



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**CHASING THE UNKNOWN: A PREDICTIVE MODEL
TO DEMYSTIFY BGP COMMUNITY SEMANTICS**

by

Joshua Werner

September 2020

Thesis Advisor:

Robert Beverly

Second Reader:

Thomas J. Krenc

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC, 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2020	3. REPORT TYPE AND DATES COVERED Master's thesis	
4. TITLE AND SUBTITLE CHASING THE UNKNOWN: A PREDICTIVE MODEL TO DEMYSTIFY BGP COMMUNITY SEMANTICS			5. FUNDING NUMBERS	
6. AUTHOR(S) Joshua Werner				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) The Border Gateway Protocol (BGP) specifies an optional communities attribute for traffic engineering, route manipulation, remotely-triggered blackholing, and other services. However, communities have neither unifying semantics nor cryptographic protections and often propagate much farther than intended. Consequently, Autonomous System (AS) operators are free to define their own community values. This research is a proof-of-concept for a machine learning approach to prediction of community semantics; it attempts a quantitative measurement of semantic predictability between different AS semantic schemata. Ground-truth community semantics data were collated and manually labeled according to a unified taxonomy of community services. Various classification algorithms, including a feed-forward Multi-Layer Perceptron and a Random Forest, were used as the estimator for a One-vs-All multi-class model and trained according to a feature set engineered from this data. The best model's performance on the test set indicates as much as 89.15% of these semantics can be accurately predicted according to a proposed standard taxonomy of community services. This model was additionally applied to historical BGP data from various route collectors to estimate the taxonomic distribution of communities transiting the control plane.				
14. SUBJECT TERMS BGP, Border Gateway Protocol, routing, exterior gateway protocols, community, BGP communities, communities, anomaly, machine learning, semantics, neural network, MLP, multi-layer perceptron, random forest, multi-class classification, AS, autonomous system			15. NUMBER OF PAGES 107	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

**CHASING THE UNKNOWN: A PREDICTIVE MODEL TO DEMYSTIFY
BGP COMMUNITY SEMANTICS**

Joshua Werner
Civilian, CyberCorps – Scholarship For Service
BS, University at Buffalo, 2016
BA, University at Buffalo, 2016

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

**NAVAL POSTGRADUATE SCHOOL
September 2020**

Approved by: Robert Beverly
Advisor

Thomas J. Krenc
Second Reader

Gurminder Singh
Chair, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

The Border Gateway Protocol (BGP) specifies an optional communities attribute for traffic engineering, route manipulation, remotely-triggered blackholing, and other services. However, communities have neither unifying semantics nor cryptographic protections and often propagate much farther than intended. Consequently, Autonomous System (AS) operators are free to define their own community values. This research is a proof-of-concept for a machine learning approach to prediction of community semantics; it attempts a quantitative measurement of semantic predictability between different AS semantic schemata. Ground-truth community semantics data were collated and manually labeled according to a unified taxonomy of community services. Various classification algorithms, including a feed-forward Multi-Layer Perceptron and a Random Forest, were used as the estimator for a One-vs-All multi-class model and trained according to a feature set engineered from this data. The best model's performance on the test set indicates as much as 89.15% of these semantics can be accurately predicted according to a proposed standard taxonomy of community services. This model was additionally applied to historical BGP data from various route collectors to estimate the taxonomic distribution of communities transiting the control plane.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

1	Introduction	1
1.1	Motivation	2
1.2	Research Questions	4
1.3	Scope	5
1.4	Summary of Major Findings	5
1.5	Thesis Structure.	7
2	Background	9
2.1	The Border Gateway Protocol	9
2.2	BGP Communities.	18
2.3	Machine Learning	25
3	Methodology	29
3.1	Data Collection	29
3.2	Developing a Model	34
3.3	Estimating Taxonomic Distribution	40
4	Results	41
4.1	Analysis of Ground-Truth Data	41
4.2	Model Performance	45
4.3	Taxonomic Distribution of Communities in Recent BGP Data.	58
5	Conclusions	73
5.1	Major Findings and Implications for BGP Community Semantics	73
5.2	Future Work	76
	List of References	79
	Initial Distribution List	85

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

Figure 2.1	Valley-Free Routing	16
Figure 2.2	Communities Taxonomy	21
Figure 4.1	Center for Applied Internet Data Analysis (CAIDA) Dictionary Dataset Clustering: January 1, 2018	43
Figure 4.2	Structure of Community Semantics in Ground-Truth Data	44
Figure 4.3	Multi-Layer Perceptron (MLP) Receiver Operating Characteristic (ROC) and Precision-Recall Curves: Subclasses	49
Figure 4.4	MLP ROC and Precision-Recall Curves: Primary Classes	51
Figure 4.5	Random Forest ROC and Precision-Recall Curves: Subclasses	55
Figure 4.6	Random Forest ROC and Precision-Recall Curves: Primary Classes	57
Figure 4.7	Conservative Estimate for Taxonomic Distribution in Recent Border Gateway Protocol (BGP) Data: Training Communities Only	62
Figure 4.8	Complete Estimate for Taxonomic Distribution of BGP Communities in Recent Data	64
Figure 4.9	Community String Length Over Time: August 10, 2020	67
Figure 4.10	Community String Length Over Time: September 1, 2020	68
Figure 4.11	Instantaneous Semantic Distributions Per Vantage: August 10, 2020	70
Figure 4.12	Instantaneous Semantic Distributions Per Vantage: September 1, 2020	71

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

Table 2.1	Subnetting	10
Table 2.2	Well-Known Communities	20
Table 2.3	Distribution of Communities	23
Table 3.1	Observed Data	30
Table 3.2	Feature Set	36
Table 4.1	CAIDA Geographic Encoding Distribution: January 1, 2018	44
Table 4.2	MLP Performance: Taxonomic Subclasses	48
Table 4.3	MLP Performance: Confusion Matrix for Subclasses	48
Table 4.4	MLP Performance: Primary Taxonomic Classes	50
Table 4.5	MLP Performance: Confusion Matrix for Primary Classes	50
Table 4.6	Random Forest Performance: Taxonomic Subclasses	54
Table 4.7	Random Forest Performance: Confusion Matrix for Subclasses	54
Table 4.8	Random Forest Performance: Primary Taxonomic Classes	56
Table 4.9	Random Forest Performance: Confusion Matrix for Primary Classes	56
Table 4.10	Control Packet Distribution	59
Table 4.11	Conservative Estimate for Taxonomic Distribution of Communities in Recent BGP Data: Training Communities Only	61
Table 4.12	Complete Estimate for Taxonomic Distribution of Communities in Recent BGP Data	63

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

ACL	Access Control List
ADASYN	Adaptive Synthetic Sampling
AI	Artificial Intelligence
AMS-IX	Amsterdam Internet Exchange
ANN	Artificial Neural Network
AS	Autonomous System
AUC	Area Under the Curve
BGP	Border Gateway Protocol
C2P	Customer-to-Provider
CAIDA	Center for Applied Internet Data Analysis
CDF	Cumulative Density Function
CDN	Content Distribution Network
CIDR	Classless Inter-Domain Routing
CNN	Convolutional Neural Network
CSV	Comma-Separated Values
DoD	Department of Defense
DoS	Denial of Service
DE-CIX	Deutscher Commercial Internet Exchange
ECIX	European Commercial Internet Exchange

EGP	Exterior Gateway Protocol
EKG	Electrocardiogram
FCN	Fully Convolutional Network
HPC	High-Performance Computing
HTML	Hypertext Markup Language
IANA	Internet Assigned Numbers Authority
IDS	Intrusion Detection System
IGP	Interior Gateway Protocol
IP	Internet Protocol
ISP	Internet Service Provider
IXP	Internet Exchange Point
LINX	London Internet Exchange
LONAP	London-based Internet Exchange Point
LSTM	Long Short-Term Memory
MANRS	Mutually-Agreed Norms for Routing Security
MED	Multi-Exit Discriminator
MLP	Multi-Layer Perceptron
MOAS	Multiple Origin Autonomous System
MRAI	Minimum Route Advertisement Interval
MRT	Multi-Threaded Routing Toolkit
NCC	Network Coordination Center
NLRI	Network Layer Reachability Information

NPS	Naval Postgraduate School
NTT	Nippon Telegraph and Telephone
NYIIX	New York International Internet Exchange
P2C	Provider-to-Customer
P2P	Peer-to-Peer
PIB	Policy Information Base
RFC	Request for Comments
RFD	Route Flap Damping
RIB	Routing Information Base
RIP	Routing Information Protocol
RIPE	Réseaux Internet Protocol (IP) Européens
RIR	Regional Information Registry
RIS	Routing Information Service
RNN	Recurrent Neural Network
ROC	Receiver Operating Characteristic
RPKI	Resource Public Key Infrastructure
RT	Route Target
RTBH	Remotely-Triggered Blackholing
RTT	Round Trip Time
S2S	Sibling-to-Sibling
SMOTE	Synthetic Minority Oversampling
TCN	Temporal Convolutional Network

TCP	Transmission Control Protocol
UDP	User Datagram Protocol
USG	United States government
USN	U.S. Navy

Executive Summary

The prefix reachability information communicated by BGP advertisements is critical for the proper operation of the global internet. These advertisements make use of the optional, transitive BGP community tags specified by the protocol for a variety of purposes including traffic steering, route manipulation, geotagging, peer-type tagging, remotely-triggered blackholing, and other services at the discretion of the particular Autonomous System (AS) whence an advertisement originates and the community target for whom it is intended. However, communities have no unifying semantic, no cryptographic protections, and often propagate much farther than intended. Consequently, AS operators are free to define their own community values, and the data collected in the course of this research suggests high variability between ASes; many communities defy the proposed standard taxonomy, while others span multiple classifications or eschew the Request for Comments (RFC) standard and encode a commonly-defined value for a different purpose. Thus, community semantics are shrouded in uncertainty, and recent research has demonstrated their potential as a vector for several different types of attack.

To demystify these semantics, this research employed a machine learning approach to prediction of community semantics; it sought to quantitatively analyze the semantic predictability between different AS semantic schemata. Ground truth community semantics data were collated from existing datasets and web scraping of known AS semantics using publicly available BGP community definitions. These data were manually labeled according to a common taxonomy. As there is a distinct class imbalance between the services in the available data, various class imbalance correction techniques were utilized on the training data, including Adaptive Synthetic Sampling (ADASYN), Synthetic Minority Oversampling (SMOTE), and Random Oversampling. Various classification algorithms, including a feed-forward MLP and Random Forest, were used as estimators for a One-vs-All (binary relevance) multi-class model and trained from vectors according to a feature set engineered from this data. Of these, the Random Forest was the most successful model developed. In all cases, hyperparameters, including layer sizing, were tuned using a grid search.

Model performance on the test set indicates varying degrees of predictability between classes, with overall accuracy indicating as much as 90.64% of community semantics can

be accurately predicted by such a model. Additionally, this model was applied through a Multi-Threaded Routing Toolkit (MRT) parser to recent BGP data at various vantages to estimate the taxonomic distribution of communities transiting the control plane.

Given the uncertainty in community semantics, and their potential for misuse, this research seeks to inform future work in detection of anomalous community usage to deter, detect, and potentially filter community-based threats or misconfiguration events. It further seeks to bring attention to the variability of semantic encodings between different ASes, especially those which contradict RFC standards or span multiple classifications.

Acknowledgments

I would like to thank my advisors, Dr. Robert Beverly and Dr. Thomas Krenc, for their constant support and enthusiasm for this research. Their knowledge in the field of networking and computer security has been inspiring, and their patient guidance as I developed my understanding of these subjects was always greatly appreciated.

I would also like to thank my family for their love, support, and encouragement through my time at the Naval Postgraduate School. Finally, I would like to thank the Scholarship for Service program for giving me this opportunity.

This material is based upon activities supported by the National Science Foundation under Agreement No 1565443. Any opinions, findings, and conclusions or recommendations expressed are those of the author and do not necessarily reflect the views of the National Science Foundation.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:

Introduction

The Border Gateway Protocol (BGP) is the *de facto* standard for inter-domain routing in the Internet today. Its primary purpose is to communicate network prefix reachability information between Autonomous Systems (ASes). It prescribes a specific path of ASes through which traffic may travel to reach a particular prefix. For this reason, BGP represents a “hop-by-hop” paradigm to routing between ASes in which the path traffic takes is dictated by its transfer from one AS to the next. This routing information can be used to map connectivity between ASes, detect routing loops, and, crucially, to implement policy decisions [1].

Network administrators at any given AS can thus utilize BGP to dictate a consistent policy for routing through their domain. Unfortunately, as the internet has grown and network complexity has increased exponentially, routing policy requirements have become increasingly complex and fine-grained [2], [3].

To help implement such policies and facilitate communication between ASes, BGP “communities” were introduced in 1996 by Request for Comments (RFC) 1997 [4]. These communities are encoded as 32-bit numbers in a variable-length attribute, allowing multiple communities to be applied to a particular route. The original intent of BGP communities was to “facilitate and simplify the control of routing information” [4] by grouping destinations such that routing decisions can be applied at a group level. A community is thus a group of destinations which share some common property, and to which a common routing policy can be applied.

This proposal further defined well-known communities and specified a standard format for their interpretation, to wit, that the first two octets (the most significant word) should specify an AS number, and the last two octets (the least significant word) should specify a community value defined by that AS. For example, AS 2914 is free to specify community values in the range `0x0B620000` to `0x0B62FFFF`, such that the first two octets always represent 2914 (`0x0B62`), while the latter two octets can encode any community defined by the network operators of AS 2914 according to any schema they devise [4]. That is to say, AS operators are free to define a value in the last two octets—a range that contains 2^{16}

possible values—as a signal or tag for a particular community service. This system can be used both for active signalling and passive encoding of information [5].

