

Linking Cryptoasset Attribution Tags to Knowledge Graph Entities: An LLM-based Approach

Régnier Avice¹[0000-0002-7320-7868], Bernhard Haslhofer²[0000-0002-0415-4491],
Zhidong Li¹[0000-0002-0784-157X], and Jianlong Zhou¹[0000-0001-6034-644X]

¹ University of Technology Sydney, Australia

regnier.avice@student.uts.edu.au, {zhidong.li, jianlong.zhou}@uts.edu.au

² Complexity Science Hub, Vienna, Austria

haslhofer@csh.ac.at

Abstract. Attribution tags form the foundation of modern cryptoasset forensics. However, inconsistent or incorrect tags can mislead investigations and even result in false accusations. To address this issue, we propose a novel computational method based on Large Language Models (LLMs) to link attribution tags with well-defined knowledge graph concepts. We implemented this method in an end-to-end pipeline and conducted experiments showing that our approach outperforms baseline methods by up to 37.4% in F1-score across three publicly available attribution tag datasets. By integrating concept filtering and blocking procedures, we generate candidate sets containing five knowledge graph entities, achieving a recall of 93% without the need for labeled data. Additionally, we demonstrate that local LLM models can achieve F1-scores of 90%, comparable to remote models which achieve 94%. We also analyze the cost-performance trade-offs of various LLMs and prompt templates, showing that selecting the most cost-effective configuration can reduce costs by 90%, with only a 1% decrease in performance. Our method not only enhances attribution tag quality but also serves as a blueprint for fostering more reliable forensic evidence.

Keywords: Forensics · LLM · Cryptoassets

1 Introduction

Attribution tags, which link pseudo-anonymous cryptoasset addresses to identifying information about real-world entities and services (e.g., cryptoasset exchanges), form the foundation of modern cryptoasset forensics [26]. Over the past decade, a multibillion-dollar industry has emerged, providing blockchain tracing tools for law enforcement investigations. These tools allow users to “follow the money”, ultimately leading to the identification and potential conviction of perpetrators. As cryptoassets have gained increasing relevance across various crime sectors, these tools are now used in investigations related to ransomware [19, 30], sextortion [31], and malware [13]. Attribution tags have also

been employed to train machine learning models for automatically categorizing service providers, such as exchanges [13, 15, 25, 44], miners [15, 25, 44], and ICO wallets [25, 44]. Moreover, these models are used to classify addresses or transactions as illicit [1, 4, 18, 22, 25, 42, 44]. Consequently, the accuracy of attribution tags becomes a critical success factor in all these application domains.

Imprecise or incorrect tags can mislead investigations, result in inaccurate model predictions, or, in the worst case, lead to false accusations. As a result, the concept of *attribution tag quality* is gaining increased attention, especially as the scientific validity of crypto-tracing techniques is being more frequently questioned and scrutinized [3]. As with any digital forensic investigation that forms the basis for legal decisions, the collected evidence must be reliable [10], since poor data quality can lead to incorrect conclusions. Ensuring high data quality also guarantees that the methodologies used in digital forensics can be consistently tested and verified.

Since attribution tags originate from and are often shared among multiple stakeholders, data consistency becomes a critical aspect of data quality. A key challenge, for instance, is the inconsistent referencing of real-world entities across different parties [10, 13]. For example, one party might refer to a specific exchange as *btc-e*, while another uses *btc-e.com*. While a human can easily recognize that both tags refer to the same entity, different tracing tools may interpret them as referring to two distinct entities. To harmonize their representation, reuse, and interpretation in forensic investigations or machine learning tasks, a suitable data format and a computational approach are needed to automatically eliminate potential data inconsistencies.

Given the semantic nature of this practical problem, knowledge graphs offer an effective solution. They represent entities as identifiable semantic concepts rather than plain strings and are widely used in search engines [14] and large knowledge bases [39]. The process of linking data to these well-defined entities is known as record linkage [9] or entity linking [36]. However, existing approaches often rely on heuristics or heavily labeled, domain-specific training datasets, which are not readily available in the context of cryptoasset attribution tags.

Therefore, in this paper, we present a novel computational approach based on Large Language Models (LLMs) that enables the linking of attribution tag datasets to well-defined knowledge graph concepts without requiring domain-specific fine-tuning. We implemented this approach by integrating several cutting-edge techniques — such as filtering, blocking, and LLMs — into an end-to-end pipeline. Our experiments, conducted on three datasets, demonstrate that:

1. Our end-to-end entity linking approach outperforms baseline methods across all three datasets, achieving up to a 37.4% improvement in F1-scores.
2. BM25₃ blocking and related-concept filtering reduce the number of candidates per tag to 5, with a recall of 93%, without requiring labeled data.
3. GPT-4o-based candidate selection achieves a 94% F1-score, while Mistral 7B-Instruct, which can run locally, achieves a 90% F1-score.
4. Using the most cost-effective prompt template reduces costs by 90%, with only a 1% drop in performance.

Our method not only enhances the quality of attribution tags but also seeks to inspire broader efforts toward improving data quality, ensuring accurate and reliable evidence in forensic investigations. To ensure reproducibility, we have made our code and datasets publicly available in the following GitHub repository: <https://github.com/ravice234/cryptoasset-attribution-tag-linker>.

2 Background

2.1 Attribution Tags

Attribution tags link pseudo-anonymous blockchain objects, such as addresses or transactions, to real-world actors or events. They provide additional context, such as the name of a service controlling an address (e.g., `btc-e`), some form of categorization (e.g., exchange), and any other information that might be useful in forensic investigations. Figure 1 illustrates an example in which two distinct attribution tags originating from different sources reference the same cryptoasset address `0x123` and describe the same real-world actor, a well-known cryptoasset exchange. One can observe that, despite describing the same entity, the attribution tag data records the name (`btc-e` vs. `btc-e.com`) and categorize the exchange differently (`Exchange` vs. `Service`).




Fig. 1. Attribution Tag Example. Two attribution tags referencing the same cryptoasset address `0x123` owned by the real-world entity *BTC-e*.

Attribution tag data quality issues can occur at various levels [16]: technical heterogeneities, such as different data formats, can impede uniform processing; syntactic heterogeneities, like the use of different encoding schemes, can hinder uniform interpretation; and semantic heterogeneities, such as the use of different names to denote the same real-world concept (synonyms, homonyms, hypernyms, etc.), can lead to inconsistent interpretations. In this paper, we primarily focus on resolving semantic interoperability issues, assuming that technical and syntactic issues can be addressed using common data preprocessing procedures.

Introducing knowledge graphs to data management environments has become a common strategy to deal with semantic interoperability issues. A knowledge graph defines nodes representing real-world entities of interest and semantic relationships between these entities [17]. An example is eBay’s product knowledge

graph, which allows them to identify if two sellers sell the same products, or if the products are related otherwise.

