

PASSPORT: Enabling Accurate Country-Level Router Geolocation using Inaccurate Sources

Muzammil Abdul
Rehman
Northeastern University
muzammil@ccs.neu.edu

Sharon Goldberg
Boston University
goldbe@cs.bu.edu

David Choffnes
Northeastern University
choffnes@ccs.neu.edu

Abstract

When does Internet traffic cross international borders? This question has major geopolitical, legal and social implications and is surprisingly difficult to answer. A critical stumbling block is a dearth of tools that accurately map *routers* traversed by Internet traffic to the *countries* in which they are located. This paper presents PASSPORT: a new approach for efficient, accurate country-level router geolocation and a system that implements it. PASSPORT provides location predictions with limited active measurements, using machine learning to combine information from IP geolocation databases, router hostnames, whois records, and ping measurements. We show that PASSPORT substantially outperforms existing techniques, and identify cases where paths traverse countries with implications for security, privacy, and performance.

1. INTRODUCTION

Determining when Internet traffic crosses international borders is of significant interest to both lawmakers and the general public. A nation’s laws usually apply even to foreign traffic that transits that nation. The US, UK, and other nations aggressively surveil traffic that crosses their national borders [46, 39, 18]. Foreign traffic that traverses countries with aggressive censorship policies can be filtered before it arrives at its destination [37]. Countries have debated (*e.g.*, Brazil [22]), or enacted (*e.g.*, China [4], Russia [52]) “data residency” laws that require their citizens’ data to remain on domestic soil. In Europe, the Data Protection Directive [15, 17, 16] forbids the movement of citizens’ data to countries that do not provide “adequate” data protection. In the US, the laws relating to the government’s ability to search and seize their own citizens’ data depend on whether that data was intercepted on US soil or abroad [8].

Despite increasing interest from the public, networking researchers still lack the tools needed to accurately determine where Internet traffic crosses international borders. While traceroute trivially reveals the IP addresses of the routers traversed by network traffic, geolocating these *router* IPs to *countries* remains a key open problem. There are a number of commercial databases [27, 41, 9, 2, 26, 21, 1] and research projects [53, 45, 54, 30, 35]

that accurately geolocate IP addresses belonging to *eye-ball* networks, Internet exchange points (IXPs), and large enterprise datacenters. However, the accuracy of these sources on *routers* is poor [20, 48]. We find that when mapping routers to countries, even the best geolocation databases achieve more than 90% accuracy for only 46.5% of the countries traversed by paths in our traceroute dataset. Worse, naïvely mapping router IPs to countries using BGP paths and AS-level geolocation, has this accuracy for only 5% of countries. Research projects that rely on active measurements can require large numbers of probes to be accurate, and their online deployments either fail to resolve many router locations [35, 54], or is not unavailable.

To address this problem, we present PASSPORT, a system that accurately maps router IPs to countries. PASSPORT identifies the set of countries traversed by a traceroute path by combining small numbers of round-trip time (RTT) latencies with a machine learning classifier that uses IP geolocation databases. PASSPORT first trains a machine-learning classifier on a set of (individually inaccurate) geolocation sources, and uses the classifier to map the IP in a path to a set of countries. These results are filtered through speed-of-light (SoL) constraints imposed by RTT latencies. This enables PASSPORT to filter out incorrect predictions and maps each IP to a single country 88.1% of the time.

Our main contributions are as follows.

- We build a dataset of 11,626 ground-truth router IP geolocations from 1,244 autonomous systems (ASes) in 120 countries, and use them to evaluate the accuracy of existing approaches for router geolocation. These locations come from well-established data sources such as IXP locations, crowdsourced labels from operators, and reverse DNS entries, and they are all cross-validated.
- We design and evaluate an empirically informed approach to aggregating a suite of machine-learning classifiers for predicting a router’s country, using our ground truth labels for training. More specifically, we demonstrate that it is possible to achieve high country-level geolocation accuracy from individually inaccurate data sources by combining independent classifiers trained on different subsets of a possibly-biased dataset.

- We build and evaluate techniques that incorporate active measurements and iterative learning that improve classifier precision and accuracy.
- We use our system to analyze the geopolitical properties of targeted intra-country and international paths, identifying interesting cases and their implications. For example, BRICS (Brazil, Russia, India, China, and South Africa) nations, among others, often wish to avoid transiting traffic through the US, where it can be subject to surveillance. However, we find that paths between Brazil and Russia, as well as those between China and India, transit the US. Russia, which reportedly meddled in recent European elections [10, 11], transits traffic for several paths between European countries. Even “purely domestic” traffic detours outside of country where sources and destinations are located, including a case where domestic Philippines traffic transits Hong Kong.
- In addition to providing an online tool to provide country-level IP geolocation, we will make all of our code and data publicly available.

2. RELATED WORK AND MOTIVATION

This section discusses prior IP geolocation work. We exclude the use of GPS, which is now commonly accessible when measuring from mobile devices, because such information is generally *not* available for router IPs.

Constraint-based schemes. These schemes [53, 45, 54, 30] primarily use speed of light (SoL) constraints based on RTT measurements from given landmark locations to identify regions where a router can feasibly be located. Each approach uses some source of ground truth to further tighten the constraints on where a host may be geolocated. Eriksson *et al.* [14] use RTTs and a Naive Bayesian classifier to identify router cities and counties (not *country*), but use an inaccurate geolocation database [41] to validate their results. Posit [13] uses statistical analysis on latency measurements and landmarks to identify possible locations.

Hostname Parsing schemes. The hostnames associated with router IPs often encode location information, *e.g.*, `ae-4-90.edge5.frankfurt1.level3.net` encodes the city of Frankfurt, Germany. Previous work leverages this to provide mappings from router hostnames to their locations via simple matching [19] and machine learning [25]. These approaches generally provide accurate IP geolocation at the regional level, but can be inaccurate when there is location ambiguity (*e.g.*, `san` can indicate the airport code for San Diego, California, or the start of several cities starting with “San” such as San Juan). Further, this technique only works if location-encoded hostname information is available for the IP address. The IXMap project [29] uses a combination of hostname parsing and next/previous hop RTT latency to assign router locations.

Geolocation Databases. Several free and paid services offer databases that map IP addresses to locations [27, 41, 9, 2, 26, 21]. However, recent work indicates that they have limited accuracy for geolocating infrastructure IP addresses [20]. In part to address this, OpenIPMap [44] maintains a crowdsourced list of router IP geolocations, but is limited to covering IPs provided by contributors. Alidade [7] uses a collection of databases and measurements to identify router locations. In contrast, we develop an adaptive strategy for incorporating unreliable information to provide reliable predictions of the *country* where a router IP is located.

Geopolitical routing implications. When Internet traffic traverses national borders, it may be subject to surveillance [49, 8, 18, 39, 46] and censorship regimes [37]. As a result, the privacy and integrity of users’ Internet traffic depends not only on endpoint locations, but also on the location of intermediate hops. In addition, large geographical detours can turn into path inflation that substantially impacts end-to-end performance [33].

Summary. Previous approaches rely on independent, complementary approaches to geolocating IP addresses, none of which alone has sufficient accuracy to reliably geolocate the country where a router IP resides. As a result, researchers are currently in the dark regarding the important question of which countries Internet traffic traverses. In our work, we leverage the observation that each of these approaches has different strengths and weaknesses, so there is an opportunity to combine them to provide greater accuracy as a whole than any individual part. The next section describes how we leverage this observation by using machine learning principles to reliably predict router-IP countries.

3. GOALS AND DESIGN

3.1 Goals

Our primary goal is to build a system that can *accurately identify the countries of Internet routers that respond to active measurements probes (e.g., pings)*. Our ancillary goals are to provide an online system that makes predictions quickly so that it can map both transient and long-lived paths, and to ensure that the system quickly adapts to changes in Internet topologies that affect router-location predictions. The system should be accurate enough to understand the *geopolitical* properties of Internet paths regarding security and privacy.

Geolocations of *router IPs* have received little attention compared to “eyeball” IPs; in fact, naïve approaches to using geolocation databases for such infrastructure IPs leads to high inaccuracy (*e.g.*, 29% of IP geolocations are incorrect in our labeled dataset for MaxMind, shown in Table 3). When looking at the accuracy on a per-country level, existing approaches achieve 90% accuracy or better only for 5–46.5% of countries (Figure 4). Such inaccuracy can signif-

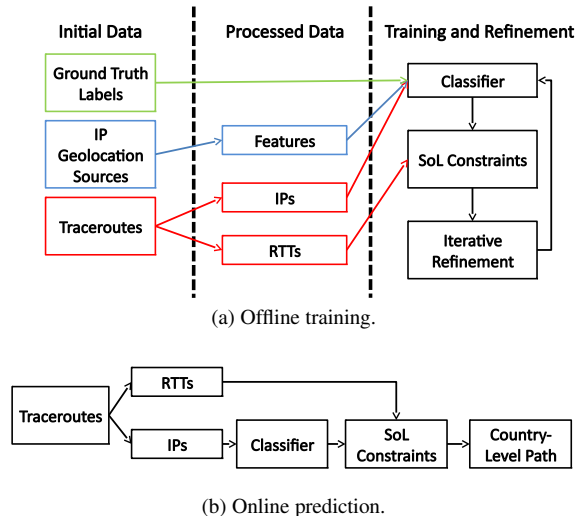


Figure 1: **Key system components.** The offline training module (top) builds a machine-learning classifier that predicts router country; the online prediction module (bottom) provides interactive prediction of router countries for an input set of traceroute measurements.

icantly impact our ability to correctly interpret geopolitical implications of paths traversing those routers.

Thus, instead of relying on any fixed set of geolocation data sources to predict a router location, we take an alternative approach that relies on machine learning. Specifically, our hypothesis is that the set of geolocation sources that will reliably predict a router’s country varies according to properties of the router (*e.g.*, IP, BGP prefix, *etc.*), and that we can build a machine-learning classifier that reliably predicts router locations based on an ensemble of individually unreliable sources. We test this hypothesis and show that a machine-learning classifier, combined with active measurement probes, can achieve substantially higher accuracy and precision compared to previous approaches (*e.g.*, at least 90% accuracy for 96.5% of countries, as shown in Fig. 4).

