail.

3.1 Feature Extraction for Source Code

As LLMs typically only accept inputs that are smaller than a project's source code, we only provide extracted features from the source code as input. Since we aim to recover component names from the source code, we do not need to maintain all the information.

In object-oriented programming languages like Java, the package structure can provide valuable information about the architecture of a system [55]. Thus, we extract the package structure from the source code. ExArch is not limited to this feature, but the feature extraction in this paper focuses only on providing a list of non-empty source code packages.

3.2 Prompting Strategies

This section describes the prompting strategies we use to generate the SAM. We consider three modes for this paper: First, we extract the component names only from the SAD. Second, we extract the component names only from the source code. Third, we incorporate both to generate the simple SAMs.

Extract Component Names from Documentation. We use Chain-of-thought prompting to generate the simple SAM from the architecture documentation. We use two prompts:

Prompt 1: Documentation to Architecture (1) Your task is to identify the high-level components based on the software architecture documentation. In a first step, you shall elaborate on the following documentation: {Software Architecture Documentation}

Prompt 2: Documentation to Architecture (2) Now provide a list that only covers the component names. Omit common prefixes and suffixes in the names in camel case. Output format:
 - Name1
 - Name2

Prompt 1 queries the LLM to identify high-level components based on the architecture documentation. We also instruct the LLM to elaborate on the architecture documentation. Thus, the LLM is not restricted to generate output in a defined format. In the second step, we use **Prompt 2** to generate a list of component names. Here, we instruct the model to provide only a list of component names without common prefixes and suffixes, like *Component*. We aim to reduce complexity and facilitate component name retrieval by defining an output format.

Generate Component Names from Source Code. We also use Chain-of-thought prompting to generate the simple SAM from the source code in two prompting steps.

Prompt 3: Code to Architecture (1) You get the {Features} of a software project. Your task is to summarize the {Features} w.r.t. the high-level architecture of the system. Try to identify possible components. {Features}; {Content}

Prompt 3 queries the model to summarize the features of a software project w.r.t. the system's architecture. Here, we instruct the model to identify possible components based on the features extracted from the source code. In this paper, we used "Packages" as the feature. This feature names all non-empty packages. We then re-use **Prompt 2**.

Generate Component Names from SAD and Source Code. Lastly, we consider a combination of the SAD and the source code to generate the simple SAM. Here, we decided to use two modes for the combination.

First, if we extract the SAM from the documentation and the source code, we can aggregate the results by using LLMs. For this purpose, we use the following prompt:

Prompt 4: Aggregation You get a list of possible component names. Your task is to aggregate the list and remove duplicates. Omit common prefixes and suffixes in the names in camel case. {Output Format (cf. Prompt 2)} Possible component names: {Possible Component Names}

In **Prompt 4**, we ask the model to aggregate the list of possible component names and remove duplicates. Afterward, we use the same statements regarding prefixes, suffixes, and output format as in the other prompts.

We calculate the similarity between the component names generated from the documentation and the source code. To merge the component names, we process them sequentially, starting with the component names from the documentation. The normalized Levenshtein distance is already known from TLR tasks [17, 33]. We use this Levenshtein distance [37] to calculate the similarity between the possible component names. If the similarity is above a certain threshold, we consider the component names as equal and omit the new one. Since we assume that the SAD is closer to the actual SAM than the source code, we start aggregation with the extracted component names from SAD. In doing so, we want to complement the component extracted from SAD with component names extracted from the source code. Moreover, this also merges similar component names.

Interpretation of Responses. Among the most challenging aspects of using LLMs is interpreting the generated responses because there is no guarantee that they adhere to the requested output format. In ExArch, the final output of the LLM should be a list of component names. This output is determined by the final prompt in each mode (cf. **Prompt 2**). To parse the responses, we consider every line of the final response. First, we check that the trimmed line starts with the character '-'. If this is the case, we consider the line as a component name. We remove the occurrences of 'components' and 'component' from the component names. Thus, we aim to only get the actual component names. Third, we remove any spaces from the component's name, as we requested the component names to be in camel case. After the initial experiments without ArTEMiS ([19]), we modified the algorithm to not only remove spaces, but to generate proper camel case by changing the capitalization. Finally, we remove any duplicates from the list of component names.

3.3 TransArC Approach

To be able to analyze the effect of LLM-extracted intermediate models on transitive TLR between SAD and code, we make use of the TransArC approach by Keim et al. [32]. TransArC links SAD to source code by using SAMs as intermediate artifacts to bridge the semantic gap. TransArC consists of two linking phases: linking documentation to SAMs, and the models to source code. The remainder of this section gives a brief overview of them.

In the first phase, TransArC uses the existing ArDoCo approach [33] to create trace links between SAD and SAMs. ArDoCo employs NLP techniques to analyze the SAD to identify architectural elements such as components. ArDoCo uses various similarity measures and heuristics to link the sentences in the SAD and components in the SAM.

In the second phase, TransArC uses ARchitecture-to-CODE Trace Linking (ArCoTL) to establish trace links between SAMs and the source code. ArCoTL first transforms the input artifacts into




Fig. 2. Overview of the ArTEMiS Approach for TLR Between SAD and SAM. ArTEMiS takes the SAD and a SAM as input. It uses two LLM prompts (blue): the first identifies architecturally relevant entities in the SAD text, with SAM entity names provided as examples (dashed arrow); the second converts the plain-text output to structured JSON. Line mapping then locates each identified entity in the original SAD text using exact and similarity-based matching. Finally, entity matching links the identified SAD entities to SAM entities using string similarity and vector embeddings, producing trace links between SAD and SAM. Artifacts are shown in orange, heuristic-based steps in white, and LLM prompting steps in blue.

intermediate representations. For SAMs, this includes identifying components, while for code, it involves extracting elements like classes, methods, and packages. ArCoTL then applies a series of heuristics to identify correspondences between architectural elements and code entities. A computational graph combines these heuristics to aggregate the confidence levels of candidates to form the final trace links.

Combining the results from both phases, TransArC generates transitive trace links between SAD and source code. It uses the intermediate SAMs to enhance the accuracy of trace link recovery, effectively reducing the semantic gap between the documentation and code. The evaluation of TransArC (cf. [32]), demonstrated its high performance in recovering trace links (weighted average F_1 -score of 0.87).

We use the TransArC approach to create trace links between the SAD and the source code. Therefore, we use the component names generated by the approach as SAM to enable trace linking between architecture documentation and source code without needing a manually created component model.

4 ArTEMiS: An Approach to Extract Architecture Entities from Software Architecture Documentation

In this section, we describe our approach ArTEMiS (Architecture Traceability with Entity Matching via Semantic inference) that takes SADs and SAMs as inputs to identify trace links between both artifacts. ArTEMiS adapts the ideas of two previous approaches of ours: SoftWare Architecture Text Trace link Recovery (SWATTR) [34], which provides trace link recovery between SAD and SAM that is also used in TransArC, and ExArch, which offers techniques for component identification in text. Accordingly, ArTEMiS uses a two-step approach: It first uses LLMs to detect architecturally relevant entities in SADs through a specialized Named Entity Recognition (NER) approach. Once identified, these entities are matched against actual architecture entities from SAMs to establish trace links.

Figure 2 provides an overview of ArTEMiS.

For the first step, the identification of architectural entities, ArTEMiS uses a two-prompt strategy.

The first prompt provides the task description (see Appendix A for details). This prompt instructs the LLM to identify explicitly named software components that are architecturally relevant. For each component, the LLM extracts: (1) a primary name, (2) any alternative names or abbreviations, and (3) complete lines where the component appears, either directly or through clear context. The

prompt includes specific rules defining what to include and exclude. The benefit of this prompt, unlike traditional heuristics or information retrieval approaches, is the capability of LLMs to identify coreferences and indirect references within their context. Additionally, ArTEMiS takes the names of architecture entities from the SAM as input for this prompt and appends them as positive examples to look out for. To ensure reliable parsing, the prompt asks the LLM to output results in structured plain-text format.

The second prompt (see [Appendix B](#)) converts the plain-text output into JSON format. We separated these tasks because combining them into a single prompt frequently produced malformed JSON. This two-step approach ensures more reliable and parseable results.

After obtaining the JSON output, ArTEMiS processes the results through several steps. First, it maps each identified component’s mentioned lines back to the original text. The mapping process begins by searching for exact matches between the mentioned lines and the original text. If none are found, it calculates string similarity scores to select the best match. We initially attempted to have the LLM return line numbers directly, but found that LLMs struggle with numerical accuracy even when line numbers are provided. Through this processing, ArTEMiS produces a structured set of components with their names, alternative names, and line locations.

Finally, ArTEMiS establishes connections between the SAM entities and the identified SAD entities. The matching process uses both string similarity metrics (Jaro-Winkler [65] and Levenshtein [38]) and vector embeddings similarity. For vector embeddings similarity, we use OpenAI’s model “text-embedding-3-large” to embed each of the name and alternative names of the identified named architecture entities from the SAD on the one hand as well as each name of the SAM entities on the other hand; we then use cosine similarity to calculate the similarity score between each pair of embedding from the SAD and embedding from the SAM. When one of the similarity scores exceeds a defined threshold, ArTEMiS creates a link between these entities. These links enable trace link generation between the various text locations where architectural entities are mentioned and their corresponding model elements.

ArTEMiS’s key advancement is replacing SWATTR’s traditional heuristics with modern LLMs for architectural entity identification. For this, ArTEMiS adopts ExArch’s idea to use LLMs to identify components in text, but maintains stronger text relationships by tracking both entity occurrences and alternative names throughout the documentation.

ArTEMiS can be deployed in three different settings. First, the approach was designed specifically for SAD-SAM TLR. Second, ArTEMiS can be integrated into TransArC to replace its existing SAD-SAM TLR approach SWATTR. Third, ArTEMiS can be combined with ExArch (using the extracted component names as SAM) and used within TransArC.

5 Experimental Evaluation

In this section, we present the evaluation setup ([Section 5.1](#)), including the dataset, evaluation metrics, and the LLMs used. Additionally, we present the baselines to which we compare our approach ([Section 5.2](#)). Afterward, [Section 5.3](#) discusses the initial experimental evaluation of ExArch by Fuchß et al. [19]. In [Section 5.4](#), we describe our full evaluation of ArTEMiS and ExArch. This section concludes in [Section 5.6](#) with our threats to validity.

5.1 Evaluation Setup: Data and Metrics

This section describes the datasets and evaluation metrics that we use for our evaluation. We use the same setup as in the original TransArC publication [32], i.e., we use the same dataset and evaluation metrics to compare the results. As ExArch generates all intermediate information needed, we do not rely on manually created SAMs for configurations that use ExArch.