This has created extreme variation in community definitions between ASes. For example, the value 666 is used for blackholing services by AS 2914; AS 3356 defines the same value as a route tag to indicate that the route was learned from a peer, rather than a customer [6], [7]. Many communities are intended for internal use only, and thus their definitions are known only to that AS’s operators. Moreover, many ASes do not publicly disclose their community definitions except to their own customers, further shrouding the semantic space in mystery.

The primary intent of this thesis is to pierce this mystery—to the extent possible—by quantitatively analyzing community semantics. To this end, this research seeks to build a predictive model capable of accurately classifying BGP communities according to an established taxonomy of community services proposed by [8]; this taxonomy identifies the three major categories and ten subcategories of passive and active community semantics used among ASes. Our research seeks to understand how much of the community value space can be identified according to this taxonomy through common features between AS-defined schemata. It further seeks to apply the best model created to recent BGP data to profile the taxonomic distribution of communities in the wild.

1.1 Motivation

BGP communities are of interest to the research community because they represent a potential attack vector; identifying “anomalous” communities—caused either through mis-configuration or malicious intent—is a challenging problem, particularly considering their lack of unifying semantics. Additionally, communities which propagate farther than their intended target may publicly leak information about the configuration and activities of a particular AS. Yet, their adoption and use is increasing.

In fact, the number of ASes and prefixes has grown considerably since the introduction of BGP communities in 1996. Either as a result of this, broader adoption of communities generally, or some combination thereof, the number of unique communities observed at route collectors has increased as well. Between 2010 and 2018, the number of unique communities increased by 296% [9].

Indeed, AS operators use BGP communities extensively to control routing policy within their domain, and even offer value-enhancing services to their customers [3], [5]. They offer a simple and convenient mechanism to engineer traffic, manage AS policy, and even mitigate attacks. This includes Denial of Service (DoS) deterrence via Remotely-Triggered Blackholing (RTBH), announcement filtering, local preference adjustments, influencing peer selection through path prepending, selective advertisement, and tagging routes by geographic ingress points and origins. In fact, the amount of information they signal is increasingly used to encode other information, such as round trip times (RTTs) [5].

Consequently, AS operators configure their infrastructure to take different actions depending on community tags [5], but these actions are implementation-dependent. Furthermore, vendor implementations sometimes vary in the way well-known communities are processed and manipulated, sometimes resulting in inconsistent behaviors that are difficult to identify and resolve [10].

At the same time, the BGP lacks built-in cryptographic protections, allowing any BGP router to announce any arbitrary route, with arbitrary community values [11]. Although secure versions of the protocol, such as BGPsec, have been proposed, it has thus far proved impractical to deploy due to its computational complexity. An alternate scheme of prefix filtering has been proposed but is difficult to implement due to the difficulty in designing effective filters [12]. To make matters worse, even these solutions offer no protections for optional attributes such as communities [5].

Similarly, BGP often suffers from misconfiguration events, particularly origin and export misconfiguration, which may become globally visible and contribute to global connectivity issues [13]. Understanding the community string in such events may be of use to operators attempting to pinpoint their cause, malicious or otherwise, and additionally aid in the creation of intelligent filters. When communities leak in such events, or otherwise propagate beyond their intended targets, they may also leak information about the activities and configuration of a specific AS.

More importantly, BGP communities have been demonstrated to be a viable attack vector. According to the proposed standard, an AS should scrub inbound communities used internally—i.e., with their AS in the first two octets—but forward any foreign communities, as these may be needed for customers to communicate with upstream providers [14].

However, if communities propagate farther than intended, they may trigger effects multiple hops beyond the direct peer to which they were originally announced, and there is no way of knowing if such behavior is intended [9].

In fact, prior research has demonstrated that they are often propagated farther than a single routing hop, at least partially due to poor understanding of their semantics. As much as 14% of transit providers propagate received communities to their peers. As the Internet's topology continues to flatten, this implies communities are continually being propagated widely through the internet [5], [9].

This enables a number of attack scenarios, including manipulation of communities for malicious interception of traffic, imposition of additional cost for a target AS, and impairment (including denial of service) or improvement of network performance. This can be mediated through manipulation of community services, optionally in conjunction with prefix hijacking, that control RTBH, path prepending, and local preference [5].

To better understand how to build effective detectors for such attacks, a researcher or operator must first be able to understand the communities an AS is propagating, particularly foreign communities. This is the primary motivation for this thesis.

1.2 Research Questions

The primary research questions this thesis seeks to answer are:

1. Is there structure to the assigned communities values within an AS such that the meaning of an unknown community can be accurately predicted?
2. Given that BGP community semantics are AS-specific, do their semantic schemata present any common features or occupy similar ranges according to the guidelines set forth by [4]? That is, is there common structure to community values between different ASes such that the meaning of community value not previously observed for a given AS can be accurately predicted?
3. Are there features that allow the meaning of a community value that has never been observed in the past to be accurately predicted?
4. Can a linear or non-linear function be used as an effective estimator for the classification boundaries between the different BGP communities according to a unifying

- taxonomy? Which technique will prove most efficacious for this purpose?
5. According to this model, assuming it can obtain reasonable accuracy, what percentage of community attribute usage in recent BGP data is attributable to policy decisions? What percentage is used for traffic engineering? What percentage is used for DoS prevention?

1.3 Scope

This thesis is limited to research into the semantics of 32-bit BGP communities. While extended and large communities have been defined, principally because many ASes now have AS numbers greater than will fit in the most significant word of a 32-bit community, their adoption is generally low compared to their 32-bit predecessors. This is, however, likely to change in the future [15]. Furthermore, this research limits itself to consideration of Internet Protocol (IP) version 4 prefixes in the context of BGP.

1.4 Summary of Major Findings

The principal findings of this research are summarized below. Chapter 4 presents these findings in greater detail:

1. The unified taxonomy proposed by [8] is broadly applicable according to our findings, however, its use as a classifying scheme is limited by the presence of communities which span multiple categories (i.e., communities which are effectively “multi-class”) and the fact that some communities have different meanings when applied alone or together.
2. A small number of communities—approximately 0.25% of observed data—do not fit into any particular subcategory of this taxonomy.
3. Many ASes use a nonstandard syntax with respect to RFC 1997 [4]; to offer specific Outbound services to a particular AS, many ASes use the most significant word to encode the service (e.g., the number of prepends) and the least significant word to specify a particular AS to which that behavior should apply. This demonstrates a fundamental limitation of the standard syntax. 118 such communities are present in collected data.

4. This research develops a feed-forward Multi-Layer Perceptron (MLP)-based multi-class model; although it was able to obtain high accuracy—see Chapter 4—it was not as successful as a Random Forest-based model.
5. The most successful model developed achieves 90.64% aggregate accuracy in classifying eight subcategories of the standard taxonomy when presented with novel testing data. Further, this model is able to categorize the principal three categories with 97.49% accuracy. This indicates that, at least among the ASes studied, there is enough information to predict community values that have not been seen before in the majority of cases.
6. The prediction accuracy varies between classes, as measured by F_1 score; low F_1 score in a given subcategory may be due to the inherent class imbalance observed in the data, and thus inability of the model to generalize from relatively few samples. In particular, this affects the Inbound: AS subcategory. Inbound: Local Preference and Blackhole were the classes with the next-lowest F_1 score. It is unclear if this indicates greater variance between ASes for specific subcategories, or if more data would increase predictive accuracy.
7. The BGP community string is a variable-length attribute; a given route announcement may contain an arbitrary number of communities in its community string. A Cumulative Density Function (CDF) for the length of this string in recent BGP data suggests that 83% of community strings include eight or fewer communities; however, community strings with 30-40 communities were observed to comprise as much as 2% of communities collected at these vantages.
8. Application of the best model to recent BGP data at three route collectors as Réseaux IP Européens (RIPE) indicates that the majority (between 94% and 98%) encode passive signaling semantics, particularly Inbound: Geographic tags. Communities encoding active Outbound semantics, such as path prepending and announcement policy, were predicted to compose between 1.08% and 2.46% of data over the two capture intervals during which BGP control packets were studied. This was an unexpected result, and may indicate a period of relative stability in the global routing tables as seen from the three vantages under study. A significant difference between capture intervals was an increase in the proportion of Blackhole communities observed, from 0.49% to 2.04%. According to our model's prediction of instantaneous distribution of community semantics, this may have arisen from RTBH events.

1.5 Thesis Structure

The remainder of this thesis will adhere to the following outline:

1. **Chapter 2** provides background information on the BGP and the common uses of the BGP communities attribute as exemplified by the taxonomy described in [8], as well as the prior work which fueled this research. It additionally introduces the machine learning methods used in the creation and selection of the predictive models created by this research, and the statistical methods in support thereof.
2. **Chapter 3** explains the data collection process, the data manipulation needed for the training set, model selection, and the verification process used. This chapter also discusses the creation of the feature set used by the model, as well as various limitations and extreme outliers discovered during the process.
3. **Chapter 4** presents the key findings of this research, with particular regard to the accuracy of the model as described by precision, recall, and F_1 score over each class, a confusion matrix, the accuracy of the model in aggregate, Receiver Operating Characteristic (ROC) curves, and Precision-Recall curves broken down by class. This chapter also presents the model's estimation of taxonomic distribution in recent BGP data.
4. **Chapter 5** offers an analysis of the key findings of this research and discusses their implications for the semantics of BGP communities and community-based threats. Furthermore, it suggests future work that can build upon this research in pursuit of a detector for anomalous community usage.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 2: Background

Routing between domains functionally enables the internet today, but our current standard, the Border Gateway Protocol, is imperfect and thus open to abuse. This chapter will introduce the foundational concepts underlying inter-domain routing generally, the role of BGP communities in this process, and the machine learning concepts applied by this research to study them.

2.1 The Border Gateway Protocol

The internet is composed of many ASes, a generic term describing any entity that administers a collection of routers within a particular domain. An AS could be an Internet Service Provider (ISP), a Content Distribution Network (CDN), an Internet Exchange Point (IXP), a university, a research network, a factory, etc. Each AS is assigned an AS number: the University of California, San Diego, for example, is AS 7377, while the ISP Comcast is AS 7922 [16]. These entities apply routing policies within their domain through an Interior Gateway Protocol (IGP) (or IGPs) and some set(s) of common metrics, and use an Exterior Gateway Protocol (EGP) to direct traffic that must be sent to another AS [17].

The BGP is the standard EGP in use in the internet today. It provides a mechanism in support of Classless Inter-Domain Routing (CIDR) by announcing reachability information for networks within a domain. More specifically, an AS router, called a speaker, will announce a set of destinations that can be reached on its domain as an IP prefix [1]. This is the mechanism by which ASes learn how to route traffic on their networks to a destination prefix in a foreign AS.

2.1.1 Prefixes

Prefixes in CIDR conveniently break up the 32-bit¹ IP address space by designating far greater variety in network sizes than the classful system of the early internet. In this earlier

¹The 32-bit address space is governed by the IPv4, but in 1995 it was already known this space would eventually face the issue of address exhaustion. CIDR can ameliorate but not solve this problem. 128-bit addresses, governed by the then-new IPv6 protocol, were introduced in 1998 [18], [19].

implementation, the IP address space was divided into five address classes of fixed size.

CIDR makes networks classless by dividing them using subnetting. Whereas in the classful implementation, the number of network bits may be only 8, 16, or 24 (classes A, B, and C, respectively), subnetting allows any number of the bits in the 32-bit space to be network bits². See Table 2.1 for more detail. It does this by application of a bitwise mask called a subnet mask; in CIDR notation, an IP address with an associated subnet mask is called a prefix [20].

As a practical example, the prefix 10.128.240.50/30 indicates that the first 30 bits are network bits. Thus, its subnet mask is 255.255.255.252, which means this subnet can effectively address two hosts. The network address is the result of a bitwise AND operation between the subnet mask and the IP address, resulting in 10.128.240.48. Thus, two hosts can be addressed at 10.128.240.49 and 10.128.240.50. Finally, its broadcast address is 10.128.240.51.

Table 2.1. Subnetting allows subdivision of networks, ameliorating the problem of address exhaustion by obviating the need to upgrade to a class B network—65,534 possible hosts—to effectively address 255 hosts, which is beyond the capacity of a class C network. Through subnetting, a class B can be divided as below for more granular network size options [20].

Network Bits	Subnet Mask	Number of Subnets	Number of Hosts
/16	255.255.0.0	0	65,534
/17	255.255.128.0	0	32,766
/18	255.255.192.0	2	16,382
/19	255.255.224.0	6	8,190
...			
/28	255.255.255.240	4,094	14
/29	255.255.255.248	8,190	6
/30	255.255.255.252	16,382	2

For this thesis, it should be noted that many BGP communities (Section 2.2), particularly blackholing communities, can only be applied to an announced prefix of particular size. For

²In practice, subnets commonly have between 8 (equivalent to a class A network) and 30 (capable of addressing two hosts) network bits [20].

example, AS 7922 defines the community 7922:666, a blackholing community which can only be applied to a prefix with 32 network bits (i.e., a single host) [6]; this is typical among blackholing communities. In general, blackholed prefixes are as specific as possible to allow normal service to non-affected hosts. Blackholing services will be negotiated between AS operators [21]. Similarly, AS 209 defines the community 209:999, an Outbound: Announcement community specific to /24 or greater routes.

Thus, a prefix allows BGP to communicate the destinations available on a particular AS; a particular BGP speaker advertises that a prefix is in its domain, and the ASes that hear this announcement know where traffic destined for a host on that prefix must go. To get there, however, BGP routers must know a path through the internet topology.