Assumptions. Our system takes as input a set of traceroute measurements (along with round-trip times to each responsive router) and a collection of IP geolocation data sources. After processing this initial data, we assign countries to each IP address¹ along a path that we observe. We assume that the errors affecting each data source are not random and that data sources achieve high accuracy for at least some networks. We further assume that our ground truth IP geolocations are correct. Though we focus on country-level geolocation, our input geolocation sources can use finer-grained precision.

To bootstrap our traceroute-based country-level geolocation analysis, we assume that fixed-line end-host geolocations can be predicted accurately. We also assume that each IP observed along a path is assigned to exactly one cor-

responding router,² and that neither the router nor the IP change locations substantially during each measurement and analysis round, which is currently one day. While it is certainly the case that IPs can be reassigned or reused arbitrarily within an ISP, we expect it to rarely affect our conclusions.

Non-goals. This work does not focus on geolocating eyeball IP addresses; rather, we assume that geolocation databases provide high accuracy for such IPs and thus use them as “anchors” for geolocating routers. We do not map routers to city-level (or finer-grained) locations; rather we focus only on the country where it resides. This is sufficient to inform several important security and privacy analyses. We do not attempt to provide perfect accuracy or coverage of Internet paths; however, our approach should geolocate most Internet paths most of the time. If a router on a path does not respond to a traceroute (or a ping) probe, then we cannot use SoL to geolocate it; however, we can use SoL constraints to other responding hops on the path to identify the set of countries such absent routers might be located in.

3.2 Design

PASSPORT is designed around two high-level components, *offline training* and *online prediction* (Figure 1). For offline training, we begin with traceroutes, ground truth location labels, and IP geolocation databases. We use these as input to train a machine-learning classifier that predicts the country where a router is located.

For online prediction, we accept as input (from a user, or via an API call) a traceroute measurement and return the set of countries in which each router is predicted to be located. The result can be zero or more countries, but as we show in § 5, we predict exactly one country for the vast majority (88.1%) of cases. This section provides a high-level overview of the system design; we cover implementation details in § 4. Table 1 presents a roadmap for the remainder of the paper.

3.2.1 Offline Training

The purpose of the offline training component of PASSPORT is to build a machine-learning classifier that reliably predicts the country-level location of routers appearing in arbitrary traceroute measurements. Offline training consists of three phases (Fig. 1a): data collection, feature selection, and training and refinement. Data collection entails gathering router IPs and RTTs from traceroutes, geolocation hints from available (unreliable) sources, and ground truth geolocations when available.

In the training and refinement phase, we first use ground truth labels (§4.1.1) and features to train an initial classifier. We then use RTTs from traceroute data to rule out any classifier predictions that violate speed of light (SoL) constraints. As a result, there may be cases where at least one router has no predicted country. In the iterative refinement phase,

¹Hereafter referred to simply as *IP*.

²As such, we will use the terms *router* and *IP* interchangeably.

§	Topic	Dataset	Key results
4.1.3	Classifier Selection	Controlled exp.	We develop an <i>ensemble</i> of 4 classifiers, trained on different subsets of labeled data.
5.2.1	Overall Accuracy	Ground truth	We show that our approach is more accurate than single classifiers and geolocation services.
5.2.3	Ensemble Construction	Ground truth	Adding more classifiers in the <i>ensemble</i> provides diminishing returns for the increase in accuracy while increasing the number of countries predicted (decreasing precision).
5.3	Constraint-based Refinement	Traceroute IPs	We increase the precision of the <i>ensemble</i> classifier in PASSPORT using SoL constraints (PASSPORT maps 88.1% IPs to a single country and 95.5% to at most two countries).
5.5	Comparison: Alternatives	Traceroute IPs	PASSPORT has a high consistency with EdgeScape and IP2Location, but lower with other geolocation databases. Most inconsistencies occur when geolocation databases predict a country that violates SoL constraints, and these affect accuracy for large fractions of paths.
6	Geopolitical Case Studies	Traceroute IPs	PASSPORT identifies many cases of international detours, reverse-forward path asymmetry, and circuitous paths.

Table 1: Roadmap for key topics covered in §4, §5 and §6.

we use the geolocated routers on a path and SoL constraints to help locate such routers. We then use the results of these analyses to *retrain* the classifier. This phase terminates when the set of predicted router locations becomes stable (*e.g.*, no more than 1% of router locations change from one iteration to the next). In practice our system converges within two or three rounds. We run this analysis periodically, currently using a (configurable) period of one week.

Note that our design incorporates a single logical classifier, but can be (and is) composed of an ensemble of classifiers. We refer to this single logical classifier as *ensemble*. As we discuss in the next section, we use a custom ensemble of classifiers, with each sub-classifier selected based on empirical analysis to achieve high accuracy (see Appendix C, to make router-country predictions).

3.2.2 Online Prediction

The online prediction component makes router-country predictions interactively (Fig. 1b). The input to this component is a set of traceroutes to be geolocated at the country level. We use the classifier trained in the offline phase to make initial predictions for router locations. Much like the refinement phase in the offline training component, we impose SoL constraints to identify feasible country locations for routers. The output of online prediction is a traceroute path labeled with router countries. See Appendix E for an example.

4. IMPLEMENTATION

We now discuss the implementation of each component of the PASSPORT design.

4.1 Offline Training

The offline training component currently runs once per week, a tunable parameter. It takes as input both ground truth and unreliable location data, extracts features for prediction, and uses them to build a machine-learning classifier. We now discuss the implementation details.

Source	IPs	Countries	ASes
<i>Union of all sources</i>	11,626	120	1,244
OpenIPMap	8,973	117	1,037
Manually Labeled	2,422	58	668
IXP	231	11	7

Table 2: Ground truth IP geolocations. Our dataset covers a substantial number of IPs worldwide.

4.1.1 Data Sources

We use the data below to train classifiers in PASSPORT. We identify the location of 11,626 IP addresses, in 120 countries and 1,244 ASes (Table 2).

Reliable location data. We assume the following datasets to be reliable and use them as ground truth. First, we use the lists of IXP addresses and locations provided by Packet Clearing House [24]. Such prefixes and locations are well known and tend to change rarely.

Second, we use crowdsourced labels from OpenIPMap [44] gathered between May 2016 and March 2017, comprising 19,257 IP addresses. OpenIPMap is a publicly available database that uses crowdsourcing to geolocate IP addresses, and as such may contain incorrect labels. We used RTT measurements from April, 2017, to throw out any infeasible IP geolocations from OpenIPMap as well as those that were mapped to a region spanning multiple countries. After this filtering step, we are left with 8,973 addresses. We cross-validated the OpenIPMap dataset and found that using it improved our coverage (by 62 countries) and country-level accuracy (nearly doubling the number of countries where we have 80% or higher accuracy) even if a significant number of labels are incorrect. In fact, our analysis showed that even when 10% or 20% of OpenIPMap locations are incorrect, the accuracy of PASSPORT only decrease by 0.6% and 4.1%, respectively. The details of this analysis are available in Appendix A.

Selecting IPs for manual labeling. To seed our system with useful measurement data spanning a diverse set of commonly traversed countries and ISPs, we conducted an initial set of traceroute measurements for the purpose




Figure 2: **Top 20 countries in terms of # of routers in the ground truth dataset.** There is bias toward the US and European countries, which we need to account for when training classifiers.

of manually labeling routers with their locations. These measurements were conducted in February, 2017. We selected 67 PlanetLab vantage points as sources, covering 20 countries. For each source, we picked the corresponding country-specific Alexa Top-100 websites as destinations. For each traceroute, we removed the endpoints and identified the countries where each intermediate hop IP was registered according to *whois*.

This set of 13,744 IPs covered a large set of countries around the world, but there were too many to manually label. We thus sampled the IPs using a set of heuristics to cover IPs in distinct prefixes across a large number of ASes globally. Specifically, we identified the top 20 largest ASes³ in each country (in terms of the customer cone size based on CAIDA AS Rank[5] data), then picked a destination IP address from up to five routable prefixes in the AS.⁴ We did this for the 50 countries with the largest number of *whois*-registered IPs, yielding 2,653 IPv4 addresses. Of these, 231 belonged to IXPs with public locations [24], leaving 2,422 unlabeled routers (from 668 ASes) for manual labeling.

Note that this approach was heuristically designed to cover a set of diverse networks worldwide with a limited amount of probes, and we do not claim to have maximized coverage or diversity. Rather, the point of this dataset is to provide useful initial data for training our classifier. Approaches that improve the scale and diversity of labeled routers should only further improve our system.

Ground truth labels. We manually labeled the location of 2,422 IP addresses. First, we identify the country for each unlabeled router using targeted traceroutes toward the router from multiple vantage points (first from a set located in different continents, then if necessary using ones in the same continent as the unlabeled router), and incorporate geolocation data inferred from corresponding RTT-based SoL constraints and router hostnames. For routers with

³In countries with 20 or fewer ASes registered, we selected all of the ASes.

⁴For ASes with fewer than five prefixes, we selected one IP for every prefix.

Information Source	Cost	Precision	Selected Features	Accuracy
IP address	N/A	N/A	IP address	N/A
CAIDA AS Rank	Free	N/A	ISP size (# ASes), customer cone (# of IPs & prefixes)	N/A
WhoIS	Free	N/A	AS: name, number, country	67%
			ISP: name, city, region, country	69%
DDec	Free	City	Country	83%
DB-IP	Free	City	Country	69%
Maxmind GeoLite2	Free	Country	Country	71%
APIgurus	Paid	City	Country	70%
IP2Location	Paid	Country	Country	81%
IP Info	Paid*	City	Country	71%

* provided for free for our research

Table 3: **Classifier features and their individual accuracy for predicting router country.** While some data sources offer high overall accuracy, we show in Fig. 4 that this accuracy is limited to a small number of countries that are overrepresented in our dataset.

hostnames that encode geolocation information (e.g., via city names or airport codes) we use both DRoP [25] and our own manual analysis to identify each router’s country.

Figure 2 shows a CDF of the fraction of IPs covered by the top 20 countries⁵. Interestingly, the top 3.3% (4 out of 120) countries cover 48.1% of the IPs. We discuss techniques to account for this bias below.

Unreliable geolocation sources. Our input dataset includes several unreliable geolocation sources: IP geolocation services (IPInfo [27], Maxmind Geolite2 [41], DB-IP [9], APIgurus [2], IP2Location [26]), hostname parsers (DDec [6], which combines DRoP [25] and undns [19]) and AS information (*whois* entries and AS Rank [5]).