Table 1. Number of Artifacts per Artifact Type and Number of Trace Links in the Gold Standard for each Project Based on [32].

Artifact Type		MS	TS	TM	BBB	JR
Sw. Architecture Documentation [17, 32]	# Sentences	37	43	198	85	13
Sw. Architecture Model	# Components	14	11	8	12	6
Source Code	# Files	97	205	832	547	1,979
SAD-Source Code [32]	# Trace links	50	707	7,610	1,295	8,240
SAD-SAM [17]	# Trace links	29	27	51	52	18

Benchmark Dataset. The TransArC approach uses a benchmark dataset from Fuchß et al. [17] comprising SAD to SAM trace links. Keim et al. [32] extended the dataset with trace links between SAD and source code. The dataset consists of five open-source projects, each differing in size and domain. The projects are MediaStore (MS), TeaStore (TS), TEAMMATES (TM), BigBlueButton (BBB), and JabRef (JR). The dataset contains the artifacts themselves and the ground truth for trace links. Table 1 provides an overview of the dataset. Every project has, at most, 14 components. The number of source code files ranges from around 100 to roughly 2000, and the SADs comprise 13 up to 198 sentences.

Evaluation Metrics. We use commonly used metrics for TLR tasks [7, 25]: precision, recall, and their harmonic mean F_1 -score. This way, we can compare our approach's performance to the reported results of other state-of-the-art approaches.

$$\text{Precision} = \frac{TP}{TP + FP}, \text{ Recall} = \frac{TP}{TP + FN}, F_1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

We define the following: True positives (TPs) are found trace links between the SAD and the source code that are also contained in the gold standard. False positives (FPs) are found trace links that are not contained in the gold standard. False negatives (FNs) are trace links contained in the gold standard but not identified by the approach.

Additionally, we use two different averaging methods. First, we report the overall (macro) average across all projects without considering their size. This average provides useful insights into the expected performance on a per-project basis. Second, we calculate a weighted average based on the number of expected trace links in the gold standard [32]. This weighting offers more in-depth insights into the anticipated effectiveness of the approach for each trace link.

Large Language Models (LLMs). We use various LLMs to generate the simple SAMs. We decided to use both closed-source models by OpenAI and locally deployed open-source models. For OpenAI, we use the following models: *GPT-4o mini*, *GPT-4o*, *GPT-4 Turbo*, *GPT-4*, and *GPT-3.5 Turbo*. As local models, we use *Codellama 13b*, *Meta AI Llama 3.1 8b*, and *Meta AI Llama 3.1 70b*.

5.2 Baselines

We reuse the baseline approaches of Keim et al. for TransArC [32]. Thus, we can directly compare their results to ours. The descriptions of the baselines are based on [32].

TAROT [22] and FTLR [28] are both recent, state-of-the-art IR-based solutions designed for linking requirements to code. CodeBERT [16] is an LLM trained to find the most semantically related source code for a given natural language description. Consequently, all three methods show promising results for similar TLR problems.

Table 2. Precision, Recall, and F_1 -score for TLR Between SAD and Code (ExArch: SAM Derived From SAD). In the Baselines section, TransArC uses manually created SAMs; all other baselines (TAROT, FTLR, CodeBERT, ArDoCode, LiSSA) operate without SAMs. ExArch generates component names from SAD using the listed LLMs and passes them to TransArC as simple SAMs, removing the need for manually created SAMs. Best result per project within each section is highlighted in bold.

Project	Met.	Baselines						ExArch							
		TAROT	FTLR	CodeBERT	ArDoCode	LiSSA	TransArC	GPT-4o mini	GPT-4o	GPT-4 Turbo	GPT-4	GPT-3.5 Turbo	Codellama 13b	Llama 3.1 8b	Llama 3.1 70b
MS	Pre.	.09	.15	.29	.05	.06	1.0	.49	.49	.45	.49	.49	.81	.49	.48
	Rec.	.24	.26	.12	.66	.72	.52	.52	.52	.40	.52	.52	.52	.52	.50
	F_1	.13	.19	.17	.09	.11	.68	.50	.50	.43	.50	.50	.63	.50	.49
TS	Pre.	.19	.19	.26	.20	.34	1.0	.95	.96	.96	.96	.96	.96	.21	.62
	Rec.	.44	.25	.57	.74	.33	.71	.67	.67	.71	.71	.67	.67	.71	.80
	F_1	.27	.21	.36	.31	.34	.83	.78	.79	.82	.82	.79	.79	.33	.70
TM	Pre.	.06	.06	.09	.37	.07	.71	.71	.71	.71	.71	.71	.66	.66	.62
	Rec.	.32	.30	.22	.92	.03	.91	.90	.90	.90	.90	.90	.49	.34	.31
	F_1	.11	.10	.12	.53	.04	.80	.80	.80	.80	.80	.80	.56	.45	.41
BBB	Pre.	.07	.04	.07	.07	.20	.77	.71	.74	.71	.63	.75	.05	.73	.62
	Rec.	.18	.42	.49	.57	.21	.91	.64	.77	.66	.84	.54	.32	.47	.43
	F_1	.10	.07	.12	.13	.21	.84	.68	.75	.69	.72	.63	.08	.57	.51
JR	Pre.	.32	.32	.49	.66	.51	.89	.89	.89	.89	.89	.89	.89	.89	.89
	Rec.	1.0	.93	.83	1.0	.01	1.0	1.0	1.0	1.0	.99	1.0	.99	1.0	1.0
	F_1	.49	.48	.61	.80	.03	.94	.94	.94	.94	.94	.94	.94	.94	.94
Avg	Pre.	.15	.15	.24	.27	.23	.87	.75	.76	.75	.74	.76	.67	.60	.65
	Rec.	.44	.43	.45	.78	.26	.81	.75	.77	.73	.79	.73	.60	.61	.61
	F_1	.22	.21	.28	.37	.14	.82	.74	.76	.73	.76	.73	.60	.56	.61
w. Avg	Pre.	.19	.19	.28	.47	.29	.81	.80	.81	.80	.80	.81	.73	.75	.74
	Rec.	.63	.59	.53	.92	.05	.94	.92	.93	.92	.93	.91	.72	.67	.66
	F_1	.29	.28	.36	.62	.06	.87	.85	.86	.86	.86	.85	.71	.68	.68

Keim et al. [32] also introduced the ArDoCode approach that uses heuristics to recover trace links between SAD and source code without using intermediate artifacts, that is based on the SWATTR approach [33, 34]. They reported that, on average, ArDoCode was the best-performing approach that does not need SAMs. Thus, this approach is one of the important baselines. Also, we compare to the original TransArC approach [32] using (manually created) SAMs.

Moreover, we compare our approaches to the LiSSA approach [18] that uses retrieval-augmented generation and LLMs to recover trace links directly. This approach has been designed to be applicable to various TLR tasks and was also recently applied to the same dataset.

Table 3. Results of a One-Sided Wilcoxon Signed-Rank Test Regarding If ExArch (GPT-4o) Using SAD for Recovery Performs Better Than the Baseline Approaches (Significance Level $\alpha = 0.05$, P-values With * Cannot Be Calculated Exactly.)

Approach / Hypothesis	Requires SAM	p-value	Significant
TAROT	No	.031	Yes
FTLR	No	.031	Yes
CodeBERT	No	.029*	Yes
ArDoCode	No	.031	Yes
LiSSA	No	.031	Yes
TransArC (ours better)	Yes	.970*	No
TransArC (ours worse)	Yes	.091*	No

5.3 ExArch and TransArC

In this section, we present the results of our evaluation of how to extract component names from SAD and/or source code using ExArch and how we integrated it with TransArC [19].

5.3.1 Extracting Component Names from SAD. This section presents the results of our approach ExArch with TransArC if we only use the SAD to generate the component names of a SAM. Table 2 provides a detailed overview of the results.

The table shows the precision, recall, and F_1 -score for each project and the average values. The table consists of two sections, one for the baseline approaches and one for our approach using different LLMs. We highlight the overall best results per project in each section. In the first section, we took the results of the approaches as presented in the work of Keim et al. [32]. Here, TransArC uses the manually created SAM, while the other baseline approaches only use the SAD and source code. Overall, the best-performing baselines are ArDoCode and TransArC.

The second section of the table presents the results of ExArch using different LLMs. We can see that the performance varies across the projects. We observe that for the weighted average F_1 -score, the models by OpenAI perform particularly well. Furthermore, we can see that they perform similarly to TransArC without the need for manually created SAMs. The best model w.r.t. weighted average F_1 -score is *GPT-4o* with a weighted average F_1 -score of 0.86 compared to 0.87 of TransArC. The other OpenAI models perform similarly. According to the classification of Hayes et al. [25], ExArch *excellently* recovers trace links between SAD and code (GPT-4o). We highlight that all models outperform the baseline approaches that also do not use manually created SAMs, and thus, require the same input as ExArch with TransArC.

Significance Tests. We use Wilcoxon's signed-rank test (one-sided) to calculate the statistical significance of ExArch's F_1 -score compared to the other baselines. Since ExArch works best using GPT-4o as LLM, we only compare the results of this configuration to the baselines. We present the results in Table 3. We mark those p-values with an asterisk that cannot be calculated exactly due to ties in the data. The table shows that ExArch significantly outperforms all baselines that only use SAD and code. Yet, our approach does not outperform TransArC using manually created SAM ($p = 0.97$). At the same time, TransArC is also not significantly outperforming our LLM-based approach ($p = 0.09$). This shows that ExArch performs comparably to TransArC without the need for manually created SAMs.

Table 4. F_1 -score for TLR Between SAD and Code (ExArch: SAM Derived From Source Code). ExArch generates component names from source code package names using the listed LLMs and passes them to TransArC as simple SAMs. Best result per column is highlighted in bold.

Approach	MS	TS	TM	BBB	JR	Avg.	w. Avg.
GPT-4o mini	.16	.45	.79	.23	.51	.43	.60
GPT-4o	.28	.64	.78	.12	.69	.50	.68
GPT-4 Turbo	.55	.28	.34	.00	.94	.42	.59
GPT-4	.31	.32	.61	.00	.74	.40	.61
GPT-3.5 Turbo	.22	.68	.72	.16	.32	.42	.49
Codellama 13b	.49	.00	.03	.10	.94	.31	.46
Llama 3.1 8b	.39	.45	.04	.00	.94	.37	.47
Llama 3.1 70b	.11	.64	.75	.46	.94	.58	.81

Conclusion RQ1.1: Applying TransArC with LLM-extracted SAMs by ExArch produces similar results as with manually created SAMs (ExArch (GPT-4o) weighted avg. F_1 -score: 0.86 vs. TransArC 0.87; not significantly different at $\alpha = 0.05$; see [Table 2](#), [Table 3](#)).