2.1.2 Autonomous System Path

After a BGP speaker announces a prefix is reachable on its domain, the news propagates to other ASes. To effectively route traffic to that prefix, downstream ASes must know how to reach that domain in the global topology of AS interconnections. In BGP, this is communicated by the mandatory “AS_PATH” and “ORIGIN” attributes of the protocol. When making an announcement, a given AS will set the “ORIGIN” and “AS_PATH” to its own AS number. When an external peer receives it, the peer prepends its own AS number to the “AS_PATH” attribute before sending it to the next external peer. This allows BGP to detect and prevent routing loops.

For example, if a speaker at AS 7322 advertises prefix 198.51.100.0/24 as reachable in its domain, it will set the “ORIGIN” and “AS_PATH” to 7322 and send the announcement to a peer AS; e.g., AS 2914. AS 2914 will not alter the “ORIGIN” attribute, but will prepend its own AS number to the “AS_PATH” attribute so it now equals the sequence 2914 7322. ASes along an announcement’s path will continue to do so, such that an AS arbitrarily-many hops away knows it needs first to traverse the ASes whose paths were prepended before it can reach the “ORIGIN” AS rightmost in the “AS_PATH” [1].

If an announcement reaches an AS with a path value equal to 3320 5459 2914 7322, then that AS knows a host within its domain trying to reach a web server at 198.51.100.21 can reach the AS where that host is located by sending its traffic first to AS 3320. AS 3320 can then forward this traffic to AS 5459, and so on until it reaches AS 7322. However, a real AS

often has several paths from which to choose, and their numbers are continually changing as paths are withdrawn or as policy dictates.

Indeed, choice of a particular path is a multi-faceted problem in BGP, not least because in the real world links are continually going down, prefixes become unreachable on a given AS, the internet topology shifts, and so on. An important metric used in path selection is the so-called “AS_PATH”; see Section 2.1.5 for more detail. Furthermore, the “AS_PATH” is used extensively for traffic engineering; for example, an AS may prepend its AS number multiple times to the path so routes through it are de-preferred. This provides a mechanism by which to align routing policies with economic constraints and preferences [1], [5].

2.1.3 Peering Sessions

BGP sessions are Transmission Control Protocol (TCP) connections between BGP routers on which BGP messages are passed. Such sessions consist of the exchange of four different types of control packets:

1. OPEN
2. KEEPALIVE
3. NOTIFICATION
4. UPDATE

To initiate a peering session, a speaker sends an OPEN packet and waits for a peer to respond in kind. In these OPEN messages, the BGP routers exchange desired values for the hold time³; both routers will use the lower value. A KEEPALIVE interval is established as, at maximum, a third of the hold time. The speaker responds with a KEEPALIVE packet, and the session is officially established.

UPDATE messages are now sent which contain a set of path attributes and Network Layer Reachability Information (NLRI), namely, the prefixes that are advertised as reachable. Previously advertised prefixes that are no longer reachable will be sent as a withdrawn route. As the BGP routers learn routes through these peering sessions, they will update their internal Routing Information Base (RIB), which stores information about inbound UPDATES, local

³The hold time is the amount of time a peer will wait for incoming BGP messages before assuming the peer has gone down. A peer can reject a hold time and thus refuse to establish the session [1], but this never happens in practice [22].

routing information, and outbound routes the router will send as an UPDATE to its peer.

If a BGP router in such a session has no UPDATE messages to send, it will send a KEEPALIVE message to inform its peer that it is still active and that the session is still live. As a form of error handling - if, for example, a peer receives an unacceptable parameter - NOTIFICATION messages may be sent and the peering attempt will fail or the session will close.

It is additionally important to note that BGP sessions may be to external peers, called eBGP, or internal to routers within an AS, called iBGP. The latter type is a means to forward eBGP advertisements through a network [1], [22].

2.1.4 Path Exploration

The BGP operates in real time, and once a peering session is established, new reachability information propagates via UPDATE messages. ASes that receive this new route must include it in their decision process and, if the route is now preferred, update downstream ASes. Due to the interconnected nature of the topology, combined with natural delays in processing and propagating new reachability information, many messages may be exchanged as better paths are iteratively discovered. This message exchange is known as “path convergence” and continues until the network has converged to adopt the new information.

To reduce instability, the BGP has a Minimum Route Advertisement Interval (MRAI) with a suggested value of 30 seconds: if, for example, there is an UPDATE for the same prefix every five seconds, nothing is sent for that prefix until the interval elapses. This ensures only the most recent information is sent to peers. Additionally, BGP implements a Route Flap Damping (RFD) mechanism to penalize BGP speakers which announce a large number of updates in a short interval, called “flapping.” Routes are suppressed when the penalty for this exceeds a set threshold; if the flapping behavior stops, the route will eventually be un-suppressed over time [23].

However, when a route is withdrawn because a prefix is no longer reachable, BGP routers will enter a phase of path exploration in which they search for longer and longer paths. In the Routing Information Protocol (RIP), this results in the “count to infinity” problem,

wherein paths continue to grow infinitely⁴ By contrast, in BGP, the MRAI allows routers time to withdraw these routes and the “AS_PATH” attribute allows BGP to immediately ignore paths with loops. Therefore, while this bad news will travel more slowly as BGP continues to explore longer paths, it will not do so infinitely [1], [24].

2.1.5 Path Selection

Path selection in the BGP is a multi-faceted problem, and is accomplished through the BGP consecutive decision process. The results of the process update the contents of the local RIB. Crucially, each AS may apply their own criteria for evaluating routes.

Generally, the most important factors in the decision process - which can vary by implementation - are the values of the following:

1. **The Path Length (“AS_PATH”)** is how BGP distinguishes shorter paths; note that this may be influenced by route manipulation due to path prepending. Prepending communities are often used to influence this value. Paths with detected loops are immediately ignored.
2. **Local Preference (“LOCAL_PREF”)** is an attribute which propagates only in internal BGP sessions; its default value is 100. Lower values are less preferred. It is commonly used for route redistribution for business purposes; see Section 2.1.6.
3. **The Multi-Exit Discriminator (MED)** provides a mechanism to manipulate another AS to select a certain route when there are multiple entry points; lower values are preferred. Some ASes use communities to influence this value [6].
4. **The IGP Metric** is used by operators that administer multiple contiguous but independent ASes and want to utilize a minimizing distance metric for path selection similar to the way an IGP works⁵.

The algorithm that evaluates these metrics, among other factors, is implementation-dependent. In general, a BGP router applies its local Policy Information Base (PIB) to inbound routes stored in its RIB. The results of this process are the routes that it will advertise to peers and are stored in the outbound routes section of its RIB [1], [25]–[28].

⁴In RIP, fortunately “infinity” is equal to 16, or one more than the number of hops it allows [24].

⁵BGP cannot generally use such a metric, as it would create significant scalability and coordination problems [25].

2.1.6 Economic Factors

Routing in BGP is heavily influenced by economic factors. Indeed, as ASes are heterogeneous by nature, routes exchanged between them will generally reflect the dynamics in their relationship through application of specific policies that affect path selection, such as path prepending to steer traffic and changes to parameters used as input to the consecutive BGP decision process [1], [29]. The “LOCAL_PREF” attribute is commonly used for this purpose [27], and is thus also a common BGP community target in signaling semantics [3], [8].

Economic relationships between ASes can be broadly categorized into four principal types:

1. Customer-to-Provider (C2P)
2. Provider-to-Customer (P2C)
3. Peer-to-Peer (P2P)
4. Sibling-to-Sibling (S2S)

C2P and P2C relationships are characterized by one AS, the customer, paying the provider to obtain connectivity to the larger internet. A P2P relationship is one in which two ASes have agreed to exchange traffic between their customers, usually without charging the other. This is called “settlement-free,” meaning that each AS merely retains the revenue from their own customers. Finally, S2S, a pair of ASes, siblings, obtain connectivity to the rest of the internet through one another. Typically, sibling ASes belong to the same organization. For this reason, a provider generally has a greater node degree⁶ than its customers. Similarly, any two peers are typically of comparable degree [29], [30].

Consequently, these relationships are a large factor in routing decisions. This is possible because a multi-homed AS can refuse to act as a transit AS for other ASes, or for a restricted set of its peers, and preferentially select certain ASes through which to route traffic [31]. This therefore suggests that, absent route leakage or BGP misconfiguration, AS operators will tend to implement policies which create “valley-free” routes. That is, once traffic has traversed a P2P or P2C link, it should not traverse a C2P or P2P link. Figure 2.1 displays an example of the dynamics of inter-AS relationships.

⁶The degree of a node is the number of edges incident upon it; that is, the number of connections an AS has to its neighbors.




Figure 2.1. If traffic transits from Customer A to ISP A, policies will tend to enforce a traffic flow that reflects the economic reality. Traffic should transit only through the peering link to ISP B, thence to Customer B, or from ISP A through the Transit AS to ISP B before reaching Customer B. Traversing through the Dual-Homed Customer creates an undesirable routing “valley”; the Dual-Homed Customer is paying for traffic through its link to ISP B, and therefore wants to avoid sending another AS’s traffic through their link.

2.1.7 BGP and Security

Security in BGP is a chief concern among researchers. RFC 7454 [14] states that BGP speakers should apply an Access Control List (ACL) to disregard all traffic to TCP port 179 from an unknown or forbidden peer. However, as the current implementation of the protocol lacks intrinsic cryptographic protections, this is insufficient. It is thus susceptible to abuse of both the confidentiality and integrity of a particular UPDATE.

Furthermore, BGP has featured in a number of attacks and misconfiguration events that have affected traffic globally. In 1997, as the result of a misconfiguration, AS 7007 announced itself as the origin of the best path to most of the internet, disrupting reachability for several hours. In 2008, Pakistan Telecom (AS 17557) announced a path to the prefix 208.65.153.0/24; its provider propagated this fake announcement throughout the internet. This hijacked global traffic to YouTube for over two hours [11], [12].

Broadly, a hacker’s objectives may include prefix blackholing, traffic redirection, traffic subversion, or creation of instability in the global routing tables [11]. Mitseva et al. [32]

identified the following principle attack types in BGP:

1. **Data Falsification Attacks:** A malicious AS may inject false data into BGP messages.

Vectors include:

- (a) **Prefix Hijacking:** This is an attack whereby an AS falsely claims to originate a prefix for which it is not responsible. This can lead to a Multiple Origin Autonomous System (MOAS) conflict. Attackers can avoid telltale MOAS conflicts by announcing prefixes that are part of the 20% of the global prefix space which is not announced.
 - (b) **Subprefix Hijacking:** Attackers can also avoid MOAS conflicts by announcing a subnetwork of an existing prefix that it has not been delegated. This is also called a de-aggregation attack.
 - (c) **AS Path Forgery:** In such an attack, attackers alter the “AS_PATH” to announce a fake link between their AS and the victim AS to avoid a MOAS conflict and induce a “one-hop” prefix hijack. This can also be done with subprefixes.
 - (d) **Interception Attack:** If an attacker has a valid route to the victim’s AS, s/he can induce traffic to be intercepted before reaching its intended destination without disturbing connectivity.
 - (e) **Suppression/Replay Attack:** A malicious AS can replay or suppress withdrawal of a route that has previously been announced.
 - (f) **Collision Attack:** This attack occurs when two colluding, non-neighbor ASes create a virtual tunnel and run a BGP session through it to announce forged routes without risk of a MOAS.
2. **Protocol Manipulation Attacks:** These are attacks in which the malicious AS seeks to manipulate properties of the BGP itself. Vectors include:
- (a) **MED modification:** a malicious AS may tamper with parameters that are used as input in the BGP consecutive decision process.
 - (b) **Exploiting the RFD and MRAI Timer:** Malicious ASes may artificially withdraw a route and subsequently re-announce it. AS see this as flapping and suppress it; meanwhile, the MRAI timer delays distribution of UPDATES, making this route seem unreachable. See Section 2.1.4 for more detail on RFD and MRAI.
3. **Data Misuse Attacks:** This is a type of attack in which ASes use correct routing data for a malicious purpose. Vectors include:

- (a) DoS: Attackers may induce heavy congestion on routes carrying BGP traffic, causing peering sessions to fail. After sessions are reestablished, the routers need to exchange full routing tables, increasing load on the devices and additionally causing convergence delays.
- (b) Route Leakage: This attacks occurs when an AS propagates routes to ASes that were not intended to receive such routes under the terms of economic agreements between them, violating valley-free export rules (see Section 2.1.6).

Thus, attacks on or misconfigurations of the BGP induce significant connectivity problems in the internet, leading to denial of service, instability, or interception of traffic. For this reason, proposals have been made for a secure version of the protocol. The most promising is BGPsec, but thus far it has proved too computationally expensive to deploy in the real world [12]. By contrast, Resource Public Key Infrastructure (RPKI), which provides a mechanism by which to cryptographically validate—via certificates—association between a specific AS numbers or prefixes with the holders of those number resources, and is slowly being deployed [33]. Additionally, more than 100 ISPs have agreed to the Mutually-Agreed Norms for Routing Security (MANRS) [34].

Unfortunately, all implementations of secure versions of the BGP fail to enforce a protection mechanism for BGP communities, which Streibelt et al. [5] recently demonstrated is an additional source of misconfiguration events and a potential attack vector.

2.2 BGP Communities

In the 90s, the internet was growing explosively, and requirements for inter-domain routing were becoming increasingly complex. The BGP communities attribute, first specified in 1996, is intended to provide a mechanism to help implement such requirements by offering a means to effect fine-grained control of routing policy to an aggregate of routes [2].

Communities are particularly useful for applying routing policies to groups of destinations which share a common property, hence their name. This allows them to encode passive semantics about a route, such as its geographic origin, as well as active semantics such as a signal to blackhole traffic for a prefix.