While geolocation services have been used as reliable sources for end host IPs, our analysis in §5 shows that they provide low accuracy for router IPs in our dataset. Hostname parsers can have high accuracy for routers, but they also offer low coverage and have false positives when hostnames are ambiguous. *whois* data contains the country where an AS or IP prefix is registered, but that often differs from where a corresponding router is located. Next, we describe how we extract features from these unreliable sources and use them as input to a machine-learning classifier.

4.1.2 Feature Selection

This section describes the features we select from unreliable data sources for training our classifier. We summarize the selected features and their individual accuracy for geolocating routers in Table 3.

The feature we use from IP geolocation services and hostname parsers (IPInfo [27], Maxmind Geolite2 [41], DB-IP [9], APIgurus [2], IP2Location [26], and DDec [6]) is

⁵The complete list of countries in terms of the number of routers in the ground truth is available at <https://goo.gl/umsbjz>

Classifier	Accuracy(%)
Decision Tree	88.45 \pm 5.11
Random Forest	84.38 \pm 5.4
Extremely Randomized Trees	84.41 \pm 3.91
1-Nearest Neighbor	58.19 \pm 7.79
AdaBoost	32.05 \pm 9.39
Linear SVM	15.83 \pm 8.67
Naive Bayes	12.07 \pm 5.09

Table 4: **Classifier Accuracy using 10-fold cross validation.** Decision trees provide high accuracy.

the country corresponding to the geolocation for each IP. Further, we use the registered country/countries for each IP’s ASN and ISP (which may own multiple ASNs) based on *whois* data.

We include the size of each AS as features, under the hypothesis that larger ISPs are more likely to span multiple countries, while smaller ones are likely to be within a single country. We define size using the number of ASes in each ISP, the number of IP addresses, and the number of routable prefixes contained in an ISP’s customer cone (using [5]).

We evaluated the *feature importance* (using Gini importance [38]) of the features in determining a router’s country, and found that IP2Location [26] was the primary determinant of the predicted country for 71.0% decisions. This was followed by the IP address (6.70%), and country predicted by DDec [6] (6.44%). The ordered list of features by importance is available in Table 8 in Appendix B.

4.1.3 Classifier Selection and Training

When implementing our machine-learning approach to predicting countries where routers are located, we identified two key challenges: selecting the type of classifier(s) to train and determining the information provided for training.

Classifier Selection. Our primary goal is to identify a classifier that can provide high accuracy and precision for predicting router country. A secondary goal is to use a classifier that is human-interpretable so we can ascertain why it performs well.

We evaluated all available classifiers in *scikit-learn* [47] to find ones with high accuracy. These include tree-based classifiers and ensembles (*e.g.*, Decision Trees and Random Forest), clustering algorithms (K-Nearest Neighbors), Support Vector Machines (SVMs), Naive Bayes, and adaptive approaches (*e.g.*, AdaBoost) as listed in Table 4.

We evaluate accuracy on the entire labeled dataset using 10-fold cross validation; *i.e.*, we sort the labeled data randomly and evenly divide it into 10 folds (disjoint groups). We test the accuracy of the classifier by training it on nine folds and testing on the one that was not in the training set. This is repeated 10 times, one for each fold to exclude from the training dataset, and the average accuracy (as well as standard deviation) are presented in the second column of Table 4. We refer to such classifiers, trained on the entire training dataset, as *default* classifier.




Figure 3: **Classifier Comparison.** CDF of the fraction of accurately predicted router locations per country, for each classifier type that performed well on the overall dataset. When training the classifiers using (a) randomly sampled number of instances, and (b) equal number of instances in each country, decision trees and forests provide the best per-country accuracy.

There are clear winners and losers. Decision trees and their ensembles perform well (with similar accuracy), while clustering approaches and AdaBoost do not. However, this finding is potentially biased by the sample size per country in our dataset, and may simply indicate that the classifiers do well only for the countries with a large sample size.

To investigate this, we compare the top-performing classifiers from Table 4 in terms of the accuracy *per country*. For each classifier type, we train one instance using data from a uniformly *random sample* of training examples, and a separate instance using an *equal number of training examples per country* (*i.e.*, by oversampling from countries with few routers and undersampling from those with large numbers of routers)⁶. Figure 3 shows our per-country accuracy results. We found that Decision Trees and Random Forests substantially outperform all other classifiers, regardless of sampling method (for details, see Appendix C.1). Given that decision trees perform well and are easy to interpret,⁷ we selected decision-tree based classifiers for our implementation. *Note, however, that no individual classifier performs well for all countries.* This motivates the need for ensemble approaches as described below.

Training an ensemble of classifiers. We have established that the training dataset used for classifiers can have a significant impact on accuracy, particularly when there is bias in the dataset. Thus, it is important to take these biases into account in PASSPORT to provide high accuracy across a wide range of countries that Internet paths may visit.

A key challenge is that we cannot know *a priori* how biased our training dataset is relative to an arbitrary set of traceroutes needing router-country predictions. Thus, instead of attempting to provide a “one-size-fits-all” classifier trained with some subset of training data, we chose to implement an *ensemble* of classifiers, each of which is trained

⁶Note that we evaluate additional training dataset selection approaches to understand classifier sensitivity in Section 5.

⁷As expected, we find that the decision-tree nodes encode information about which database is most accurate for a given IP address.

using complementary approaches to sampling our dataset. In this model, each classifier in the ensemble predicts router countries independently and the ensemble returns a set that is the union of these countries. As we discuss in the following paragraphs, we use RTT latency with SoL constraints from traceroute data to eliminate infeasible countries from this set. If the set still contains more than one country, PASSPORT can optionally return the set or report that it was unable to isolate a single country for the router. In our evaluation, we show that PASSPORT returns exactly one country 88.1% of the time, and zero countries only 3.1% of the time.

Our *ensemble* consists of multiple classifiers.⁸ One classifier simply uses the entire labeled dataset. For the remaining classifiers, we train each instance using a different subset of data to mitigate bias in the dataset. One of the classifiers uses a fixed number of training instances per country, as we found that doing so will reduce the impact of country-level bias in the dataset. The remaining two classifiers are trained using empirically derived sets of training instances. For this, we investigate the marginal impact on accuracy from adding individual training instances for each country. One approach uses this to find the minimal set of training instances that maximizes the accuracy for each country (which we refer to as *maximum accuracy*). The other approach finds the point at which adding more instances has diminishing returns on accuracy (*i.e.*, the second derivative of the accuracy vs. number of instances curve is zero), which we refer to as *knee* (of the curve).

SoL constraints and iterative refinement. The *ensemble* classifier predicts one or more countries for each router. To improve the accuracy and precision of the ensemble, we use traceroute data in the following way.

First, we use RTT latencies associated with each traceroute path that includes the targeted router, r . For each source s in the set of sources S with a path that includes targeted router r , we find the minimum RTT delay between s and r . We conservatively find the distance d between s and r using the empirically derived propagation delay from Laki *et al.* (0.47c) [34], then identify the set of permissible regions that r can be located in using the geodesic circle c_s centered at s (geolocated using a geolocation service) with radius d . When there are paths from multiple sources containing r , we identify the set of permissible regions for r as bounded by the intersection of circles c_s for all s in S .

After identifying the permissible region where r may be located, we identify the countries intersected by region. Finally, we adjust the set of predicted countries by finding their intersection with the permissible countries. At the end of this step, a router is assigned to zero⁹ or more countries.

⁸We tested 23 classifiers, not shown due to space limitations, and used the 4 that provide highest accuracy gain in our *ensemble*.

⁹The intersection of permissible regions c_s may be an empty region, or permissible regions may not include any predicted countries.

We incorporate two optimizations to improve coverage and efficiency. First, we use lists of router aliases [32] to ensure that we can intersect paths visiting different IP addresses on the same router. Second, we ignore large latency measurements (which lead to intersection regions larger than a country). Namely, we use only measurements with an RTT less than 100 milliseconds to ensure that the diameter of c_s is less than 40% of the circumference of the earth in any direction.

For cases where the latency from r to all s in S is large (*i.e.*, our vantage points are far from the targeted router), we use an iterative refinement approach as follows. Given a path measured from s containing r , we find all *landmark routers* l with IPs that can be located with high precision (*e.g.*, reliable geolocation sources from §4.1.1). We then estimate the latency between l and r as $\frac{1}{2} * |rtt_{s-r} - rtt_{s-l}|$, where rtt_{s-r} is the RTT between s and r and rtt_{s-l} is the RTT between s and l . In other words, we assume that the path from l back to s is the same as r back to s , and find the average additional delay from l to r and r to l (which is equal to or larger than the smallest unidirectional delay on the round-trip path). Because the above path assumption may be violated in practice, we use this approach only when no other information is available to reliably isolate a router country.

After this step, each router r is mapped to a permissible set of countries (or the empty set). We use this as a new label to retrain the *ensemble* classifier to improve its predictions. We repeat the above process until the set of countries mapped to routers stabilizes.

4.2 Online Prediction

A user submits a traceroute (or an IP address) to the PASSPORT web server using an online submission form or API request. PASSPORT parses the traceroute to extract IP addresses and RTTs, then uses its classifier to predict the countries where these IPs are located, and applies SoL constraints to eliminate any infeasible countries. The output is the filtered list of IP-to-country mappings. The online system predicts each IP address to *up to two* countries or returns an error for that IP address. The detailed predictions of countries from the *ensemble* classifier, the list of countries filtered using SoL constraints, and a complete list of countries predicted by PASSPORT can be accessed using the public API.

The online prediction system, user guide, and the API are publicly available at <https://passport.ccs.neu.edu>

5. EVALUATION

This section addresses key questions regarding the effectiveness and efficiency of our approach. First, we conduct microbenchmarks on our classifier design and implementation. Next, we demonstrate the advantage of incorporating SoL constraints and iterative refinement. We then show that our approach incurs reasonably low overhead to train

and is suitable for online prediction. Last, we compare our approach with alternatives.

5.1 Methodology

We evaluate our classifier design and implementation using the same dataset and methodology presented in Section 4.1. This contains a set of router IP addresses for which we have ground truth labels. To evaluate our system against IPs not in our training dataset, we conduct additional traceroutes (using a similar methodology) as specified in Section 5.3.