Conclusion RQ1.2: ExArch significantly outperforms state-of-the-art TLR approaches between SAD and code that do not use SAMs as intermediates (e.g., ExArch (GPT-4o) weighted avg. F_1 -score 0.86 vs. ArDoCode 0.62; $p < 0.05$; see [Table 2](#), [Table 3](#)).

Conclusion RQ1.3: On average, OpenAI’s closed-source LLMs outperform the open-source Llama-based models in this task (best closed-source: GPT-4o, weighted avg. F_1 -score 0.86; best open-source: Codellama 13b, 0.71; see [Table 2](#)).

5.3.2 Generate Component Names from Source Code. This section focuses on **RQ1.4** and presents the results of ExArch with TransArC if we only use the source code to generate the component names. [Table 4](#) provides a detailed overview of the results. As described in [Section 3](#), we provide a list of all non-empty packages as features for the prompts. We argue that the packages can be a good representation of the high-level structure of a software project. Nevertheless, comparing the results to ExArch using the SAD to extract the component names, the average performance using only source code is lower. Especially, the performance for the project *BBB* is bad, i.e., for many models, the F_1 -score is 0 or close to 0. The overall best model is *Llama 3.1 70b* with a weighted average F_1 -score of 0.81. Notably, this performance is better than its performance when using the SAD-extracted SAMs. Nevertheless, the overall average F_1 -score is worse.

Since, on average, all other models perform worse than with SAD, we conclude that the packages alone are insufficient to generate SAMs for TLR. Moreover, this mode’s performance is worse than the mode that only considers documentation. Further, the results vary even more across projects and LLMs.

5.3.3 Generate Component Names from SAD and Code. This section presents the results of ExArch with TransArC if we use both the SAD and the source code to generate the component names

Table 5. F₁-score for TLR Between SAD and Code (ExArch: SAM Derived From Both SAD and Source Code). Component names extracted from both sources are merged using either LLM-based prompt aggregation (top) or Levenshtein-distance-based similarity aggregation ($t = 0.5$, bottom). Best result per column and type of aggregation is highlighted in bold.

Approach	MS	TS	TM	BBB	JR	Avg.	w. Avg.
<i>Prompt-Aggregation</i>							
GPT-4o mini	.05	.62	.79	.34	.51	.46	.62
GPT-4o	.06	.58	.78	.42	.94	.56	.82
GPT-4 Turbo	.43	.82	.80	.63	.94	.72	.85
GPT-4	.31	.80	.62	.53	.94	.64	.77
GPT-3.5 Turbo	.44	.64	.74	.44	.93	.64	.80
Codellama 13b	.44	.56	.42	.13	.94	.50	.65
Llama 3.1 8b	.43	.00	.00	.00	.94	.27	.44
Llama 3.1 70b	.11	.64	.71	.22	.94	.52	.78
<i>Similarity-Aggregation ($t = 0.5$)</i>							
GPT-4o mini	.27	.51	.79	.49	.93	.60	.82
GPT-4o	.28	.64	.78	.54	.89	.63	.81
GPT-4 Turbo	.43	.82	.80	.69	.94	.73	.86
GPT-4	.31	.80	.61	.70	.94	.68	.78
GPT-3.5 Turbo	.22	.70	.72	.44	.93	.60	.80
Codellama 13b	.63	.56	.56	.10	.94	.56	.70
Llama 3.1 8b	.63	.34	.46	.44	.94	.56	.68
Llama 3.1 70b	.13	.64	.78	.45	.94	.59	.82

of the SAM. Here, we consider the two different aggregation strategies described in Section 3.2: the LLM-based strategy and the similarity-based one. We use a similarity threshold of $t = 0.5$ in the experiments for the normalized Levenshtein distance. Nevertheless, this threshold can also be adjusted to the specific needs of a project.

Aggregation via Prompting. We present our results in F₁-score regarding the aggregation via prompting in the first part of Table 5. The data shows that the best-performing model is *GPT-4 Turbo* with a weighted average F₁-score of 0.85. The results are slightly worse in macro average, compared to the mode that only uses the SAD but mostly better than the mode that only uses the source code.

Aggregation via Similarity. Next, we present the results of our evaluation of the aggregation via similarity. As described in Section 3.2, we use the normalized Levenshtein distance to aggregate common names.

We present our results in the second part of Table 5. The aggregation is better than the results when only code is used. The best-performing model is *GPT-4 Turbo* with a weighted average F₁-score of 0.86. Nevertheless, the aggregation via similarity also performs not better than the mode that only uses the SAD, particularly when considering the non-weighted average F₁-score. However, on average, the performance is often better than the aggregation via prompt.

Conclusion RQ1.4: On average, using only the SAD to generate the simple SAM performs best (best avg. F_1 -score: SAD-only 0.76 (GPT-4o), code-only 0.58 (Llama 3.1 70b), both-aggregated 0.72–0.73 (GPT-4 Turbo); see Table 2, Table 4, Table 5).

Overall, we conclude that to extract performing component names for TLR, the SAD is a better source than the source code. LLMs promise to have great language comprehension capabilities. Thus, we assume that the task of summarizing and extracting component names based on SAD is easier for the LLMs than inferring component names from source code.

5.3.4 Discussion & Error Analysis. This section discusses our results and analyzes some errors of the LLMs. First, we discuss exemplary differences between the extracted SAMs from the SAD using different LLMs. Second, we analyze exemplary differences between extracted SAMs from SAD and source code. Finally, we discuss the traceability from the SAD to the generated SAM.

Differences between extracted SAMs from SAD using different LLMs. As seen in Section 5.3.1, the performance of the different LLMs varies across the projects. To provide insights into the reasons for that, we perform a detailed analysis on the *MediaStore* project with a particular focus on the difference between the results of *Codellama 13b* and *GPT-4o*. We picked this example because they show a rather large gap in precision (*Codellama 13b* with 0.81 vs 0.49 of *GPT-4o*) while having the same recall. The first column of Figure 3 shows the SAM extracted by *Codellama 13b*, the second column the component names of the original, manually created SAM, and the third column the SAM extracted by *GPT-4o*. While the SAMs are not the same, they share similar component names. The connections between the names indicate that we treat them as being similar. The highlighting of nodes indicates whether they are matched (green), not matched (red), or related (yellow). Both generated SAMs are missing components like *Cache* that are part of the manually created SAM but not described in the SAD. Besides that, the difference is that *Codellama 13b* generates a component *PersistenceTier* and *GPT-4o* generates a component called *DataStorage*. This small difference causes the drop in precision from 0.81 to 0.49 in the performance for the TLR task. Therefore, we also emphasize that according to model theory, the purpose of a model (e.g., an SAM) is important [57]. *Codellama 13b* creates component names that are more suitable for the considered TLR task. However, this does not mean that they are better for other tasks, and we do not assume that these SAMs can be directly used in reverse engineering tasks, as these tasks would require more details.

Differences between extracted SAMs from SAD and Code. In particular, we are interested in the actual differences between the extracted SAMs from the SAD and the source code. Since this requires manual analysis, we only focus on the projects with fewer components according to the manually created SAMs (cf. Table 1), *JabRef* and *TEAMMATES*.

First, we analyze the differences between the extracted SAMs for the project *JabRef* using *Llama 3.1 70b*. Our evaluation shows that the performance regarding TLR is the same as *TransArC* using the manually created SAM. In Figure 4a, we show the component names *Llama 3.1 70b* extracts. The first column defines the components extracted using only SAD, the second column the components of the manually created SAM, and the third column the components extracted from the source code. In the third column, we only show the elements that are extracted as main components. The LLM also generates a list of sub-components with no effects on performance. We use the same color scheme as in the previous sections. The names for the matching components of the modes only differ in their letter case, *globals* is missed by both modes, the SAD-extracted SAM additionally includes *EventBus*, and the code-extracted adds *Networking*. The similar results in TLR suggest that these differences do not matter for the specific TLR task.

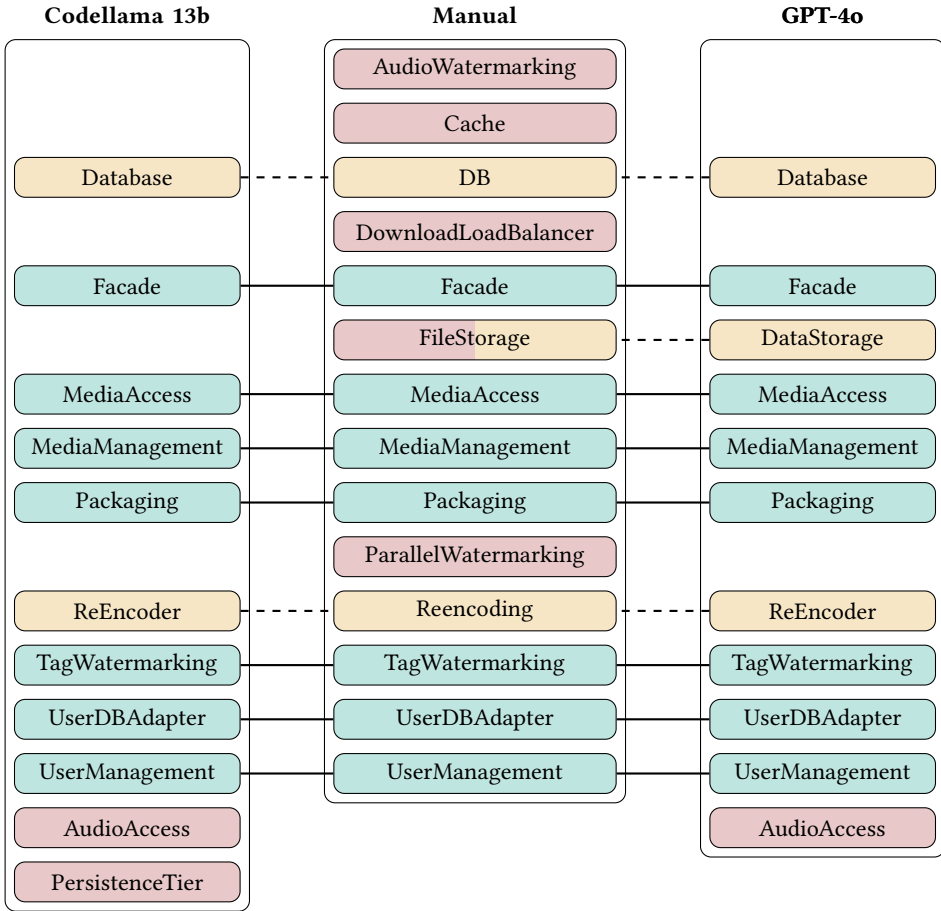
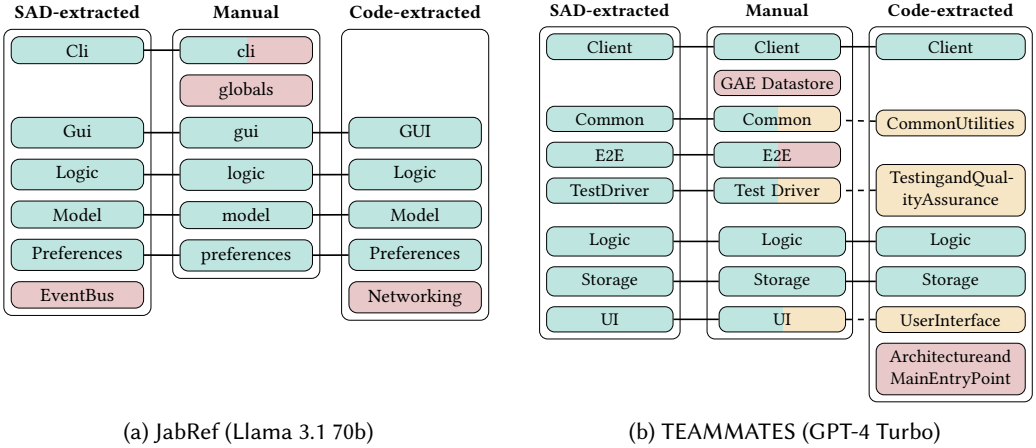


