UNARI: An Uncertainty-aware Approach to AS Relationships Inference

Guoyao Feng
Carnegie Mellon University
gfeng@cs.cmu.edu

Srinivasan Seshan
Carnegie Mellon University
srini@cs.cmu.edu

Peter Steenkiste
Carnegie Mellon University
prs@cs.cmu.edu

ABSTRACT

Over the last two decades, several algorithms have been proposed to infer the type of relationship between Autonomous Systems (ASes). While the recent works have achieved increasingly higher accuracy, there has not been a systematic study on the uncertainty of AS relationship inference. In this paper, we analyze the factors contributing to this uncertainty and introduce a new paradigm to explicitly model the uncertainty and reflect it in the inference result. We also present UNARI, an exemplary algorithm implementing this paradigm, that leverages a novel technique to capture the interdependence of relationship inference across AS links.

CCS CONCEPTS

• **Networks** → Network architectures; Network measurement.

KEYWORDS

AS relationships, routing policies, measurement

ACM Reference Format:

Guoyao Feng, Srinivasan Seshan, and Peter Steenkiste. 2019. UNARI: An Uncertainty-aware Approach to AS Relationships Inference. In *The 15th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '19)*, December 9–12, 2019, Orlando, FL, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3359989.3365420>

1 INTRODUCTION

The Internet inter-domain routing topology consists of thousands of independent interconnected autonomous systems (AS). To route traffic in the Internet, these ASes connect to each other through the Border Gateway Protocol (BGP) and establish routing relationships, which are based on their private business agreements. Knowing these relationships is essential to understanding how packets are routed on the network today and, therefore, critical to evaluating new protocol designs. Unfortunately, these business agreements are typically kept private, which has motivated numerous past efforts [7–9, 15, 20, 25, 28] to infer the relationships based on observations of network routing. Past efforts classify the relationship between two ASes into one of the three types: customer-to-provider

(c2p) ¹, peer-to-peer (p2p), and sibling-to-sibling (s2s). These algorithms have focused on improving the accuracy of these labels, but inference algorithms inherently incorporate some form of uncertainty and error.

For example, many prior works rely on the valley-free principle [9], which is, by itself, inadequate for deterministically inferring the relationship of every link in the AS-level topology [8]. As a result, these works use additional heuristics to classify links. Past efforts do consider the reliability of these heuristics “implicitly” by rank ordering them, but the relative uncertainty that is associated with these heuristics is never conveyed in the final output. In addition, inference results depend on the coverage of AS paths received by the route collectors. For example, some algorithms [15] rely on ASes’ transit/node degrees to determine the provider in a p2c relationship. However, newly deployed collectors may uncover more AS paths, causing ASes’ transit/node degrees and inferred relationships to change.

Motivated by the challenges above, we propose a novel framework that allows the algorithmic implementation to explicitly express and model uncertainty associated with inferred relationships. Our framework features two steps: *principle-based* and *uncertainty-aware* inference. The principle-based inference step applies only dependable rules (principles) to make deterministic relationship inferences and ensures the inference results remain stable under varying coverage of AS paths. Next, for each undecided link, the probabilistic inference step uses less reliable heuristics and keeps track of the level of confidence in the inferences. This information is then conveyed to the users. This can be done in a variety of ways, ranging from binary confidence tag (high/low) to a numeric score. Thus, our paradigm embraces uncertainty and makes it explicit when there is no strong evidence in favor of a particular type of relationship.

We have also designed a concrete algorithm, UNARI, that illustrates the value of this framework. UNARI uses two principles for the principle-based step, the existence of a clique of well-known Tier-1 ASes interconnected by p2p links [20] and the valley-free principle. For the uncertainty-aware step, we develop a novel technique to capture the interdependence between undecided links. Our evaluation results show that UNARI’s principle-based inference achieves higher accuracy and is more stable than prior algorithms. Moreover, UNARI’s uncertainty-aware inference outperforms a naïve approach to expressing uncertainty. We also evaluate the efficacy of the share metric by comparing it against other metrics such as transit and node degrees as features for learning. We show that the share metric is not only effective by itself but also complementary to the transit and node degrees.

The contributions of our work are summarized as follows:

¹Inferring an AS link $u - v$ as c2p is equivalent to inferring the link $v - u$ as p2c.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CoNEXT '19, December 9–12, 2019, Orlando, FL, USA
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6998-5/19/12...\$15.00
<https://doi.org/10.1145/3359989.3365420>

- (1) We identify 4 sources of uncertainty in AS relationship and analyze their impact on the inference results.
- (2) We propose a novel framework to model inference uncertainty explicitly and reflect it in the results.
- (3) We propose UNARI, an algorithm that implements this framework and leverages a novel technique to capture the interdependence between undecided links. Our evaluation shows that it achieves higher accuracy and stability than prior algorithms for deterministic inference and outperforms a naïve strategy to express uncertainty.

The paper is organized as follows. §2 describes the related works on AS relationship inference. In §3 we provide insights into the causes of uncertainty in relationship inference. §4 describes the new paradigm to explicitly reflect uncertainty in the inference output, followed by §5 that presents the UNARI algorithm as an exemplary implementation of the paradigm. We evaluate the performance of UNARI in §6. §7 concludes our work.

2 RELATED WORK AND BACKGROUND

Inference Algorithms. The topic of AS relationship inference was started by Gao’s pioneering work [9] that classified an AS link into one of the three relationship types: p2c, p2p, and s2s. In a c2p relationship, the customer AS pays the provider AS for the traffic transmitted between them to obtain global reachability. In a p2p relationship, the two parties transmit traffic between their own networks and their customers’ networks with no payment involved. In the case of s2s relationship, the two ASes are owned by the same organization, and they transmit traffic between their providers/peers/siblings for free.

New inference algorithms in follow-up research, including MVP [25], UCLA [20], and ASRANK [15], have achieved increasingly higher accuracy. But **uncertainty** arises in relationship inference when they face problems such as the incomplete coverage of AS-level topology and the use of potentially unreliable heuristics. Unfortunately, none of these algorithms explicitly address it.

ProbLink [13] is a recent attempt to address uncertainty due to noisy but useful link attributes by taking a probabilistic approach to inference. The authors of ProbLink find that the types of most links in validation datasets are easy to infer, which explains the contradiction between prior algorithms’ high accuracy and their sub-optimal performance when applied to practical applications. ProbLink seeks to improve the inference accuracy on subsets of hard-to-infer links. Unlike UNARI, ProbLink’s objective is thus not modelling or expressing uncertainty. The distinction is reflected in their inference output; ProbLink chooses a single relationship with the highest probability for each link, whereas UNARI aims to develop a new paradigm that calculates uncertainty for all links and displays multiple relationships for each link, along with the probability of each relationship. Moreover, we highlight two differences in terms of algorithm design. ProbLink simultaneously takes into account all relevant link attributes to infer a link’s type. In contrast, UNARI performs inference in multiple phases and applies more dependable inference rules first. To capture the interdependence across links on different AS paths, ProbLink continuously updates the link type inferences and iterates until it converges. On the other hand, UNARI’s uncertainty-aware inference step computes

a numeric feature for undecided links. The feature measures the number of valley-free inferences over the entire topology, under the condition that a link is inferred as one of the relationship types.

Data. This paper uses the following datasets. First, we downloaded BGP data from three archives: PCH [11], RIS [5], and Routeviews [19]. PCH (Packet Clearing House) manages route collectors deployed on more than 100 IXPs around the world. It makes publicly available the daily routing table snapshots from these collectors. RIS (Routing Information Service), developed by RIPE NCC, also collects and stores Internet routing data from several locations around the globe. Routeviews started as a tool to obtain real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet. It maintains a data archive of BGP RIBs and updates. Each collector manages multiple vantage points and receives their BGP data feeds.