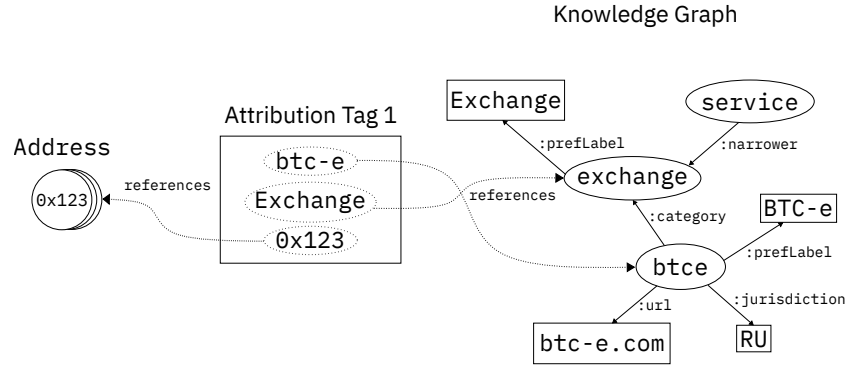


Fig. 2. Linking an Attribution Tag to the Knowledge Graph. Attribution tag instances are linked to concepts defined in the knowledge graph.

The process of linking text mentions to entities in a knowledge graph is called *entity linking* [36]. This process can be further broken down into: 1) generating a subset of entities (*candidates*) that are most likely to match, and 2) selecting which, if any, candidate matches to the text mention. Candidate set generation can be achieved by excluding implausible (*filtering*) and clustering similar (*blocking*) entities [29]. The selection of a matching candidate is typically performed using a decision function learned through machine learning. In this paper, the term *entity linking* is used to describe the process of linking tags to a knowledge graph; while our solution also encompasses techniques that stem from the closely related problem of linking database records (record linkage, entity matching), for the sake of simplicity, we will consistently refer to it as entity linking.

2.2 Related Work

Entity linking approaches where tuples are linked to knowledge base entities have utilized look-up methods [34], embedding comparisons [7], and hybrid approaches [8]. Record-based entity linking has seen advancements using classical machine learning [20], deep neural networks [27], and pre-trained language models [23]. In [38], a mixture-of-experts approach was proposed, utilizing the training results from various data integration matching tasks. For blocking, many state-of-the-art approaches use deep learning [2, 23, 37, 40]. A simple tf-idf based blocker can achieve competitive results without training and labeled data as shown in [32].

More recently, researchers started to examine LLMs on data integration tasks. Studies have evaluated and proposed various models, including GPT-3 [28], Jellyfish variants [43], and others [33], on tasks such as entity linkage, data imputa-

tion, and error detection. Different matching strategies for LLMs in detail have also been explored in [33, 41].

Within the specific domain of cryptoasset investigations, [13] described a method for resolving conflicting attribution tags. They use the edit distance to harmonize strings referring to the same entity. The application of LLMs in the context of cryptoassets and blockchains has been explored for multiple tasks, such as detecting anomalous Ethereum transactions [11], auditing smart contracts [5], and identifying discrepancies between smart contract bytecode and project documentation [12].

In this paper, we go beyond this approach, by linking attribution tag datasets to well-defined knowledge graph concepts using an LLM-based approach.

3 Data

In this section, we present the data that informed the design and implementation of our approach and was used throughout our experiments.

3.1 GraphSense TagPacks

GraphSense is an open-source cryptoasset analytics platform that provides a curated collection of over 500,000 publicly available attribution tags. A TagPack is a data structure used to package and share attribution tags³. Each tag corresponds to a blockchain address and includes a label, several optional fields, and categorization information. The categories are drawn from a subset of the *INTERPOL Dark Web and Virtual Assets Taxonomy (DWVA)*⁴, a community-driven effort to define common forms of abuse and entities representing real-world actors and services within the broader Darknet and Cryptoasset ecosystems.

In addition, GraphSense provides a curated list of 2,862 actors, each representing a well-defined real-world entity within the cryptoasset ecosystem. These actors encompass a wide range of roles and types in the industry such as centralized exchanges (e.g., Binance), decentralized finance platforms (e.g., Aave), and mixing services (e.g., Tornado Cash). Each actor is assigned a unique ID, a label, one or more categories from the DWVA taxonomy, and optional fields such as a URL or jurisdiction. This enables GraphSense to partially implement a knowledge graph, offering explicit links between entities, which we use as a ground-truth dataset. In total, 378,550 attribution tags contain such actor links.

Since many attribution tags share the same labels and actor links and we are only interested in unique records, we filtered out duplicates, leaving 2,570 unique linked attribution tags. These were split into training, validation, and test datasets in a 1:30:69 ratio. The training set was used to create few-shot examples, while the validation set was employed in experiments 1 and 2 to optimize the individual components. The test set was exclusively used in experiment 3 to evaluate our approach end-to-end.

³ <https://github.com/graphsense/graphsense-tagpacks>

⁴ <https://misp-galaxy.org/interpol-dwva/>

3.2 WatchYourBack Attribution Tags

We utilize the dataset published in [13]. Structurally, these attribution tags are similar to the GraphSense TagPack, as they also include categories and subcategories. However, their taxonomy is not harmonized with the GraphSense Actor Taxonomy.

The dataset consists of tags aggregated from various sources, including GraphSense and other shared datasets. To avoid duplicating tags, we filter out any that are linked to addresses already present in the GraphSense TagPack database. From this filtered set, we manually selected and annotated 126 records, 67 of which contain an actor link.

3.3 DeFi Rekt Database

The DeFi Rekt database [6] contains over 3,500 events related to crimes involving cryptoassets. Many of these events involve actors within the ecosystem. Each event includes a title and optional fields such as the date of the event, and funds involved. For our experiments, we randomly sampled 100 records from events where the loss of funds exceeded 100,000 USD. We manually annotated these sampled events with actor links. Out of the 100 sampled records, only 32 contained an actor link.

3.4 Dataset summary

Table 1 summarizes the data we used for design and implementation of our approach and the experiments we conducted. The total number of distinct actors is lower than the sum of individual datasets due to overlapping actors between datasets.

Table 1. Overview of the three attribution tag datasets used in this study.

Dataset	Samples	Actor Links	Distinct Actors
GraphSense Tag Pack	2,570	2,570	484
Train	25	25	20
Validation	771	771	198
Test	1,774	1,774	361
WatchYourBack	126	67	58
Defi Rekt	100	32	32
Total	2,794	2,669	520

4 Approach

Our approach for linking attribution tags to knowledge graph concepts comprises two main components, as illustrated in Figure 3: the *candidate set generator* and

the *candidate selector* modules. The candidate set generator reduces the pool of potential actors in the knowledge graph that may correspond to an attribution tag by applying filtering and blocking techniques. Next, the candidate selector module determines which, if any, of the proposed entities match the attribution tag. The following sections offer a detailed explanation of each module.




Fig. 3. Approach Overview. The *candidate set generator* filters potential entities and the *candidate selector module* identifies the matching entity.

4.1 Candidate Set Generator

The goal of the candidate set generator is to reduce the pool of candidates in the knowledge graph that potentially match a given attribution tag. Since LLM inference is expensive and pairwise comparison is of order $O(nm)$, where n is the number of records and m is the number of entities in the knowledge graph, an approach that minimizes the comparisons is required.

To reduce the overall cost, we limit the entities in each prompt to k candidates by applying *filtering* and *blocking* techniques, thereby reducing the problem complexity to $O(nk)$, where $k \ll m$. Filtering is the process of eliminating incorrect candidates, while blocking refers to similarity-based clustering that identifies likely and unlikely candidates [29].