Fig. 3. Comparison of SAMs Extracted from SAD for MediaStore: manually created SAM (center) vs. ExArch with Codellama 13b (left) and GPT-4o (right). Green nodes are exact matches between extracted and manual components (solid lines); orange nodes are approximate matches where names differ but refer to the same concept (dashed lines); red nodes have no match.

Second, we analyze the differences between the extracted SAMs for the project *TEAMMATES* using *GPT-4 Turbo*. Here, our evaluation shows that using SAD, the performance is comparable to TransArC, while using the source code, the performance is worse.

In Figure 4b, we show the component names *GPT-4 Turbo* extracts. We can see that our extracted SAM from SAD is only missing the *GAE Datastore* component. The extraction from the source code diverges more. Also, many components are correctly extracted, some components are abbreviated in the manually created SAM, while the LLM extracts the full name from the source code (e.g., *UserInterface*). However, the LLM also extracts additional components that are not present in the manually created SAM. The component *ArchitectureandMainEntryPoint* originates from the LLM's output "Architecture and Main Entry Point". Since we instructed the LLMs to provide component names in camel case, the final component name is as described. While there are packages *teammates.architecture* and *teammates.main* in the source code, the LLM could not identify that these packages should not be mapped to a new component. The first package contains an



(a) JabRef (Llama 3.1 70b)

(b) TEAMMATES (GPT-4 Turbo)

Fig. 4. Comparison of SAMs Extracted from SAD and Source Code vs. manually created SAMs for JabRef (Llama 3.1 70b) and TEAMMATES (GPT-4 Turbo). Each subfigure shows the SAD-extracted SAM (left), the manually created SAM (center), and the code-extracted SAM (right). Green nodes are exact matches (solid lines); orange nodes are approximate matches where names differ but refer to the same concept (dashed lines); red nodes have no match. For JabRef, only the main components are shown in the code-extracted column.

architecture test and the second package contains the application’s main entry point. The component *TestingandQualityAssurance* is also absent in the manually created SAM. While the source code contains packages for test cases, these are not part of the manually created SAM. We mark the components that are related to the manually created SAM with dashed lines. We assume that the testing-related components extracted via code belong to the *Test Driver* component from the manually created SAM. This example shows that if we consider source code for extracting simple SAMs, distinguishing between code for production and code for testing is challenging.

Traceability from SAD to generated SAM. To discuss the performance of ExArch, we briefly analyze the performance of the TLR approach from SAD to the generated SAM. Since we do not have a gold standard for this task, we decided to analyze our results using the following steps: First, we identify the components generated by the LLMs. Second, we manually match the components that are also present in the original SAM. Those components are assigned the same unique identifier as in the original SAM. All other components are assigned a new unique identifier. By doing so, we can use the evaluation of ArDoCo [33] to analyze the TLR from SAD to the generated SAM. Since we want to do this to get insights into the overall performance from SAD to source code, we only selected some projects and LLMs for analysis, based on our evaluation results in Table 2.

We first analyze the SAD-generated SAM for *MediaStore* by *GPT-4o* (see Figure 3). We selected this model because the LLM is the overall best performing. Compared to the original SAM, the generated SAM contains 11 components, while the original SAM contains 14 components. It also contains two components that are not present in the original SAM: *AudioAccess* and *DataStorage*. Running ArDoCo with the generated SAM increases recall for the TLR between SAD and SAM from 0.62 to 0.79 compared to running with the manually created SAM. The precision decreases from 1.0 to 0.68. Our results for TLR from SAD to code also reflect this.

We then analyze the generated SAM for *BigBlueButton* by *Codellama 13b*. We selected this model because it is the best-performing open-source LLM, except for *BBB*. Comparing the generated SAM

Table 6. Number of Artifacts per Artifact Type and Number of Trace Links in the revised Gold Standard for each Project. Underline marks the changed values compared to Table 1.

Artifact Type		MS	TS	TM	BBB	JR
Sw. Architecture Documentation	# Sentences	37	43	198	<u>87</u>	13
Sw. Architecture Model	# Components	14	11	8	12	6
Source Code	# Files	97	205	832	547	1,979
SAD-Source Code	# Trace links	<u>59</u>	707	<u>8,097</u>	<u>1,529</u>	<u>8,268</u>
SAD-SAM	# Trace links	<u>31</u>	27	<u>57</u>	<u>62</u>	18

to the original SAM, the generated SAM contains 11 components, while the original SAM contains 12. The LLM additionally generates a component called *BigBlueButton*. Also, it slightly renamed the components *Apps* and *FSESL* to *AppsAkka* and *FSESLAkka*. The LLM could not generate the components *BBB web* and *Recording Service*. It also generated names like *Joiningavoicereference* that represent tasks but not a component name. Thus, the recall of the TLR from SAD to the generated SAM is 0.19 and the precision is 0.27. With the original SAM, the recall is 0.83 and the precision is 0.88. If we consider the package names of the code, one problem can be the generation of the *BigBlueButton* component. Since all Java packages of the project start with *org.bigbluebutton*, the component can lead to many false-positive trace links that decrease precision.

5.4 The Impact of ArTEMiS

This section covers our evaluation of ExArch and ArTEMiS. First, we describe how we revised the benchmark after the initial experiments of ExArch [19]. Afterward, we discuss the results of our evaluation of ArTEMiS for SAD-SAM TLR. Finally, we evaluate TransArC and ExArch with the integrated ArTEMiS approach.

5.4.1 Revision of the Benchmark. After the initial experiments with ExArch [19], we re-assessed trace links present in the benchmark and created a revised version of it. We identified some missing trace links (mostly coreferences) and added them. Also, we identified false positives and removed them. Finally, we fixed a part of the documentation of *BigBlueButton* where a few sentences were wrongly terminated. The characteristics of the revised benchmark are summarized in Table 6. The revision of trace links between SAD and SAM also affected the amount of trace links between SAD and the code. We use this revised version of the benchmark for the following evaluations.

In Table 7, we present the results of our ExArch approach for the revised benchmark. We only consider the extraction of component names from the SAD, as our experiments showed that this is the best-performing mode. In comparison to the initial version, ExArch now tries to recover camel-cased component names by removing spaces and enforcing camel-casing in post-processing. Moreover, we included recent versions of the LLMs, i.e., *GPT-4.1* and *GPT-5*, to provide more insights into the performance of the most recent LLMs. As you can see, the results are similar to the initial experiment. Also, the newer models perform similarly, but the best-performing model is *GPT-4o* with an average F_1 -score of 0.76.

5.4.2 ArTEMiS for SAD-SAM TLR. In this section, we present the results of our new ArTEMiS approach for TLR between SAD and SAM. To compare our approach with the state of the art, we use the component of TransArC that creates the trace links between SAD and SAM — namely SWATTR [33, 35]. We execute ArTEMiS with three recent LLMs: *GPT-4o*, *GPT-4.1*, and *GPT-5*. Since ArTEMiS does not produce deterministic results, we run each configuration five times and

Table 7. Precision, Recall, and F_1 -score for TLR Between SAD and Code with ExArch on the Revised Benchmark (SAM Derived From SAD). Results are shown for three LLMs (GPT-4o, GPT-4.1, GPT-5). The revised benchmark contains corrected and additional trace links compared to the original (cf. Table 6). Best result per metric is highlighted in bold.

Project	Metric	GPT-4o	GPT-4.1	GPT-5
MediaStore	Precision	.51	.51	.51
	Recall	.46	.46	.46
	F_1 -score	.48	.48	.48
TeaStore	Precision	.96	.59	.62
	Recall	.67	.69	.80
	F_1 -score	.79	.64	.70
Teammates	Precision	.75	.72	.75
	Recall	.90	.90	.90
	F_1 -score	.82	.80	.82
BigBlueButton	Precision	.80	.67	.70
	Recall	.73	.74	.84
	F_1 -score	.77	.70	.76
JabRef	Precision	.89	.89	.89
	Recall	1.0	1.0	1.0
	F_1 -score	.94	.94	.94
Average	Precision	.78	.68	.69
	Recall	.75	.76	.80
	F_1 -score	.76	.71	.74
w. Average	Precision	.83	.79	.80
	Recall	.92	.92	.93
	F_1 -score	.87	.85	.86

report average results with standard deviation. Table 8 shows the results of our evaluation. The best ArTEMiS configuration w.r.t. F_1 -score uses *GPT-5*. On average, this configuration is slightly better than SWATTR. Notably, ArTEMiS achieves a better average recall than SWATTR. This is mainly due to the increased recall for MediaStore and TeaStore. For the other projects, the recall is decreased. However, the standard deviation is different across the projects and LLMs. In particular, for Teammates with *GPT-5*, the standard deviation of the precision is comparably high with 0.10. Still, we can summarize that ArTEMiS offers an alternative to traditional heuristics-based approaches like SWATTR for TLR between SAD and SAM.

Conclusion RQ2.1: ArTEMiS (GPT-5) outperforms the state-of-the-art heuristics-based approach SWATTR in average recall (0.85 vs. 0.77) and achieves a slightly higher average F_1 -score (0.81 vs. 0.80; see Table 8).

5.4.3 ArTEMiS and ExArch in TransArC. In this section, we present the results of combining our approaches ExArch, ArTEMiS, and TransArC. For the SAD to code TLR task, we consider three configurations that are compared to the original TransArC approach and the baseline ArDoCode. First,

Table 8. Comparison of ArTEMiS to its heuristic-based baseline SWATTR: Average Precision, Recall, and F₁-score (\pm Std. Dev.) of SAD to SAM TLR over five runs. ArTEMiS is executed with different LLMs (GPT-4o, GPT-4.1, GPT-5). Best result per metric is highlighted in bold.