The community attribute is an optional, transitive field of variable length; this means, unlike

some other attributes, the communities attached to a route announced by a speaker may be passed on from one AS to the next. In fact, this is often required so customers can signal upstream providers [4]. However, communities are generally intended for use between two direct AS neighbors; they are not supposed to propagate widely in the internet, but recent investigations have revealed this is often the case. In fact, as many as 14% of transit ASes forward communities that they receive, potentially inducing unknown effects in ASes far beyond the intended target [5], [9]. Since the field is of variable length, an announcement often contains multiple communities values.

Unfortunately, community semantics are very poorly defined. This is a major contributing factor to their improper propagation in the internet. RFC 1997, which originally specified the attribute, defines a common syntax for 32-bit communities, whereby the first two octets designate an AS number and the second two octets represent a specific community value.

It further defines three so-called “well-known” communities; that is, communities of global significance. As of September 24, 2019, the Internet Assigned Numbers Authority (IANA) recognizes 14 well-known communities proposed in subsequent RFCs [4], [5], [35]:

Table 2.2. RFC 1997 defines three well-known communities and a common syntax for reading them and defining new ones. As of September 24, 2019, the IANA recognizes the following 14 well-known communities. Source: [4], [10], [35]–[39].

Purpose	Hex Value	Standard Syntax	Meaning
GRACEFUL_SHUTDOWN	0xFFFF0000	65535:0	Reduce traffic loss: inform peer of planned maintenance.
ACCEPT_OWN	0xFFFF0001	65535:1	Accept routes with router's same ORIGIN AS (for route reflectors).
ROUTE_FILTER_TRANSLATED_v4	0xFFFF0002	65535:2	Attach translated Route Targets (RTs) for VPNv4 route filtering.
ROUTE_FILTER_v4	0xFFFF0003	65535:3	Attach RTs for VPNv4 route filtering as-is.
ROUTE_FILTER_TRANSLATED_v6	0xFFFF0004	65535:4	Attach translated RTs for VPNv6 route filtering.
ROUTE_FILTER_v6	0xFFFF0005	65535:5	Attach RTs for VPNv6 route filtering as-is.
LLGR_STALE	0xFFFF0006	65535:6	Mark stale routes retained (long-lived graceful restart).
NO_LLGR	0xFFFF0007	65535:7	This route not suitable for long-lived graceful restart.
ACCEPT_OWN_NEXTHOP	0xFFFF0008	65535:8	Accept routes with router's same NEXT_HOP IP address (for route reflectors).
BLACKHOLE	0xFFFF029A	65535:666	Drop any traffic to this prefix.
NO_EXPORT	0xFFFFF001	65535:65281	This route should not leave the BGP confederation.
NO_ADVERTISE	0xFFFFF002	65535:65282	Do not advertise this route to other peers.
NO_EXPORT_SUBCONFED	0xFFFFF003	65535:65283	Do not advertise this route to external peers, including those in confederation.
NOPEER	0xFFFFF004	65535:65284	Do not advertise this prefix in a bilateral peering session.

Aside from the communities listed in Table 2.2 and several reserved ranges of values, AS operators are free to define their own communities according to the common syntax defined in RFC 1997⁷. This allows them to encode passive and active semantics in the final two octets of the field. This means they are free to designate any value in a space equal to 2^{16} as a particular community service [4]. For example, Level 3 (AS 3356) allows customers to set local preference to 70 with the community 3356:70 [6].

This has produced variable schemata for community definitions between ASes. Additionally, many ASes do not publish their community definitions publicly; they are available only to their customers [5], and the number of unique communities almost tripled between 2010 and 2018 [9], exacerbating the problem. This research is an attempt to quantitatively study these semantics, to the extent possible.

2.2.1 Taxonomy

To better categorize BGP community semantics, Bonaventure et al. [8] developed the first unified taxonomy for communities. This taxonomy is displayed in Figure 2.2:




Figure 2.2. This first community taxonomy, developed in 2008, divides communities by their point of application or when they are added: i.e., when a route is received (Inbound), for routes that are or are not advertised to other peers (Outbound), and when traffic should be dropped to the advertised prefix altogether (Blackhole). Adapted from [8].

The Bonaventure taxonomy designates three principal categories of communities based on where they are relevant in BGP sessions. These are subdivided based on the most common meanings these communities encode or the services they signal:

⁷This research has in fact found that the semantics of many ASes are in conflict with this paradigm.

1. **Inbound:** Inbound communities are added or used by a router when a route is received by a router during an eBGP session.
 - (a) Local Pref: These communities are used to influence the “LOCAL_PREF” attribute, as has been discussed in Sections 2.1.5 and 2.1.6 in more detail. It is an important factor in the BGP consecutive decision process. Thus, it is an example of an active community semantic. For example, the community 174:10 tells AS 174 to set the local preference for that route to 10.
 - (b) Route Tagging: These communities encode passive information about a route; they are “tagged” with this information.
 - i. IXP: indicates the route was learned at an IXP; e.g., 4589:14901 informs AS 4589 that it learned the route at the IXP Deutscher Commercial Internet Exchange (DE-CIX).
 - ii. Type of Peer: indicates the relationship the receiving AS has with the peer; e.g., 3356:123 tells AS 3356 that it learned that route from a customer [6]. This is to help distinguish the relationship between ASes for learned routes.
 - iii. Geographic: indicates the geographic origin where the route was received; e.g., 3356:500 tells AS 3356 that the route was received in the United Kingdom.
 - iv. AS: indicates the AS from which a route was learned; e.g., 15997:1080 is used by AS 15997 to mark a route from Tata Communications.
2. **Blackhole:** Blackhole communities are designed to tell the receiving AS to drop all traffic to the prefix it advertises; these are thus an active semantic. The taxonomy subdivides this by the scope and locality of the effect.
 - (a) Global Blackhole: indicates that traffic is blocked everywhere within the receiving peer (generally an ISP).
 - (b) Border Blackhole: indicates that traffic should be blocked at the peer’s border routers.
 - (c) Upstream Blackhole: indicates that traffic be blocked by the peer’s upstream transit provider before it enters the peer’s network.
3. **Outbound:** Outbound communities are designed for traffic engineering. That is, they have active semantics. In particular, they are used for redistribution of routes at routers downstream.
 - (a) Announcement: marks a route that should not be announced, typically to a

specific peer or group of peers. For example, 2914:4029 tells AS 2914 not to advertise the route to any peer in North America.

- (b) Path Prepending: induces prepending of the AS number n times, where typically n is a number between one and four, when propagating this prefix (see Section 2.1.2). This is usually applied to specific peers or groups of peers. For example, 2914:4023 tells AS 2914 to prepend 2914 to the path three times when announcing it to peers in North America [6].

This taxonomy functionally enables this research by providing classifications to attach to particular community values. As many of these encode active signaling semantics, understanding them *in situ* is a valuable stepping stone to creating a functional anomaly detector for BGP communities.

Furthermore, the current distribution of these communities according to this taxonomy is a secondary research question this work seeks to study. Bonaventure et al. [8] were able to estimate this distribution: see Table 2.3:

Table 2.3. A 2008 estimation of the distribution of passive and active BGP community semantics. Note that this is the distribution of defined semantics, rather than a proportion of the communities seen, for example, at a route collector. Adapted from [8].

Category	Subcategory	Type	Proportion (%)	Totals (%)
Blackhole			0.22%	0.22%
Inbound	Local Pref		1.62%	26.25%
	Routes Tagging	IXP	1.82%	
		Type of Peer	6.33%	
		Geographic Location	9.11%	
AS	6.08%			
	Unknown		1.26%	
Outbound	Route Redistribution	Announcement	25.9%	73.52%
		Path Prepending	47.43%	
	Unknown		0.16%	

Table 2.3 indicates that the majority of communities defined are for route redistribution

purposes; i.e., for path prepending and announcement.

2.2.2 BGP Extended and Large Communities

The communities presented thus far have been 32-bit communities. However, RFC 4360 [40] introduced 64-bit communities, called “Extended Communities,” to the BGP in 2006. According to the specification, extended communities are intended to provide an extended range for community values to alleviate the fear of overlap and add structure to the community space. More specifically, extended communities provide structure by inclusion of a specified “Type” field which designates the type of application for which a given community value will be used. Additionally, RFC 8092 introduced so-called “Large Communities” in 2017; these large communities are 96 bits, or 12 octets, in length.

Unfortunately, for purposes of this research, extended and large communities have not seen widespread use and thus are not yet a significant factor. However, this is likely to change soon. As of March 25, 2019, as many as several hundred large communities were viewed at route collectors at the Routing Information Service (RIS) administered by RIPE .

Additionally, and more importantly, there are nearly 20,000 4-byte AS numbers currently visible in the global routing table, and Routing Information Registries (RIRs) are currently at the point where they are mostly assigning 4-byte AS numbers. This indicates large community usage will continue to increase generally [9], [15].

2.2.3 Communities and Security

Given the inherent lack of security in the BGP, communities currently offer a vector for potential attackers. Indeed, as discussed in Section 2.1.7, even current proposals for securing the BGP make no provisions for enforcing the security of communities. Streibelt et al. [5] has demonstrated practical scenarios in which communities can be leveraged to deny service, improve service, impose additional cost, and intercept traffic. In fact, this study was the inspiration for this thesis.

These attacks include use of community services for:

1. **RTBH**: by using blackholing communities for a particular AS, a malicious actor can effectively launch a DoS attack that cripples its reachability. This can be done in

- conjunction with prefix hijacking to drop traffic at the destination AS, denying service universally.
2. **Traffic Steering:** though harder to initiate since ASes typically respond only to steering attempts from their customers, it is possible to influence both the local preference and path prepending behavior in providers immediately upstream. However, AS relationships make it harder to use these communities against ASes farther away, even when the communities propagate. This may be most useful in conjunction with prefix hijacking. See Section 2.1.6 for more detail on AS relationships and their effects on path selection.
 3. **Route Manipulation:** an attacker can induce a redistribution of routes using specific communities; e.g., by sending a community to redistribute the route for a prefix to a particular AS along with a community to suppress that announcement. Notably, this requires knowledge about the community evaluation order for the intermediate AS.

2.3 Machine Learning

Machine learning is a method of data analysis that attempts to automate the process of building a model. Many models are classifiers for a given problem: for example, classifying different species of flowers on the basis of a set of observed features present in the data. In a supervised method, each species will be labeled, and each unique label is an example of a class. Simply put, a model is an attempt to train a classifier to correctly predict the probability that a particular flower belongs to a particular class. At its core, such a model is a solution to an optimization problem that attempts to minimize a loss function; i.e., the probability that the model incorrectly classifies an instance of a class [41]. This research attempts to use such a model to classify particular BGP communities.

2.3.1 Artificial Neural Networks (ANNs)

An ANN is one of the most common types of machine learning models. They arose from attempts to recreate the form and function of the human brain in the field of Artificial Intelligence (AI), and though they do this poorly, they function well as classifiers. More specifically, the brain is composed of individual neurons, cells which receive input signals from surrounding cells and act if the cumulative input of these cells exceeds a given

threshold. ANNs are based on a simplified form of neurons and their interconnections, wherein individual nodes with an activation function are connected in layers.

The simplest form of ANN is a single-layer perceptron. A perceptron is only capable of solving linear problems; that is, it can only fit the data if the classes are linearly separable.

It does this by taking a number of inputs; i.e., features of the data, and assigning weights to them. A bias is also applied to the activation function at each node. The node then transforms its input into an output signal with respect to a threshold just as in a real neuron (the bias adjusts this calculation). Like a real neuron, a perceptron learns by adjusting the weights and biases until it optimally reduces its error rate.

An MLP is simply a perceptron with multiple layers; the input layer is connected to one or more fully-connected “hidden” layers which enable it to solve non-linear problems. Such a network learns by first computing the output given some set of inputs, evaluating the result, and then propagating the error backward in a process appropriately called backpropagation. Since, in an MLP, input to one layer is a function of the input to the preceding layer⁸, backpropagation is merely an application of the chain rule of derivatives. Other types of neural architectures are commonly derived by manipulating the way neurons in the hidden layers are interconnected.

The amount that weights are updated during the training process is called the learning rate, and is adjustable⁹: the higher the value, the more rapidly the model changes. In stochastic gradient descent, the prototypical optimization algorithm for such a network, the network is iteratively attempting to guess, at each epoch, the error gradient. That is, it tries to guess how far the loss function is away from its minimum, and adjusts the weights and biases (the parameters) accordingly. The learning rate determines how much the parameters are adjusted at each epoch.

This process continues until the model converges at the minimum of the loss function, at which point the model has ideally become a competent classifier for the data. At this point, it can be used against novel testing data to evaluate its accuracy; if it has high performance

⁸For this reason, an MLP is called a “feed forward” network [42].

⁹Adjustable values that control the learning process are called “hyperparameters,” as they do not change between epochs. They are distinct from parameters like weights and biases, which the model learns on its own during training. Hyperparameters can be tuned through a grid search of possible combinations of values.

on the training set, but poor performance on the testing set, the model is likely overfitting the training data. This occurs when a model attunes itself too strongly to the training data, and thus fails to generalize well. Some simple solutions to overfitting include setting an early stopping point for the optimization algorithm, using k -many folds to train the data iteratively, gathering more data, removing features, etc. [41], [42].

2.3.2 Random Forests

A Random Forest is another classification algorithm that relies on recursive partitioning of the feature space. This is accomplished through classification trees which infer classification boundaries through the separation of nodes in the tree. This can also be represented as a rectangular partition of the feature space. Therefore, this approach will produce every possible combination that can be derived through recursive splitting. An estimate of class probabilities can be inferred from their relative frequencies in each partition (i.e., a terminal node of the tree).

However, individual classification trees are susceptible to instability due to small changes in the data. The boundary defined at each partition, sometimes called the “cutpoint,” strongly depends on the distribution of observations present in the training data. Additionally, an individual tree’s prediction will be piecewise constant; that is, it may vary too much for small changes in the values of the predictor variables (the feature set).