The BGP dataset used by our experiments contained the AS paths from the daily routing table snapshots between April 1, 2012 and April 5, 2012. We chose this time range to match the timestamp of the ground truth dataset for evaluation, as we describe below. We combine AS paths from all collectors into the single *aggregate* path set, Q , which is then used as the input to the inference algorithms. We use I to denote the set of all collectors. I captures 79 route collectors with at least one valid AS route, over 128k AS links and more than 41k ASes.

We also use the *partial* ground truth dataset released by [15]. It is derived from community attributes [6] and RPSL [1] from April 2012. It covers 16,248 p2p links and 31,886 p2c/c2p links.

3 UNCERTAINTY IN AS RELATIONSHIP INFERENCE

In this section, we present 4 factors (§3.1-3.4) that make it impossible to deterministically infer a single relationship for certain links and describe how these factors can influence an algorithm’s output. We also provide some insights into possible strategies to mitigate their impact, paving the way for an uncertainty-aware solution. At the end, we discuss the potential benefits of exposing inference uncertainty.

3.1 Dependence on Route Collectors’ Coverage

The ability to deploy route collectors can place limits on our visibility of the AS-level Internet topology. The incompleteness of the observable AS-level Internet topology has been studied by past research efforts. For example, [20, 21] showed that single-period snapshots of public BGP datasets only reveal a small portion of the links connecting Tier-2 ASes; combining long period snapshots and historical data of BGP updates leads to significant improvement. Moreover, the public BGP data misses many peer links at Tier-2 and below, as a result of the valley-free property and the poor coverage of monitors at stub ASes. In this section, we analyze how these coverage limitations influence the results of inference algorithms. Specifically, §3.1.1 illustrates the effect through concrete examples drawn from prior algorithms: Gao’s, ASRANK, and MVP. §3.1.2 presents an empirical study using public BGP datasets.

Let us define the key terms used in the rest of the paper. The *vantage points* of a collector are identified by the set of first-hop ASes along the AS paths it receives. In Figure 1, vantage point A

is feeding BGP routing data (i.e., path $A - B - C - D$) to its route collector RC .

We quantify a route collector's *coverage* using the number of links and vantage points it is able to observe². We refer to the undirected graph formed by the set of AS paths from a set of vantage points as the *Original* AS topology whereas the one with additional vantage points as the *Augmented* AS topology. Given an AS link shared by both, an inference algorithm is deemed *unstable* if it infers different relationships under the two topologies, which implies inference uncertainty due to route collectors' limited coverage.

3.1.1 Instability Causes. Algorithms can be unstable when paths from new vantage points becomes available as input, if they infer relationships based on measures that involve aggregates such as the total number of adjacent ASes.

We use Gao's algorithm as an example to illustrate the problem. The algorithm starts by computing the node degree of each AS in the topology. Next, for each AS path, one of the ASes with the highest node degree is selected as the peak of the valley-free path. It then increments the c2p score of links preceding the peak and the p2c score of links following the peak. The third step infers links as p2c/c2p if their scores exceed a user-defined threshold. Finally, the algorithm revisits each AS path, filters out links that are impossible to have peering relationships, and assigns p2p relationship to links whose end points have similar node degrees.

We now use Figure 1 to show how incomplete coverage affects the output of Gao's algorithm. Consider the AS path $A - B - C - D$ in the Original topology (step 1). Both B and C have a node degree of 2. Gao's algorithm chooses one of the ASes with the highest node degree as the peak of the valley-free path (step 2); ASes on both sides of the peak are classified as either peers or customers. Without loss of generality we assume it selects B as the peak of the AS path and $B - C$ is classified as either p2c or p2p. Now consider a new path (step 3) in the Augmented topology, $X - C - Y$, which leads to C 's node degree to increase from 2 to 4. C is now chosen as the peak in the first AS path and the link $B - C$ changes from p2c/p2p to c2p (step 4).

Note that the addition of vantage points impacts other algorithms in similar ways. For example, vantage point addition impacts transit degree, used by the ASRANK algorithm, and vantage-point rank, used by the MVP algorithm.

3.1.2 Empirical Study of Stability. We present an empirical study to evaluate the stability of Gao's, MVP, UCLA and ASRANK algorithms.

Varying Topology Coverage. We evaluate the stability of relationship inference by progressively excluding route collectors from the set of all collectors, I . This approach allows us to study the behavior of inference algorithms when existing collectors become unavailable or new collectors are introduced. We show in App. A that individual router collectors' coverage is unique and complementary to each other. Hence, excluding a subset of route collectors reduces the collective coverage of remaining route collectors.

²We choose not to measure by the number of unique paths because paths traversing the top of the hierarchy tend to share the same segments and contribute little new information in terms of coverage. For example, the same path announcement received by two vantage points constitute two unique paths even though the only difference is the ASN prepended by the vantage points.

Figure 1: An example illustrating the impact of incomplete coverage on Gao's algorithm's stability. In the Original topology, either B or C is the top provider. C becomes the top provider after a new path is introduced by the Augmented topology

Suppose H denotes a set of indices of excluded collectors, the remaining collectors are therefore $\bar{H} = I \setminus H$. For example, if $H = \{rrc15\}$, $\bar{H}$ includes all collectors but "rrc15". We refer to this step as *shrinking*. We bootstrap H by generating a small random sample from the set of all collectors and gradually grows its by adding more random samples from the remaining collectors. Throughout this process, a collector never returns to $\bar{H}$ once it is moved to H . In total, we randomly produce 20 series of $\bar{H}$ for the results below.

We measure the stability of inference algorithms as the number of *relationship transitions*, or the number of links whose inference result changes due to shrinking.

Stability Results. Figure 2a and 2b show that the number of links and ASes removed increases as we progressively shrink $\bar{H}$. As $|\bar{H}|$ declines from 79 to 7, the average loss of ASes rises from 0 to 2301 (5.6%). However, the number of links experiences a dramatic drop, losing coverage of more than 38% (on average) of the links observed by I . It is also important to realize that the number of links removed is not proportional to the number of excluded datasets because link coverage by different route collectors may overlap, especially at the top of the network hierarchy.

Figure 2c shows the number of relationship transitions by UCLA as a result of shrinking $\bar{H}$. At each step of shrinking $\bar{H}$, we measure the number of links whose inferred relationship is different under I and $\bar{H}$. For instance, the curve labeled 'p2p_p2c' represents links whose inferred relationship changes from p2p to p2c with respect to I . In the case of UCLA, the curve of 'p2c_p2p' is trending upwards as more collectors are excluded and eventually reaches a mean of 7.8k transitions. In contrast, the curves of 'p2c_c2p' and 'p2p_p2c' stay mostly flat. Judging from the rightmost column, one is tempted to think that an average of 8k total transitions does not seem substantial with respect to the 128k links in I . But we note that it only requires the loss of coverage over 38% of the links and the transitions affect 10% of the remaining links, as shown in Figure 2a. We further study the number of relationship transitions for Gao's,

Figure 2: Impact of reducing coverage.

MVP, and ASRANK. Similar to UCLA, we observe thousands of links whose inferred relationship changes from p2c to p2p with respect to I . The details are presented in App. B.

The observations illustrate the level of stability for different algorithms. While experimental results [15] show the best-performing algorithms such as ASRANK and UCLA achieve high precision and recall ($> 90\%$), they are sensitive to the number of route collectors available.