Project	Metric	SWATTR (<i>baseline</i>)	ArTEMiS		
			GPT-4o	GPT-4.1	GPT-5
MediaStore	Precision	.94	1.0 (\pm .00)	.99 (\pm .01)	.87 (\pm .01)
	Recall	.55	.85 (\pm .03)	.92 (\pm .02)	.97 (\pm .02)
	F ₁ -score	.69	.92 (\pm .02)	.96 (\pm .01)	.92 (\pm .01)
TeaStore	Precision	1.0	.65 (\pm .02)	.63 (\pm .02)	.59 (\pm .01)
	Recall	.74	.86 (\pm .04)	.96 (\pm .00)	.99 (\pm .02)
	F ₁ -score	.85	.74 (\pm .01)	.76 (\pm .01)	.74 (\pm .01)
Teammates	Precision	.60	.89 (\pm .05)	.58 (\pm .07)	.78 (\pm .10)
	Recall	.86	.55 (\pm .04)	.68 (\pm .04)	.65 (\pm .07)
	F ₁ -score	.71	.68 (\pm .04)	.62 (\pm .04)	.70 (\pm .03)
BigBlueButton	Precision	.90	.84 (\pm .05)	.83 (\pm .04)	.86 (\pm .03)
	Recall	.71	.37 (\pm .08)	.63 (\pm .06)	.62 (\pm .03)
	F ₁ -score	.79	.51 (\pm .09)	.71 (\pm .03)	.72 (\pm .02)
JabRef	Precision	.90	.93 (\pm .02)	.87 (\pm .02)	.97 (\pm .04)
	Recall	1.0	.90 (\pm .05)	1.0 (\pm .00)	1.0 (\pm .00)
	F ₁ -score	.95	.91 (\pm .03)	.93 (\pm .01)	.98 (\pm .02)
Average	Precision	.87	.86	.78	.81
	Recall	.77	.71	.84	.85
	F ₁ -score	.80	.75	.80	.81
w. Average	Precision	.83	.86	.76	.81
	Recall	.76	.62	.77	.77
	F ₁ -score	.77	.69	.75	.77

the configuration *TransArC + ArTEMiS*: Here, we replaced the SWATTR component of TransArC that recovers trace links between SAD and SAM with our new ArTEMiS approach using its best LLM (i.e., *GPT-5*). Here, the comparison to the original TransArC approach is interesting as it also needs existing SAMs. Second, the configuration *TransArC + ExArch*: This is the same configuration that has been evaluated in the initial experiment [19]. Here, we use ExArch with its best performing LLM, i.e., *GPT-4o*, to extract the component names from the SAD (cf. Table 7). Third, we analyze the configuration *TransArC + ExArch + ArTEMiS*: This configuration combines both approaches, i.e., we use ExArch to extract the component names from the SAD, then use ArTEMiS to recover the trace links between SAD and SAM, and finally use TransArC to recover the trace links between SAD and code. For both configurations, the interesting comparison is mainly ArDoCode, as it is the best baseline that does not require SAMs for the SAD to code TLR task.

We present the results of our evaluation on SAD to code TLR in Table 9. Comparing the results of TransArC with ArTEMiS to the original TransArC, we can see that the average recall is increased from 0.77 to 0.87. This is due to the increased recall for MediaStore and TeaStore, analogous to the results of ArTEMiS vs. SWATTR. However, we can also see that the weighted average recall

Table 9. Combining Approaches: Average Precision, Recall, F_1 -score (\pm Std. Dev.) of SAD to Code TLR over five runs. TransArC marks the baseline of TransArC + ArTEMiS. ArDoCode is the baseline for TransArC + ExArch and TransArC + ArTEMiS + ExArch. We use the best-performing LLMs for our approaches: For ExArch, we use GPT-4o; for ArTEMiS, we use GPT-5.

Project	Metric	TransArC (<i>baseline</i>)	TransArC + ArTEMiS	TransArC + ArTEMiS + ExArch	TransArC + ExArch	ArDoCode (<i>baseline</i>)
Needs SAMs		✓	✓	✗	✗	✗
MediaStore	Precision	.96	.86 (\pm .05)	.53 (\pm .00)	.51	.05
	Recall	.42	.96 (\pm .03)	.51 (\pm .00)	.46	.59
	F_1 -score	.59	.91 (\pm .02)	.52 (\pm .00)	.48	.10
TeaStore	Precision	1.0	.59 (\pm .00)	.54 (\pm .00)	.62	.20
	Recall	.71	.99 (\pm .01)	1.0 (\pm .00)	.80	.74
	F_1 -score	.83	.74 (\pm .00)	.70 (\pm .00)	.70	.31
Teammates	Precision	.75	.96 (\pm .06)	.94 (\pm .07)	.75	.39
	Recall	.90	.67 (\pm .05)	.68 (\pm .05)	.90	.92
	F_1 -score	.82	.79 (\pm .03)	.79 (\pm .02)	.82	.55
BigBlueButton	Precision	.82	.87 (\pm .04)	.78 (\pm .01)	.70	.08
	Recall	.84	.74 (\pm .03)	.64 (\pm .01)	.84	.56
	F_1 -score	.83	.80 (\pm .02)	.70 (\pm .00)	.76	.14
JabRef	Precision	.89	.98 (\pm .04)	.98 (\pm .04)	.89	.66
	Recall	1.0	.99 (\pm .01)	1.0 (\pm .00)	1.0	1.0
	F_1 -score	.94	.99 (\pm .02)	.99 (\pm .02)	.94	.80
Average	Precision	.88	.85	.75	.69	.28
	Recall	.77	.87	.77	.80	.76
	F_1 -score	.80	.85	.74	.74	.38
w. Average	Precision	.83	.95	.93	.80	.48
	Recall	.93	.83	.83	.93	.92
	F_1 -score	.87	.88	.87	.86	.62

is decreased from 0.93 to 0.83, reflecting the decreased recall for projects with more trace links between SAD and code. In terms of F_1 -score, we can see that the new ArTEMiS approach slightly increases the average and weighted average F_1 -score of TransArC. Considering the combination of all approaches, i.e., TransArC with ExArch and ArTEMiS, we can see that the results are similar to the results of TransArC with ExArch. This configuration outperforms the baseline approach ArDoCode in terms of F_1 -score. While the results are similar to the first experiment without ArTEMiS, we can see some interesting performances. Notably, the results for JabRef are nearly perfect. Also, the combination of ExArch and ArTEMiS leads to a better F_1 -score in weighted average, meaning that it currently is the best-performing SAD to code TLR approach that does not require SAMs.

Conclusion RQ2.2: ArTEMiS (GPT-5) improves both average recall (0.87 vs. 0.77) and average F_1 -score (0.85 vs. 0.80) when integrated into TransArC. The combination TransArC + ExArch (GPT-4o) + ArTEMiS (GPT-5) achieves the best weighted avg. F_1 -score among approaches without SAMs (0.87 vs. ArDoCode 0.62; see Table 9).

5.5 LLMs in Architectural Element Extraction

In our work, we used two LLM-based approaches, ExArch and ArTEMiS, to bridge the semantic gap between Software Architecture Documentation (SAD) and code. We further combine them with TransArC and enable SAD-to-code-TLR without manually created Software Architecture Models. In doing so, they collectively answer our overarching research question:

Conclusion RQ0: Using LLMs to extract architectural elements yields TLR performance between SAD and code comparable to approaches that rely on manually created SAMs. Integrating ExArch into TransArC yields a weighted average F_1 -score of 0.86 (vs. 0.87 with manual SAMs), and integrating ArTEMiS raises average recall from 0.77 to 0.87. The combination TransArC + ExArch + ArTEMiS attains a weighted average F_1 -score of 0.87, well above the strongest baseline without SAMs (ArDoCode, 0.62).

Besides that, the design and evaluation of ExArch and ArTEMiS yield broader lessons about using LLMs in software engineering tasks. Both ExArch and ArTEMiS follow the same design principle: They decompose the overall TLR problem and apply LLMs only to the subtasks where they excel (ExArch extracts component names from natural language or code; ArTEMiS identifies architectural entities in SAD text) while delegating the final trace link recovery to the heuristics of TransArC. Both also adopt a multi-step prompting strategy. ExArch uses two prompts per mode: one to elaborate on the architecture, one to produce a rather basic list of names. A simple list format is easier and more reliable to parse into compared to structured data formats, as minor formatting deviations do not break parsing. ArTEMiS requires richer output (entity names, alternative names, and line references), making a structured JSON format necessary but also harder to produce reliably in a single prompt; hence, the second prompt is dedicated solely to converting the plain-text output into JSON. As ArTEMiS operates on longer prompts (containing the full SAD text, the SAM entity names, and detailed task instructions) and requires a structured JSON formatting step, we run each configuration multiple times and report average results with standard deviation.

Why does SAD-based extraction work better? We assume that the semantic gap between SAD and component names is smaller than between code and component names, enabling LLMs to act like named-entity recognizers on SAD. In contrast, component boundaries and names are often implicit in code. Still, simple SAMs produced even by smaller LLMs can be sufficient to enable TransArC to outperform many state-of-the-art baselines on SAD-to-code TLR. The purpose of these SAMs matters: the names best for TLR are not necessarily suited for reverse engineering tasks requiring richer architectural detail.

5.6 Threats to Validity

In this section, we address potential threats to validity, drawing upon the guidelines provided by Runeson and Höst [52]. We also consider the work of Sallou et al. [54], which discusses threats specific to experiments involving LLMs in the context of software engineering. Unless stated otherwise, the following threats apply to both ExArch and ArTEMiS, as well as to their combination with TransArC.

Regarding *construct validity*, one potential threat is dataset bias. To mitigate this, we use the same datasets as those employed in state-of-the-art approaches for TLR. By doing so, we ensure coverage of various domains and project sizes. Moreover, our study does not emphasize prompt engineering, meaning we do not extensively modify or compare different prompts. As a result, the selection of prompts potentially threatens the validity of our findings regarding our research questions. Another concern is the possible bias in metric selection. To address this, we employ commonly used metrics in TLR research. Overall, we strive to minimize confounding factors that could hinder our ability to address our research questions effectively.

Concerning *internal validity*, there is a risk that other factors influence our evaluation, potentially leading to incorrect interpretations or conclusions. We adhere to established practices for risk mitigation, utilizing datasets and projects from state-of-the-art approaches. We also clearly document the origins of these projects and the associated ground truths. Project quality and consistency variations may still impact the performance of our approaches. For the combined pipeline (TransArC + ExArch + ArTEMiS), errors may propagate and compound across stages: imprecise component names produced by ExArch may cause ArTEMiS to miss or misattribute entity matches, which in turn affects the final SAD-to-Code trace links. We mitigate this by evaluating each component individually as well as in combination, allowing us to isolate the contribution of each stage.