5.2 Classifier Analysis

In this section, we answer the following key questions:

- *How well do simple approaches work at predicting IP location?* None is particularly accurate, especially when considering *per-country* accuracy.
- *Are all the unreliable geolocation sources similarly bad, or are there cases where at least one correctly predicts locations most of the time?* Usually at least one unreliable source can correctly predict location, opening the door to classifier-based prediction.
- *Is any one classifier sufficient to optimize accuracy?* No, an *ensemble* provides the best results.
- *Is the machine-learning approach both accurate and precise?* Yes, in the vast majority of cases our approach identifies one country for a router.

The following subsections describe which ML classifiers work best individually and how we combine them to provide higher accuracy. All these results present average statistics using 10-fold cross validation.

5.2.1 Overall accuracy

We start with the accuracy of our final *ensemble* classifier as described in Section 4.1.3, using the labeled data from Section 4.1.1. Accuracy is defined in terms of the fraction of router-to-country predictions that are correct according to ground truth labels. We exclude cases where multiple countries are predicted.

We compare the performance of our *ensemble* with different geolocation services in terms of the fraction of routers in each country that are correctly predicted (singleton set). This is plotted in Figure 4 using a CDF where the y-axis represents countries and x-axis is the fraction of a country’s routers predicted correctly. Curves closer to the bottom and right edges indicate higher accuracy.

Our *ensemble* substantially outperforms all other approaches—it achieves 90% or better accuracy for 96.5% of countries in our dataset. By comparison, individual decision trees can achieve the same level of per-country accuracy only for 61% of countries. Worse, the most accurate IP geolocation service, IP2Location achieves this only for 46.5% of countries and *whois* registry information for only 5% of countries.

In terms of implications, it is clear that machine-learning approaches are able to synthesize individually inaccurate geolocation sources to more reliably predict router geolo-

cations. Further, the accuracy for IP geolocation services and *whois* data do not extend beyond a small fraction of countries, *meaning any conclusions about geopolitical properties of Internet paths using such data (e.g., [12, 51]) are highly likely to be incorrect.*

5.2.2 Microbenchmarks and Classifier Sensitivity

We now investigate the individual components of our *ensemble* classifier and their sensitivity to the training dataset. For each country, we select the number of training instances that either (1) provide maximum per-country accuracy (*maximum accuracy*), or (2) at the point where increasing instances provide diminishing returns in accuracy (*knee*), as discussed in §4.1.3.

The outcome depends on how we allocate training instances to *other countries*, so we investigate three schemes using three schemes: using the same number of samples for other countries as in the target country (*balanced*), double the instances in other countries as in the target country (*unbalanced*), and a random number of samples for other countries (*random*).

Overall, we find that *random* performs systematically worse, while *maximum accuracy* and *knee* approaches on both *balanced* and *unbalanced* sampling methods yield similar results. The details on how we select training instances to achieve high accuracy are further described in Appendix C.2.

5.2.3 Ensemble Construction

In Section 4.1.3, we described several techniques for training classifiers in addition to the *default* classifier. We now investigate how to combine these classifiers to optimize for accuracy and precision.

Figure 5 shows the change in accuracy for all size- n sets of classifiers with the number of classifiers in the set (x -axis) against the average increase in the accuracy (y -axis) and the number of countries predicted by the set of classifiers (y_2 -axis). The figure shows a diminishing return on accuracy improvement as more classifiers are added. At $x = 3$, the increase in average accuracy is roughly 1.1% while it decreases to 0.3% for $x = 4$ and beyond this point the increase in accuracy is negligible.

We pick the first 4 classifiers to form our *ensemble* classifier. It contains *default*, *balanced* dataset bias with *knee* sampling approach, *unbalanced* with *maximum accuracy* approach, and one classifier with equal number of training instances for all countries.

5.2.4 Precision

A key question is whether the high accuracy in the *ensemble* classifier is due to one correct country being predicted for a router, or rather that the correct country is one of a large set of countries returned. To evaluate the precision, we computed the distinct number of countries and continents returned by the classifier for each router in our




Figure 4: **Comparison of our ensemble approach with alternatives**, using a CDF of per-country router location-prediction accuracy. The *ensemble* curve (bottom-right) has a high accuracy for most countries, while traditional approaches have high accuracy only for a few countries.

dataset. The *ensemble* predicts a single country for 65.7% of router IPs, while it predicts a single continent for 82.4% of router IPs (see Appendix C.4). Thus, the *ensemble* alone is accurate but not precise; further, due to the spread in continents predicted, this precision is not sufficient to reason about geopolitical implications of Internet paths. In the next section, we demonstrate how we use SoL constraints to improve precision such that the vast majority of the time (88.1% of IPs) a router is predicted to be located in one country.

5.3 Constraint-based Refinement

In this section, we discuss how we use RTT latencies measured from traceroute paths to substantially improve the precision of the *ensemble* classifier. Specifically, we use SoL constraints as described in Section 4.1.3 to rule out infeasible countries from the set predicted by the *ensemble* classifier. Further, we leverage these observations to train a classifier on the subset of feasible countries predicted by the *ensemble* to improve precision for routers without sufficiently narrow SoL constraints imposed by RTT latency.

Dataset. To analyze the impact of SoL constraints, we need a large set of traceroute data from which we can impose substantial numbers of constraints. We thus conducted an additional measurement campaign as follows. We performed forward traceroutes and reverse traceroutes [31] in April, 2017 from 172 PlanetLab nodes in 32 countries to up to four government websites for 190 countries¹⁰ (generating 30,248 and 10,477 *successful* traceroutes for forward and reverse traceroutes, respectively). To add diversity in terms of sources and destinations, we also used 442,862 traceroutes (311,684 UDP and 131,178 ICMP) from RIPE Atlas from 9,553 probes in 176 countries during the same period. These

¹⁰We used fewer than four websites if we could not locate four. The limit of four was set due to API limitations by unreliable geolocation sources




Figure 5: **Change in average accuracy for all countries for each new classifier.** We find that four classifiers are sufficient for high accuracy. By limiting the classifiers in the ensemble, the predictions are more precise (i.e., include fewer countries).




Figure 6: **Impact of using speed of light (SoL) constraints.** By incorporating SoL constraints to rule out infeasible router locations, then using this information to retrain our classifiers, our system predicts exactly one country 88.1% of the time, compared to the initial classifier that does so only 63.7% of the time.

measurements yielded at least two intersecting regions for SoL analysis for 63% of routers.

The following paragraphs compare the *precision* of several schemes for using SoL constraints.¹¹ First, we compare the precision of our *ensemble* classifier to that of SoL constraints alone, then evaluate the combination of the two techniques. Since SoL constraints can offer additional ground truth labels (e.g., by ruling out predicted countries or reinforcing correct predictions), we further evaluate the impact of retraining our classifiers with the combined information. We plot a histogram summarizing our precision results in Figure 6, with each bar indicating the fraction of router IPs (y-axis) that are predicted to be located in x countries.

Precision of SoL constraints. The leftmost two bars indicate the precision for the *ensemble* classifier and SoL constraints. We find that SoL constraints predict substantially fewer countries than the *ensemble*, but neither predicts a single country more than 64% of the time. The combination of these two approaches (“Ensemble + SoL” bar), where we use SoL constraints to identify which of the predicted countries is feasible, substantially improves precision with almost all of the routers predicted to be located in 1–2 countries.

Iterative refinement using classifier retraining. To further improve precision, we leverage two cases from the previous analysis to retrain our classifier: instances where the SoL constraints identified exactly one country for router (11,308 addresses spanning 110 countries), and instances where the *ensemble* classifier was incorrect. We retrain the *ensemble* classifier, and using the combination of SoL constraints and the predicted countries by this retrained *ensemble*, plotted “Retrained Ensemble + SoL”. This step substantially improves precision, with 79.4% of routers be-

¹¹Due to lack of ground truth in this dataset, we cannot evaluate accuracy.




Figure 7: **System performance.** Training time is reasonably low and increases linearly with the number of training instances; the online prediction component returns more than 1200 predictions in less than 300 ms.

ing located to a single country and only 5.3% of predictions containing more than one country.

The last optimization for precision incorporates cases where the *ensemble* predicted multiple countries while the SoL constraints predicted only one country and there was no intersection between the two (showing that the *ensemble* was incorrect in predicting countries). In this case, we can ignore the classifiers and *only* use the SoL constraints to predict the country. This is depicted as the rightmost “PASSPORT” bar in Figure 6. Putting it all together, our final classifier is able to predict exactly one country for a router 88.1% of the time.

5.4 Efficiency

We now evaluate PASSPORT in terms of time to train the classifiers in the offline component and time to predict in the online component.¹² All experiments used a 4-core *i7* processor (4.2 GHz) with 16 GB of RAM.

Figure 7a plots the time required to train our classifier (y-axis) as a function of the number of training instances (x-axis), both for the *ensemble* classifier and for an individual decision tree. The graph shows that an individual classifier can be trained with 100,000 instances in approximately 50 seconds. When looking at training the *ensemble* of 4 classifiers used in our system, we find that training time takes approximately 4.8 seconds when trained in parallel (for a combined set of 31,603 instances). Thus, re-training classifiers in the offline phase is not a bottleneck in our system, and if necessary, retraining can be done on the scale of 10s of seconds.

For the online prediction component, we expect users to submit traceroutes and obtain the corresponding countries on the path interactively (*i.e.*, within a small number of seconds or less). Figure 7b plots the response time as a function of the number of parallel requests for predictions where measurements are available. The plot shows that under low load the system can make predictions within 10s of milliseconds, serve 350 parallel requests in less than 100 ms and can serve up to 2,000 parallel requests in less than half a second. Thus, PASSPORT is sufficiently fast to provide interactive predictions for traceroutes.

¹²We do not include the variable times necessary to conduct active measurements, though they are on the order of seconds.

5.5 Comparison with Alternatives

We now compare our approach with geolocations provided by other commonly used services. Unlike the analysis in §5.2, we use the dataset described in §5.3, which is larger but does not have ground truth labels for router locations.

Geolocation Databases While comparing public databases, instead of determining whether each service is *accurate*, we evaluate other approaches in terms of whether they are consistent with PASSPORT results *for the cases where exactly one country is predicted*. While we cannot guarantee that all these cases are accurate, we have high confidence in them due to SoL constraints and classifier predictions.