3.2 Uncertainty in p2p Relationship Inference

Figure 3: An example illustrating the challenge of inferring p2p relationships with certainty.

It is hard to unambiguously infer p2p relationships because commonly used techniques, such as the valley-free constraint, are incapable of determining that a link must be p2p and completely eliminate the possibility of it being p2c [7, 8]. Consider a valid path with at least one AS link. If all links are c2p/p2c, there must exist a top provider adjacent to one or two customers under the valley-free constraint. The path remains valid if we classify a p2c link to the top provider as p2p. Conversely, given a valid path with a p2p link, replacing the p2p relationship with p2c or c2p will still leave the path valid. In general, this makes it more difficult to infer p2p than

p2c relationships and the inference algorithms tend to leave the inference of p2p relationships as the last step after all p2c inferences are made.

Figure 3 illustrates the challenge of inferring p2p relationship with certainty. Consider a path $A - B - C - D$. In the first scenario, suppose that UCLA, or any algorithm that prioritizes p2c/c2p inference over p2p [8, 9, 15], is able to infer $C - D$ as p2c and $A - B$ as c2p.³ Algorithms typically infer the undecided link $B - C$ as p2p. But it could in fact be inferred as any one of the three relationships without violating the valley-free constraint. However, algorithms pick a single relationship, despite the uncertainty inherent to the scenario. Now consider another pathological scenario where the algorithm is only able to infer $C - D$ as p2c, leaving two consecutive links undecided. Inferring both $A - B$ and $B - C$ as p2p violates the valley-free property and assign a single relationship, so they must speculate whether there exists a p2p link on this path and if yes, which of $A - B$ and $B - C$ should be inferred as p2p.

Prior algorithms have developed techniques to achieve high accuracy in p2p inference, but they are not as effective in handling such scenarios. For example, a common approach is to accurately infer as many p2c/c2p relationships as possible and the remaining links are then inferred as p2p. We focus our discussion on c2p/p2p/p2c because s2s relationships are uncommon. The accuracy of p2p inference now hinges on the performance of p2c/c2p inference. However, the lack of evidence to support p2c/c2p inference is not a guarantee that a link is p2p, as the above example shows. As another example, UCLA assumes that a provider wants to announce paths from its customers to its providers and the monitors at Tier-1 ASes should be able to reveal all the downstream p2c relationships over time [20]. A link hidden from all the Tier-1 ASes is less likely to be p2c, and thus should be inferred as p2p. Again, this rule only reduces rather

³The UCLA algorithm obtains a clique of all Tier-1 ASes from external sources and infer as p2c the links following these ASes along the AS paths. The remaining links are inferred as p2p.

than eliminates the possibility of incorrect p2p inference. The growing number of region-specific p2c relationships observable only below the provider AS [15] makes the technique even less reliable. We conclude that we need a new strategy to express this inherent uncertainty in the inference results.

3.3 Algorithmic Paradigms

Broadly speaking, prior algorithms follow one of the two design paradigms: greedy [9, 15, 20, 25, 28] and optimization-based [7, 8]. An algorithm's design paradigm is a deciding factor in its ability to identify the most probable relationship inference from multiple valid candidates.

Using a collection of rules and heuristics, a greedy algorithm typically computes relationship inference over AS links in a sequential order. For a given link, its inference result might depend on another link's result computed in a previous step. The greedy approach to AS relationship inference has a critical limitation: it does not support backtracking and branching. When making inference on an AS link, the algorithm must commit to one of the feasible relationships as governed by the rules and heuristics. If a mistake is made, a greedy algorithm cannot recover from it and the error can impact the inference results for the remaining links. Returning to the example above, we suppose that the algorithm has already inferred $v_1 - v_2$ as p2c and $v_2 - v_3$ as p2p, before it encounters the path $v_1 - v_2 - v_3$. Clearly the path now violates the valley-free constraint. When the p2p inference was made on $v_2 - v_3$, however, it was possible that both p2p and p2c were feasible according to the valley-free constraint. Then a heuristic was used to break the tie and gave precedence to p2p over p2c. The violation could have been avoided if the algorithm had chosen p2c instead. Unfortunately, greedy algorithms lack the capability to maintain the uncertainty of relationship inference across steps.

In contrast to the greedy approach, an optimization-based algorithm explores all feasible inferences simultaneously for all links. It usually encodes the heuristics and rules as constraints and optimize for a certain objective. For example, DPP [7] uses boolean variables to represent the orientation of links (i.e., p2c v.s. c2p) and a collection of 2-literal disjunctive clauses to express the valley-free constraint. The objective is set to minimize the number of unsatisfiable clauses, or equivalently, the number of invalid paths. But [8] identifies the discrepancy that maximizing the number of valley-free paths does not necessarily translate to higher inference accuracy. In general, the objective of optimization-based algorithms is not fully aligned with the ideal objective to maximize the accuracy because the true relationships are not known and it is difficult to quantify the difference between the two objectives. It is this discrepancy in the two objectives that affects our level of confidence in the inference output and leads to uncertainty.

The main takeaway is that a new paradigm should be able to overcome the limitations of greedy and optimization-based algorithms. We consider ProbLink as a hybrid of both paradigms; it starts with an initial classification of links produced by a greedy algorithm and then uses a probabilistic inference model to continually update the inferred type of links until convergence. The second phase allows it to mitigate the limitations of the greedy algorithm. As we will show in §4, UNARI takes a different approach. It uses dependable

rules to make deterministic inference followed by unreliable rules to make uncertainty-aware inference.

3.4 Unreliable Assumptions and Rules

In addition to the valley-free constraint, inference algorithms must rely on additional assumptions to derive rules and heuristics. These assumptions vary in terms of reliability but prior algorithms rely on simple prioritization to differentiate between them in terms of confidence level when making inferences.

For example, [20] and [15] assume there exists a clique of Tier-1 providers interconnected through p2p links. It is easy to obtain a list of Tier-1 providers and the rule derived from the assumption achieves near-perfect inference accuracy when validated using a partial ground truth dataset [15]. However, these algorithms also use assumptions that depend on a predefined parameter or "magic constant". For example, Gao's algorithm compares the ratio of node degrees against a constant to determine the provider and customer roles for adjacent ASes. Links that are categorized using this parameter are not marked any differently than those using near-perfect rules. As another example, [15] assumes that vantage points reporting routes to less than 2.5% of ASes are not announcing provider routes to the route collector. It is unclear why the 2.5% threshold is appropriate and how a different value can affect the accuracy. We believe unreliable assumptions can introduce uncertainty into the inference process and the labels generated by the corresponding rules should convey the uncertainty of the results.

3.5 Potential Benefits of Exposing Inference Uncertainty

The results of AS relationship inference have supported numerous research works [4, 12, 14, 16, 24]. Exposing the uncertainty in inference opens up the opportunity to improve the robustness of their findings. We use two prior works to illustrate the potential benefits. [12] used AS relationships to evaluate the impact of peer-assisted video-on-demand (VoD) on ISPs. The authors showed that ISP-friendly peer-assisted VoD achieved 50% savings compared to solutions with no P2P. They also stated that the inference by CAIDA had been conservative because ISPs were unwilling to share their sibling and peering relationships. It is possible that quantifying the uncertainty in inference helps develop a bound rather than a single value on the estimated savings. As another example, [16] relied on AS relationships to understand BGP misconfigurations. The export policies, derived from the relationships, enabled the authors to identify export misconfigurations. The authors acknowledged that incorrect inference could lead to failure in identifying misconfigurations or mistake legitimate configurations as erroneous. Based on the uncertainty in inference, it might be possible to estimate the uncertainty in identifying misconfigurations.