There are also potential threats concerning *external validity*. First, our evaluation is limited to a small number of projects and TLR tasks, which may affect the generalizability of the results to other projects and tasks. We utilize an established dataset covering various domains and project sizes to reduce this threat. However, these datasets mainly consist of open-source projects, which may differ from closed-source projects in key characteristics. Second, we evaluate the approach using only a limited set of LLMs, which may limit the robustness of our conclusions. We further evaluated the approach using only the configuration and the most promising models. Nonetheless, we evaluated our approaches using various LLMs to show the performance of fixed prompts using different LLMs. A third concern is the use of closed-source models in evaluation. Since the training data of these models is unknown, we cannot entirely rule out data leakage. To ensure transparency, we provide our code, prompts, and LLM responses in our replication packages [20, 21]. Furthermore, we also use locally deployed models for ExArch. Finally, the non-determinism of LLMs poses another threat. To mitigate this, we set the LLMs' temperature to zero and used the same fixed seed for each run. By doing so, we make sure that future research can replicate our results. Setting the temperature to zero is only possible for non-GPT-5 models, as GPT-5 does not allow setting the temperature. Moreover, for ArTEMiS, the longer prompts and the structured JSON formatting step can cause non-deterministic behavior with effects on the results even at temperature zero. Therefore, we run each ArTEMiS configuration multiple times and report average results with standard deviation to account for this variability.

By conducting our evaluation on established datasets and using widely accepted metrics, we address the *reliability* of our research. Therefore, we do not introduce a threat regarding the manual creation of the gold standards or datasets. However, the existing gold standards might be imperfect.

6 Conclusion and Future Work

We presented two complementary LLM-based approaches, ExArch and ArTEMiS, to bridge the semantic gap between SAD and code. ExArch generates simple SAMs (component names) from SAD and/or code to enable TransArC without manually created SAMs. ArTEMiS recovers trace links between SAD and SAM as a modern alternative to heuristics. Together with TransArC, these approaches enable SAD-to-code-TLR without manual SAMs and improve average performance over state-of-the-art baselines.

First, on generating component names, we studied three modes (from SAD, from code, and both). Consistent with Section 5.3, generating names from SAD performs best for the investigated TLR task. On the original benchmark, ExArch with GPT-4o combined with TransArC achieves a weighted average F_1 -score of 0.86, comparable to TransArC with manually created SAMs (0.87), and significantly outperforms baselines that rely only on SAD and code. On the revised benchmark, ExArch alone (SAD-only) achieves an average F_1 -score of 0.76 with GPT-4o. Extraction from code alone underperforms on average; combining SAD and code via aggregation helps but does not surpass SAD-only in macro average. Across models we evaluated, closed-source OpenAI models on average outperform the open-source Llama-based models.

Second, on linking SAD to SAM, ArTEMiS provides a modern LLM-based alternative to heuristics. It improves over SWATTR in average recall and is slightly better in average F_1 -score (with GPT-5, recall 0.85 vs. 0.77; F_1 -score 0.81 vs. 0.80). When integrated into TransArC, ArTEMiS increases average recall for SAD to code TLR from 0.77 to 0.87 and raises average F_1 -score (0.80 to 0.85). In contrast, the weighted recall decreases from 0.93 to 0.83 (driven by projects with many links). The combination TransArC + ExArch + ArTEMiS attains the best weighted average F_1 -score (0.87) among approaches that do not require manually created SAMs, outperforming the strongest baseline without SAMs (ArDoCode, 0.62).

Taken together, these results show that LLMs are a viable means to extract the architectural elements needed for TLR between SAD and code, enabling performance similar to approaches that rely on manually created SAMs.

In the future, we want to advance code-based SAM extraction by leveraging architecture recovery techniques; information-fusion methods [69] appear particularly promising. We also want to enrich prompts with additional project context (e.g., ADRs, architecture diagrams, and other SAD sections). Furthermore, it would be interesting to determine the capabilities of LLMs regarding trace link identification. To reduce noise, we aim to separate production and test code when deriving names and refine pre-processing (e.g., package selection and filtering). We want to explore code summarization to provide more context to LLMs and study models with integrated Chain-of-thought prompting (e.g., *OpenAI o1*). Finally, we intend to assess per-project calibration (such as similarity thresholds) and devise strategies to improve weighted recall for projects with many links.

Data Availability Statement

We provide our code, baselines, evaluation data, and results in our replication packages [20, 21].

Acknowledgments

This work was funded by Core Informatics at KIT (KiKIT) of the Helmholtz Assoc. (HGF), the German Research Foundation (DFG) - SFB 1608 - 501798263, the Topic Engineering Secure Systems of the Helmholtz Association (HGF), and supported by KASTEL Security Research Labs, Karlsruhe.

References

- [1] Aakash Ahmad, Muhammad Waseem, Peng Liang, Mahdi Fahmideh, Mst Shamima Aktar, and Tommi Mikkonen. 2023. Towards Human-Bot Collaborative Software Architecting with ChatGPT. In *Proceedings of the 27th International Conference on Evaluation and Assessment in Software Engineering (Oulu, Finland) (EASE '23)*. Association for Computing Machinery, New York, NY, USA, 279–285. doi:10.1145/3593434.3593468
- [2] Amarjeet and Jitender Kumar Chhabra. 2017. Improving modular structure of software system using structural and lexical dependency. *Information and Software Technology* 82 (2017), 96–120. doi:10.1016/j.infsof.2016.09.011
- [3] G. Antoniol, G. Canfora, G. Casazza, A. De Lucia, and E. Merlo. 2002. Recovering Traceability Links between Code and Documentation. *IEEE Transactions on Software Engineering* 28, 10 (Oct. 2002), 970–983. doi:10.1109/TSE.2002.1041053

- [4] Thazin Win Win Aung, Huan Huo, and Yulei Sui. 2020. A Literature Review of Automatic Traceability Links Recovery for Software Change Impact Analysis. In *Proceedings of the 28th International Conference on Program Comprehension* (Seoul, Republic of Korea) (*ICPC '20*). Association for Computing Machinery, New York, NY, USA, 14–24. doi:10.1145/3387904.3389251
- [5] YuXuan Chen, Jianwei Ding, Dashuang Li, and Zhouguo Chen. 2021. Joint BERT Model based Cybersecurity Named Entity Recognition. In *Proceedings of the 2021 4th International Conference on Software Engineering and Information Management* (Yokohama, Japan) (*ICSIM '21*). ACM, New York, NY, USA, 236–242. doi:10.1145/3451471.3451508
- [6] Choongki Cho, Ki-Seong Lee, Minsoo Lee, and Chan-Gun Lee. 2019. Software Architecture Module-View Recovery Using Cluster Ensembles. *IEEE Access* 7 (2019), 72872–72884. doi:10.1109/ACCESS.2019.2920427
- [7] Jane Cleland-Huang, Orlena Gotel, Andrea Zisman, et al. 2012. *Software and systems traceability*. Vol. 2. Springer. doi:10.1007/978-1-4471-2239-5
- [8] Anna Corazza, Sergio Di Martino, Valerio Maggio, and Giuseppe Scanniello. 2011. Investigating the use of lexical information for software system clustering. In *2011 15th European Conference on Software Maintenance and Reengineering*. 35–44. doi:10.1109/CSMR.2011.8
- [9] Javier Cámara, Lola Burgueño, and Javier Troya. 2024. Towards standardized benchmarks of LLMs in software modeling tasks: a conceptual framework. *Software and Systems Modeling* (Sept. 2024). doi:10.1007/s10270-024-01206-9
- [10] Javier Cámara, Javier Troya, Lola Burgueño, and Antonio Vallecillo. 2023. On the assessment of generative AI in modeling tasks: an experience report with ChatGPT and UML. *Software and Systems Modeling* 22, 3 (May 2023), 781–793. doi:10.1007/s10270-023-01105-5
- [11] Javier Cámara, Javier Troya, Julio Montes-Torres, and Francisco J. Jaime. 2024. Generative AI in the Software Modeling Classroom: An Experience Report With ChatGPT and Unified Modeling Language. *IEEE Software* 41, 6 (2024), 73–81. doi:10.1109/MS.2024.3385309
- [12] Souvick Das, Novarun Deb, Agostino Cortesi, and Nabendu Chaki. 2023. Zero-shot Learning for Named Entity Recognition in Software Specification Documents. In *2023 IEEE 31st International Requirements Engineering Conference (RE)*. 100–110. doi:10.1109/RE57278.2023.00019
- [13] R. Dhar, K. Vaidhyathan, and V. Varma. 2024. Can LLMs Generate Architectural Design Decisions? - An Exploratory Empirical Study. In *2024 IEEE 21st International Conference on Software Architecture (ICSA)*. IEEE Computer Society, Los Alamitos, CA, USA, 79–89. doi:10.1109/ICSA59870.2024.00016
- [14] J. Andrés Díaz-Pace, Antonela Tommasel, and Rafael Capilla. 2024. Helping Novice Architects to Make Quality Design Decisions Using an LLM-Based Assistant. In *Software Architecture*. Springer Nature Switzerland, Cham, 324–332.
- [15] Tobias Eisenreich, Sandro Speth, and Stefan Wagner. 2024. From Requirements to Architecture: An AI-Based Journey to Semi-Automatically Generate Software Architectures. In *Proceedings of the 1st International Workshop on Designing Software* (Lisbon, Portugal) (*Designing '24*). ACM, New York, NY, USA, 52–55. doi:10.1145/3643660.3643942
- [16] Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, and Ming Zhou. 2020. CodeBERT: A Pre-Trained Model for Programming and Natural Languages. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 1536–1547. doi:10.18653/v1/2020.findings-emnlp.139
- [17] Dominik Fuchß, Sophie Corallo, Jan Keim, Janek Speit, and Anne Kozirolek. 2023. Establishing a Benchmark Dataset for Traceability Link Recovery Between Software Architecture Documentation and Models. In *Software Architecture. ECSA 2022 Tracks and Workshops*, Thais Batista, Tomáš Bureš, Claudia Raibulet, and Henry Muccini (Eds.). Springer International Publishing, Cham, 455–464. doi:10.1007/978-3-031-36889-9_30
- [18] Dominik Fuchß, Tobias Hey, Jan Keim, Haoyu Liu, Niklas Ewald, Tobias Thirolf, and Anne Kozirolek. 2025. LiSSA: Toward Generic Traceability Link Recovery through Retrieval-Augmented Generation. In *Proceedings of the IEEE/ACM 47th International Conference on Software Engineering* (Ottawa, Canada) (*ICSE '25*). Institute of Electrical and Electronics Engineers (IEEE). doi:10.1109/ICSE55347.2025.00186
- [19] Dominik Fuchß, Haoyu Liu, Tobias Hey, Jan Keim, and Anne Kozirolek. 2025. Enabling Architecture Traceability by LLM-based Architecture Component Name Extraction. In *2025 IEEE 22nd International Conference on Software Architecture (ICSA)*. Institute of Electrical and Electronics Engineers (IEEE). doi:10.1109/ICSA65012.2025.00011
- [20] Dominik Fuchß, Haoyu Liu, Sophie Corallo, Tobias Hey, Jan Keim, Johannes von Geisau, and Anne Kozirolek. 2025. Replication Package: Who's Who? LLM-assisted Software Traceability with Architecture Entity Recognition. doi:10.5281/zenodo.18957376
- [21] Dominik Fuchß, Haoyu Liu, Tobias Hey, Jan Keim, and Anne Kozirolek. 2024. Replication Package: Enabling Architecture Traceability by LLM-based Architecture Component Name Extraction. doi:10.5281/ZENODO.14506935
- [22] Hui Gao, Hongyu Kuang, Kexin Sun, Xiaoxing Ma, Alexander Egyed, Patrick Mäder, Guoping Rong, Dong Shao, and He Zhang. 2023. Using Consensual Biterms from Text Structures of Requirements and Code to Improve IR-Based Traceability Recovery. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering (ASE '22)*. Association for Computing Machinery, New York, NY, USA, 1. doi:10.1145/3551349.3556948