Filtering The candidate set can be narrowed down when both the attribution tag and the corresponding knowledge graph entity are associated with categorization information from the same controlled vocabulary. For example, an attribution tag might be categorized as “exchange” and a knowledge graph entity as “service”, where the latter represents a semantically broader concept than the former. When such information is available, as in the GraphSense TagPack dataset (see Section 3.1), it can be leveraged for filtering. We define two possible filtering methods:

- *Same-Concept Filtering*: This method excludes actors that belong to a category different from the one specified in the attribution tag.

- *Related-Concept Filtering*: This approach is more flexible than same-concept filtering, as it leverages the taxonomy structure to exclude all actors associated with concepts unrelated to the attribution tag. Related concepts include both ancestors and descendants of the original concept, but exclude descendants of ancestor concepts that are not directly related.

Blocking The basic idea behind blocking is to avoid comparing entities that are unlikely to match, significantly reducing the number of comparisons. This is done by partitioning the data into smaller subsets (*blocks*), where entities share some similarities or common attributes. Inspired by [32], who demonstrate that simple blocking methods — without requiring machine learning or pre-trained models — can achieve strong results, we apply two straightforward methods:

BM25₃ This method is based on the Okapi BM25 [35] that is part of the family of term frequency-inverse document frequency (tf-idf) scoring functions. We tokenize the document and query strings into trigrams as proposed in [2] to allow for approximate string matching. The overall score for a tokenized document D with respect to a tokenized query Q is given by:

$$\text{BM25}(D, Q) = \sum_{i=1}^n \text{idf}(q_i) \cdot \frac{f(q_i, D) \cdot (k_1 + 1)}{f(q_i, D) + k_1 \cdot (1 - b + b \cdot \frac{|D|}{\text{avgdl}})} \quad (1)$$

where $f(q_i, D)$ is the frequency of token q_i of query q in document D and avgdl is the average document length in the corpus. The inverse document frequency (IDF) is calculated as:

$$\text{idf}(q_i) = \log \left(\frac{N - n(q_i) + 0.5}{n(q_i) + 0.5} + 1 \right) \quad (2)$$

We use the standard parameters $k_1 = 1.5$ and $b = 0.75$ defined in the implementation of the `rank_bm25` library⁵.

Overlap₃ This method measures the similarity between two strings based on the overlap $\frac{|A \cap B|}{\min(|A|, |B|)}$ of their trigram sets A and B .

4.2 Candidate Selector

The candidate selector module takes a set of candidates for each attribution tag and selects the best matching entity. Technically, this step is implemented using an LLM. In the first stage, the module constructs a batch of prompts, where each prompt corresponds to an attribution tag and includes all associated candidates. Optionally, the prompts can include examples (few-shot prompts) showcasing both matching and non-matching cases. This prompt batch is then fed to the LLM, which is tasked with either selecting the candidate that best

⁵ https://github.com/dorianbrown/rank_bm25

matches the attribution tag or indicating that none of the candidates correspond to the tag.

Figure 4 shows the template used for this prompting task. It consists of an extended system message (SYS, SYS+), the few-shot examples (FEW-SHOT), a task description (TASK), a domain statement (DOMAIN), the input data (INPUT), the selection question (QUEST), and an extended output format reminder (OUT, OUT+).

```

SYS You are an AI assistant that follows instruction extremely well.
SYS+ User will give you a question. Your task is to answer as faithfully as you can.
      Determine if the attribution tag label refers to any of the listed entities and, if so, to which one.
      The attribution tag describes an entity that is related to a blockchain address.
      Attribution Tag Label: compound_Pair_ETH_USDC
      [0] Compound
FEW-SHOT [1] Morpho Aave/Compound
           [2] Tether (USDT)
           [3] BitLaunder.com
           [4] CryptoBounty.com
           [5] None of the entities above
           0
DOMAIN The attribution tag describes an entity that is related to a blockchain address.
TASK Determine if the attribution tag label refers to any of the listed entities and, if so, to which one.
      Attribution Tag Label: curvefinance_UST-StableSwapUST
      [0] TrustSwap
      [1] Curve Finance
INPUT [2] PancakeSwap (AMM, v2, Stableswap)
           [3] Zyberswap (AMM, Stableswap)
           [4] StableKoi
           [5] None of the entities above
QUEST Choose your answer from: [0, 1, 2, 3, 4, 5].
OUT Does the attribution tag refer to any of the listed entities and, if so, to which one?
OUT+ Before answering, make sure your answer only contains a number.

```

Fig. 4. Prompt template. The template used for prompting candidate selection. It consists of several parts (e.g., SYS, FEW-SHOT, INPUT, etc.) and defines the instructions provided to an LLM to guide its response or generated output.

5 Experiments

5.1 Setup

Models The candidate selector module (see Section 4.2) leverages LLMs, integrating both remote models via the OpenAI API and locally hosted models. For our experiments, we utilized two remote and six local models, all compatible with consumer-grade hardware:

- *GPT-4o, GPT-3.5 Turbo*: These models, hosted by OpenAI, were employed in their specific versions: gpt-3.5-turbo-0125 and gpt-4o-2024-05-13.
- *Jellyfish-7B, Jellyfish-13B-AWQ*: Developed and fine-tuned for data pre-processing tasks [43], the Jellyfish models include a 7B version, based on Mistral-7B-Instruct-v0.2, and a 13B version, based on OpenOrca-Platypus2-13B. The 13B model was quantized to fit on consumer-grade GPUs using Activation-aware Weight Quantization (AWQ) [24].

- *Mistral 7B/7B-Instruct*: Both the base and instruction fine-tuned versions of the Mistral 7B model, developed by Mistral AI, were evaluated in their v0.3 release.
- *Meta LLaMA 3 8B / 8B-Instruct*: These models, developed by Meta, include both the base and instruction fine-tuned versions of the Meta LLaMA 3 8B model.

Prompt Configurations Prompt engineering can significantly influence LLM performance in tasks such as entity linking, as demonstrated by [33]. The authors show that different template structures yield better results with certain models, while other templates are more effective with different models. Building upon the structure of our base prompt template (see Figure 4), we are testing nine additional shorter template configurations, as detailed in Table 2. In all ten configurations, we are considering few-shot examples (FEW-SHOT) as optional and place them before the task description (TASK).

Table 2. Prompt template configurations. Each template represents a variation of the structural elements from our base prompt template in Figure 4.

Template	Configuration
0	SYS, SYS+, TASK, DOMAIN, QUEST, OUT, OUT+
1	SYS, TASK, DOMAIN, QUEST, OUT, OUT+
2	SYS, SYS+, TASK, DOMAIN, QUESTION, OUT
3	SYS, SYS+, TASK, QUEST, OUT, OUT+
4	SYS, SYS+, TASK, DOMAIN, QUESTION
5	TASK, DOMAIN, QUEST, OUT, OUT+
6	SYS, TASK, QUEST, OUT
7	TASK, QUESTION, OUT, OUT+
8	SYS, TASK, QUEST
9	TASK, QUEST

Hardware All experiments are conducted on an AWS g5.xlarge instance equipped with 16 GB of RAM, 4 vCPUs powered by a second-generation AMD EPYC processor, and an NVIDIA A10G Tensor Core GPU with 24 GB of VRAM. For all models, we set the temperature to 0. Local models are run using vLLM [21], an LLM serving system, with the `max_num_of_seq` parameter set to 128 and `enable_prefix_caching` to true.