Table 5 presents the results for several geolocation services.¹³ We find that all of the geolocation sources yield locations that are inconsistent with our data (including SoL constraints), ranging from 7.3–21.5% of routers in our dataset for April 2017. We found that this number increases to 9.1–23.1% for the same geolocation services with data from June 2016 (not shown), indicating that using slightly older geolocation databases (as is the case for Maxmind GeoIP2 [40] and IPLigence [28]) yields similar results.

We evaluate the impact of these inconsistencies when it comes to mapping traceroutes to country-level paths and find that a substantial fraction (37.4–66.7%) of paths are affected. We find that 0.1–14% of the inconsistencies occur due to SoL violations, and thus are not only inconsistent but *incorrect*.

We found several patterns behind the observed SoL violations. IPInfo [27] returns EU as a country for 42% of their SoL violations (instead of a single country), while Maxmind [41] incorrectly returns either Switzerland or Sweden for 48% of the violations. Similarly, IP2Location [26], and EdgeScope [1] return United States as the most commonly mislabeled country, corresponding to 22% and 11% of incorrect cases, respectively. IPLigence fails to predict a country for over 27% of violations.

For cases where PASSPORT and other databases were inconsistent, we sampled and analyzed a subset of cases for manual analysis. In the vast majority of cases, we find that PASSPORT is correct. We sampled and labeled 596 IP addresses from 519 ASes in 146 countries in a manner similar to §4.1.1. Our results showed that PASSPORT had the highest accuracy with 96.47% IP addresses being labeled correctly, followed by IP2Location at 92.3% and Edgescape at 90.9%. PASSPORT outperformed other sources in all continents except Europe, where Edgescape had the highest accuracy (95.35%, vs. 93.4% for PASSPORT). Appendix D provides an in-depth comparison of the geolocation sources with PASSPORT.

Comparison with active measurement systems. A key advantage of PASSPORT is that it can accurately geolocate

¹³We attempted to compare with Alidade [7], but the service was not running at the time of writing. The authors instead provided us with locations provided by Alidade’s geolocation services, which are included in the table.

Source	SoL (%)	Viol.	Inconsistent IPs (%)	Affected Paths (%)
EdgeScape	0.1		7.6	37.4
IP2Location	1.4		7.3	31.2
DB-IP	6.4		14.7	60.5
Maxmind	10.0		21.2	64.3
GeoIP2*				
Maxmind	13.5		21.1	64.1
GeoLite2				
IP Info	14.0		21.5	65.6
IPligence*	13.2		23.1	66.7

Table 5: **Comparison with other approaches.** While many geolocation databases are highly consistent with PASSPORT, their inconsistencies affect large fractions of paths in our dataset. Further, several databases exhibit SoL violations for a substantial fraction of IPs in our dataset; we know in these cases that the databases are incorrect. * from June 2016

router countries with few (or no) additional measurements beyond a traceroute. To demonstrate this, we investigated the number of vantage points required to issue ping measurements to determine whether a predicted country is feasible according to SoL constraints. We explored a greedy approach (use VPs with the lowest RTT to a router) and a random approach (select VPs randomly).

Both approaches are highly efficient at achieving high precision. When using a greedy approach, one VP provides the same result as using all VPs. The random approach, which requires no a priori knowledge, can predict 87.41% of routers to a single country (compared to the optimal 88.1% in Fig. 6). Further, at most 6 randomly selected VPs are as good as using all of them.

By comparison, Wang *et al.* [53] and other geolocation approaches that use active measurements [13, 54] perform measurements from *all* available VPs (see Eriksson *et al.* [13], which evaluates such approaches using a minimum of 25 VPs). Wong *et al.* [54] found that Octant can locate 80% of target addresses using only 10 VPs, but this is still an order of magnitude larger than our approach. In short, PASSPORT provides high accuracy without needing a large distributed set of vantage points or large numbers of measurements, thus reducing the barrier to deployment.

6. CASE STUDIES

We now use PASSPORT and the dataset of §5.3, to study paths with interesting properties. We focus on *detours*—paths that traverse at least one country that is not the source or destination country. Paths that start and end in the same country (*e.g.*, Canada) but detour through another (*e.g.*, the US) are especially interesting, because such “purely domestic traffic” is subject to the surveillance and censorship regime in the detoured country. We are also interested in paths that transit multiple continents. Table 6 lists sources, destinations, detour countries for several paths in our dataset. Each path described in Table 6 has been manually validated using hostnames, pings from multiple vantage points, and

Source	Dest.	Interesting Detours	Cases (Detours/ Total Traceroutes)
BR	RU	US, FR, DE	12 / 12
PH	PH	HK	1 / 12
CA	CA	US	23 / 457
CZ	CZ	PL	24 / 132
LV	FR	RU	2 / 11
AM	BG, LU	RU	8 / 20
CN	LB, SA	US, UK, FR	12 / 12
CN	SG, IN	US	326 / 326
CN	QA	US, FR	8 / 8
SG, JP, CN	PK	US, FR	22 / 22
SG	CN, PH	US	5 / 5
PH	PK, LK	US, UK	5 / 5
PH	CN, TH	US	11 / 11
PH	LB	FR	9 / 9
TW	TH	US	3 / 3
GH	MW	UK, FR	312 / 312
GH	TN	UK	58 / 58
ZA	RE	UK, FR	17 / 17
AR, JP	RE	US, FR	11 / 11
NZ	RE	BR, US, FR	7 / 7

Table 6: **Interesting Cases.** Path starting and ending in Asia/Africa detour through the US and/or EU. Paths starting and ending in the EU detour through Russia. Some Canadian paths detours the US, a Philippines path detours through Hong Kong, and paths starting and ending in Oceania take a circuitous trip around the globe.

review of RTTs between hops and next-hop location. We discuss the a few of these paths in more detail below.

BRICS. Around late 2014, the BRICS countries (Brazil, Russia, India, China, and South Africa) were reportedly planning to build a undersea fiber cable that would interconnect them, while avoiding the US and Europe [36]. Thus far, however, this effort has come to nought. In our dataset, all the traceroutes from Brazil to Russia transited through the United States and France. We also saw US detours in all paths from China to India.

Russia. Paths that start and end in Europe but detour through Russia are interesting, given reports of potential Russian meddling in European elections [10, 11]. That said, Russia is home to one of the world’s largest IXPs (MSK-IX) [42], which could explain why we saw several detours from within the EU to Russia. Some of the paths between Latvia and France are carried by a Russian Telecom company through Russia. In addition, a set of the paths between Armenia and Luxembourg or Bulgaria travel through Russia.

Asia to the US. We identified several cases where traffic originating and terminating in different Asian countries detours to the US. In fact, we found detours to the US in *all* the traceroutes in our dataset from China to its neighbors (India, Pakistan, and Philippines) and from China to the Middle East (Saudi Arabia, Lebanon, and Qatar). This is significant because foreign communications that transit the US are subject to warrantless surveillance under Section 702 of the Foreign Intelligence Surveillance Act (FISA) [46].

Canada to the US. We identified a path that started and ended in Canada, and took a detour through the US, corrob-

orating evidence provided by the IXmaps project [29]. The US intelligence community recognizes Canadian citizens as “second party” persons, and thus requires some additional approval (beyond the standard FISA Section 702 authorizations) before they can be “targeted” for surveillance [43].

Phillipines to Hong Kong on the reverse path. One particularly interesting path started and ended in the Phillipines; while its forward path stayed inside the Phillipines, its reverse path passed through Hong Kong. This path highlights the importance of using reverse traceroute [31] to measure international detours.

Long paths. Some of the longest (geographically) paths we observed were between Oceania and Africa. For example, we found a path originating in New Zealand that traversed Brazil, the US, France, and then eventually arriving at the destination in Reunion Island. We also found detours through the US, UK, and France for traffic originating and terminating in *different* African countries. The inflated paths between African countries was initially reported by Gupta *et al.* [23], and tends to result from poor connectivity at regional IXPs.

US traffic transiting abroad. Of particular interest are cases where traffic starting and ending the US transits through a foreign country, because US surveillance law applies fewer restraints when American’s Internet traffic intercepted abroad, rather than inside the US [3]. Our dataset, however, suggests that such cases are rare. We did find some cases (7 out of 32,609 traceroutes) where a traceroute might have exited and entered the United States via a Level 3 router in Toronto. However, the router labeled as Toronto was unresponsive to probes upon validation and so we cannot confirm this finding.

7. DISCUSSION

Generalizability. We used a large set of traceroute measurements to inform the design, implementation, and evaluation of our system. However, our results apply necessarily only to the data that we collected. We believe that our results will be favorable for other datasets using similar vantage points; however, we cannot make claims about how well it will perform in networks and countries that were not in our dataset. Our classifiers were limited by the ground truth labels made available to them, and we expect that additional labels will improve our results.

Caveats for machine learning. We used relatively simple machine-learning classifiers in large part due to their already-high accuracy, but it is possible that more advanced techniques would improve the system. An advantage to our decision-tree-based approach is that one can inspect the trees to determine whether the classifier is learning something intrinsic to data sources. Such analysis can provide confidence in the ability to perform well when provided with different labeled data.

Improving datasets. We focused on country-level router geolocation for unidirectional paths based on single snapshots of paths between source/destination pairs. As part of ongoing work, we are investigating how geopolitical properties of paths change over time. We will incorporate real-time BGP feeds to investigate suspicious transient geographic detours. We will also expand the set of measured paths to provide greater coverage. Finally, we will investigate how to include crowdsourced data from our online PASSPORT tool.

Implications. We highlighted a number of scenarios of “interesting” paths in terms of the countries they traversed and whether they were inflated. In many of these cases, it is possible that this behavior is normal and even intentional. While we tried to highlight cases that we thought were poignant, in general we leave such decisions to individuals with sufficient knowledge to draw strong conclusions about implications.

Certain stakeholders may wish to avoid undesirable paths. We will investigate how to use the PEERING testbed [50] to make BGP announcements that cause routes to avoid certain countries.

8. CONCLUSION

This paper showed that one can reliably predict the countries visited by routers along a traceroute path using collection of unreliable geolocation sources, when paired with speed-of-light constraints and machine learning. We designed and built a system, PASSPORT, that does this, and demonstrated that it is accurate, precise, and efficient enough to provide information for submitted traceroutes interactively. We showed that its accuracy is substantially better than standard geolocation sources. We also used our system to evaluate the implications of the geopolitical paths our system identifies, revealing potential security, privacy, and performance issues.