- [23] Joshua Garcia, Daniel Popescu, Chris Mattmann, Nenad Medvidovic, and Yuanfang Cai. 2011. Enhancing architectural recovery using concerns. In *2011 26th IEEE/ACM International Conference on Automated Software Engineering (ASE 2011)*, 552–555. doi:10.1109/ASE.2011.6100123
- [24] Jameleddine Hassine. 2024. An LLM-based Approach to Recover Traceability Links between Security Requirements and Goal Models. In *Proceedings of the 28th International Conference on Evaluation and Assessment in Software Engineering (Salerno, Italy) (EASE '24)*. ACM, New York, NY, USA, 643–651. doi:10.1145/3661167.3661261
- [25] Jane Huffman Hayes, Alex Dekhtyar, and Senthil Karthikeyan Sundaram. 2006. Advancing Candidate Link Generation for Requirements Tracing: The Study of Methods. *IEEE TSE* 32, 1 (Jan. 2006), 4–19. doi:10.1109/TSE.2006.3
- [26] Jane Huffman Hayes, Alex Dekhtyar, Senthil Karthikeyan Sundaram, E Ashlee Holbrook, Sravanthi Vadlamudi, and Alain April. 2007. REquirements TRacing On target (RETRO): improving software maintenance through traceability recovery. *Innovations in Systems and Software Engineering* 3 (2007), 193–202.
- [27] Guntur Budi Herwanto, Gerald Quirchmayr, and A. Min Tjoa. 2024. Leveraging NLP Techniques for Privacy Requirements Engineering in User Stories. *IEEE Access* 12 (2024), 22167–22189. doi:10.1109/ACCESS.2024.3364533
- [28] Tobias Hey, Fei Chen, Sebastian Weigelt, and Walter F. Tichy. 2021. Improving Traceability Link Recovery Using Fine-grained Requirements-to-Code Relations. In *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME) (2021-09)*, 12–22. doi:10.1109/ICSME52107.2021.00008
- [29] Tobias Hey, Dominik Fuchß, Jan Keim, and Anne Kozirolek. 2025. Requirements Traceability Link Recovery via Retrieval-Augmented Generation. In *Requirements Engineering: Foundation for Software Quality*. Springer, Cham. doi:10.1007/978-3-031-88531-0_27
- [30] Tobias Hey, Jan Keim, and Sophie Corallo. 2024. Requirements Classification for Traceability Link Recovery. In *2024 IEEE 32nd International Requirements Engineering Conference (RE)*, 155–167. doi:10.1109/RE59067.2024.00024
- [31] Xinyi Hou, Yanjie Zhao, Yue Liu, Zhou Yang, Kailong Wang, Li Li, Xiapu Luo, David Lo, John Grundy, and Haoyu Wang. 2024. Large Language Models for Software Engineering: A Systematic Literature Review. *ACM Trans. Softw. Eng. Methodol.* (Sept. 2024). doi:10.1145/3695988 Just Accepted.
- [32] Jan Keim, Sophie Corallo, Dominik Fuchß, Tobias Hey, Tobias Telge, and Anne Kozirolek. 2024. Recovering Trace Links Between Software Documentation And Code. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 215, 13 pages. doi:10.1145/3597503.3639130
- [33] Jan Keim, Sophie Corallo, Dominik Fuchß, and Anne Kozirolek. 2023. Detecting Inconsistencies in Software Architecture Documentation Using Traceability Link Recovery. In *2023 IEEE 20th International Conference on Software Architecture (ICSA)*, 141–152. doi:10.1109/ICSA56044.2023.00021
- [34] Jan Keim, Sophie Schulz, Dominik Fuchß, Claudius Kocher, Janek Speit, and Anne Kozirolek. 2021. Trace Link Recovery for Software Architecture Documentation. In *Software Architecture*, Stefan Biffl, Elena Navarro, Welf Löwe, Marjan Sirjani, Raffaella Mirandola, and Danny Weyns (Eds.). Springer International Publishing, Cham, 101–116.
- [35] Jan Keim, Sophie Schulz, Dominik Fuchß, Claudius Kocher, Janek Speit, and Anne Kozirolek. 2021. Trace Link Recovery for Software Architecture Documentation. In *Software Architecture*, Stefan Biffl, Elena Navarro, Welf Löwe, Marjan Sirjani, Raffaella Mirandola, and Danny Weyns (Eds.). Springer International Publishing, Cham, 101–116. doi:10.1007/978-3-030-86044-8_7
- [36] Hongyu Kuang, Patrick Mäder, Hao Hu, Achraf Ghabi, LiGuo Huang, Jian Lü, and Alexander Egyed. 2015. Can method data dependencies support the assessment of traceability between requirements and source code? *Journal of Software: Evolution and Process* 27, 11 (2015), 838–866. doi:10.1002/smr.1736
- [37] Vladimir I Levenshtein et al. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, Vol. 10. Soviet Union, 707–710.
- [38] Vladimir I Levenshtein et al. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, Vol. 10. Soviet Union, 707–710.
- [39] Jinfeng Lin, Yalin Liu, Qingkai Zeng, Meng Jiang, and Jane Cleland-Huang. 2021. Traceability Transformed: Generating more Accurate Links with Pre-Trained BERT Models. In *Proceedings of the 43rd International Conference on Software Engineering (ICSE '21)*. IEEE Press, Madrid, Spain, 324–335. doi:10.1109/ICSE43902.2021.00040
- [40] Thibaud Lutellier, Devin Chollak, Joshua Garcia, Lin Tan, Derek Rayside, Nenad Medvidović, and Robert Kroeger. 2018. Measuring the Impact of Code Dependencies on Software Architecture Recovery Techniques. *IEEE Transactions on Software Engineering* 44, 2 (2018), 159–181. doi:10.1109/TSE.2017.2671865
- [41] Garima Malik, Mucahit Cevik, Swayami Bera, Savas Yildirim, Devang Parikh, and Ayse Basar. [n. d.]. Software requirement specific entity extraction using transformer models.
- [42] Niklas Meissner, Sandro Speth, and Steffen Becker. 2024. Automated Programming Exercise Generation in the Era of Large Language Models. In *2024 36th International Conference on Software Engineering Education and Training*, 1–5. doi:10.1109/CSEET62301.2024.10662984