A Random Forest, by contrast, is an ensemble method which computes a prediction not from a single tree, but from a set of regression trees using a bootstrap sample of the learning data. The prediction of all trees can thus be combined, or “bagged,” to create smoother decision boundaries and improve the overall accuracy of the forest. A particular facet of Random Forests in general is that they introduce additional diversity by randomly restricting the set of predictor variables used at each split of the tree. For this reason, Random Forests are particularly good at inferring even weak interactions between predictor variables [43].

2.3.3 Binary Relevance for Multi-Class Problems

Multi-class data, such as the BGP communities that are the subject of this thesis, are difficult for many classifiers to handle natively. Instead, categorical labels are typically one-hot encoded such that every individual label becomes an array of binary variables equal in

number to the total number of classes. For example, for the labels “red,” “green,” and “blue,” red becomes [1, 0, 0], green becomes [0, 1, 0] and blue becomes [0, 0, 1]. To effectively train the model, an individual classifier is built for each label. For this reason, this method is called “Binary Relevance” or “One-vs-All.” The classifier with the highest output for an instance determines the class prediction for the model [41].

2.3.4 Class Imbalance

Real-world data often has a class imbalance. That is, the training data has a minority class with insufficient samples for training. This is called a minority class, and a classifier may struggle to learn how to recognize it. In fact, it is likely to never predict it at all in sufficiently skewed data. As demonstrated in Table 2.3 and observed in the data collection process, the BGP community taxonomy has such an imbalance.

Fortunately, there are multiple methods to train models from such data:

1. **Random Oversampling:** Random Oversampling simply takes samples of the minority class and duplicates them until it is better represented proportionally.
2. **Synthetic Minority Oversampling (SMOTE):** generates synthetic samples from a minority class. These samples are linear combinations of two samples from the minority class, increasing the number of samples present in the training data; this linear combination is chosen by selecting a random nearest neighbor from a k -Nearest Neighbors algorithm and created at a shared point between them [44].
3. **Adaptive Synthetic Sampling (ADASYN):** uses the density distribution to determine the number of synthetic instances to generate for each minority class; it focuses particularly on generating samples next to samples that were classified incorrectly using a k -Nearest Neighbors unsupervised approach [45].

CHAPTER 3: Methodology

This chapter presents the methods and procedures used in the course of this research into BGP communities. In particular, this chapter discusses the data collection process, feature engineering, model selection and validation, and application of such models to recent BGP data. Note that, throughout this chapter, for the sake of simplicity, communities will be presented in the standard syntax proposed by RFC 1997 as discussed in Section 2.2. Further, discussion of community semantics reflects the taxonomy presented in Section 2.2.1.

3.1 Data Collection

Collecting ground-truth data was the longest and most challenging part of this research. The major source of data for this research was the Center for Applied Internet Data Analysis (CAIDA) Community Dictionary Dataset [46]; web scraping was used as an additional source of data.

3.1.1 The CAIDA Dictionary Dataset

The BGP Community Dictionary Dataset, provided by CAIDA, was a useful but limited source of data for this research. This dataset contains only Inbound communities, particularly geographic and IXP tags; as described in Section 2.2.1, this represents only two out of the ten classification subcategories this research sought to classify. Thus, although it offers 7,311 community definitions, it was incomplete for purposes of this research. See Table 3.1 for more detail.

As a preliminary feasibility study, this dataset was examined for clustering or patterns within the semantic space in a three-dimensional plot examining data by community value, AS number, and the geographic location which they encode, as inferred from the standard community syntax. The results of this are presented in Chapter 4.

Table 3.1. The distribution of BGP communities used as training data in this research according to the standard taxonomy. In total, 10,027 community definitions were used from a pool of 848 unique ASes. The CAIDA Dictionary Dataset contributed 72.91% of these communities, while the remaining 27.09% were collected via web scraping. Note that, given the disparity of sources, these data are not expected to represent the distribution of defined semantics in the wild. This imbalance was adjusted for modeling as explained in Section 3.2.3.

Category	Subcategory	Type	Proportion (%)	Totals (%)
Blackhole			3.26%	3.26%
Inbound	Local Pref		0.89%	87.36%
	Routes Tagging	IXP	45.65%	
		Type of Peer	8.62%	
		Geographic Location	32.41%	
AS	0.44%			
	Other		0.24%	
Outbound	Route Redistribution	Announcement	3.17%	9.71%
		Path Prepending	6.53%	
	Other		0.01%	

3.1.2 Web Scraping

Given the limitations of the Dictionary Dataset, significant additional data—totalling 2,716 defined communities—were manually collected to gather information on the other community classes in the standard taxonomy.

To this end, web scraping was employed to automate the workload of collecting data. Web scraping is the process of programmatically extracting useful data from Hypertext Markup Language (HTML) pages on the internet. A significant number of community definitions were thus extracted from the community guide at One Step Consulting [6] using the BeautifulSoup library in Python. One Step is a technical consulting company that gathers data from provider source material; e.g., published guides or queries to major registry databases such as that of RIPE. This information was spot-checked where possible with public guides from the ASes under study, for example the publicly available communities

guide at Nippon Telegraph and Telephone (NTT), a Tier-One provider AS [7]. Communities noted as legacy definitions were not added to the dataset.

Due to the disparity of formats for the presentation of data on the website, even with web scraping this required manual adjustment unique to each HTML page containing a particular AS's community definitions; the programmatic equivalent was more time-consuming than simply doing so manually.

Additionally, many such definitions would require significant natural language processing; the explanations for communities are sometimes presented in very diverse formats, with critical information written in plain English before a particular section, or use letter variables to represent potential values in a definition. For example, NTT defines several communities with the form `65442:nnn`, which means prepend twice to peer “nnn” in Asia, where “nnn” is meant to be replaced with the AS number of the peer in question [7]. Similarly, TeliaSonera community prepending definitions are of the form `1299:252x`, where x can be the number of prepends (typically any number between zero and three) or nine, which signifies “do not announce” to a certain peer.

For this reason, there was insufficient time to use all of the communities defined at [6]; ASes of particularly high rank according to CAIDA's database¹⁰ were preferentially chosen as they are expected to represent communities of higher interest to the research community; however, some communities were also taken from ASes of lower rank for classes which were under-represented. Future work would greatly benefit from additional time or resources spent in data collection; e.g., via Amazon Mechanical Turk.

This data was stored in Comma-Separated Values (CSV) format and a manual classification was added to each definition, to serve as the true label for the supervised models developed to demystify this semantic space.

3.1.3 The Trouble with Taxonomy

Unfortunately, as there is no unifying semantic for most communities, a small but not insignificant number of communities do not conform to the taxonomy presented by [8]; for

¹⁰The rank of a particular AS, as defined by CAIDA, is a topological estimation of its customer cone size; i.e., the number of customers or indicated customers an AS has [16].

example, European Commercial Internet Exchange (ECIX), a major IXP, defines Round Trip Time (RTT) communities for granular application of policy with respect to latency information. The community 65030:ms means, “do not announce this route to peers with an RTT greater than or equal to the value ms.” Similarly, 65011:0 is an Inbound tag community that signifies that the RTT is between one and five milliseconds [47].

For purposes of this study, the former were classified by their principal effect (i.e., Outbound: Prepending); for data preprocessing, as with all communities that follow this alternate syntax (see Section 3.1.4), the first two octets were taken as the community value and the AS number was inferred separately. The latter communities can only be labeled as a general “Other” class; although they certainly fall under the domain of Inbound: Route Tagging, a more specific class does not exist in the Bonaventure taxonomy, and thus these communities fall outside the scope of this research. Extremely few such communities were observed in collected data (see Table 3.1).

Furthermore, since the BGP community field is of variable length, it is often the case that a route will have multiple communities. In at least one observed case, these communities have different meanings in different combinations: AS 209 uses 209:888 to signify a peer route when by itself. Ergo, this is clearly an Inbound: Type of Peer community. However, when used in conjunction with 209:64740, it means “do not announce to Deutsche Telekom”; that is, it modifies an Outbound: Announcement community. By contrast, 209:64740 alone means to announce that route to Deutsche Telekom, and 209:64743 means to prepend 209 three times to the AS path (i.e., it is an Outbound: Prepending community). Even more confusingly, the community string could be 209:888 209:64520 209:64749. Taken together, this tells AS 209 not to announce to any peers *except* Deutsche Telekom [48]. The model developed by this research does not capture this semantic complexity; fortunately, this was only observed in the semantic schema of one AS. However, there is no regulation of the community space to preclude it in other ASes from which data were not gathered.

Other communities are effectively multi-label with respect to the Bonaventure taxonomy. For example, Global Crossing, AS 3549, defines Inbound: Type of Peer communities with the format 3549:TCCC, where “T” is a type code with value one, two, four, or five to specify the relationship with the originating peer and “CCC” is a country code (e.g., 3549:4840 is a customer route learned in the United States). This means it tags both Type of Peer and

Geographic origin; it is multi-label with respect to our taxonomy. A more sophisticated model would need to be built to effectively predict such complexities; that is, the model developed herein is expected to perform poorly with respect to such communities. In total, 523 communities such communities were collected. Most of these were from AS 6461, which defines separate peer and customer (i.e., Inbound: Type of Peer) communities for specific locations of geographic ingress (i.e., Inbound: Geographic). For this research, these communities were labeled Inbound: Type of Peer; however, the model’s inability to capture this complexity is a fundamental limitation that should be addressed in future work.

Additionally, although the standard taxonomy proposed three separate categories of blackholing based on the scope and locality of their effect, this was very difficult to discern from ground-truth data; the observed blackholing community definitions rarely offer enough detail to distinguish between them. For this reason, the developed model treats blackholing generally. The reason for this may be that this is a question of provider implementation of community application, while these guides are intended for customer use. This is a fundamental limitation of this study, and perhaps of semantic classification in general. Further research into distinguishing between them may require testing community effects in the wild.

3.1.4 Syntactic Limitations

Additionally, many communities do not follow the syntax proposed in RFC 1997 [4], which states that “...community attribute values shall be encoded using an autonomous system number in the first two octets. The semantics of the final two octets may be defined by the autonomous system.” For example, the IXP DE-CIX, AS number 6695, uses the community 0:6695 to indicate that a route should not be redistributed, and 0:nnn to indicate that a route should not be redistributed to a specific peer, where “nnn” is the AS number of said peer.

Conversely, 6695:nnn is used to signify that a route *should* be redistributed to that particular peer. Similarly, Level 3, AS 3356, uses 65001:0, 65002:0, 65003:0, and 65004:0 to signify prepending once, twice, three times, or four times to any peer, respectively. The last two octets may also be occupied by the AS of a particular peer to specify to whom this prepending behavior should occur [6]. Fundamentally, this is a limitation of the syntax, as

such granular encoding information cannot be done easily with the last two bytes of a 32-bit community. Use of extended or large communities (Section 2.2.2) is a potential solution to this.

In considering communities of the form `xxx:0` or `xxx:nnn`, where “xxx” is *not* the AS number of the defining AS, and “nnn” is the AS number of a target peer, “xxx” was taken as the community value and the AS number inferred separately during preprocessing. In total, 118 such communities were collected.

3.2 Developing a Model

The intent of this research was to attempt to demystify the semantics of BGP communities. To this end, several One-vs-All models were created with the Python library Scikit-Learn [49]. Choosing the correct classification algorithm as the internal estimators for these was a crucial part of designing an effective predictive model.

Fortunately, Van Efferen et al. demonstrated that MLPs are useful in flow-based anomaly detection in Intrusion Detection Systems (IDSes) with appropriate feature extraction due to their fault tolerance, ability to adapt to changes in information, and resilience to noisy signals in the data. Their research compared the results of such a model with a J48 Decision Tree algorithm, which has proven high precision when working with such data [42]. Given the similarities of the classification problem studied in this thesis—for example, the need to choose appropriate features in one-dimensional data¹¹—MLPs were chosen as the estimators in the first model developed adapted. Its efficacy as an estimator was compared with use of Random Forests, which are “ensembles” of decision trees (see Section 2.3.2).

¹¹The BGP communities in the dataset created for this research do not have a multi-dimensional component; features were engineered such that the feature vectors used to train the model are one-dimensional. Data with additional dimensions (e.g., temporal), such as would occur with packet-based anomaly detection, typically require more complex neural architectures than a feed-forward MLP. Fully Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) with Long Short-Term Memory (LSTM) architectures are potential solutions to classification problems involving time-series data [50]. This has implications on future work in anomalous community detection, and will be discussed further in Section 5.2.

To train these models, the collated data—collected as described in the preceding section—were read into a data structure called a DataFrame using the Pandas library. The data were separated into training and testing datasets and stripped of their labels. Values were then scaled according to unit variance in each feature according to the equation

$$z = \frac{x - u}{s} \quad (3.1)$$

where u is the mean of the training samples and s is the standard deviation. This scaler was fitted only to the training data, and applied separately to the testing data. As this was a multiclass model, labels were one-hot encoded, and the chosen classifier was used in a Binary Relevance paradigm as explained in Section 2.3.3. Thus, a separate, binary classifier was optimized to predict each class versus all others.

3.2.1 Feature Engineering

In the wild, operators will typically have only the community string itself from which to draw features. The AS Rank, as determined by CAIDA [16], was determined to be unsuitable as a feature for this reason, as well as any number of other quantitative values calculated from historical data, such as a community’s frequency. Ground truth data would be difficult to collect for such communities as well, and vantage-specific bias would limit the model’s ability to generalize. Thus, in order to build an effective model, qualitative analysis of this data alone was used to identify a feature set capable of differentiating individual community values.