4 TOWARDS UNCERTAINTY-AWARE AS RELATIONSHIP INFERENCE

As we demonstrate in Section §3, uncertainty is inherent to AS relationship inference but largely ignored by prior algorithms. To bridge this gap, we propose a new framework that explicitly reflects uncertainty in the inference results. An exemplary algorithm is presented in §5.

We start by deriving the design requirements based on the insights from §3. First, when it comes to single-label (deterministic) inference, the algorithm should not rely on heuristics whose output may vary depending on the amount of information available, such as heuristics based on node or transit degrees (§3.1). This allows us to minimize the uncertainty caused by an incomplete view of the AS topology. Second, it should expose uncertainty by associating more than one type of relationship with an AS link if necessary⁴ (§3.2). Third, it must differentiate unreliable assumptions and heuristics from the more trustworthy rules, such as the valley-free constraint (§3.4). Fourth, it should mitigate the limitations of greedy and optimization-based paradigms (§3.3).

Our framework consists of two major steps. After extracting a set of AS paths from the dataset, we identify potentially erroneous AS paths and remove or sanitize them. The remaining AS paths, collectively forming the AS topology, are then passed to the principle-based inference step where we attempt to apply *principles*, such as the valley-free constraint and other reliable rules, to deterministically infer one of the three possible relationships for each link in the topology (Step ①). If the principles manage to eliminate all but one relationship, the inference is complete for this link; otherwise, we leave it to the next step. Note that Step ① is intentionally conservative and incorporates rules that commit to an inferred relationship only if the confidence level is very high. The final step (Step ②) uses potentially less reliable assumptions and heuristics to perform uncertainty-aware inference and compute a *tag*. The tag can be numerical or categorical and represents our confidence in the inference for each remaining link. The separation of principle-based and uncertainty-aware inference allows us to expose the level of uncertainty associated with each link. Since misclassifications made in Step ① can impact the tags calculated in Step ②, the conservative selection of rules in Step ① prioritizes accuracy over coverage and thus minimizes such mistakes. We defer describing how our framework also addresses limitations of the optimization-based approach to the concrete implementation in §5 where we exploit a partial ground truth dataset and formulate the objective to maximize inference accuracy directly.

4.1 A Flexible Framework

The proposed framework does not prescribe the specific rules to be used in Step ① and ②. We demonstrate the framework's flexibility by illustrating the very diverse choices that can be made in the two steps.

Principle selection for Step ①. To attain high certainty, Step ① could for example rely on the valley-free constraint. It can also incorporate external knowledge such as the set of all Tier-1 ASes.

Heuristic selection for Step ②. Step ② is the major enabler of uncertainty-aware inference, allowing not only a wide variety of heuristics but also different tag formats to express the confidence in the inference output.

First, the output of Step ② can be a categorical tag indicating “high” or “low” confidence. For example, building on prior work,

⁴The Gao's algorithm may infer multiple relationships on one link but it seems to be an artifact of not enforcing mutual exclusion in the algorithm design. In fact, the total number of inferences exceed the number of links in the topology in the evaluation results [9].

recall that in Figure 3, B and C can be inferred as any one of the three relationships without violating the valley-free constraint. Now suppose we further observe that link $B - C$ is not seen in any AS path announced to B or C 's providers. It is an indication that $B - C$ is more likely to be p2p. However, the heuristic exploits a “negative observation”; it is grounded in the lack of paths that the vantage points expect to observe and we have seen impact of incomplete coverage in §3.1. Thus, labeled links can receive a “high” or “low” confidence tag depending on whether they were classified based on the valley free constraint or based on negative observation heuristic. Optionally, we can also document the reason for the low confidence tag, $B - C$ would be labeled as (p2p, “low”, “negative observation”). We note that [15, 20] explore similar “negative observations”.

As another example, the tag could be numeric, e.g., a scalar. Recall from §3.4 that Gao's algorithm compares the ratio of node degrees against a constant to determine the provider and customer roles for adjacent ASes. The intuition is that the extent to which the ratio deviates from 1 is an indication of how likely that one of the ASes is transiting traffic for the other. Thus, we may simply use this ratio as the numerical tag along with the c2p/p2c label. As for links inferred as p2p, we may set the tag value to be inversely proportional to the difference in transit/node degree of the end points.

Figure 4: An example to implement the uncertainty-aware AS relationship inference framework.

We can also introduce richer tags and new heuristics. Consider a heuristic that outputs as a tag a *vector of scores*, where each score represents our confidence in inferring the link as one of the three relationships. For example, it may assign equal score values to all feasible relationships according to the valley-free constraint. We provide an example in Figure 4 where the AS topology is formed by two paths. Suppose that $X - F$ and $X - E$ are inferred as p2p and p2c respectively in Step ①. At this point, the other 2 links have two or more feasible relationships under the chosen principles. In this example, we use the aforementioned heuristic that assigns equal score values to all feasible relationships. Observe that if we infer $X - Y$ as any one of c2p/p2p/p2c, there always exists an inference of $Y - Z$ such that all paths are valley-free. Thus, the heuristic assigns $\frac{1}{3}$ to all elements in the score vector of $X - Y$.

5 UNCERTAINTY-AWARE AS RELATIONSHIP INFERENCE ALGORITHM (UNARI)

In this section, we describe the Uncertainty-aware AS Relationship Inference (UNARI) algorithm that leverages our framework from §4. §5.1 to §5.3 describe in detail the implementation of Step ① and ② of the framework.

5.1 Preprocessing AS Paths

Similar to the filtering and sanitizing step described in [15], UNARI starts by preprocessing the set of AS paths. Using the list of Tier-1 ASes obtained from [27], it first removes AS paths with loops and paths where two Tier-1 ASes are separated by one or more non-Tier-1 ASes. Both of these path types are signs of path poisoning usage. It also removes paths with unassigned ASes. For each path, if it identifies an AS used to operate IXP route servers, UNARI removes it and establishes a link between its preceding and following neighbors.

5.2 Principle-based Inference

UNARI's principle-based inference step is summarized in Algorithm 1. It adopts as principles two rules commonly used by the recent inference algorithms [15, 20]. For the first principle, UNARI leverages the observations that there exists a clique of Tier-1 ASes interconnected with p2p links at the top of the hierarchy and that the AS members of the clique are known. We can either identify the clique through preprocessing [15] or obtain the list of top-tiers from public sources [20]. The second principle imposes the constraint that the relationships assigned to links along an AS path must be valley-free.

After extracting a collection of valid AS paths from the dataset, UNARI infers the links connecting a pair of Tier-1 ASes as p2p based on the first principle. Then, UNARI infers a set of p2c links, with high confidence, following the second principle. The inference of p2c links has two phases. Phase I exploits the knowledge of Tier-1 ASes. Given an AS path $p_j : v_1 - v_2 - \dots - v_h - v_{h+1} - \dots - v_m \in Q$, UNARI searches for the last Tier-1 AS starting from v_1 . Suppose v_h is the last Tier-1 AS along the path, we can deduce that $v_h - v_{h+1}$ is of type p2p or p2c because a Tier-1 AS has no provider. Moreover, the type of every link on the path segment $v_{h+1} \rightsquigarrow v_m$ must be p2c by the valley-free principle. UNARI applies this procedure on each path and ends up with a set of p2c inferences. An idea similar to Phase I of p2c inference was proposed in [20]. However, UNARI takes it one step further and introduces a Phase II of p2c inference. It takes advantage of the outputs from Phase I and infers as many p2c links as possible by repeatedly iterating through the list of paths. Within each iteration, it applies the valley-free principle based on links that have been inferred in the previous iterations and makes p2c inference on undecided links accordingly. For example, consider another AS path $p_i : v_{h+1} - v_{h+2} - v_w \in Q$ where none of the nodes is Tier-1 AS. While Phase I cannot infer p2c links on p_i directly, one of its links, $v_{h+1} - v_{h+2}$, is inferred as p2c after p_j is processed. The first iteration of Phase II is able to apply the valley-free principle and infer $v_{h+2} - v_w$ as p2c. The succeeding iterations can now leverage the inference result on $v_{h+2} - v_w$. Phase II terminates when no new p2c links are produced in the last iteration.