9. ACKNOWLEDGMENTS

This work is partially supported by NSF awards CNS-1618955 and CNS-1405871. We thank the anonymous reviewers for their feedback. We also thank IPInfo[27] for providing us with their geolocation API.

References

- [1] Akamai. 2018. Akamai EdgeScape. <http://www.akamai.com/html/technology/products> (2018).
- [2] APIgurus. 2018. EurekaAPI/APIgurus: IP-GeoLoc IP Address Geolocation Online Service. <https://www.apigurus.com/>. (2018).
- [3] Axel Arnbak and Sharon Goldberg. 2015. Loopholes for Circumventing the Constitution: Unrestrained Bulk Surveillance on Americans by Collecting Network Traffic Abroad. *Michigan Telecommunications and Technology Law Review: MTTLR* (2015).

- [4] Courtney M. Bowman Bowman, Ying Li, and Lijuan Hou. 2017. A Primer on China’s New Cybersecurity Law: Privacy, Cross-Border Transfer Requirements, and Data Localization. *Proskauer: Privacy Law Blog* (May 9 2017). <http://privacylaw.proskauer.com/2017/05/article/2017-05-09-china-new-cybersecurity-law-privacy-cross-border-transfer-requirements-and-data-localization/>
- [5] CAIDA. 2018. AS Rank: AS Ranking. <https://as-rank.caida.org/>. (2018).
- [6] CAIDA. 2018. DDec - DNS Decoded (BETA). <http://ddec.caida.org/>. (2018).
- [7] Balakrishnan Chandrasekaran, Mingru Bai, Michael Schoenfeld, Arthur Berger, Nicole Caruso, George Economou, Stephen Gilliss, Bruce Maggs, Kyle Moses, David Duff, et al. 2015. *Alidade: IP Geolocation without Active Probing*. Technical Report. Duke University.
- [8] Jennifer C Daskal. 2015. The Un-Territoriality of Data. *125 Yale Law Journal* 326 (2015).
- [9] DBIP. 2018. DB-IP Database. <https://db-ip.com/db/>. (2018).
- [10] Mark Deen. 2017. Frances Presidential Front-Runner Says Russia Is Hacking Him Now. <https://www.bloomberg.com/politics/articles/2017-02-14/macron-urges-eu-pressure-on-russia> *Bloomberg Politics* (2017).
- [11] Janosch Delcker. 2017. Germany fears Russia stole information to disrupt election. <http://www.politico.eu/article/hacked-information-bomb-under-germanys-election/>. *Politico* (2017).
- [12] Anne Edmundson, Roya Ensafi, Nick Feamster, and Jennifer Rexford. 2016. Characterizing and Avoiding Routing Detours Through Surveillance States. *CoRR* (2016).
- [13] Brian Eriksson, Paul Barford, Bruce Maggs, and Robert Nowak. 2012. Posit: a lightweight approach for IP geolocation. *ACM SIGMETRICS Performance Evaluation Review* 40, 2 (2012), 2–11.
- [14] Brian Eriksson, Paul Barford, Joel Sommers, and Robert Nowak. 2010. A Learning-based approach for IP geolocation. In *Passive and Active Measurement*. Springer, 171–180.
- [15] EU. 1995. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data. *Official Journal of the EC* 23, 6 (1995).
- [16] EU. 2016. Data transfers outside the EU. <http://ec.europa.eu/justice/data-protection/article-2016-05-18/> (2016).
- [17] EU. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the EC* L111 (2016).
- [18] Cyrus Farivar. 2013. German NSA has deal to tap ISPs at major Internet Exchange. *ars Technica* (July 2013).
- [19] Michael J Freedman, Mythili Vutukuru, Nick Feamster, and Hari Balakrishnan. 2005. Geographic locality of IP Prefixes. In *Proc. of IMC*.
- [20] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensafi, and Christos Papadopoulos. 2017. A Look at Infrastructure Geolocation in Public and Commercial Databases. In *Proc. of IMC*.
- [21] Google. 2017. The Google Maps Geolocation API. <https://developers.google.com/maps/documentation/geolocation/> (2017).
- [22] Allison Grande. 2014. Brazil Nixes Data Localization Mandate From Internet Bill. <https://www.law360.com/articles/520198/brazil-nixes-data-localization-mandate-from-internet-bill> *law360.com* (2014).
- [23] Arpit Gupta, Matt Calder, Nick Feamster, Marshini Chetty, Enrico Calandro, and Ethan Katz-Bassett. 2014. Peering at the Internet’s Frontier: A First Look at ISP Interconnectivity in Africa. In *Proc. of PAM*.
- [24] Packet Clearing House. 2017. Internet Exchange Directory. <https://prefix.pch.net/applications/ixpdir/members/> (2017).
- [25] Bradley Huffaker, Marina Fomenkov, and kc claffy. 2014. DRoP: DNS-based router positioning. *Proc. of ACM SIGCOMM CCR* (2014).
- [26] IP2Location. 2018. IP2Location: IP Address to Identify Geolocation. <https://www.ip2location.com/>. (2018).
- [27] IPInfo. 2018. IPInfo.io. <https://ipinfo.io/>. (2018).
- [28] IPLigence. 2017. IPLigence: IP Geocator. <http://www.ipligence.com/geolocation/>. (2017).
- [29] IXMaps. 2018. IXmaps: Map Your Internet. <https://www.ixmaps.ca/>. (2018).
- [30] Ethan Katz-Bassett, John P. John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. 2006. Towards IP Geolocation Using Delay and

- Topology Measurements. In *Proc. of IMC*. <http://doi.acm.org/10.1145/1177080.1177090>
- [31] Ethan Katz-Bassett, Harsha V. Madhyastha, Vijay Adhikari, Colin Scott, Justine Sherry, Peter van Wesep, Arvind Krishnamurthy, and Thomas Anderson. 2010. Reverse traceroute. In *Proc. of USENIX NSDI*.
- [32] K. Keys, Y. Hyun, M. Luckie, and k. claffy. 2013. Internet-Scale IPv4 Alias Resolution with MIDAR. *IEEE/ACM Transactions on Networking* 21, 2 (Apr 2013), 383–399.
- [33] Rupa Krishnan, Harsha V. Madhyastha, Sridhar Srinivasan, Sushant Jain, Arvind Krishnamurthy, Thomas Anderson, and Jie Gao. 2009. Moving Beyond End-to-End Path Information to Optimize CDN Performance. In *Proc. of IMC*.
- [34] Sándor Laki, Péter Mátray, Péter Hága, István Csabai, and Gábor Vattay. 2009. A detailed path-latency model for router geolocation. In *Proc. of TridentCom*. IEEE, 1–6.
- [35] Sándor Laki, Péter Mátray, Péter Hága, Tamás Sebők, István Csabai, and Gábor Vattay. 2011. Spotter: A model based active geolocation service. In *INFOCOM, 2011 Proceedings IEEE*. IEEE, 3173–3181.
- [36] Stacia Lee. 2017. International Reactions to U.S. Cybersecurity Policy: The BRICS undersea cable. <https://jsis.washington.edu/news/reactions-u-s-cybersecurity-policy-bric-undersea-cable/>. *University of Washington, H.M. Jackson School of International Studies* (2017).
- [37] Philip Levis. 2012. The Collateral Damage of Internet Censorship by DNS Injection. *ACM SIGCOMM CCR* 42, 3 (2012).
- [38] Gilles Louppe, Louis Wehenkel, Antonio Sutera, and Pierre Geurts. 2013. Understanding variable importances in forests of randomized trees. In *Advances in neural information processing systems*. 431–439.
- [39] Ewen MacAskill, Julian Borger, Nick Hopkins, Nick Davies, and James Ball. 2013. GCHQ taps fibre-optic cables for secret access to world’s communications. *The Guardian* (June 2013).
- [40] MaxMind. 2018. MaxMind GeoIP2: IP Address to Country Database. <https://www.maxmind.com/en/geoip2-country-database> (2018).
- [41] MaxMind. 2018. MaxMind GeoLite2: IP Address to Country Database. <http://dev.maxmind.com/geoip/geoip2/geolite2/> (2018).
- [42] MSK-IX. 2018. Internet eXchange #1 in Russia; 100G peering ports and colocation. <https://www.msk-ix.ru/en/>. (2018).
- [43] Ellen Nakashima and Ashkan Soltani. 2014. Privacy Watchdog’s Next Target: the Least-Known but Biggest Aspect of NSA Surveillance. *Washington Post* (2014).
- [44] RIPE NCC. 2017. RIPE Atlas OpenIPMap database. <https://github.com/RIPE-Atlas-Community/openipmap> (2017).
- [45] Venkata N Padmanabhan and Lakshminarayanan Subramanian. 2001. An investigation of geographic mapping techniques for Internet hosts. In *Proc. of ACM SIGCOMM*.
- [46] PCLOB. 2014. Privacy And Civil Liberties Oversight Board: Report on the Surveillance Program Operated Pursuant to Section 702 of the Foreign Intelligence Surveillance Act. (July 2014).
- [47] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12, Oct (2011), 2825–2830.
- [48] Ingmar Poesse, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. 2011. IP geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review* 41, 2 (2011), 53–56.
- [49] Ronald Regan. 1981. Executive Order 12333: United States Intelligence Activities. <http://www.archives.gov/federal-register/codified-compilations/1981-12-01> (December 1981).
- [50] B. Schlinker, K. Zafiris, I. Cunha, N. Feamster, and E. Katz-Bassett. 2014. PEERING: An AS for Us. In *Proc. of HotNets*.
- [51] Anant Shah, Romain Fontugne, and Christos Papadopoulos. 2016. Towards Characterizing International Routing Detours. In *Proc. of AINTEC*. <http://doi.acm.org/10.1145/3012695.3012698>
- [52] Shaun Walker. 2015. Russian data law fuels web surveillance fears. *The Guardian* (September 12 2015). <https://www.theguardian.com/world/2015/sep/01/russian-data-law>
- [53] Yong Wang, Daniel Burgener, Marcel Flores, Aleksandar Kuzmanovic, and Cheng Huang. 2011. Towards Street-Level Client-Independent IP Geolocation. In *Proc. of USENIX NSDI*.
- [54] Bernard Wong, Ivan Stoyanov, and Emin Gün Sirer. 2007. Octant: A Comprehensive Framework for the Geolocalization of Internet Hosts. In *Proc. of USENIX NSDI*.