Consequently, the feature set presented in Table 3.2 were engineered from the community string in an attempt to recognize patterns in the diverse definitions of communities between ASes. All integers were scaled according to the above formula; the feature’s mean was calculated columnwise and subtracted for each sample, such that each feature has a mean of zero and a standard deviation of one.

Table 3.2. These features were engineered from the collected community data based on domain knowledge and observed improvement to F_1 score for a particular class on a given seed. In total, thirteen features were engineered from the community string alone.

Feature	Type	Explanation
AS Number	Integer	The number of the AS that defined the community. In the wild, this must be inferred from the standard syntax alone; i.e., the most significant word.
Community Value	Integer	The value defined in the least significant word of the community, or the most significant word in select cases (see Section 3.1.4).
Community-AS Number Product	Integer	The product of the community value and the defining AS was computed for each instance.
Ends in Zero	Boolean	Indicates whether the community value ends in zero; many Outbound: Announcement values end in a zero. Many Inbound: Geographic and Inbound: Local Preference values end in zero as well.
Ends in One, Two, Three, or Four	Boolean	Indicates whether the community value ends in one, two, three, or four; the majority of prepending communities end in one of these values as a way to encode how many prepends an AS should make, regardless of the preceding digits.
Ends in Nine	Boolean	Indicates whether the community value ends in nine. Many Outbound: Announcement communities end in this value.
Multiple of Five	Boolean	Indicates whether the community value is a multiple of five, using modular arithmetic. Most Inbound: Local Preference communities are a multiple of five.
Less Than or Equal to 200	Boolean	Indicates whether the community value is less than or equal to 200. The majority of Inbound: Local Preference values observed fit this criterion, as the default for "LOCAL_PREF" - the policy metric these communities influence - is 100 and communities rarely deviate from this value by more than 100% in observed data.
Repeating	Boolean	Indicates whether, when treated as a string, the community value is a repeating character. Many Blackhole communities have this form; e.g., 9999 and 666.
Number of Sixes	Integer	A count of the number of sixes in the community value. Many Blackhole communities in particular have more sixes, even if they are not the classic Blackhole value 666.
Number of Zeroes	Integer	Indicates the number of zeroes in the community value. Many Inbound communities in particular are a multiple of 1,000 and thus have more zeroes than other values.
Number of Digits	Integer	Indicates the number of digits in the community value.
Three Digits	Boolean	Indicates whether the community value is exactly three digits; many Inbound: Type of Peer are three characters in length.

3.2.2 Model Validation

This thesis applied the general process of model validation in machine learning. In this process, features are extracted or engineered from data and preprocessed in a context-appropriate manner. The dataset can then be separated into training and testing sets. Subsequently, an appropriate classification algorithm (e.g., an MLP) is selected, hyperparameters are chosen, and the model is trained with the training data (or k -many folds thereof) in an attempt to minimize the loss function according to the model's optimization algorithm. In this research, multiple classification algorithms were variously tested, including an MLP and a Random Forest; the comparative results for each model are presented in Chapter 4. The subset of the data reserved for testing comprised approximately 25% (2,501 communities) of the total number of labelled communities collected as described above. Training and testing sets were stratified with respect to their taxonomic label such that the proportion of instances of each class in the testing set maintains its proportional representation in all data. This stratification was applied to ensure the testing set had an adequate number of samples of each minority class to test its accuracy. Unfortunately, as some classes were underrepresented in ground-truth data, they were perforce underrepresented in testing data as well (see Section 3.2.3).

The trained model was then used on the testing data and its accuracy was determined according to various metrics of performance (see Section 3.2.4). If a model does poorly, the training phase can be repeated by introducing new data, selecting new hyperparameters, introducing new features, etc. In this research, a grid search—performed on a High-Performance Computing (HPC) cluster—was additionally applied to search for optimal hyperparameters including learning rate, optimization algorithm (“solver”), hidden layer sizes, and the activation function in use at each node. The model, its fitted training scaler, and the label encoder used to one-hot encode the multi-class labels were saved using Python's Joblib library to enable their use in other applications. Additionally, the model is implemented in a class, such that if no saved model is present, it can be instantiated and retrained in another context.

3.2.3 Dealing with Class Imbalance

A particular problem with the data collected for this research is that there is a distinct class imbalance, in terms of actual semantic definitions, present in collected data and observed

by [8] (see Table 3.1).

As a result of this, a continual problem during the model evaluation process was the model's failure to predict certain minority classes. The ideal method of fixing this imbalance would be to collect more data, but this proved difficult, as described previously. To this end, Random Oversampling, SMOTE, and ADASYN were variously applied to the training set to permit the classifier to recognize samples of these classes in the testing data. Random Oversampling proved most efficacious for this purpose by a small margin of approximately 0.5% in aggregate accuracy depending on the seed. This was an important step in building the final model.

3.2.4 Metrics of Performance

To evaluate model performance, precision, recall, and F_1 score were calculated for each class, and accuracy was calculated over all testing data. These are equations that employ true positive and false positive predictions made on the data:

1. A **true positive** (T_p) is when a model predicts that a sample is an instance of a class, and it is *correct* in its prediction.
2. A **false positive** (F_p) is when the model predicts that a sample is a member of a class, and it is *incorrect* in its prediction; the sample is *not* a member of that class. This is also called a **Type I Error**.
3. A **true negative** (T_n) is when a model predicts that a sample is *not* an instance of a class, and it is *correct* in its prediction.
4. A **false negative** (F_n) is when the model predicts that a sample is *not* a member of a class, and it is *incorrect* in its prediction. That is, it *is* a member of that class, and the model wrongly predicted it was not. This is also called a **Type II Error**.

Precision, Recall, F_1 Score, and Accuracy are defined by the following formulae:

$$Precision = \frac{T_p}{T_p + F_p} \quad (3.2)$$

$$Recall = \frac{T_p}{T_p + F_n} \quad (3.3)$$

$$Accuracy = \frac{T_p + T_n}{Total} \quad (3.4)$$

$$F_1 = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (3.5)$$

Recall is therefore the ability of a model to find all relevant data within a dataset, whereas precision is a model's ability to identify only the relevant data. A model with low recall but high precision would miss a lot of predictions for the class, but what it did predict to be in the class would really be in the class. That is, it is very selective in its predictions, but usually right when it positively identifies something. Conversely, a model with high recall but low precision predicts that many things are in the class, and many of them are, but many of them are not. It positively identifies many things, but often does so wrongly. The F_1 score is the harmonic mean between them; accuracy is simply the proportion of correct predictions.

Additionally, ROC and Precision-Recall curves were measured for each sample. ROC curves plot a model's true positive rate against its false positive rate. This uses the threshold that determines if a sample is positively identified as a member of the class; ROC curves plot how precision vs recall changes while varying this threshold. In other words, starting at a threshold of 1.0 (positively identifying no instances, truly or falsely), the threshold is decreased slowly and more positive predictions are made (true positives and false positives), until every point is positively identified at threshold 0.0. The Area Under the Curve (AUC) can be then be calculated to determine classification performance. Precision-Recall curves are similar; they measure precision vs. recall at varying thresholds. Like ROC curves, their AUC—computed as the average precision over all thresholds—can be taken and used as a measure of classification performance.

Finally, a confusion matrix was created for the testing data; this is simply a table wherein the model's total predictions for a class are separated into their true labels [41], [51], [52]. Accuracy was taken over the entire testing set as well. These metrics were used extensively in the model validation process; a final, predictive model was iteratively built in this way.

3.3 Estimating Taxonomic Distribution

As an immediate, practical application of the best model created, recent routing data was collected from the RIS database provided by RIPE [53]. Specifically, UPDATEs aggregated in Multi-Threaded Routing Toolkit (MRT) format were taken from multiple vantages, including the London Internet Exchange (LINX), the London-based Internet Exchange Point (LONAP), the RIPE Network Coordination Center (NCC) in Amsterdam, and the New York International Internet Exchange (NYIIX). All collected data for August 10, 2020 between 20:00 and 20:35 were collated and parsed on a local HPC cluster using the `Mrt-parse` and `Multiprocessing` libraries in Python. For comparison, UPDATEs were additionally collected from the same vantages on September 1, 2020 between 4:00 and 4:35.

The community string, as well as several summary statistics, were extracted from each UPDATE found in this data, and the model was applied to each community within. A feature vector was created for each instance in accordance with the feature set engineered for the training data (see Table 3.2). While ground-truth AS numbers were immediately available for training data, the defining AS number had to be inferred during this analysis as the most significant word of a community; given that the community attribute is transitive, it is difficult to deduce the true, defining AS in real-world data. Future work may apply topological considerations to increase confidence when the community contravenes the standard syntax of its foundational RFC. This is, unfortunately, a limitation of any approach that uses the AS number as a feature. Note also that well-known communities were considered separately. The best model—trained separately for both the major three categories and the eight subcategories of the standard taxonomy on all available data—was then used to predict the taxonomic distribution of community services present in this real-world data; the results of this process are presented in Chapter 4.

CHAPTER 4: Results

This chapter will present the major findings of this research. This includes a preliminary examination of semantic structure within ground-truth data, the performance of the various models developed in the course of this research, and the results of the estimation of taxonomic distribution in recent BGP data.

4.1 Analysis of Ground-Truth Data

As explained in Section 3.1.1, 72.91% of community data used in this research was contributed by the CAIDA Dictionary Dataset. One of the primary research questions set forth in this study is whether there is common structure to community definitions among ASes, thereby implying that one can make accurate predictions of unknown communities. To obtain intuition over the extant structure of community definitions among ASes, communities in this dataset were evaluated according to the standard syntax (see Section 2.2) and plotted against their encoded country as defined by CAIDA's dictionary in Figure 4.1. Specifically, the most significant word of the community was inferred to be the defining AS number, while the least significant word was taken to be the community value. For comparison, the semantic structure of all collected ground-truth communities, including those obtained through web scraping, is plotted in Figure 4.2. The most clear pattern this presents is an approximate horizontal line for Blackhole communities with the community values 666 and 9999. The least significant words for most communities, with the notable exceptions of the two types of location-encoding ingress communities, are clustered at values less than 10,000.

This analysis indicates that individual geographic encoding for a given country can occupy any value up to the value of a word (65,535); some ASes define geographic communities at most of the possible values in this range for different locations in the same country. It should additionally be noted that certain countries were disproportionately represented in this data; this is likely due to vantage-specific bias in CAIDA's data collection process and more granular geographic community needs for countries with more ASes. See Table 4.1 for a listing of countries for which many different geographic origin communities were

observed.

In fact, according to these data, some ASes in CAIDA’s dictionary have hundreds of communities which encode the same facility. For example, there are 666 recorded communities in the dataset encoding an IXP facility in Ashburn, Virginia operated by Equinix. The most significant word in all of these communities is 24115, whereas the least significant word takes almost any value possible up to the value of a word. This pattern is common within the data and is the source of the vertical series of communities in Figure 4.1.

According to CAIDA, all of these communities denote a specific IXP facility; e.g., one operated by Amsterdam Internet Exchange (AMS-IX), ECIX, Equinix, or DE-CIX [46]. It is possible that this is merely an artifact of CAIDA’s data collection process; likewise, CAIDA’s dictionary may be incorrect, and thus these communities may have different semantics altogether, including active semantics (i.e., these communities may be for route redistribution). However, several alternative sources of ground-truth data were examined for independent verification of selected examples (including vertical series attributed to ASes 24115, 19996, 39107, and 42476), and neither conflicting nor corroborating definitions for these communities could be found [6], [54].

Unfortunately, another implication of this analysis is that, at least among the ASes represented in this dataset, the standard taxonomy is often not clearly delineated by clear ranges in the least significant word. This is particularly concerning because the CAIDA Dictionary Dataset purportedly contains only “location-encoding ingress communities,” i.e., Inbound: Geographic and Inbound: IXP tags [46]. This precipitated the need for the feature engineering approach explained in Section 3.2.1.




Figure 4.1. Communities present in the CAIDA Dictionary Dataset for January 1, 2018, were plotted in a three-dimensional grid according to the standard community syntax. Country codes for the encoding region were plotted along the x-axis, while the community value and the defining AS's number—as derived from the standard community syntax—were plotted on the y-axis and the z-axis, respectively. A colormap was applied along the x-axis; it has no meaning beyond distinguishing three-dimensional data by country.

Table 4.1. These countries were encoded with the greatest frequency in analyzed data. 78 countries with fewer than 100 definitions in the dataset were aggregated for ease of reference.

Country Code	Country	Number of Communities	Percentage
US	United States	1,289	18.00%
NL	Netherlands	1,158	16.16%
DE	Germany	1,106	15.44%
GB	United Kingdom	1,006	14.04%
RU	Russia	848	11.83%
PL	Poland	300	4.19%
CH	Switzerland	258	3.60%
FR	France	163	2.28%
IT	Italy	102	1.42%
Other		934	13.04%
Total		7,164	100%




Figure 4.2. This figure scatters all 10,027 ground-truth community data according to the syntax proposed by RFC 1997 to illustrate macro-structure in known community semantics.

4.2 Model Performance

To develop the best method of community value estimation, various classification algorithms were applied to the dataset using the approach described in Section 3.2. This section will present the accuracy of the models as expressed through various metrics of performance (see Section 3.2.4) computed over a subset—2,501 communities, or approximately 25%—of the data reserved for testing. As described in Section 3.2.2, training and testing data were stratified on their labels to ensure that minority classes would appear in testing data.

However, it must be noted that several classes are underrepresented in the data; various imbalance correction strategies were applied as explained in Section 3.2.3. Iterative model development indicated that overall performance was highest with random oversampling. However, this process was only applied to the training data to avoid an overly-optimistic evaluation of model skill. Consequently, the support in the testing data reflects only real communities for each class. For this reason, performance in categories with fewer true samples is expected to be relatively lower. As seen in the “Support” column of Tables 4.2 and 4.6, this particularly affects the Inbound: AS and Inbound: Local Preference subclasses; re-sampling the dataset has only allowed the model to predict these classes in novel data. Therefore, our model is expected to be less generalizable with respect to these minority subclasses of the standard taxonomy.