Conflicts might arise from the procedure above. Specifically, a link $u - v$ might be inferred as p2c in one path but c2p in another path. The problem can be attributed to the observation of s2s links, where both end points can export their own routes and the routes of their customers, providers or peers [9]. A study [8] has shown that 31 out of 3724 verified relationships are s2s according to its survey data. UNARI labels these links as “conflict” and leaves final classification of the link as either p2c or c2p to the uncertainty-aware inference step.

Algorithm 1 Principle-based Inference

- 1: Infer p2p links between Tier-1 ASes
 - 2: Infer p2c links following Tier-1 ASes ▷ Phase I
 - 3: Gather all inference results into link set $\mathcal{J}$
 - 4: **repeat** ▷ Phase II
 - 5: Scan through the path list to infer more p2c links using $\mathcal{J}$
 - 6: Update $\mathcal{J}$ to incorporate the new p2c links
 - 7: **until** No new p2c links inferred in the last iteration
 - 8: Label conflicting p2c inferences
-

5.3 Uncertainty-aware Inference

We implemented two uncertainty-aware inference versions, following the heuristic briefly described in §4.1.

The first version generates a binary tag for each undecided link to indicate the level of confidence that it is indeed p2p. Consider an undecided link $U - V$. The heuristic searches for a set of ASes, Ω , that satisfies the following conditions. First, each $X \in \Omega$ must be an inferred provider of either U or V . Second, $X - U$ or $X - V$ must be on some path announced to a vantage point and X is closer to the vantage point. If Ω is $\emptyset$, it assigns a “high” confidence tag to the link; otherwise, a “low” confidence tag is assigned. Recall that the heuristic exploits the “negative observation”; if $U - V$ is indeed c2p/p2c, at least one of the vantage points is expected to observe the announcement through the providers of U or V .

The second version adopts a more sophisticated algorithm that generates a probability score. It starts with a procedure to compute a *share vector*, denoted by S , for every undecided link. Each element in the vector measures the number of valley-free inferences over the entire topology, under the condition that the link is inferred as one of the relationships. Then we refine the procedure with a relaxation step to handle the case where no valley-free inference can be found. Finally, UNARI outputs the score as the conditional probability that the type of relationship of a link is r given its share vector s . We choose Multinomial Logistic Regression (MLR) and leverage a partial ground truth dataset for learning the probabilistic model.

We assume there exist a correlation between the element in a link's share vector and its type of relationship. For example, over all possible valley-free inferences, if a link appears as p2p much more often than c2p/p2c, we may assume it is more likely to be p2p.

5.3.1 Computing Shares. By the valley-free principle, the inferred types of relationship of undecided links found within the same AS path are interdependent. For example, Figure 5 shows an AS topology with links $v_4 - v_5$ and $v_6 - v_7$ inferred as p2c by principle-based inference while the other links are undecided. In Path I, if

we infer any one of the undecided links as p2p, the other two links must be inferred as p2c or c2p; otherwise, it violates the valley-free principle.

Figure 5: An AS topology with undecided links and its interdependence graph ($\mathcal{G}$).

Table 1: All valley-free Inferences for the topology in Figure 5.

n_{12}	n_{23}	n_{34}	n_{36}
p2c	p2c	p2c	p2c/p2p/c2p
p2p	p2c	p2c	p2c/p2p/c2p
c2p	p2c	p2c	p2c/p2p/c2p
c2p	p2p	p2c	p2c/p2p/c2p
c2p	c2p	p2p	p2c
c2p	c2p	c2p	p2c

To capture the interdependence imposed by valley-free principle, we introduce a graph called interdependence graph ($\mathcal{G}$). Each undecided AS link corresponds to a node in $\mathcal{G}$. For each pair of nodes in the topology, we connect them with an arc iff there exists at least one AS path that contains both AS links. For example, the two paths in Figure 5 leads to four nodes $n_{12} : v_1 - v_2$, $n_{23} : v_2 - v_3$, $n_{34} : v_3 - v_4$, and $n_{36} : v_3 - v_6$. An arc is created for a pair of nodes if they appear in the same path. Note that the interdependence relationship can span across different AS paths if they share the same undecided links. $v_3 - v_4$ is the common link shared by Path I&II.

Now we associate each node n_i in $\mathcal{G}$ with three Boolean variables x_i , y_i , and z_i and use them to denote c2p, p2p, and p2c relationships respectively. The valley-free principle can be expressed as a Boolean formula over these variables. Returning to the example in Figure 5, we can use $y_{23} \Rightarrow z_{34}$ to express the following constraint: if node n_{23} is inferred as p2p, node n_{34} must be p2c. By solving for a satisfying assignment of the formula, we can identify a valid valley-free inference over the undecided links. Moreover, the valley-free principle is typically insufficient to produce a unique satisfying assignment. For instance, if we assume that Path I&II in Figure 5 are the only paths containing the four undecided links, then these nodes form a single connected component in $\mathcal{G}$ and their inference

are independent from other nodes in $\mathcal{G}$. Each satisfying assignment matches one of the valley-free inferences in Table 1.

We use C_j to label the connected components in the $\mathcal{G}$ and $F_E(C_j)$ to denote the number of unique valley-free inferences over C_j given that the predicate E evaluates to true. The subscript is dropped when no predicate is specified. Returning to Figure 5, the number of valley-free inferences is $F(C_j) = 14$ and the number of valley-free inferences given that n_{12} is inferred as c2p is $F_{x_{12}}(C_j) = 8$. We express the share associated with n_i and relationship R as a normalized count

$$S_i^R = \frac{F_{R_i}(C_j)}{F(C_j)} \quad (1)$$

where $R \in \{x, y, z\}$ and C_j is the connected component in which n_i resides. In the running example, the share associated with n_{12} and c2p is $S_{12}^x = \frac{8}{14}$. The computation of $F_E(C_j)$ can be reduced to a propositional model counting problem [2], or #SAT.

The definition above makes two assumptions. First, it enforces that the correct inference for the connected component is valley-free. It is justified because valley-free violation is uncommon [10] and the valley-free principle is fundamental. We will describe how UNARI handles the case where no valley-free inference can be found in §5.3.2. Second, it gives all possible valley-free inferences the same weight. We believe it is a reasonable assumption as it simplifies the computation of shares and leave part of the complexity of deciding the weights to the final step—computing probability estimates.

5.3.2 Relaxation for Unsatisfiable Connected Components. Some connected components might have no valley-free inference due to two reasons. First, [10] identified valley paths using BGP Community data. If C_j contains nodes representing AS links from these paths, the algorithm is unable to find valley-free inference. Second, incorrect inferences are made during the principle-based inference step and propagate to the uncertainty-aware inference step.