APPENDIX

A. OPENIPMAP ANALYSIS

We evaluated the accuracy of OpenIPMap [44] before considering it as a source of reliable data for training our classifier, because this data source uses crowdsourced labels that are not independently validated (outside of this study). As a first step in the analysis, we removed the IP addresses with wrongly labeled countries from OpenIPMap that violated SoL constraints. Then we performed a two-fold analysis to evaluate the accuracy of classifiers with and without OpenIPMap. Our results show that using filtered OpenIPMap data when training a Decision Tree classifier improves accuracy, even if a significant fraction of labels are wrong.

A.1 Data Sources

We use the two data sources as described in Section 4.1.1. These data sources are as follows.

Reliable Sources/Manually Labeled. These are the IP addresses that we manually labelled using traceroutes and ping measurements, hostname entries, and IP prefixes to responsive routers, as shown in Table 2. This includes IXP locations.

OpenIPMap. These are the IP addresses that were collected from OpenIPMap [44] and validated using SoL constraints, as shown in Table 2.

Entire Dataset. This dataset is a *union* of our *reliable data* and *OpenIPMap* datasets, as shown in Table 2.

A.2 Evaluation

To evaluate the effect of using OpenIPMap [44] as our ground truth, we looked at two aspects of the dataset; poisoning the reliable data, and evaluating the accuracy of this reliably labeled dataset using OpenIPMap. Our overall results show that using OpenIPMap as training data has a net effect of improving accuracy and coverage compared to not using it.

A.2.1 Poisoning

To understand the effect of incorrect labels provided by OpenIPMap, we evaluate the resilience of our classifier and its ability to correctly predict the true country despite given wrong training labels.

We take our manually labeled dataset (reliable data) and “poison” different percentages of the data by randomly assigning an incorrect label. Then use 10-fold cross validation on each corresponding Decision Tree classifier to identify how the poisoned data impacts accuracy.

We evaluate four categories of predictions; i) where classifier was able to predict the correct country despite being provided with a wrong label, ii) where classifier predicted to the poisoned country (for the routers where it was poisoned and the true label where it was not poisoned),

iii) where the predicted country was neither the true country nor the poisoned country and, iv) where classifier was unable to predict any country and suggested an “unknown country” label.

Figure 8 shows the effect of poisoning on the reliable data for all four categories mentioned. With 10% of the labels being incorrect, the accuracy decreased by 0.6% while with 20% the data being poisoned, the accuracy decreases by 4.1%. Even when the 40% of the dataset is poisoned, the accuracy of the classifier remains well above 80% showing that our classifiers are resilient to reasonably low levels of incorrect labels.

A.2.2 Accuracy Analysis for Training on Datasets

To further understand the affect of OpenIPMap on the accuracy, we evaluated the accuracy of the classifier with and without OpenIPMap. We evaluate the accuracy of the classifiers by training the classifiers using reliable data, OpenIPMap, and the aggregation of both OpenIPMap and manually labeled dataset. We then evaluate them and test the accuracy of the prediction for the reliable data.

We used 10-fold cross validation (where *training* and *testing* datasets were the same) with a Decision Tree classifier to evaluate the accuracy.

Table 7 shows that when our reliable dataset is *tested* on the all the other datasets, the difference in accuracy due to different training dataset is low. The second row shows that when using both OpenIPMap and reliable labels, the accuracy difference compared to using only reliable data (first row) is statistically small—within a standard deviation. We also investigate other combinations of training and testing datasets (last three rows) for completeness.

Not only is the difference in accuracy low when including OpenIPMap, the resulting coverage in terms of per-country accuracy is *substantially better*. Figure 9 shows the CDF of the country-level accuracy when we *tested* the accuracy of reliable dataset by *training* the classifier on reliable data, OpenIPMap, and the entire dataset. OpenIPMap, despite having the lowest overall accuracy in Table 7 outperforms other training datasets (curve to the bottom right), closely followed by the entire dataset. When trained on OpenIPMap and entire dataset, the classifier was able to predict 73% and 72% countries with at least 80% accuracy, respectively, while this accuracy level is achieved only for 39% of the countries when the classifier was trained on the reliable data.

The Table 7 and Figure 9 show that using the *entire dataset* rather than *reliable data* alone provides higher *per-country* accuracy with an insignificant decrease in the *overall* accuracy.

We also investigated the accuracy of OpenIPMap when trained on our reliable data, but omit the results because a significant number of countries in the OpenIPMap dataset are not included in the training dataset.

B. FEATURE IMPORTANCE

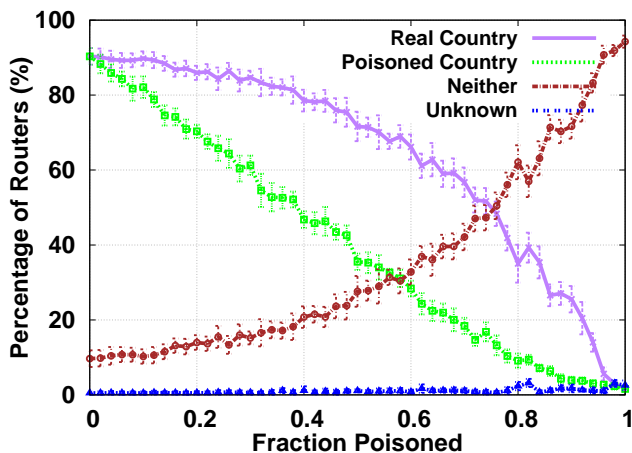


Figure 8: **Poisoning the dataset** When 10% of reliable dataset is poisoned only, the overall accuracy decrease is only 0.6% (less than half of a standard deviation) showing that our classifiers are resilient to poisoning and can filter out incorrect labels during the training phase.




Figure 9: **Country-level accuracy of the classifier** when it was trained on the reliable (manually labeled) dataset, OpenIPMap, and the entire dataset (OpenIPMap and manually labeled) and then tested using the manually labeled only. OpenIPMap has the best country-level accuracy while using manually labeled data alone decreases the per-country accuracy for the dataset.

In a Decision Tree classifier, different features have different weights associated with them. One feature is, sometimes, preferred over another to make the decision.

To understand the effect of all the features on the classifiers, and identify the most important ones in deciding the country label by our classifier, we perform feature selection analysis on the *default* classifier.

We train the classifier on our ground truth training dataset and then analyze the assigned weights. We do not enforce

Training Data	Testing Data	Accuracy (%)
Reliable Data	Reliable Data	88.91 ± 1.56
Entire Dataset	Reliable Data	87.46 ± 0.93
OpenIPMap	Reliable Data	83.39 ± 0.99
Entire Dataset	Entire Dataset	88.45 ± 1.01
OpenIPMap	OpenIPMap	87.32 ± 0.69

Table 7: **Accuracy for different training-testing datasets** with 10-fold cross validation. The decrease in accuracy is relatively insignificant when we compare the accuracy of testing the manually labeled data set trained on manually labeled data against training on the entire dataset.

Feature	Importance (%)
IP2Location Country	71.00
IP address	6.70
DDec Country	6.44
DB-IP Country	5.00
ISP Name	4.26
IPInfo Country	1.83
AS Country	0.90
IP Prefix	0.76
AS Name	0.83
ISP Country	0.46
EurekAPI	0.42
Maxmind GeoLite	0.41
AS Number	0.23
AS Registry	0.22
ISP size (# ASes)	0.17
ISP customer cone	0.17
ISP city	0.09
ISP Region	0.08

Table 8: **Feature importance for all features** in the default classifier (with all labeled data) in the ensemble. The country predicted by IP2Location is the most important feature in the classifier predictions.

any weights and let the classifier decide the power of each feature.

Table 8 summarizes the importance of each feature for the *default* (all labeled data) classifier. Our results show that IP2Location is the primary feature used by the classifier to decide a country label, with IP2Location being the primary decision factor for 71% of the rules in the classifier. The IP address and DDec [6] each contribute to 6.7% and 6.4%, respectively. The reason for low contribution by DDec is because the hostnames are available for only 7% of routers in our dataset. The feature importance for *balanced* with maximum accuracy sampling approach and *unbalanced* with knee-based sampling approach was similar but not identical to the *default* classifier, whereas, for the classifier trained with *equal number of training instances per country*, the percentage was less skewed (in favor of any specific source) giving higher weights to ISP Country and ISP size.




Figure 10: **Decision tree sensitivity to number of training instances.** More than two thirds of countries require a small number of instances to achieve 90% accuracy, but some require 100s of instances to achieve high accuracy.




Figure 11: **Decision tree sensitivity to algorithm for selecting training instances.** Comparing the accuracy of the countries under different selection methods for the training dataset.




Figure 12: **Precision.** PASSPORT is precise in that the total number of countries and continents mapped by the ensemble classifier for each unique IP address tends to be one most of the time.

C. ENSEMBLE CLASSIFIER

This section provides supplementary details of decisions behind the construction of *ensemble* classifier, details on the sensitivity analysis of the classifiers, the reasoning behind the design decisions involved in the construction of the *ensemble* and the evaluation of the precision of the *ensemble*.

C.1 Classifier Selection

A simple statistic of the *overall* classifier accuracy in Table 4 is potentially biased by the number of samples per country in our dataset and is not necessarily instructive of whether the classifier offers high accuracy across a wide range of countries.

To investigate this, we compare the top-performing classifiers from Table 4 by training them with data by *randomly sampling* training data from our dataset.¹⁴ We then find the accuracy of each classifier according to the fraction of a country’s routers that are predicted correctly. We plot this in Figure 3a as a CDF of accuracy per country, where the x-axis represents accuracy and y-axis is the fraction of countries with router locations predicted correctly at least $x\%$ of the time.

For $x = 0.9$, the figure shows that Decision Tree and Random Forest achieves an accuracy of 90% or better for 35% of the countries. The 1-Nearest Neighbor algorithm achieves the same level of accuracy for only 3% of the countries and AdaBoost never achieves 90% accuracy.

Figure 3a shows that all the classifiers fail to predict 23% of the countries. In our dataset, these countries are the ones that have only one or two training instances.

To study the impact of bias, we plot Figure 3b, with the axes similar to Figure 3a, as a CDF of accuracy per country by using an *equal number of training examples per country* (i.e., by oversampling from countries with few routers and undersampling from those with large numbers of routers).