5.2 Experiment 1: Candidate Set Generation

The goal of this experiment is to evaluate different filtering and blocking techniques (described in Section 4.1) and find an appropriate candidate set size. We run the experiment on the GraphSense Tag Pack validation dataset on candidate

set sizes (k) of 1, 5, 10, and 25. The task is to predict the correct actor for each sample within the k candidates by comparing the tag label with the actor label, e.g., `btc-e.com` with `BTC-e`. To evaluate the performance of combined filtering and blocking techniques, we compute the ratio of attribution tags for which the correct actor link is included among the k candidates. This metric is often referred to as top-k accuracy. However, since each attribution tag has exactly one correct actor link, the proportion of correctly recovered actor links, *recall*, is equivalent in this context. Following prior works [29,32] on blocking methods, we refer to this metric as recall. Note that in this scenario, precision depends solely on k and recall, and thus does not provide additional information.

Table 3. Effectiveness of candidate set generation. Recall of candidate set generation for different candidate set sizes (k), blocking (BM25₃, Overlap₃), and filtering (*same-concept*, *related-concept*) techniques.

Filter	Blocker	Recall			
		$k=1$	$k=5$	$k=10$	$k=25$
No Filtering	Overlap ₃	0.467	0.826	0.872	0.904
	BM25 ₃	<u>0.765</u>	<u>0.899</u>	<u>0.916</u>	0.957
same-concept	Overlap ₃	0.553	0.739	0.763	0.774
	BM25 ₃	0.669	0.770	0.776	0.782
related-concept	Overlap ₃	0.658	0.873	0.914	0.938
	BM25 ₃	0.811	0.933	0.944	<u>0.955</u>

The results in Table 3 demonstrate that BM25₃ consistently outperforms Overlap₃ across all candidate set sizes k . Using the *same-concept* filter degrades performance, suggesting that it is overly restrictive. In contrast, the related-concept filter produces slightly better candidate sets than using no filtering. The candidate sets improve only marginally when their size exceeds 5. Therefore, we conclude that for our subsequent experiments, $k = 5$ is the optimal candidate set size and BM25₃ is the preferred blocking method. The related-concept filter will be applied where relevant. An illustration of the performance on the different candidate set sizes can be found in Appendix 1

This demonstrates that basic blocking techniques, such as the BM25₃ blocker, are effective in generating candidate sets with as little as 5 elements, achieving 90% recall, and that related-concept filtering can further improve recall to 93%.

5.3 Experiment 2: Candidate Selection

The goal of our second experiment is to assess the performance of different LLM-based candidate selectors using various prompt template configurations (see Table 2) and few-shot examples. We run the experiment on the GraphSense

TagPack validation set, using the candidate sets generated by the previous experiment. The model’s task is to select the correct actor for each attribution tag from the candidate set if present, otherwise to predict that no candidate matches. The model selects a candidate by responding with the corresponding number, as illustrated in Figure 4. Invalid responses are classified as no match, following the guidelines of [28, 33]. We treat the task as a multiclass classification problem and evaluate the models using accuracy and macro-averages for recall, precision, and F1-score.

Table 4. Candidate Selector results on the GraphSense TagPack validation set, showing recall (**R**), precision (**P**), F1-score (**F1**), and accuracy (**Acc.**) for each model’s top-performing template (**T**). Best performance is in bold, second best is underlined.

Model	Zero-Shot					Five-Shot				
	T	R	P	F1	Acc.	T	R	P	F1	Acc.
GPT4o	1	0.907	0.919	0.910	0.953	7	0.939	0.940	0.939	0.983
GPT3.5	7	0.536	0.573	0.535	0.499	9	<u>0.929</u>	<u>0.928</u>	<u>0.927</u>	0.951
Jellyfish 7B	0	<u>0.804</u>	<u>0.802</u>	<u>0.797</u>	0.847	3	0.876	0.870	0.869	0.946
Jellyfish 13B	6	0.568	0.588	0.571	0.708	2	0.797	0.790	0.789	0.879
Llama 3 8B	0	0.744	0.743	0.738	0.827	0	0.868	0.859	0.861	0.947
Llama 3 8B-Inst	7	0.729	0.731	0.725	0.844	3	0.883	0.874	0.875	0.955
Mistral 7B	1	0.737	0.732	0.728	<u>0.850</u>	8	0.829	0.813	0.812	0.908
Mistral 7B-Inst	7	0.711	0.702	0.698	0.799	0	0.913	0.900	0.903	<u>0.958</u>

The results in Table 4 indicate that GPT-4o outperforms all other methods in both zero-shot and five-shot scenarios. Notably, while GPT-3.5 has the weakest performance with zero examples, it achieves the second-highest F1-score when five examples are included in the prompt. The Jellyfish 7B model ranks second in zero-shot F1-score, suggesting that its fine-tuning on diverse data preprocessing tasks positively impacts the entity selection task. In contrast, the Jellyfish 13B model underperforms compared to the other local models, despite having more parameters, confirming the findings of [43] that it performs worse on unseen tasks than its 7B counterpart. Among the local models, Mistral 7B-Instruct performs the best, achieving an F1 score of more than 90%.

In Figure 5, we can see that the template choice can have a significant effect on some models. For example, GPT-4o zero-shot results are close to zero on some templates, while on others, they achieve F1-scores of over 90%. The Meta Llama 3 base model has the most stable performance and is the only model that achieves more than 60% F1-score on every template with zero examples. An overview of all models can be found in Appendix 2.

Cost Analysis Choosing more extensive template configurations and introducing few-shot examples increases the size of the prompts, and thus makes inference




Fig. 5. Model Performance with different Templates. *GPT-4o*'s zero-shot results vary significantly across different templates, while *Llama-3 8B* has a more stable performance.

more expensive. For remote models, we define cost in \$USD based on the OpenAI API usage policy⁶ that charges based on the prompt size. The current rates are (\$5,\$15) per 1 million input/output tokens for GPT-4o and (\$0.5,\$1.5) for GPT-3.5 Turbo.

For local LLMs, we use inference time as the cost metric. We define inference time as the time taken by the model to process all prompt batches and generate the responses, which includes tokenization but does not include the loading of the weights or any prompt pre/post-processing. We calculate the average run time of five runs for each template in both zero- and five-shot settings.

To determine which models and what configurations provide the best cost-performance value we first define value as:

$$V_{T/S} = F1_{T/S} * (1 - \tilde{C}_{T/S}) \quad (3)$$

where T/S is the template/shot configuration and $\tilde{C} = \frac{C - \min(C)}{\max(C) - \min(C)}$ is the normalized cost. We normalize the cost for remote and local models separately.

In Table 5 we see that we can save over 90% on costs when using GPT-3.5 with template 9 while only losing 1% on the F1-score compared to the best GPT-4o configuration. The template configuration for local LLMs has less cost impact because we can reduce redundant computation to encode the large shared prefix of the prompts by using vLLM prefix caching [21]. However, using Mistral-7B-Inst with template 9 instead of template 0 still reduces inference time by more than 30%, with a performance decrease of less than 0.5%.

This experiment shows LLM-based entity linking is effective, with GPT-4o achieving a 94% F1-score. Additionally, LLM's that can run locally in a consumer-grade GPU perform well, reaching 90% F1-score. Furthermore, we show that cost-performance analysis can yield a 90% cost reduction for only 1% of performance decrease.