Inspired by the idea of maximizing the number of satisfiable clauses [7, 8], we propose a relaxation method to mitigate these problems. As mentioned in §5.3, the valley-free constraint imposed over C_j can be expressed as a Boolean formula. Our approach first leverages a MAX-SAT solver to determine an assignment maximizing the number of clauses that can be made true. For each clause that expresses an interdependence between AS links but is unsatisfied by the assignment, we remove the corresponding arc from C_j . At the end of this process, C_j is broken into one or more connected components. The same process is applied to each resulting component recursively if it remains unsatisfiable. Instead of ignoring the unsatisfied clauses, as in [7, 8], our method removes the interdependence that leads to these clauses.

Once an arc is removed from a component, the interdependence between two AS links, which must have appeared in the same AS path, is lost. The goal is to preserve as many interdependence relationships as possible. It is challenging to incorporate the existence of valley paths in UNARI as the valley-free principle is fundamental in our approach to deriving the shares.

Algorithm 2 summarizes the procedure of computing share vectors in §5.3.1 and §5.3.2. It begins by identifying a set of connected components, $\mathcal{U}$, from the interdependence graph. Then for each component it calculates the number of valley-free inferences $F(C_j)$.

If the value is 0, meaning the associated Boolean formula is unsatisfiable, it dismantles it into a set of smaller, satisfiable connected components, and merge them back to the group. Now that all connected components in the set are satisfiable, the nested for loop (line 7) proceeds to calculate the share vectors for each link in every connected component, which is parallelizable because no dependency spans across multiple components. Our evaluation on real datasets show that the while loop (line 2) terminates within two iterations and only a small percentage of links are removed due to splitting the connected components.

Algorithm 2 Computing Share Vectors

```

1: Identify a set of connected components  $\mathcal{U} = \{C_j\}$  from the
   interdependence graph  $\mathcal{G}$ 
2: while  $\exists C_j \in \mathcal{U}$  s.t.  $F(C_j) = 0$  do
3:   Remove  $C_j$  from  $\mathcal{U}$ 
4:   Compute a set of satisfiable connected components  $\{C'_j\}$ 
     from  $C_j$  using relaxation.
5:   Merge  $\{C'_j\}$  back to  $\mathcal{U}$ 
6: end while
7: for all  $C_j \in \mathcal{U}$  do
8:   for all  $n_i \in C_j$  do
9:     Compute  $F_{R_i}(C_j)$  and  $S_i^R$  where  $R \in \{x, y, z\}$ 
10:   end for
11: end for
  
```

5.3.3 Computing Probability Estimates. Finally, we apply MLR to model the conditional probability $Pr(R = r|S = s)$ using the elements of share vector as features. We decide to use MLR because it outputs not only the predicted type of relationship, but also the probability estimates. The ground truth dataset is partitioned into the training and test sets. MLR assumes that $Pr(R = r|S = s) = f(s; \theta)$ for some function f parameterized by θ . Intuitively, it aims to estimate θ^* , the value of θ that is most consistent with the true relationship from the training set. The test set will be used to evaluate the model's performance. Now consider an unlabeled link beyond the partial ground truth dataset. Based on its score vector s' , the trained model outputs $f(s'; \theta^*)$ to estimate the probability that its type of relationship is r .

Importantly, MLR assumes that the relationship of undecided links are independent; but this is not true because the valley-free constraint induces interdependence across them. Markov Logic Network [23] (MLN) allows us to explicitly express the valley-free constraint using first-order logic. But our experience with an MLN inference engine (i.e., Tuffy [18]), indicates that it faces scalability issues when applied to the AS-level Internet topology. Therefore, we consider the share vectors as a reasonable compromise and leaves the search of a more scalable MLN-based solution as future work.

6 UNARI EVALUATION

In this section, we begin by showing the results of the preprocessing step and introduce the set of tools used for the implementation. Then, we present the evaluation results for Step ① and ②. The details of the BGP data and the partial ground truth data are described in §2.

We rely on public BGP datasets in this evaluation. UNARI's preprocessing step extracts 7,153,208 out of 13,482,724 unique AS paths from these datasets. The result is a topology of 41,127 ASes and 130,076 links⁵.

In UNARI's uncertainty-aware inference, we use sharpSAT [26] for counting the number of valid assignments over each connected component and Open-WBO [17] to determine the maximum number of satisfiable clauses for unsatisfiable components.

6.1 Principle-based Inference Results

Figure 6: The upper plot shows the number of p2c (c2p) links inferred by UCLA, ASRANK, UNARI, and UNARI* (a variant of UNARI), when $|\bar{H}| = 7$. UNARI infers more p2c (c2p) links than UCLA but fewer than ASRANK. The lower plot compares the number of transitions as $|\bar{H}|$ drops from 79 to 7 and illustrates the principle-based inference step's strong stability.

We evaluate the principle-based inference step in terms of accuracy and stability. First, we show that UNARI achieves higher accuracy than prior algorithms when it infers a single relationship on an AS link. We then turn to stability evaluation and analyze how the principle-based inference step helps address the issue in §3.1. When evaluating the accuracy and stability of prior algorithms, the same public BGP datasets are used as input to our implementation of these algorithms. We focus our comparison on UNARI with UCLA and ASRANK, and leave the comparison with ProbLink as future work.

6.1.1 Accuracy. A set of 22 Tier-1 ASes is retrieved from [27]⁶. UNARI infers 151 p2p links across them. 131 of these links are found in the ground truth dataset and all labeled as p2p. One Tier-1 p2p link in the ground truth is not present in any of the paths in

⁵The topology after preprocessing has more links than the topology formed by all route collectors from I in §2 because UNARI introduces new links between the peering participants at the XPs when the ASes that are serving as route servers are removed from AS paths.

⁶For consistency, we verify that they match the Tier-1 ASes identified in [15], the source of the ground truth dataset.

the input, so it cannot be inferred by UNARI. Therefore, UNARI correctly infers all validated p2p links in the clique of Tier-1 ASes.

The principle-based inference step infers 78,689 c2p/p2c links and 383 s2s links, among which 26,637 links are found in the ground truth dataset. Its confusion matrix is shown in Table 2. Observe that UNARI correctly infers 26,339 out of 26,389 validated p2c links but it misclassifies 26 p2c links as c2p. 64 p2c links are inferred as both p2c and c2p. As described in §5.2, they represent a very small group of links experiencing conflicting inference results; each of them appears in multiple AS paths and applying the principles does not lead to the same relationship across all paths. The principle-based inference step does not infer p2p links, but it misclassifies a total of 208 p2p links as p2c, c2p, or conflict.

Table 2: Confusion matrix of principle-based inference step. Conflict links are inferred as both c2p and p2c.

True relationship	Inferred relationship			Total
	p2c	c2p	conflict	
p2c	26339	26	64	26429
p2p	110	93	5	208

The accuracy of the principle-based inference step is 98.9%. Over the same set of validated links, ASRANK and UCLA achieve accuracies of 96.9% and 96.3%, respectively. Note that they are consistent with the overall performance of ASRANK and UCLA when applied to all validated links in the topology (97% and 95.5%). It means the subset of links are not deliberately chosen to favor UNARI. To conclude, the principle-based step prioritizes reliability over coverage, achieving higher accuracy than ASRANK and UCLA on a subset of links.