¹⁴Note that we evaluate additional training dataset selection approaches to understand classifier sensitivity in Section 5.

Decision Trees and Random Forests still outperform other classifiers, achieving an accuracy of 90% or better for 76% of the countries. However, we also found that accuracy decreased for certain countries compared to the *default* classifier, thus a single classifier does not necessarily perform well for *all* countries.

C.2 Microbenchmarks and Classifier Sensitivity

We investigate the effectiveness of the individual components of our *ensemble* classifier and their sensitivity to the training dataset under different data *sampling* techniques. We focus on the empirically derived number of training instances *for a country* that provide maximum per-country accuracy (*maximum accuracy*) and diminishing returns for increasing number of training instances (*knee*) approach from Section 4.1.3.

First, we investigate the number of instances required to achieve a per-country accuracy of 60% and 90%. The outcome depends on how we allocate training instances to other countries, so we investigate three schemes: using the same number of samples as in the target country (*balanced*), double the instances in other countries (*unbalanced*), and a random number of samples (*random*).

Figure 10 shows the minimum number of training instances required to achieve a given per-country accuracy level (x-axis), as a CDF over all countries in our dataset (y-axis). Unsurprisingly, lower accuracy thresholds require fewer training instances. Focusing on the 90% accuracy threshold, about two thirds of countries need eight or fewer training instances. This is encouraging because providing ground truth labels is a time-consuming, manual process for many router IPs. Note, however, that there is a long tail to the graph, indicating that some countries, like United States, France, Germany and Russia, need large numbers (hundreds or thousands) of instances to achieve high accuracy.

We now evaluate our strategies for automatically selecting the best instances to use for training. Recall from Sec-

tion 4.1.3 that we incorporate two schemes: maximizing per-country accuracy (*maximum accuracy*) and finding the point of diminishing returns for accuracy improvement (*knee*). Figure 11 shows the per-country accuracy for each approach using a CDF of the per-country accuracy (x-axis) over all countries in our dataset (y-axis). Overall, we find that *random* performs systematically worse (not shown), while *maximum accuracy* and *knee* approaches on both *balanced* and *unbalanced* sampling methods yield similar results.

C.3 Ensemble Construction

In constructing the ensemble, we investigate how to combine the classifiers to optimize between accuracy and precision.

Figure 11 shows the approaches to achieve a maximum gain in the accuracy per country, however, it hides the overlap for the predicted countries between different approaches.

We now analyze the impact of adding each of these classifiers along with the classifiers trained using the same number of instances, to a classifier trained on the entire training dataset (*i.e.*, the *default* classifier).

We plot a figure similar to Figure 6 to evaluate this (not shown). Our study found that at least 6 classifiers with *equal number of training instances for each country* are required to achieve convergence in the marginal increase in overall accuracy, while keeping the average number of countries predicted to a minimum, as more classifiers are added to the *default* classifier. These classifiers are trained using *mean* of the *number of training instances per country* in the dataset, since for our dataset, the mean is a better representative to have a significant number of training instances as compared to the median or the mod.

We evaluate candidates for the ensemble by combining i) 6 classifiers trained using different subsets *with equal number of instances per country*, ii) all 4 classifiers from Figure 11 and, iii) *default* classifier. This gave us a total of 11 classifiers.

To find the the minimum number of classifiers to achieve convergence in the increase in accuracy, we make *n*-length combinations of using all the sets of classifiers (all 11 classifiers).

As seen in Figure 5, we achieve diminishing returns in the accuracy as accuracy more classifiers are added. At $x = 3$, the increase in average accuracy is roughly 1.1% while it decreases to 0.3% for $x = 4$ and beyond this point the increase in accuracy is negligible.

We pick the first 4 classifiers to form our *ensemble* classifier. It contains *default*, *balanced* dataset bias with *knee* sampling approach, *unbalanced* with *maximum accuracy* approach, and one classifier with equal number of training instances for all countries.

C.4 Precision

To evaluate the question of precision of our *ensemble* classifier, we plot the distinct number of countries and

Source	Accuracy (%)
Passport	96.47
EdgeScape	90.93
IP2Location	92.28
DB-IP	86.41
Maxmind GeoLite2	85.07
Maxmind GeopIP2*	81.88
IPInfo	84.39
IPLigence*	81.38

Table 9: Accuracy of PASSPORT as compared to other geolocation sources. PASSPORT outperforms all other geolocation sources for the labeled data. * from June 2016

continents returned by the classifier for each router in our dataset, using a CDF in Figure 12. A point in each curve represents the fraction of IP addresses (y-axis) for which at least x countries were predicted (x-axis).

The graph shows that a single country is predicted for 65.7% of router IPs, while a single continent is predicted for 82.4% of router IPs. Thus, the *ensemble* alone is accurate but not precise; further, due to the spread in continents predicted, this precision is not sufficient to reason about geopolitical implications of Internet paths.

D. ACCURACY COMPARISON WITH GEOLOCATION SOURCES

This section expands on our analysis at the end of Section 5.5. We sampled and labeled the inconsistent locations between PASSPORT and geolocation services in a manner similar to Section 4.1.1. Of the 80 inconsistent locations between PASSPORT and EdgeScape, PASSPORT was correct for 60 of those, while EdgeScape was correct for 18, and both were wrong for 2 of those locations. Similarly, of the 54 inconsistencies with IP2Location labeled across 41 countries, PASSPORT was correct for 39 of those, IP2Location for 13, and none of them for 2 locations. For Maxmind-GeoLite2 database, of the 78 labels across 54 countries, PASSPORT successful for 65 and Maxmind for only 10 cases with both being wrong for 3 cases. As for DB-IP, 86 instances of inconsistent locations for 56 countries was labeled, PASSPORT correctly located 68 cases while DB-IP was successful for only 11 cases. IPInfo showed similar results as DB-IP as 78 out of 97 labeled inconsistencies were correctly predicted by PASSPORT while IPInfo was correct for only 10 cases.

When all the inconsistent labels were fed PASSPORT, it was able learn and predict 3 more consistent cases with EdgeScape, 2 more for IP2Location, DB-IP and Maxmind-GeoLite 2 and 1 more case for IPInfo.

To further establish a better understanding of PASSPORT accuracy when compared to other geolocation sources, we manually label 596 IP addresses across 519 ASes and 146 countries (and independent territories) in a manner similar to

Section 4.1.1, except we chose 1 IP address per AS instead of 5. We used these labels and found the accuracy of all the geolocation sources including PASSPORT.

Table 9 summarizes the results showing that PASSPORT had the highest accuracy with 96.47% IP addresses being labeled correctly, followed by IP2Location at 92.28% and Edgescape at 90.90%.

The results showed that our diversity of probes in different countries allowed us to locate 4.02% routers correctly where all other geolocation sources failed to predict the correct country. Moreover, for 1.51% cases, atleast one of the geolocation sources was correct while PASSPORT failed to predict the correct location.

Of all incorrect by predictions by EdgeScape, it defaulted and provided United States as the predicted country for 3.02% of the cases (one-third of failed cases). No patterns were found for other geolocation sources.

Looking at the continent-level accuracy, PASSPORT had highest accuracy for all continents except Europe. EdgeScape had the highest for Europe with EdgeScape predicting 95.35% routers correctly and PASSPORT being correct for 93.4% cases

We analyzed DDec [6] as well. The DDec-interpretable hostname information was available for only 8% of the routers. Of the routers where hostname country was available, PASSPORT had a consistency of 99.8% with DDec interpretations.

We intended to include NetAcquity as a comparison point; however, the company would not offer us access to their database at a reasonable price without publication restrictions, which we refused to accept on principle.

E. ONLINE SYSTEM

PASSPORT system is publicly available for use via a web interface and a REST API. The user interface developer documentation, and the source code, are also public at <https://passport.ccs.neu.edu>

Figure 13 shows a snapshot of the output result from the website of PASSPORT for a sample traceroute measurement. While the web interface provides a visual representation of the countries visited for a traceroute (or an IP address), the API is more powerful and has the ability to provide predictions individual results by *ensemble*, the SoL system, and the overall PASSPORT prediction.

When a traceroute is provided to website or the API of PASSPORT, it's parsed for valid router IP addresses and their corresponding RTT measurements, as shown in Figure 1b. Using the IP address, the location predictions from geolocation services is collected and this information is used to predict a set of countries by the classifier. The IP address is used to issue new ping measurements from our vantage points. These are combined with the user-provided measurements to construct the SoL constraints for the router location. The classifier predictions are then evaluated using these SoL constraints and result is returned.

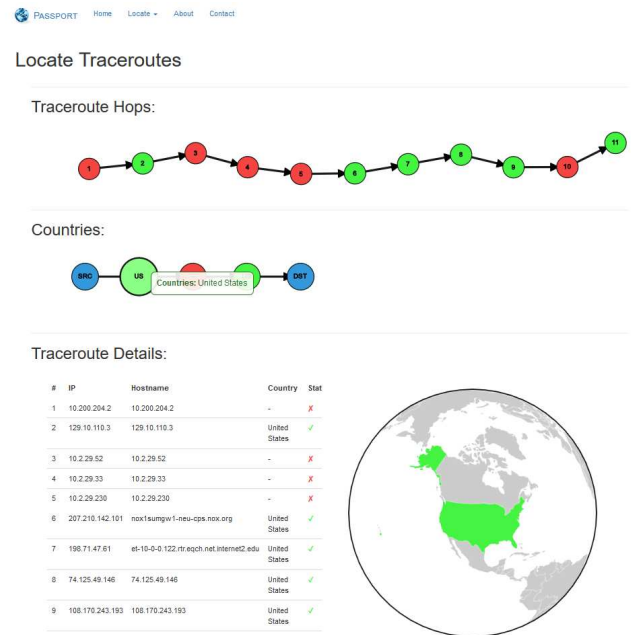


Figure 13: **Traceroute performed by PASSPORT.** The web interface provides a green circle for hops for all the hops that were located, while a yellow one representation ongoing measurement for the router and red one shows private addresses and system failures (with reasoning). For countries, the blue circles represent the source and the destination, while the map shows a visualization of the countries traversed in a path.

All this information is also stored for offline analysis and caching is performed at each step to speed up future queries with a flexible cache duration.

Since PASSPORT incorporates access to some external rate-limited APIs and shared measurement vantage points, new requests in the PASSPORT core (online system) use a job scheduler.