4.2.1 MLP Metrics

The Binary Relevance method, described in Section 2.3.3, was used to create a model with internal MLP estimators. This was the first model developed; it obtained approximately 78.65% aggregate accuracy on testing data with respect to eight subclasses of the Bonaventure taxonomy presented in Section 2.2.1. Due to the limitations of the ground-truth data from which community definitions were scraped (see Section 3.1.3), the subclasses of Blackhole were not considered during the training phase. Thus, the model was trained to predict any of the following eight taxonomic classes:

1. Blackhole Communities
2. Inbound Communities
 - (a) AS
 - (b) Geographic
 - (c) IXP
 - (d) Local Preference
 - (e) Type of Peer
3. Outbound Communities
 - (a) Announcement
 - (b) Prepending

Performance metrics, including precision, recall, and F_1 score, were computed for each class and are presented in Table 4.2; a complete confusion matrix is additionally presented in Table 4.3. As described in Section 3.2.4, ROC and Precision-Recall curves were plotted and are presented in Figure 4.3. The AUC for each class is noted in their respective legends. Note that, for Precision-Recall curves, the AUC is computed as the average precision across all thresholds.

Recall from Section 3.2.4 that ROC curves contrast the number of Type I errors (false positives) of a model as it makes more positive predictions about class membership at varying confidence thresholds. A theoretically-perfect ROC curve would pass through the point $(0, 1)$; such a model would be able to identify *all* true samples without a *single* false positive in its predictions. Realistically, most models have a commensurate number of false positives as the threshold for a positive prediction is lowered. An ROC curve thus always increases monotonically. Precision-Recall curves are fundamentally similar, but contrast

precision and recall instead of the true positive rate and the false positive rate. Recall from Section 3.2.4 that recall measures the total fraction of true positives among all true samples, and therefore will always be monotonically increasing as the threshold is lowered, whereas precision is sensitive to the number of false positives in predictions at a given threshold. Thus, it is not necessarily the case that it will decrease monotonically. Furthermore, it should be noted that, since the number of thresholds is proportional to the number of unique probability predictions made by the model, classes with fewer true samples in testing data are tested at fewer thresholds.


For comparison, an additional MLP-based multi-class model was trained to classify the three primary categories of the taxonomy, namely, Blackhole, Inbound, and Outbound communities. The performance metrics and confusion matrix of this model are presented in Tables 4.4 and 4.5, respectively; ROC and Precision-Recall curves for this model are displayed in Figure 4.4. As expected, given the decreased complexity of this classification problem, this model attained a higher overall accuracy of 93.84% on novel testing data.

Table 4.2. This table presents the major performance metrics for the eight subclasses fitted to the model during training. Performance is calculated over a testing set of data. Support reflects the number of “true” samples of that type in the testing data.


Class	Precision	Recall	F_1 Score	Support
Blackhole	0.60	0.81	0.69	81
Inbound: AS	0.26	0.91	0.41	11
Inbound: Geographic	0.78	0.73	0.75	803
Inbound: IXP	0.90	0.76	0.82	1131
Inbound: Local Preference	0.38	0.82	0.52	22
Inbound: Type of Peer	0.76	0.95	0.85	213
Outbound: Announcement	0.56	0.85	0.67	78
Outbound: Prepending	0.73	0.97	0.83	162
Accuracy			0.79	2501
Macro Average	0.62	0.85	0.69	2501
Weighted Average	0.81	0.79	0.79	2501

Table 4.3. This table presents the confusion matrix for the MLP-based model with respect to eight subclasses of the Bonaventure taxonomy. “Inbound” and “Outbound” are omitted for the sake of brevity.

		Predicted								
		Blackhole	AS	Geographic	IXP	Local Preference	Type of Peer	Announcement	Prepending	Support
Actual	Blackhole	66	1	4	6	1	0	3	0	81
	AS	0	10	0	0	0	0	1	0	11
	Geographic	24	13	585	89	16	27	26	23	803
	IXP	16	13	152	862	7	32	17	32	1131
	Local Preference	0	0	1	1	18	1	1	0	22
	Type of Peer	1	1	2	1	2	203	1	2	213
	Announcement	3	0	2	1	3	2	66	1	78
	Prepending	0	0	1	0	0	1	3	157	162
	Total Predictions	110	38	747	960	47	266	118	215	2501



(a) ROC Curves



(b) Precision-Recall Curve


Figure 4.3. ROC and Precision-Recall curves for the MLP-based model estimating eight subclasses of the Bonaventure taxonomy. Note that the “No Skill” line represents a model incapable of predicting classes better than if it were randomly guessing. Similarly, F_1 score is the harmonic mean between precision and recall: isocurves for varying F_1 scores are plotted for reference.

Table 4.4. This table presents the major performance metrics for the three primary taxonomic classes fitted to the model during training. Performance is calculated over a testing set of data. Support reflects the number of “true” samples of that type in the testing data.


Class	Precision	Recall	F_1 Score	Support
Blackhole	0.61	0.78	0.68	81
Inbound	0.99	0.94	0.96	2180
Outbound	0.72	0.97	0.82	240
Accuracy			0.94	2501
Macro Average	0.77	0.90	0.82	2501
Weighted Average	0.95	0.94	0.94	2501

Table 4.5. This table presents the confusion matrix for the MLP-based model with respect to the three primary classes of the Bonaventure taxonomy.

		Predicted			
		Blackhole	Inbound	Outbound	Support
Actual	Blackhole	63	17	1	81
	Inbound	38	2052	90	2180
	Outbound	2	6	232	240
Total Predictions		103	2075	323	2501



(a) ROC Curves



(b) Precision-Recall Curve

Figure 4.4. ROC and Precision-Recall curves for the MLP-based model estimating the primary three classes of the Bonaventure taxonomy. Note that the “No Skill” line represents a model incapable of predicting classes better than if it were randomly guessing. Similarly, F_1 score is the harmonic mean between precision and recall: isocurves for varying F_1 scores are plotted for reference.

As expected, the performance of this model varies with respect to the individual classes of the community taxonomy. The most challenging classes for the MLP-based model were Inbound: AS and Inbound: Local Preference with F_1 scores of 0.41 and 0.52, respectively. It is no coincidence that these classes are also those with the poorest representation in data; support for these classes in testing data included only 33 “true samples” between them. While this is likely the reason the model’s predictive accuracy suffered, it may also suggest that their class boundaries are harder to estimate by such a model. If the engineered feature set is insufficient for these classes, these classes will be more challenging to predict effectively.

Similarly, Blackhole communities, even without exploring subclasses within the taxonomy, suffered from the lowest F_1 score of the primary three classes at 0.68, despite having reasonably high support in testing data. This suggests that Blackholing is particularly variable among AS semantic schemata. Future work may be able to achieve higher accuracy with more extensive work in data collection and feature engineering.

Contrariwise, Inbound: Type of Peer, Inbound: IXP, and Outbound: Prepending all exhibited F_1 scores in excess of 0.80, which indicates that the model was able to more effectively estimate their classification boundaries. This was expected particularly with regard to prepending, which during the data collection and feature engineering phases was observed to be more likely to share similar structure between ASes.

It must also be noted that the model’s recall was higher than its precision for all classes save Inbound: Geographic. This indicates that the model often makes positive predictions about class membership, and thus is more likely to identify true members, but is sometimes incorrect in its assessment. This may suggest that some ASes structure their community definitions such that certain community types exhibit features indicative of a different class in another AS’s semantic schema.

For example, Level 3, AS number 3356, defines the community 3356:666 as an Inbound tag for a route learned from a peer (i.e., it is an Inbound: Type of Peer community). This is unusual because RFC 7999 defines a well-known Blackhole community, 65535:666; the value 666 has thus been widely adopted among ASes for organizationally-defined Blackhole communities [6], [39]. In fact, this model indeed misclassifies 3356:666 as a Blackhole community. Since precision accounts for false positives, this Type I error will

lower the precision for Inbound: Type of Peer communities. Recall, however, is unaffected; it accounts only for true positives and false negatives (see Section 3.2.4).

Finally, the ROC and Precision-Recall curves present an interesting dichotomy in model skill. Namely, the ROC curves indicate better performance than is observed in the Precision-Recall curves: note, for example, that the AUC (or, equivalently, the average precision) of the micro-average for all classes in the Precision-Recall curves in Figure 4.3b is 0.86, whereas the micro-averaged AUROC in Figure 4.3a is 0.96. This is because ROC curves are plots of the true positive rate versus the false positive rate at varying prediction thresholds; the false positive rate incorporates true negatives, making changes in the ROC curve insensitive to changes in class distributions. Thus, they may give an optimistic evaluation of model skill for minority classes. This is why Precision-Recall curves were plotted for comparison: they evaluate the fraction of true positives among positive predictions and thus offer a more accurate evaluation of model skill with respect to imbalanced data [55], [56].

4.2.2 Random Forest Metrics

For comparison with MLPs as a classification algorithm, One-vs-All models utilizing internal Random Forests as estimators were also trained and tested on collected BGP community data. This model was able to outperform the MLP-based model by a significant margin. It attained 90.64% aggregate accuracy with respect to eight classes of the Bonaventure taxonomy. Likewise, it was able to obtain 97.80% accuracy when trained to predict the primary three classes of the Bonaventure taxonomy.


A complete set of performance metrics and a confusion matrix for this model are presented in Tables 4.6 and 4.7, respectively; ROC and Precision-Recall curves are presented in Figure 4.5. Additionally, an evaluation of the model's performance with respect to the three primary classes of the Bonaventure taxonomy is presented in Tables 4.8 and 4.9 as well as Figure 4.6.

Table 4.6. This table presents the major performance metrics for the eight subclasses fitted to the model during training. Performance is calculated over a testing set of data. Support reflects the number of “true” samples of that type in the testing data.


Class	Precision	Recall	F_1 Score	Support
Blackhole	0.76	0.72	0.74	81
Inbound: AS	0.75	0.55	0.63	11
Inbound: Geographic	0.88	0.88	0.88	803
Inbound: IXP	0.91	0.93	0.92	1131
Inbound: Local Preference	0.89	0.73	0.80	22
Inbound: Type of Peer	0.95	0.96	0.96	213
Outbound: Announcement	0.93	0.86	0.89	78
Outbound: Prepending	0.98	0.99	0.98	162
Accuracy			0.91	2501
Macro Average	0.88	0.83	0.85	2501
Weighted Average	0.91	0.91	0.91	2501

Table 4.7. This table presents the confusion matrix for the Random Forest-based model with respect to eight subclasses of the Bonaventure taxonomy. “Inbound” and “Outbound” are omitted for the sake of brevity.

		Predicted								Support
		Blackhole	AS	Geographic	IXP	Local Preference	Type of Peer	Announcement	Prepending	
Actual	Blackhole	58	0	11	10	0	0	2	0	81
	AS	0	6	2	2	0	1	0	0	11
	Geographic	10	0	709	75	1	4	3	1	803
	IXP	3	2	73	1047	1	4	0	1	1131
	Local Preference	1	0	1	4	16	0	0	0	22
	Type of Peer	2	0	3	3	0	204	0	1	213
	Announcement	2	0	2	5	0	1	67	1	78
	Prepending	0	0	2	0	0	0	0	160	162
Total Predictions	76	8	803	1146	18	214	72	164	2501	



(a) ROC Curves



(b) Precision-Recall Curve


Figure 4.5. ROC and Precision-Recall curves for the Random Forest-based model estimating eight subclasses of the Bonaventure taxonomy. Note that the “No Skill” line represents a model incapable of predicting classes better than if it were randomly guessing. Similarly, F_1 score is the harmonic mean between precision and recall: isocurves for varying F_1 scores are plotted for reference.

Table 4.8. This table presents the major performance metrics for the three primary taxonomic classes fitted to the model during training. Performance is calculated over a testing set of data. Support reflects the number of “true” samples of that type in the testing data.


Class	Precision	Recall	F_1 Score	Support
Blackhole	0.84	0.69	0.76	81
Inbound	0.98	0.99	0.99	2180
Outbound	0.98	0.93	0.96	240
Accuracy			0.98	2501
Macro Average	0.93	0.87	0.90	2501
Weighted Average	0.98	0.98	0.98	2501

Table 4.9. This table presents the confusion matrix for the Random Forest-based model with respect to the three primary classes of the Bonaventure taxonomy.

		Predicted			
		Blackhole	Inbound	Outbound	Support
Actual	Blackhole	56	25	0	81
	Inbound	9	2166	5	2180
	Outbound	2	14	224	240
Total Predictions		67	2205	229	2501



(a) ROC Curves



(b) Precision-Recall Curve

Figure 4.6. ROC and Precision-Recall curves for the Random Forest-based model estimating the primary three classes of the Bonaventure taxonomy. Note that the “No Skill” line represents a model incapable of predicting classes better than if it were randomly guessing. Similarly, F_1 score is the harmonic mean between precision and recall: isocurves for varying F_1 scores are plotted for reference.

The performance of the Random Forest-based model is higher for every class; this difference is particularly pronounced with respect to the eight subclasses of the Bonaventure taxonomy. As with the MLP-based model, classes with relatively few true samples in the testing data fared worse than classes with better support. As was observed with the previous model, Inbound: AS had the worst F_1 score among the eight subclasses (0.63); Blackhole communities were similarly challenging for this model to classify correctly. Both of these classes saw only mild increases in F_1 score when compared to the MLP-based model. While the lack of support is most likely the proximate cause for this difficulty with Inbound: AS communities, the modest increase in predictive performance for Blackhole communities lends further credence to the class's likely variability between AS schemata. This is an unexpected result.