⁶ <https://openai.com/api/pricing/>

Table 5. Template/Shots (**T/S**) configuration and cost and performance differences (Δ **C/F1**) between models with the highest F1-score and the one with the best cost-performance value (**V**). Local and remote models are compared separately

Model	T/S	F1	C	V	Δ (%)	
					C	F1
GPT-4o	7/5	0.939	\$3.359	0.116		
GPT-3.5	9/5	0.927	\$0.256	0.876	-92.380	-1.163
Mistral-7B-Inst	0/5	0.903	33.951s	0.659		
Mistral-7B-Inst	9/5	0.899	22.777s	0.797	-32.941	-0.390

5.4 Experiment 3: End to End Entity Linking

The goal of this experiment is to test our approach of linking attribution tags to a knowledge graph end-to-end and compare it to baseline solutions. We run the experiment on the GraphSense TagPack test set, WatchYourBack, and Defi Rekt datasets. All samples pass through the candidate generator, and the resulting batch of candidate sets are then fed to the candidate selector. For the candidate set generator, we employ the BM25₃ blocker. Additionally, we apply the *related-concept* filtering for the GraphSense dataset, while no filtering is used for the other datasets, because their categories are not linked to knowledge graph concepts. For the candidate selector, we use for each LLM the template that achieved the best performance in the previous experiment. We follow experiment 2’s evaluation method, with the difference that a miss by the candidate set generator is counted as an error. We compare the LLMs with the following baseline methods:

BM25₃: We use the top ranking candidate of our BM25₃ blocker, and decide based on a threshold if it is a match or not. The threshold of *15.7238* was determined by evaluating the model’s precision and recall on all top-candidate scores in the GraphSense TagPack validation set, and selecting the one that maximizes the F1-score.

UnicornPlus, UnicornPlusFT: UnicornPlus [38] is a DeBERTa-based mixture-of-expert model that is fine-tuned for data integration matching tasks. For a fair comparison, we apply the same candidate set generation process and perform pairwise matching between attribution tags and their candidates. In case multiple candidates match, we apply the softmax function on each prediction and choose the one with the highest probability. Furthermore, we create UnicornPlusFT, a fine-tuned version of the model. For this, we reshuffle our GraphSense train and validation sets with a 80/20 training/validation split and train the model for 10 epochs using the settings proposed in [38].

Table 6 demonstrates that our end-to-end linking approach outperforms baseline methods across all datasets. Using BM25₃ blocking and GPT-4o as candidate selector, we achieve F1-scores of 79-85%. With Mistral 7B-Instruct, our

Table 6. Performance of various models on different datasets

Model	GraphSense		WatchYourBack		DeFi Rekt	
	F1	Acc.	F1	Acc.	F1	Acc.
BM25 ₃	0.718	0.789	0.495	0.516	0.393	0.510
UnicornPlus	0.667	0.445	0.558	0.651	0.352	0.670
UnicornPlusFT	0.783	0.873	0.542	0.603	0.419	0.610
GPT4o	0.853	0.927	0.801	0.873	0.793	0.930
GPT3.5	0.810	0.881	<u>0.756</u>	<u>0.825</u>	<u>0.691</u>	<u>0.890</u>
Llama 3 8B-Inst	0.814	<u>0.921</u>	0.625	0.746	0.516	0.710
Mistral 7B-Inst	<u>0.821</u>	0.918	0.692	0.786	0.547	0.770

best-performing local candidate selector, F1-scores range from 55-82%. Detailed results for all models and error type analysis are available in Appendix 3.

This experiment shows that our approach outperforms baseline methods on all three datasets by up to 37.4% in F1-score. Furthermore, it demonstrates the generalization capabilities of our approach by achieving F1-scores of over 79% on all datasets.

6 Discussion and Conclusions

In this paper, we addressed the issue of *attribution tag quality*, with a particular focus on data inconsistencies that arise when attribution tags are shared among different parties. We argue that data quality issues can mislead forensic investigations and even result in false convictions if addresses are labeled incorrectly. To solve this, we proposed a novel computational approach based on Large Language Models (LLMs) that automatically links attribution tags to well-defined concepts in knowledge graphs, addressing the semantic nature of the problem. We implemented our approach in an end-to-end pipeline and demonstrated that, when combined with filtering and blocking techniques, it outperforms existing methods. Additionally, we showed that pre-trained LLMs running locally on consumer-grade hardware achieve performance comparable to remote models. Furthermore, we demonstrated that carefully designed prompts can significantly reduce costs with only a marginal decrease in performance. Overall, we believe our approach not only addresses the pressing issue of inconsistent attribution tags but also has the potential to inspire broader efforts to improve data quality in other forensic investigation tools and platforms.

One limitation of our approach is its binding to specific application domains; so far, we lack evidence that our method is generally applicable to all record linkage problems. However, we believe that a data- and measurement-driven approach would be valuable for assessing the broader suitability of this method. Another limitation is the assumption that different parties (e.g., exchanges, investigators) use the same knowledge graph when exchanging attribution tag

records. If this is not the case, our approach does not harmonize the data but merely shifts the problem to a different abstraction level. However, significant efforts are being made to harmonize and adopt shared knowledge graphs within the field. For example, the Darkweb and Virtual Assets taxonomy developed by INTERPOL has been integrated into the Malware Information Sharing Platform (MISP) Galaxy⁷. This taxonomy helps categorize and enrich threat intelligence, with MISP Galaxies organizing related data clusters to describe higher-level concepts such as adversary groups, malware families, and vulnerabilities, simplifying complex data analysis for organizations.

Future work could include fine-tuning pre-trained LLMs to the cryptoasset domain, improving performance by recognizing semantically related terms that are syntactically different. Additionally, leveraging relationships like hypernyms and hyponyms in knowledge graphs could refine candidate matching. Extending the approach to automatically categorize cryptoasset addresses based on well-defined categories would also provide a more comprehensive solution.

Acknowledgments. Avice gratefully acknowledges funding from the Digital Finance CRC which is supported by the Cooperative Research Centres program, an Australian Government initiative. This research was partially funded by the Austrian security research program KIRAS of the Federal Ministry of Finance (BMF) under the project DeFiTrace (grant agreement number 905300) and the FFG BRIDGE project AMALFI (grant agreement number 898883).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

⁷ <https://misp-galaxy.org/interpol-dwva/>

Appendix 1

A visual illustration of the candidate set generation performances of the different blocking and filtering techniques can be found in Figure 6. We can see that the recall gain from $k = 1$ to $k = 5$ is significant, while subsequent increases in candidate set size have only a marginal impact.