6.1.2 Stability. The principle-based inference step not only attains higher accuracy than UCLA and ASRANK, but also features greater stability when faced with shrinking visibility over the AS topology. The upper plot of Figure 6 compares the number of p2c inferences made by UCLA, ASRANK and UNARI when $|\bar{H}|$ (the set of active route collectors) is 7. UNARI infers more p2c (c2p) links than UCLA but fewer links than ASRANK. Despite their similarity, UNARI is able to infer more p2c links than UCLA. Recall that the principle-based inference step, as described in §5.2, starts by inferring p2c links that follow the p2p links connecting two Tier-1 ASes. Once it gathers a set of p2c links, it repeatedly scans the AS paths to expand the p2c link set based on inferences made in the prior iterations. It is this iterative procedure that allows UNARI to infer more p2c links. At the same time, its accuracy is not compromised for two reasons. First, UNARI follows a more elaborate approach to filtering and sanitizing AS paths, which should generally reduce the number of erroneous paths and improve inference accuracy. Second, UNARI avoids committing to conflicting inference results that can lead to misclassification. If a link is inferred as p2c in one path and c2p in the other, UNARI leaves it undecided while UCLA chooses either p2c or c2p. We believe ASRANK infers more p2c links than UNARI because it incorporates a larger pool of rules, some of which are less reliable than the ones chosen by UNARI.

The lower plot of Figure 6 depicts the number of p2c_c2p transitions as $|\bar{H}|$ drops from 79 to 7, following the procedure presented in Figure 2. We focus on p2c_c2p transitions because the principle-based inference step only infers p2p relationship for links connecting Tier-1 ASes. Observe that UNARI experiences many fewer transitions in comparison to UCLA and ASRANK, exhibiting stronger stability. We also evaluate UNARI*, a variant of UNARI that omits the iterative p2c inference process. Surprisingly, the transitions can be further reduced to 0 at the expense of a minor drop in the number of p2c inferences. UNARI's advantage can be attributed to the decision to always categorize links with conflicting inferences as anomalies (i.e., links inferred as different types of relationship on multiple AS paths) instead of committing to one of them.

6.2 Uncertainty-aware Inference Results

Experiments in this section only involve links that remain unclassified after UNARI's principle-based inference step. We evaluate the two versions of uncertainty-aware inference separately.

For the first version that generates binary tags, we evaluate it against a baseline implementation that assigns p2p to all undecided links. Note that the baseline essentially follows the UCLA algorithm, which assigns p2p relationship to all remaining AS links, after it completes the inference of all c2p/p2c links. The heuristic above assigns "high" confidence tag to 99.6% of all p2p links in the validation dataset and achieves a higher precision than the baseline at (93.4% vs. 92.5%).

For the second version that generates probability scores, we identify the subset of these links that are labeled by the ground truth dataset and further divide this subset into training (70%) and test (30%) sets. For training, we use the MLR implementation (LogisticRegressionCV) from scikit-learn [22], with cross-validation and L2 regularization to lessen the amount of overfitting. We report the performance of uncertainty-aware inference by evaluating the probability estimates computed by the trained model over the test set. Unless stated otherwise, the training and test error rates, measured by cross-entropy loss, are within 5% difference.

In order to compare the outputs of uncertainty-aware inference and the deterministic algorithms, we convert every probability distribution into a single-relationship assignment. Specifically, we designate the type of relationship with the maximum probability as the single-relationship output of uncertainty-aware inference. For example, consider the distribution $(0.06, 0.04, 0.9)^T$ where the elements are interpreted as the probability that the relationship is c2p, p2c, and p2p. The single-relationship output is p2p. It follows that conventional metrics such as accuracy becomes applicable. To evaluate an algorithm's performance, we measure the cross-entropy between the true and inferred probability distributions. The true distribution is obtained by applying 1-of- K encoding [3] to the single-relationship assignment. For example, if a link is assigned p2p, we can associate it with the distribution $(0.0, 0.0, 1.0)^T$.

Below, we examine UNARI's uncertainty-aware inference in terms of its accuracy of single-relationship assignment, how well its final output probabilities reflect confidence and whether its share vectors are better than other heuristics for estimation of probabilities.

Figure 7: Correlation between the relationship with the maximum probability estimate (Max) and the true relationship ($True$) for UNARI.

Figure 8: Average cross-entropy between the true and inferred probability distributions for 4 heuristics.

6.2.1 Single-relationship Assignment Accuracy. For single-relationship assignment, UNARI’s uncertainty-aware inference is comparable to the state-of-the-art algorithms. UNARI achieves an accuracy of 93%, which is slightly lower than UCLA (94%) and ASRANK (96%) but higher than Gao’s (65%). Note that we present the accuracy of single-relationship assignment for completeness. The two key metrics to evaluate uncertainty-aware inference are presented next.

6.2.2 Probability Estimate Accuracy. One of UNARI’s advantages over prior algorithms is its ability to effectively express a varying level of confidence in the inference output that is consistent with the ground-truth data. We illustrate it through the correlation between the relationship with the maximum probability estimate (Max) and

the true relationship ($True$), of links in the test set. The upper plot of Figure 7a shows the distribution of links with $Max = p2p$, binned by the maximum value of probability distribution. The majority of links fall within the range of 0.9 and 1.0.⁷ Meanwhile, we observe that a significantly higher number of matches ($Max = p2p; True = p2p$) than mismatches ($Max = p2p; True = p2c$) in the lower plot for the same range, indicating low uncertainty when the probability of p2p is high. In particular, in the rightmost bar (around 1.0), the number of matches at that point (lime bar) is nearly two orders of magnitude higher than that of mismatches (orange bar). As the maximum probability value decreases from 0.9 to towards 0.5, the percentage of mismatches increases but the number of links also decreases. We conclude that for the majority of links with $Max = p2p$, the maximum value of probability distribution effectively expresses the algorithm’s level of confidence in the inference output. As for links with $Max = p2c$ or $c2p$, we recognize a similar pattern within the range of 0.73 and 0.75 in Figure 7b but the number of matches is less dominant.

6.2.3 Efficacy of Share Vectors. The share vectors are not the only heuristic that can be used to compute probability estimates. UNARI’s key strength is that it opens the exploration of such algorithm designs. To illustrate how well share vectors work, we consider three alternative for comparison: Naïve, Degrees and Hybrid. For any given link, the Naïve heuristic always outputs the relative frequency distribution of the true labels. For example, the probability of p2p is set to 0.92 if 92% of the links from the input have p2p as the true label. The Degrees heuristic is a variant of UNARI’s uncertainty-aware inference step that replaces the link’s share vector with the node/transit degrees of its endpoints as the link’s feature. The last heuristic, Hybrid, includes both share vectors and node/transit degrees as features.

⁷Note the log scale on y-axis.

Figure 8 measures the average cross-entropy between the true and inferred probability distributions for the above four heuristics, over 5 random splits of training and test sets. A lower value implies smaller difference between the two distributions. We use the Naïve heuristic as a baseline for comparison. Note that all variants of UNARI achieve lower average cross-entropy than Naïve. Degrees outperforms UNARI with share vectors, showing that node/transit degrees is a useful feature. Hybrid combines share vectors and node/transit degrees, two complementary features, and it outperforms using either alone. We believe Hybrid’s advantage can be attributed to share vectors’ ability to encapsulate the interdependence across neighbor links as imposed by the valley-free constraint, which cannot be easily captured by the local properties such as node/transit degrees.

7 CONCLUSION

In this paper, we identify several causes of uncertainty in AS relationship inference and demonstrate why the problem deserves more attention by the research community. We propose a novel framework to inferring AS relationships that explicitly models the uncertainty and reflect it in inference output. We propose an exemplary algorithm (UNARI) that implements this design paradigm and leverages a novel technique to capture the interdependence between undecided links. Our evaluation indicates it achieves higher accuracy and stability compared to prior algorithms for deterministic inference and outperforms a naïve strategy to express uncertainty.

ACKNOWLEDGMENTS

We would like to thank our shepherd, Jelena Mirkovic, and the anonymous reviewers for their valuable feedbacks. This research was funded in part by NSF under award number CNS-1565343 (https://www.nsf.gov/awardsearch/showAward?AWD_ID=1565343).