- [43] Kevin Moran, David N. Palacio, Carlos Bernal-Cárdenas, Daniel McCrystal, Denys Poshyvanyk, Chris Shenefiel, and Jeff Johnson. 2020. Improving the Effectiveness of Traceability Link Recovery Using Hierarchical Bayesian Networks. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (ICSE '20)*. Association for Computing Machinery, New York, NY, USA, 873–885. doi:10.1145/3377811.3380418
- [44] Kevin Moran, David N. Palacio, Carlos Bernal-Cárdenas, Daniel McCrystal, Denys Poshyvanyk, Chris Shenefiel, and Jeff Johnson. 2020. Improving the effectiveness of traceability link recovery using hierarchical bayesian networks. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (ICSE '20)*. Association for Computing Machinery, New York, NY, USA, 873–885. doi:10.1145/3377811.3380418
- [45] Kazuki Nishikawa, Hironori Washizaki, Yoshiaki Fukazawa, Keishi Oshima, and Ryota Mibe. 2015. Recovering transitive traceability links among software artifacts. In *2015 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 576–580. doi:10.1109/ICSM.2015.7332517
- [46] Marc North, Amir Atapour-Abarghouei, and Nelly Bencomo. 2024. Code Gradients: Towards Automated Traceability of LLM-Generated Code. In *2024 IEEE 32nd International Requirements Engineering Conference (RE)*. 321–329. doi:10.1109/RE59067.2024.00038
- [47] A. Panichella, C. McMillan, E. Moritz, D. Palmieri, R. Oliveto, D. Poshyvanyk, and A. De Lucia. 2013. When and How Using Structural Information to Improve IR-Based Traceability Recovery. In *2013 17th European Conference on Software Maintenance and Reengineering*. 199–208. doi:10.1109/CSMR.2013.29
- [48] Patrick Rempel and Parick Mäder. 2017. Preventing Defects: The Impact of Requirements Traceability Completeness on Software Quality. *IEEE Transactions on Software Engineering* 43, 8 (2017). doi:10.1109/TSE.2016.2622264
- [49] Alberto D. Rodriguez, Jane Cleland-Huang, and Davide Falessi. 2021. Leveraging Intermediate Artifacts to Improve Automated Trace Link Retrieval. In *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. 81–92. doi:10.1109/ICSM.2021.00014
- [50] Alberto D. Rodriguez, Katherine R. Dearstyne, and Jane Cleland-Huang. 2023. Prompts Matter: Insights and Strategies for Prompt Engineering in Automated Software Traceability. In *2023 IEEE 31st International Requirements Engineering Conference Workshops (REW)*. 455–464. doi:10.1109/REW57809.2023.00087
- [51] Satrio Adi Rukmono, Lina Ochoa, and Michel Chaudron. 2024. Deductive Software Architecture Recovery via Chain-of-thought Prompting. In *Proceedings of the 2024 ACM/IEEE 44th International Conference on Software Engineering: New Ideas and Emerging Results (Lisbon, Portugal) (ICSE-NIER)*. Association for Computing Machinery, New York, NY, USA, 92–96. doi:10.1145/3639476.3639776
- [52] Per Runeson and Martin Höst. 2008. Guidelines for conducting and reporting case study research in software engineering. *Empirical Software Engineering* 14, 2 (2008), 131. doi:10.1007/s10664-008-9102-8
- [53] Daniel Russo, Sebastian Baltes, Niels van Berkel, Paris Avgeriou, Fabio Calefato, Beatriz Cabrero-Daniel, Gemma Catolino, Jürgen Cito, Neil Ernst, Thomas Fritz, Hideaki Hata, Reid Holmes, Maliheh Izadi, Foutse Khomh, Mikkel Baun Kjærgaard, Grischa Liebel, Alberto Lluch Lafuente, Stefano Lambiase, Walid Maalej, Gail Murphy, Nils Brede Moe, Gabrielle O'Brien, Elda Paja, Mauro Pezzè, John Stouby Persson, Rafael Prikladnicki, Paul Ralph, Martin Robillard, Thiago Rocha Silva, Klaas-Jan Stol, Margaret-Anne Storey, Viktoria Stray, Paolo Tell, Christoph Treude, and Bogdan Vasilescu. 2024. Generative AI in Software Engineering Must Be Human-Centered: The Copenhagen Manifesto. *Journal of Systems and Software* 216 (2024), 112115. doi:10.1016/j.jss.2024.112115
- [54] June Sallou, Thomas Durieux, and Annibale Panichella. 2024. Breaking the Silence: the Threats of Using LLMs in Software Engineering. In *Proceedings of the 2024 ACM/IEEE 44th International Conference on Software Engineering: New Ideas and Emerging Results (Lisbon, Portugal) (ICSE-NIER'24)*. Association for Computing Machinery, New York, NY, USA, 102–106. doi:10.1145/3639476.3639764
- [55] Zipani Tom Sinkala and Sebastian Herold. 2022. Hierarchical Code-to-Architecture Mapping. In *Software Architecture*, Patrizia Scandurra, Matthias Galster, Raffaella Mirandola, and Danny Weyns (Eds.). Springer International Publishing, Cham, 86–104.
- [56] Skander Soltani and Elias Limouni. 2025. LLM Based Data Annotation and Augmentation for NER and Relationship Extraction Models Enhancement. In *Artificial Intelligence for Global Security*, Dominique Verdejo and Eunika Mercier-Laurent (Eds.). Springer Nature Switzerland, Cham, 153–160.
- [57] Herbert Stachowiak. 1973. *Allgemeine Modelltheorie*. Springer Verlag, Wien.
- [58] Chao Sun, Mingjing Tang, Li Liang, and Wei Zou. 2020. Software Entity Recognition Method Based on BERT Embedding. In *Machine Learning for Cyber Security*, Xiaofeng Chen, Hongyang Yan, Qiben Yan, and Xiangliang Zhang (Eds.). Springer International Publishing, Cham, 33–47.
- [59] Jeniya Tabassum, Mounica Maddela, Wei Xu, and Alan Ritter. 2020. Code and Named Entity Recognition in Stack-Overflow. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 4913–4926. doi:10.18653/v1/2020.acl-main.443

- [60] Mingjing Tang, Tong Li, Wei Gao, and Yu Xia. 2022. AttenSy-SNER: software knowledge entity extraction with syntactic features and semantic augmentation information. *Complex & Intelligent Systems* 9, 1 (June 2022), 25–39. doi:10.1007/s40747-022-00742-5
- [61] Fangchao Tian, Tianlu Wang, Peng Liang, Chong Wang, Arif Ali Khan, and Muhammad Ali Babar. 2021. The impact of traceability on software maintenance and evolution: A mapping study. *Journal of Software: Evolution and Process* 33, 10 (2021), e2374. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/smr.2374 doi:10.1002/smr.2374
- [62] MM Tikhomirov, NV Loukachevitch, and BV Dobrov. 2020. Recognizing named entities in specific domain. *Lobachevskii Journal of Mathematics* 41, 8 (2020), 1591–1602.
- [63] Vassilios Tzerpos and Richard C Holt. 2000. Accd: an algorithm for comprehension-driven clustering. In *Proceedings Seventh Working Conference on Reverse Engineering*. IEEE, 258–267.
- [64] M. Veera Prathap Reddy, P. V. R. D. Prasad, Manjunath Chikkamath, and Sarathchandra Mandadi. 2019. NERSE: Named Entity Recognition in Software Engineering as a Service. In *Service Research and Innovation*, Ho-Pun Lam and Sajib Mistry (Eds.). Springer International Publishing, Cham, 65–80.
- [65] William E Winkler. 1990. String comparator metrics and enhanced decision rules in the Fellegi-Sunter model of record linkage. (1990).
- [66] Han Wu, Xiaoyong Li, and Yali Gao. 2020. An Effective Approach of Named Entity Recognition for Cyber Threat Intelligence. In *2020 IEEE 4th Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)*, Vol. 1. 1370–1374. doi:10.1109/ITNEC48623.2020.9085102
- [67] C. Xiao and V. Tzerpos. 2005. Software clustering based on dynamic dependencies. In *Ninth European Conference on Software Maintenance and Reengineering*. 124–133. doi:10.1109/CSMR.2005.49
- [68] Deheng Ye, Zhenchang Xing, Chee Yong Foo, Zi Qun Ang, Jing Li, and Nachiket Kapre. 2016. Software-Specific Named Entity Recognition in Software Engineering Social Content. In *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, Vol. 1. 90–101. doi:10.1109/SANER.2016.10
- [69] Yiran Zhang, Zhengzi Xu, Chengwei Liu, Hongxu Chen, Jianwen Sun, Dong Qiu, and Yang Liu. 2023. Software Architecture Recovery with Information Fusion. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (San Francisco, CA, USA) (ESEC/FSE 2023)*. Association for Computing Machinery, New York, NY, USA, 1535–1547. doi:10.1145/3611643.3616285
- [70] Zejun Zhang, Zhenchang Xing, Xiaoxue Ren, Qinghua Lu, and Xiwei Xu. 2024. Refactoring to Pythonic Idioms: A Hybrid Knowledge-Driven Approach Leveraging Large Language Models. *Proc. ACM Softw. Eng.* 1, FSE, Article 50 (July 2024), 22 pages. doi:10.1145/3643776
- [71] Cheng Zhou, Bin Li, and Xiaobing Sun. 2020. Improving software bug-specific named entity recognition with deep neural network. *Journal of Systems and Software* 165 (2020), 110572. doi:10.1016/j.jss.2020.110572

A ArTEMiS Prompt (Task)

```

1 Identify all architecturally relevant software components that are explicitly named in the
  following text.
2 For each identified component, provide:
3 - The primary name (as it appears in the text)
4 - All alternative names or abbreviations used for the same component in the text (case-
  insensitive)
5 - All full lines where the component is mentioned (directly or via clear context)
6
7 Rules for identifying components:
8 1. Only include explicit modular software components with distinct technical
  responsibilities. These may include:
9   - services (e.g., UserService)
10  - APIs (e.g., PaymentAPI)
11  - adapters, handlers, managers, routers, engines
12  - infrastructure components (e.g., Media Server, Presentation Conversion Pipeline)
13  - client-side or server-side subsystems (e.g., electron client, backend server)
14 2. Exclude domain-level entities, even if capitalized - such as business data objects, file
  types, or general functionalities - unless used as part of a named technical unit.
15 Do not include non-technical concepts even if mentioned with verbs like "convert", "
  generate", or "store" - these are often subject-side actions unless framed as
  components.
16
17 Examples of domain terms (do not include):
18 - image - "Each item includes an image."
19 - recommendation - "Recommendations are generated..."
20 - file - "Uploads include a JSON file."
21 - session - "Each session is stored separately."
22 - presentation - "Uploaded presentations go through conversion..."

```

```

23
24   Include only when wrapped in named software components that perform active, modular
25   responsibilities (if explicitly named and described).
26 3. DO include technical subsystems described with proper software roles, and clearly scoped
27   :
28   - (Web) server - if described as a component implementing client-server communication or
29   event dispatching
30   - (Web) client - if described as rendering or subscribing to events/data
31   - Media Server / MS - as a media streaming component implementing SFU/MCU
32 4. Do not include:
33   - Package, class, or namespace names (e.g., common.util, x.y.z)
34   - Interfaces (unless directly implemented and deployed)
35   - General use of technologies or third-party tools like "React" or "Spring" unless
36   internally wrapped as system components
37 5. A component is valid if:
38   a) Its name includes a functional suffix or architecture-relevant term (Service, Client,
39   Engine, Manager, Adapter, Server, Router, Converter, etc.)
40   OR
41   b) The text clearly describes it as implementing a technical function within the system
42   (e.g., routing requests, synchronizing state, managing media streams)
43 6. Reverse pronoun references are allowed only when strongly tied to a previously named
44   component across adjacent lines.
45 Do not infer vague or implied components through generic phrases like:
46   - it handles the process
47   - this system
48   - the module
49 7. Do not create implied components from action nouns (e.g., "conversion", "delivery", "
50   recommendation") unless these are mentioned as named, distinct architectural elements.
51 8. If an external technology (e.g., MongoDB, Redis, etc.) is used in a custom component (e.
52   g., our RedisPublisher, or MongoSyncService), include that named component - not the
53   technology itself.
54
55 Return the results in a clearly structured, unambiguous plain-text format that enables
56 straightforward conversion to JSON (e.g., using key-value sections per component).

```

B ArTEMiS Prompt (Formatting)

```

1 Given the last answer (see below), for each component, return a JSON object containing:
2 - "name": the primary name of the component (use the most descriptive name).
3 - "type": "COMPONENT"
4 - "alternativeNames": a list of alternative or ambiguous names, if applicable.
5 - "occurrences": a list of lines where the component appears or is referenced.
6
7 Output should be a JSON array (and nothing else!), like:
8 [
9   {
10    "name": "...",
11    "type": "COMPONENT",
12    "alternativeNames": [...],
13    "occurrences": [...]
14  },
15  ...
16 ]
17
18 Example:
19 [
20   {
21     "name": "AuthenticationService",
22     "type": "COMPONENT",
23     "alternativeNames": ["service"],
24     "occurrences": ["The AuthenticationService handles login requests.", "It forwards
25     valid credentials to the UserDatabase.", "The service logs each attempt."]
26   },
27   {
28     "name": "UserDatabase",
29     "type": "COMPONENT",
30     "alternativeNames": ["DB"],
31     "occurrences": ["It forwards valid credentials to the UserDatabase.", "The DB then
32     validates the credentials."]
33   }
34 ]

```