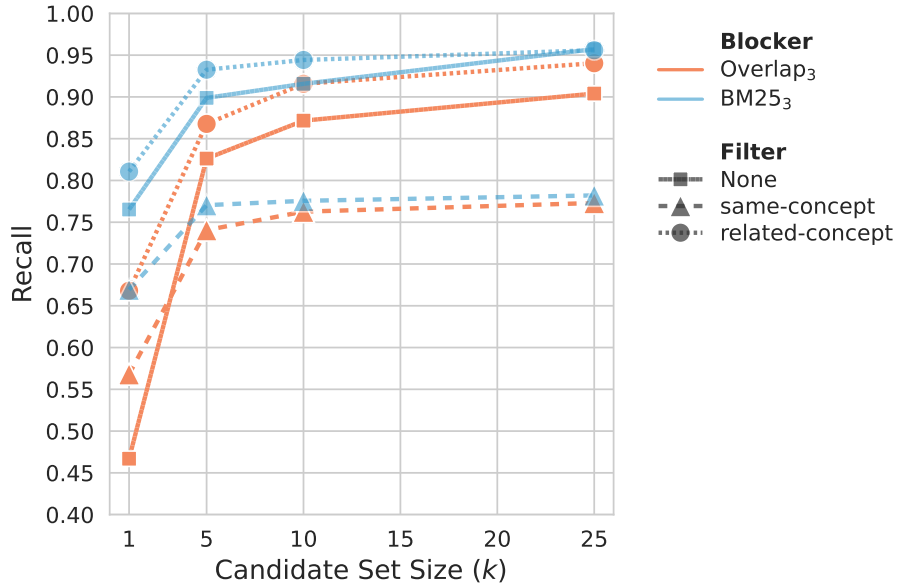


Fig. 6. Choosing Candidate Set Size. Recall of candidate set generation for different candidate set sizes (k). The performance gain for all methods visibly slows down for $k > 5$.

Appendix 2

A comprehensive overview of all candidate selector model results across each template can be found in Figure 7. We can see that the zero-shot performance of the models varies significantly across templates; however, introducing five examples reduces this variance considerably. Interestingly, while GPT-3.5 is ineffective without examples, with five examples it achieves the second-highest F1-score. Also notable is that the zero-shot performance for templates 4, 8, and 9 is lowest, signaling that the output format reminder (OUT) is crucial for good results without examples.

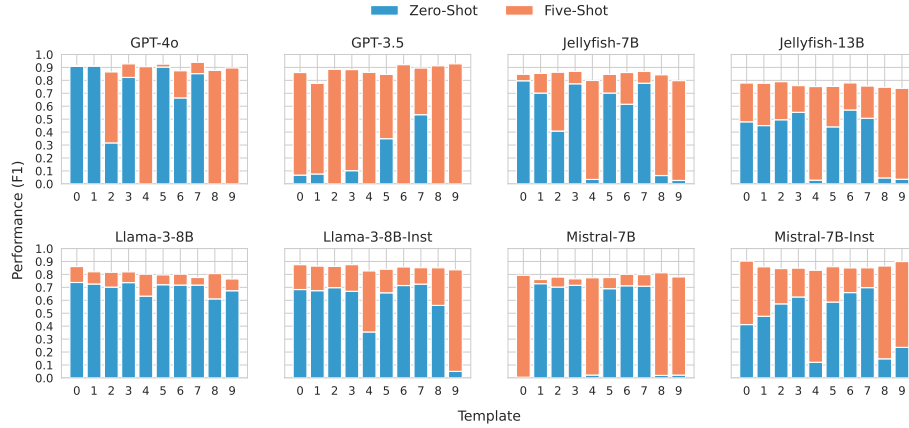


Fig. 7. *Candidate Selection Experiment:* Model performances (**F1**) across each template using no (*Zero-Shot*) and five (*Five-Shot*) examples in the prompt.

Appendix 3

In our experiments, we treat entity linking as a multi-classification problem and measure the performance with the macro F1 score and the accuracy, see Table 7. For the macro F1-score, this means that we calculate each entity’s individual F1 score and average them. Consequently, any prediction that is different from the ground truth is treated equally as an error. In a practical setting, we argue that errors predicting no entity matches are less problematic than errors linking the attribution tag to the wrong entity.

Figure 8 shows us the error composition of the different models across all three datasets. We distinguish between the *Missed Entity* error from a wrong no-match prediction and the *Wrong Entity* error with a wrong entity predicted. We can see that GPT-4o not only has the best performance in terms of accuracy and macro F1-score but also has less than 1% of wrong entities predicted, with the majority of the errors coming from missed entities. In contrast, the local LLMs all have more *Wrong Entity* than *Missed Entity* errors. Our best-performing local LLM, Mistral 7B-Instruct, has, despite better overall accuracy, more *Wrong Entity* errors than GPT-3.5.

In Table 8, we identify the three actors most frequently misclassified by each model. The results are similar across different models, with a few exceptions. For example, `maker` appears in the top three only for UnicornPlus, despite being present in 135 different tags. Notably, none of the models correctly classify any of the 20 tags associated with `0x` as an actor.

Table 7. Full overview of end-to-end entity linking performance including the results of all models the three datasets

Model	GraphSense		WatchYourBack		DeFi Rekt	
	F1	Acc.	F1	Acc.	F1	Acc.
BM25 ₃	0.718	0.789	0.495	0.516	0.393	0.510
UnicornPlus	0.667	0.445	0.558	0.651	0.352	0.670
UnicornPlusFT	0.783	0.873	0.542	0.603	0.419	0.610
GPT4o	0.853	0.927	0.801	0.873	0.793	0.930
GPT3.5	0.810	0.881	<u>0.756</u>	<u>0.825</u>	<u>0.691</u>	<u>0.890</u>
Jellyfish 7B	0.799	0.914	0.634	0.746	0.432	0.640
Jellyfish 13B	0.716	0.842	0.595	0.730	0.478	0.720
Llama 3 8B	0.785	0.917	0.618	0.683	0.440	0.680
Llama 3 8B-Inst	0.814	<u>0.921</u>	0.625	0.746	0.516	0.710
Mistral 7B	0.745	0.887	0.467	0.452	0.327	0.330
Mistral 7B-Inst	<u>0.821</u>	0.918	0.692	0.786	0.547	0.770



Fig. 8. Error analysis. The number of missed and wrongly predicted actor links for each model in the end-to-end entity linking experiment. The experiment contained 2000 samples from three different datasets.

Table 8. Overview of the most frequently misclassified actors for each model

Model	Actor	# Miss Classified	# Total Occurrences
BM25 ₃	aaave	121	133
	curve	38	128
	binance	33	37
UnicornPlus	aaave	133	133
	maker	128	135
	synthetix	122	198
UnicornPlusFT	huobi	51	53
	aaave	35	133
	0x	20	20
GPT4o	0x	20	20
	aaave	18	133
	feiprotocol	10	27
GPT3.5	synthetix	30	198
	aaave	22	133
	0x	20	20
Jellyfish-7B	aaave	29	133
	0x	20	20
	feiprotocol	10	27
Jellyfish-13B	aaave	52	133
	0x	20	20
	synthetix	19	198
Llama 3 8B	0x	20	20
	aaave	20	133
	feiprotocol	10	27
Llama 3 8B-Inst	aaave	24	133
	0x	20	20
	feiprotocol	11	27
Mistral 7B	aaave	23	133
	0x	20	20
	curve	19	128
Mistral 7B-Inst	aaave	21	133
	0x	20	20
	feiprotocol	10	27