A ROUTE COLLECTOR COVERAGE

Following our analysis in §3.1 we provide evidence to illustrate the variation in route collectors’ coverage.

We start by profiling the coverage of the AS topology by route collectors. The histogram in Figure 9a groups the AS links by the number of collectors that are able to observe these links. We note that a significant portion of the links are seen by no more than 2 collectors. They are likely connected to stub ASes near the edge of the AS topology. The spike at 23 can be attributed to the high variance in the number of AS links observed by individual collectors. In particular, 53 collectors observe fewer than 30,000 links while 22 collectors observe over 50,000 links. A similar distribution is found in Figure 9b which captures the collectors’ coverage of vantage points and a minor spike takes place at 21.

The variation in route collectors’ coverage has a key implication. While individual collectors might cover some links or vantage points in common, each of them still contributes complementary information to form a more complete view of the AS topology. Moreover, topological information is lost if a route collector is not deployed or becomes unavailable because the majority of links and vantage points are only seen by a single collector.

(a) Coverage of AS links. 32,322 of them are observable by only one route collector. (b) Coverage of vantage points. 1,511 of them are observable by only one route collector.

Figure 9: Route collectors’ coverage of the AS topology.

Figure 10: Relationship transitions as a function of $\bar{H}$ for Gao’s.

B MORE STABILITY RESULTS

As mentioned in §3.1.2, we now show the stability results of Gao’s, MVP, and ASRANK. We observe a surge in the ‘p2c_p2p’ curve of ASRANK (12), which is similar to Figure 2c but with a smaller magnitude. Its ‘p2p_p2c’ curve experiences a gradual increase as $|\bar{H}|$ shrinks from 79 to 23 and witnesses a minor drop as the collector count decreases from 23 to 15. We also notice a small number of transitions from p2c to c2p. In contrast, the ‘p2c_p2p’ curve of Gao’s (10) follows the shrinking size of $\bar{H}$ and trends downward. What is surprising, however, is the large number of transitions at $|\bar{H}| = 71$, caused by a minor drop in the number of route collectors. MVP (11) seems more stable than the rest as measured by the number of transitions. The results support our speculation that these algorithms are sensitive to the number of route collectors available and the level of sensitivity may vary depending on the algorithm.

Figure 11: Relationship transitions as a function of $\bar{H}$ for MVP.

Figure 12: Relationship transitions as a function of $\bar{H}$ for ASRANK.

REFERENCES

- [1] Cengiz Alaettinoglu, Curtis Villamizar, Elise Gerich, David Kessens, David Meyer, Tony Bates, Daniel Karrenberg, and Marten Terpstra. 1999. *Routing policy specification language (RPSL)*. Technical Report.
- [2] Armin Biere, Marijn Heule, and Hans van Maaren. 2009. *Handbook of satisfiability*. Vol. 185. IOS press.
- [3] Christopher M. Bishop. 2006. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg.
- [4] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. 2013. Mapping the expansion of Google's serving infrastructure. In *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 313–326.
- [5] RIPE Network Coordination Center. 2017. Routing Information Service Raw Data. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data>. [Online; accessed 12-October-2017].
- [6] R Chandra, P Traina, and T Li. 1996. *BGP communities attribute*. Technical Report.
- [7] Giuseppe Di Battista, Maurizio Patrignani, and Maurizio Pizzonia. 2003. Computing the types of the relationships between autonomous systems. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies*. IEEE, 156–165.
- [8] Xenofontas Dimitropoulos, Dmitri Krioukov, Marina Fomenkov, Bradley Huffaker, Young Hyun, George Riley, et al. 2007. AS relationships: Inference and validation. *ACM SIGCOMM Computer Communication Review* 37, 1 (2007), 29–40.
- [9] Lixin Gao. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking (ToN)* 9, 6 (2001), 733–745.
- [10] Vasileios Giotsas and Shi Zhou. 2012. Valley-free violation in Internet routing—Analysis based on BGP Community data. In *Communications (ICC), 2012 IEEE International Conference on*. IEEE, 1193–1197.
- [11] Packet Clearing House. 2017. Packet Clearing House Daily Routing Snapshots. https://www.pch.net/resources/Routing_Data/. [Online; accessed 12-October-2017].
- [12] Cheng Huang, Jin Li, and Keith W Ross. 2007. Can internet video-on-demand be profitable? *ACM SIGCOMM Computer Communication Review* 37, 4 (2007), 133–144.
- [13] Yuchen Jin, Colin Scott, Amogh Dhamdhere, Vasileios Giotsas, Arvind Krishnamurthy, and Scott Shenker. 2019. Stable and Practical AS Relationship Inference with Problink. In *Proceedings of the 16th USENIX Conference on Networked Systems Design and Implementation (NSDI'19)*. USENIX Association, Berkeley, CA, USA, 581–597. <http://dl.acm.org/citation.cfm?id=3323234.3323282>
- [14] Matthew Luckie, Amogh Dhamdhere, David Clark, Bradley Huffaker, et al. 2014. Challenges in inferring internet interdomain congestion. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. ACM, 15–22.
- [15] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, et al. 2013. AS relationships, customer cones, and validation. In *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 243–256.
- [16] Ratul Mahajan, David Wetherall, and Tom Anderson. 2002. Understanding BGP misconfiguration. In *ACM SIGCOMM Computer Communication Review*, Vol. 32. ACM, 3–16.
- [17] Ruben Martins, Vasco Manquinho, and Inês Lynce. 2014. Open-WBO: A modular MaxSAT solver. In *International Conference on Theory and Applications of Satisfiability Testing*. Springer, 438–445.
- [18] Feng Niu, Christopher Ré, AnHai Doan, and Jude Shavlik. 2011. Tuffy: Scaling up statistical inference in markov logic networks using an rdbms. *Proceedings of the VLDB Endowment* 4, 6 (2011), 373–384.
- [19] University of Oregon. 2017. University of Oregon Route Views Project. <http://www.routeviews.org/>. [Online; accessed 12-October-2017].
- [20] Ricardo Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. 2010. The (in) completeness of the observed internet AS-level structure. *IEEE/ACM Transactions on Networking (ToN)* 18, 1 (2010), 109–122.
- [21] Ricardo V Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. 2008. In search of the elusive ground truth: the internet's as-level connectivity structure. In *ACM SIGMETRICS Performance Evaluation Review*, Vol. 36. ACM, 217–228.
- [22] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [23] Matthew Richardson and Pedro Domingos. 2006. Markov logic networks. *Machine learning* 62, 1-2 (2006), 107–136.
- [24] Devavrat Shah and Tauhid Zaman. 2011. Rumors in a network: Who's the culprit? *IEEE Transactions on information theory* 57, 8 (2011), 5163–5181.
- [25] Lakshminarayanan Subramanian, Sharad Agarwal, Jennifer Rexford, and Randy H Katz. 2002. Characterizing the Internet hierarchy from multiple vantage points. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, Vol. 2. IEEE, 618–627.
- [26] Marc THURLEY. 2006. sharpSAT: Counting models with advanced component caching and implicit BCP. *Lecture notes in computer science* (2006), 424–429.
- [27] tier-one 2017. Tier 1 network. https://en.wikipedia.org/wiki/Tier_1_network. [Online; accessed 12-October-2017].
- [28] Jianhong Xia and Lixin Gao. 2004. On the evaluation of AS relationship inferences [Internet reachability/traffic flow applications]. In *Global Telecommunications Conference, 2004. GLOBECOM'04. IEEE*, Vol. 3. IEEE, 1373–1377.