References

1. Bartoletti, M., Carta, S., Cimoli, T., Saia, R.: Dissecting ponzi schemes on ethereum: Identification, analysis, and impact. *Future Generation Computer Systems* **102**, 259–277 (Jan 2020). <https://doi.org/10.1016/j.future.2019.08.014>, <http://dx.doi.org/10.1016/j.future.2019.08.014>
2. Brinkmann, A., Shraga, R., Bizer, C.: SC-Block: Supervised Contrastive Blocking Within Entity Resolution Pipelines, p. 121–142. Springer Nature Switzerland (2024). https://doi.org/10.1007/978-3-031-60626-7_7, http://dx.doi.org/10.1007/978-3-031-60626-7_7
3. Tell congress: Stop the use of chainalysis services, <https://www.change.org/p/tell-congress-stop-the-use-of-chainalysis-services>
4. Chen, W., Zheng, Z., Ngai, E.C.H., Zheng, P., Zhou, Y.: Exploiting blockchain data to detect smart ponzi schemes on ethereum. *IEEE Access* **7**, 37575–37586 (2019). <https://doi.org/10.1109/access.2019.2905769>, <http://dx.doi.org/10.1109/ACCESS.2019.2905769>
5. David, I., Zhou, L., Qin, K., Song, D., Cavallaro, L., Gervais, A.: Do you still need a manual smart contract audit? (Jun 2023). <https://doi.org/10.48550/arXiv.2306.12338>, <http://arxiv.org/abs/2306.12338>, arXiv:2306.12338 [cs]
6. De.Fi: Top Crypto Hacks - REKT Database. <https://de.fi/rekt-database> (2024), <https://de.fi/rekt-database>
7. Deng, X., Sun, H., Lees, A., Wu, Y., Yu, C.: Turl: table understanding through representation learning. *Proceedings of the VLDB Endowment* **14**(3), 307–319 (Nov 2020). <https://doi.org/10.14778/3430915.3430921>, <http://dx.doi.org/10.14778/3430915.3430921>
8. Efthymiou, V., Hassanzadeh, O., Rodriguez-Muro, M., Christophides, V.: Matching Web Tables with Knowledge Base Entities: From Entity Lookups to Entity Embeddings, p. 260–277. Springer International Publishing (2017). https://doi.org/10.1007/978-3-319-68288-4_16, http://dx.doi.org/10.1007/978-3-319-68288-4_16
9. Elmagarmid, A.K., Ipeirotis, P.G., Verykios, V.S.: Duplicate record detection: A survey. *IEEE Transactions on Knowledge and Data Engineering* **19**(1), 1–16 (Jan 2007). <https://doi.org/10.1109/tkde.2007.250581>, <http://dx.doi.org/10.1109/TKDE.2007.250581>
10. Fröwis, M., Gottschalk, T., Haslhofer, B., Rückert, C., Pesch, P.: Safeguarding the evidential value of forensic cryptocurrency investigations. *Forensic Science International: Digital Investigation* **33**, 200902 (Jun 2020). <https://doi.org/10.1016/j.fsidi.2019.200902>, <http://dx.doi.org/10.1016/j.fsidi.2019.200902>
11. Gai, Y., Zhou, L., Qin, K., Song, D., Gervais, A.: Blockchain Large Language Models (Apr 2023). <https://doi.org/10.48550/arXiv.2304.12749>, <http://arxiv.org/abs/2304.12749>, arXiv:2304.12749 [cs]
12. Gan, R., Zhou, L., Wang, L., Qin, K., Lin, X.: DeFiAligner: Leveraging Symbolic Analysis and Large Language Models for Inconsistency Detection in Decentralized Finance. *LIPICs, Volume 316, AFT 2024* **316**, 7:1–7:24 (2024). <https://doi.org/10.4230/LIPICs.AFT.2024.7>, <https://drops.dagstuhl.de/entities/document/10.4230/LIPICs.AFT.2024.7>
13. Gomez, G., Moreno-Sanchez, P., Caballero, J.: Watch your back: Identifying cybercrime financial relationships in bitcoin through back-and-forth exploration. In: *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, p. 1291–1305. CCS '22, ACM (Nov 2022). <https://doi.org/10.1145/3548606.3560587>, <http://dx.doi.org/10.1145/3548606.3560587>

14. Google: Introducing the knowledge graph: Things, not strings (2012), <https://blog.google/products/search/introducing-knowledge-graph-things-not/>, accessed: 2024-08-30
15. Harlev, M., Sun Yin, H., Langenheldt, K., Mukkamala, R., Vatrappu, R.: Breaking bad: De-anonymising entity types on the bitcoin blockchain using supervised machine learning. In: Proceedings of the 51st Hawaii International Conference on System Sciences 2018. pp. 3497–3506. Proceedings of the Annual Hawaii International Conference on System Sciences, Hawaii International Conference on System Sciences (HICSS), United States (2018), <http://www.urbanccd.org/events/2018/1/3/hawaii-international-conference-on-system-sciences-hicss-51>, the 51st Hawaii International Conference on System Sciences. HICSS 2018, HICSS2018 ; Conference date: 03-01-2018 Through 06-01-2018
16. Haslhofer, B., Klas, W.: A survey of techniques for achieving metadata interoperability. *ACM Computing Surveys* **42**(2), 1–37 (Feb 2010). <https://doi.org/10.1145/1667062.1667064>, <http://dx.doi.org/10.1145/1667062.1667064>
17. Hogan, A., Blomqvist, E., Cochez, M., D'amato, C., Melo, G.D., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., Ngomo, A.C.N., Polleres, A., Rashid, S.M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., Zimmermann, A.: Knowledge graphs. *ACM Computing Surveys* **54**(4), 1–37 (Jul 2021). <https://doi.org/10.1145/3447772>, <http://dx.doi.org/10.1145/3447772>
18. Hu, S., Zhang, Z., Luo, B., Lu, S., He, B., Liu, L.: Bert4eth: A pre-trained transformer for ethereum fraud detection. In: Proceedings of the ACM Web Conference 2023. WWW '23, ACM (Apr 2023). <https://doi.org/10.1145/3543507.3583345>, <http://dx.doi.org/10.1145/3543507.3583345>
19. Huang, D.Y., Aliapoulios, M.M., Li, V.G., Invernizzi, L., Bursztein, E., McRoberts, K., Levin, J., Levchenko, K., Snoeren, A.C., McCoy, D.: Tracking ransomware end-to-end. In: 2018 IEEE Symposium on Security and Privacy (SP). IEEE (May 2018). <https://doi.org/10.1109/sp.2018.00047>, <http://dx.doi.org/10.1109/SP.2018.00047>
20. Konda, P., Das, S., Suganthan G. C., P., Doan, A., Ardalan, A., Ballard, J.R., Li, H., Panahi, F., Zhang, H., Naughton, J., Prasad, S., Krishnan, G., Deep, R., Raghavendra, V.: Magellan: toward building entity matching management systems. *Proceedings of the VLDB Endowment* **9**(12), 1197–1208 (Aug 2016). <https://doi.org/10.14778/2994509.2994535>, <http://dx.doi.org/10.14778/2994509.2994535>
21. Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C.H., Gonzalez, J., Zhang, H., Stoica, I.: Efficient memory management for large language model serving with pagedattention. In: Proceedings of the 29th Symposium on Operating Systems Principles. SOSP '23, ACM (Oct 2023). <https://doi.org/10.1145/3600006.3613165>, <http://dx.doi.org/10.1145/3600006.3613165>
22. Li, S., Gou, G., Liu, C., Hou, C., Li, Z., Xiong, G.: Ttag: Temporal transaction aggregation graph network for ethereum phishing scams detection. In: Proceedings of the ACM Web Conference 2022. WWW '22, ACM (Apr 2022). <https://doi.org/10.1145/3485447.3512226>, <http://dx.doi.org/10.1145/3485447.3512226>
23. Li, Y., Li, J., Suhara, Y., Doan, A., Tan, W.C.: Deep entity matching with pre-trained language models. *Proceedings of the VLDB Endowment* **14**(1), 50–60 (Sep 2020). <https://doi.org/10.14778/3421424.3421431>, <http://dx.doi.org/10.14778/3421424.3421431>
24. Lin, J., Tang, J., Tang, H., Yang, S., Chen, W.M., Wang, W.C., Xiao, G., Dang, X., Gan, C., Han, S.: Awq: Activation-aware weight quantization

- for on-device llm compression and acceleration. In: Gibbons, P., Pekhimenko, G., Sa, C.D. (eds.) *Proceedings of Machine Learning and Systems*. vol. 6, pp. 87–100 (2024), https://proceedings.mlsys.org/paper_files/paper/2024/file/42a452cbafa9dd64e9ba4aa95cc1ef21-Paper-Conference.pdf
25. Liu, X., Tang, Z., Li, P., Guo, S., Fan, X., Zhang, J.: A graph learning based approach for identity inference in dapp platform blockchain. *IEEE Transactions on Emerging Topics in Computing* **10**(1), 438–449 (Jan 2022). <https://doi.org/10.1109/tetc.2020.3027309>, <http://dx.doi.org/10.1109/TETC.2020.3027309>
 26. Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G.M., Savage, S.: A fistful of bitcoins: characterizing payments among men with no names. In: *Proceedings of the 2013 conference on Internet measurement conference*. IMC'13, ACM (Oct 2013). <https://doi.org/10.1145/2504730.2504747>, <http://dx.doi.org/10.1145/2504730.2504747>
 27. Mudgal, S., Li, H., Rekatsinas, T., Doan, A., Park, Y., Krishnan, G., Deep, R., Arcaute, E., Raghavendra, V.: Deep learning for entity matching: A design space exploration. In: *Proceedings of the 2018 International Conference on Management of Data*. SIGMOD/PODS '18, ACM (May 2018). <https://doi.org/10.1145/3183713.3196926>, <http://dx.doi.org/10.1145/3183713.3196926>
 28. Narayan, A., Chami, I., Orr, L., Ré, C.: Can foundation models wrangle your data? *Proceedings of the VLDB Endowment* **16**(4), 738–746 (Dec 2022). <https://doi.org/10.14778/3574245.3574258>, <http://dx.doi.org/10.14778/3574245.3574258>
 29. Papadakis, G., Skoutas, D., Thanos, E., Palpanas, T.: Blocking and filtering techniques for entity resolution: A survey. *ACM Computing Surveys* **53**(2), 1–42 (Mar 2020). <https://doi.org/10.1145/3377455>, <http://dx.doi.org/10.1145/3377455>
 30. Paquet-Clouston, M., Haslhofer, B., Dupont, B.: Ransomware payments in the bitcoin ecosystem. *Journal of Cybersecurity* **5**(1) (Jan 2019). <https://doi.org/10.1093/cybsec/tyz003>, <http://dx.doi.org/10.1093/cybsec/tyz003>
 31. Paquet-Clouston, M., Romiti, M., Haslhofer, B., Charvat, T.: Spams meet cryptocurrencies: Sextortion in the bitcoin ecosystem. In: *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*. AFT '19, ACM (Oct 2019). <https://doi.org/10.1145/3318041.3355466>, <http://dx.doi.org/10.1145/3318041.3355466>
 32. Paulsen, D., Govind, Y., Doan, A.: Sparkly: A simple yet surprisingly strong tf/idf blocker for entity matching. *Proceedings of the VLDB Endowment* **16**(6), 1507–1519 (Feb 2023). <https://doi.org/10.14778/3583140.3583163>, <http://dx.doi.org/10.14778/3583140.3583163>
 33. Peeters, R., Bizer, C.: Entity matching using large language models (Feb 2024), <https://arxiv.org/abs/2310.11244v2>
 34. Ritze, D., Lehmborg, O., Bizer, C.: Matching html tables to dbpedia. In: *Proceedings of the 5th International Conference on Web Intelligence, Mining and Semantics*. WIMS '15, ACM (Jul 2015). <https://doi.org/10.1145/2797115.2797118>, <http://dx.doi.org/10.1145/2797115.2797118>
 35. Robertson, S.E., Walker, S.: Some Simple Effective Approximations to the 2-Poisson Model for Probabilistic Weighted Retrieval, p. 232–241. Springer London (1994). https://doi.org/10.1007/978-1-4471-2099-5_24, http://dx.doi.org/10.1007/978-1-4471-2099-5_24
 36. Shen, W., Wang, J., Han, J.: Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineer-*

- ing **27**(2), 443–460 (Feb 2015). <https://doi.org/10.1109/tkde.2014.2327028>, <http://dx.doi.org/10.1109/TKDE.2014.2327028>
37. Thirumuruganathan, S., Li, H., Tang, N., Ouzzani, M., Govind, Y., Paulsen, D., Fung, G., Doan, A.: Deep learning for blocking in entity matching: a design space exploration. *Proceedings of the VLDB Endowment* **14**(11), 2459–2472 (Jul 2021). <https://doi.org/10.14778/3476249.3476294>, <http://dx.doi.org/10.14778/3476249.3476294>
 38. Tu, J., Fan, J., Tang, N., Wang, P., Li, G., Du, X., Jia, X., Gao, S.: Unicorn: A unified multi-tasking model for supporting matching tasks in data integration. *Proceedings of the ACM on Management of Data* **1**(1), 1–26 (May 2023). <https://doi.org/10.1145/3588938>, <http://dx.doi.org/10.1145/3588938>
 39. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Communications of the ACM* **57**(10), 78–85 (2014)
 40. Wang, R., Li, Y., Wang, J.: Sudowoodo: Contrastive self-supervised learning for multi-purpose data integration and preparation. In: *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. vol. 14, p. 1502–1515. IEEE (Apr 2023). <https://doi.org/10.1109/icde55515.2023.00391>, <http://dx.doi.org/10.1109/ICDE55515.2023.00391>
 41. Wang, T., Chen, X., Lin, H., Chen, X., Han, X., Wang, H., Zeng, Z., Sun, L.: Match, compare, or select? an investigation of large language models for entity matching (2024), <https://arxiv.org/abs/2405.16884>
 42. Wu, J., Yuan, Q., Lin, D., You, W., Chen, W., Chen, C., Zheng, Z.: Who are the phishers? phishing scam detection on ethereum via network embedding. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **52**(2), 1156–1166 (Feb 2022). <https://doi.org/10.1109/tsmc.2020.3016821>, <http://dx.doi.org/10.1109/TSMC.2020.3016821>
 43. Zhang, H., Dong, Y., Xiao, C., Oyamada, M.: Jellyfish: A large language model for data preprocessing (Mar 2024), <https://arxiv.org/abs/2312.01678v4>
 44. Zhou, J., Hu, C., Chi, J., Wu, J., Shen, M., Xuan, Q.: Behavior-aware account de-anonymization on ethereum interaction graph. *IEEE Transactions on Information Forensics and Security* **17**, 3433–3448 (2022). <https://doi.org/10.1109/tifs.2022.3208471>, <http://dx.doi.org/10.1109/TIFS.2022.3208471>