An interesting trend observed with this model when contrasted with the MLP-based model is that, unlike the latter, this model generally has higher precision than recall for each class. Wherever this is not the case; e.g., for the Inbound: IXP class, precision and recall are within a few hundredths of one another. This indicates that the Random-Forest model is better overall than the MLP model; it often makes positive predictions for members of a class and is usually correct in its assessment. This is the desired behavior of any predictive model, and thus the Random Forest-based model—trained on all collected community semantic data—was selected for use in the subsequent section.

4.3 Taxonomic Distribution of Communities in Recent BGP Data

Recent BGP data in MRT format were collected from three RIPE vantages in New York, Amsterdam, and London. In total, 9.0 million communities were observed over 2.2 million UPDATEs received on August 10, 2020, across all vantages studied. For comparison, BGP data for the same three vantages were collected over a different capture interval; an additional 7.7 million communities were collected on September 1, 2020. Table 4.10 displays the distribution of BGP control packet types observed between the two collection intervals.

Table 4.10. This table represents the types of control packet observed during the two intervals studied; see Section 2.1.3 for more detail. Only UPDATE packets were further analyzed.

Date	Type	Count	Percent
August 10, 2020	OPEN	396	0.02%
	UPDATE	2,231,412	99.22%
	NOTIFICATION	190	0.01%
	KEEPALIVE	16,883	0.75%
	Total	2,248,881	100%
September 1, 2020	OPEN	394	0.80%
	UPDATE	2,101,134	99.17%
	NOTIFICATION	197	0.02%
	KEEPALIVE	16,940	0.80%
	Total	2,118,665	100%

As an immediate application of the best model—the Random Forest-based model—predictions for the taxonomic description of each class were made on BGP communities observed by parsing and extracting all community strings collected. For the capture interval on August 10, 2020, 1,355,391 observed communities (15.09%) were present in training data; however, many of the communities collected were repeated in subsequent UPDATES. In fact, only 10,425 unique communities (0.12% of the 9.0 million communities collected) were observed at these vantages during the capture interval; only 950 (9.11%) of these unique communities were present in training data. Similarly, of the communities collected on September 1, 2020, 1,202,982 (15.59%) were present in training data. Only 10,409 (0.13% of the 7.7 million communities collected) of these communities were unique, however. 963 (9.25%) of these unique communities were present in training data.

A conservative estimate for the taxonomic distribution of communities during both capture intervals was made over the communities observed in training data; see Table 4.11 and Figure 4.7 for both capture intervals. Ergo, this estimate accounts for 15.09% and 15.59%

of communities seen on August 10, 2020, and September 1, 2020, respectively. Additionally, an estimate for all communities observed over these periods—including the remaining, novel inputs—is presented in Table 4.12 and Figure 4.8 using only the primary three categories of the Bonaventure taxonomy for the most accurate possible result. Thus, this estimation accounts for 100% of communities in both capture intervals, though is necessarily a less confident result. Future work in expanding the pool of ground-truth predictions used as inputs for training this model will increase its accuracy and generalizability with respect to novel data, as new ground-truth training data will expand the model’s coverage of extant semantic structure among ASes.

Table 4.11. The best multi-class model developed predicted the following taxonomic distribution of BGP communities in UPDATES collected August 10, 2020, between 20:00 and 20:35 at various RIPE vantages. This distribution is compared with UPDATES collected at the same three vantages on September 1, 2020, between 4:00 and 4:35. Note that the proportions shown constitute only those communities which were present in training data, and therefore those for which we have higher confidence in model skill. This accounted for approximately 15% of all communities observed during both capture intervals. Note that this also implies that this estimate is slightly biased by the proportion of communities present in training data (see Table 3.1).

Date	Category	Subcategory	Type	Proportion (%)	Totals (%)
August 10, 2020	Blackhole			0.49%	0.49%
	Inbound	Local Pref		0.46%	98.43%
		Routes Tagging	IXP	8.34%	
			Type of Peer	18.71%	
			Geographic Location	70.86%	
AS	0.06%				
Outbound	Route Redistribution	Announcement	0.66%	1.08%	
		Path Prepending	0.42%		
September 1, 2020	Blackhole			2.04%	2.04%
	Inbound	Local Pref		0.42%	96.71%
		Routes Tagging	IXP	5.70%	
			Type of Peer	18.39%	
			Geographic Location	72.13%	
AS	0.07%				
Outbound	Route Redistribution	Announcement	0.97%	1.24%	
		Path Prepending	0.27%		




Figure 4.7. These figures represent a conservative estimate for the taxonomic distribution of communities observed in recent RIPE BGP data, as predicted by the Random Forest-based model described in Section 4.2.2. In this estimate, only communities which were present in training data were considered, to afford the highest possible confidence. Therefore, these estimates account for approximately 15% of all communities observed during both capture intervals. Note that the scale for the y-axis is logarithmic.

Table 4.12. The best multi-class model developed predicted the following taxonomic distribution of BGP communities in UPDATEs collected August 10, 2020 between 20:00 and 20:35 at various RIPE vantages. This distribution is compared with UPDATEs collected at the same three vantages on September 1, 2020 between 4:00 and 4:35. Note that, while this accounts for 100% of communities collected across both capture intervals, the lack of ground-truth data for as many as 85% of these communities necessarily reduces our confidence in these predictions.

Date	Category	Proportion (%)
August 10, 2020	Blackhole	2.65%
	Inbound	94.88%
	Outbound	2.46%
	Well-Known	0.01%
September 1, 2020	Blackhole	2.73%
	Inbound	94.94%
	Outbound	2.31%
	Well-Known	0.01%




Figure 4.8. A complete accounting of the taxonomic distribution of all communities observed in recent RIPE BGP data. This distribution was estimated by the Random Forest-based model described in Section 4.2.2, and focuses on the primary three subcategories of the Bonaventure taxonomy for the most accurate result. This estimate is necessarily a less confident result, given the inclusion of novel communities not seen during training and for which ground-truth information may not be available.

This process estimated that the majority of the communities studied—in both applications of the model—could be classified as Inbound communities, a fact which did not vary significantly between capture intervals. In fact, the model predicted that the majority of communities are Inbound: Geographic, Inbound: Type of Peer, and Inbound: IXP; together, they accounted for between 96.2% and 97.9% of the communities for which we have ground-truth data collected at these vantages. It should be noted that this result is slightly biased by the proportion of communities present in ground-truth data, but the less-confident estimate including all communities in Table 4.8 suggests that this distribution is largely

accurate across the entirety of communities received during these periods. This was an expected result; it corroborates the findings of Bonaventure et al. who were able to use their community database to identify between 22.80% and 23.05% of communities in BGP data across eight vantages, and found that the majority—55% on average—were Inbound communities. However, as they were simply using a database of communities, they were unable to obtain predictions on between 76.95% and 78.20% of the communities they observed [8]. This suggests that the majority of communities transiting the control plane have passive semantics.

By contrast, Outbound communities were predicted to comprise a small fraction of the communities collected during both capture intervals. This is a somewhat surprising result, and may be indicative of a period of general stability in the global routing tables with respect to the short duration during which UPDATES were studied at these specific vantages. It may also suggest that Outbound communities are simply used more sparingly because they have the potential to influence routing globally. Furthermore, it is possible that Inbound communities are more likely to propagate improperly through the Internet, and thus are over-represented in community strings. Interestingly, the predicted proportion of Blackhole communities in ground-truth communities was much higher during the second capture interval (increasing from 0.49% to 2.04%). It is unknown if this corresponds to a spike in malicious DoS traffic during this period, or if the model is simply misclassifying Inbound or Outbound communities, given its relatively lower F_1 score for Blackhole communities as described above. Besides this increase, the taxonomic distribution was very similar between capture intervals. Additionally, our estimate for the remaining 85% of communities is proportional during both capture intervals to the conservative estimate using only ground-truth communities. However, our confidence in this result is necessarily lower, given the lack of ground-truth labels for many of these communities.


Finally, it should be noted that well-known communities (see Table 2.2) accounted for less than a tenth of a percent of the communities studied during both capture intervals (approximately 0.01%). This appears to indicate that AS-defined communities see much wider use given that they typically offer far greater granularity of signaling semantics, though this result may be partially due to vantage-specific bias.

To complement this work, the length of community strings observed at each route collector


was measured at these vantages during the two capture intervals. In total, 1,147,308 community strings were collected on August 10, 2020; 814,104 (36.48%) UPDATEs did not utilize the optional community attribute. This information was summarized as a histogram and a CDF in Figure 4.9. This analysis reveals approximately 85% of community strings include ten or fewer communities. The average number of communities per string across all vantages was approximately six communities. This analysis indicates that, for the time interval under study, community string length remains relatively stable; UPDATEs with more than ten communities in their community string appear only intermittently.

Interestingly, 6,819 (0.5%) of these community strings contained thirty or more communities. 585 (0.04%) extreme outliers with more than 50 communities were observed; of these, the five UPDATEs with the longest string contained 386 communities. These extreme values were omitted from Figure 4.9 to ensure its readability. It is unclear if these strings are anomalous—i.e., the result of runaway propagation, misconfiguration, or malicious intent—or if this behavior is intended.

To corroborate this result, 1,290,167 community strings were extracted from BGP data at the same vantages during the second capture interval on September 1, 2020; 810,967 (38.60%) UPDATEs did not utilize the optional community attribute. This information is presented in Figure 4.10. This analysis reveals very little variation from the previous capture interval; the average number of communities per UPDATE was approximately 6 communities, and the proportion of outliers is similar. For example, 669 community strings—0.05%—with more than 50 communities were observed, although the longest of these contained only 173 communities in contrast with the prior capture interval.




(a) Community String Lengths Over Time: August 10, 2020




(b) Community String Length: Histogram and CDF for August 10, 2020

Figure 4.9. The length of the BGP community string observed in UPDATES at three route collectors at RIPE was measured over time on August 10, 2020 between 20:00 and 20:35. For legibility, Figure 4.9a sampled the length of one in every ten community strings.



(a) Community String Lengths Over Time: September 1, 2020




(b) Community String Length: Histogram and CDF for September 1, 2020

Figure 4.10. To corroborate the results obtained in the prior capture interval, the length of the BGP community string observed in UPDATES at three route collectors at RIPE was additionally measured on September 1, 2020 between 4:00 and 4:35. For legibility, Figure 4.10a sampled the length of one in every ten community strings.


To further analyze variation in community string length and predicted semantics, stacked area plots were created to illustrate the instantaneous distribution of community semantics seen at each vantage as predicted by the best model; see Figures 4.11 and 4.12. To ensure the highest confidence in these results, these predictions were restricted to ground-truth communities present in training data.

This analysis suggests that the majority of UPDATES at any given moment—at least for these vantages, and during the two capture intervals—are location-encoding ingress communities, and particularly Inbound: Geographic communities. Often, these strings appear to contain Type of Peer communities as well, presumably to annotate economic relationships between ASes as described in Section 2.1.6; according to the model, the NCC vantage appears to collect more Inbound: Type of Peer communities than either of the other two vantages. However, this may be due simply to the presence of these specific communities in collected ground-truth data. Interestingly, the community strings observed at the NCC vantage during September 1, 2020 (Figure 4.12c) appear to show brief spikes in Blackhole communities, particularly at 4:03 and 4:23; this may be indicative of RTBH events. This correlates with the increased proportion of Blackhole communities observed between the two capture intervals in Table 4.11, and may be its proximate cause.




(a) NYIIX

(b) LINX and LONAP




(c) NCC

Figure 4.11. The instantaneous distribution of community semantics on August 10, 2020 was analyzed for the period between 20:00 and 20:35 at three RIPE vantage points. For the greatest degree of accuracy, only communities for which ground truth were available were considered during this period.



(a) NYIIX

(b) LINX and LONAP



(c) NCC

Figure 4.12. The instantaneous distribution of community semantics on September 1, 2020 was analyzed for the period between 4:00 and 4:35 at three RIPE vantage points. For the greatest degree of accuracy, only communities for which ground truth were available were considered during this period.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 5: Conclusions

This research sought to demystify the semantics of BGP communities using a machine learning approach to classification of communities according to a unified taxonomy of community services. To this end, ground-truth community definitions were collected using web scraping and CAIDA’s Community Dictionary Dataset for use as training data for a predictive model capable of classifying new communities. This research therefore serves as a proof-of-concept for the application of machine learning methods to semantic classification of BGP communities.

This thesis was motivated by a desire to identify the taxonomic distribution of communities, allow AS operators and researchers to identify unknown communities as they transit the control plane, and to inform future work in the creation of a community anomaly detection system designed to detect and even combat community-based threats.

5.1 Major Findings and Implications for BGP Community Semantics

During the data collection process, our research revealed that the Bonaventure taxonomy does not fully account for all types of communities; approximately 0.25% of the 10,027 communities used in training and testing data could not be classified according to any particular taxonomic subclass. One example of this is RTT communities such as are defined by ECIX [47]. However, all communities in collected data can be classified by one of the three primary taxonomic classes (Inbound, Outbound, or Blackhole).

Similarly, many BGP communities observed in ground-truth data also do not conform to the standard syntax proposed by their foundational RFC; that is, the defining AS’s number is not encoded in the most significant word in all 32-bit communities. 118 (1.18%) such communities were collected, and there may be significantly more in ASes from which community data was not collected. Fundamentally, this appears to have arisen from a limitation of the syntax whereby ASes were unable to encode sufficiently granular information in the least significant word of a community. For example, the community 65402:nnn, as defined

by AS 2914 (NTT), means to prepend its AS number twice to peer “nnn” in North America. It would not be possible to encode both the active semantic and the peer’s AS number in the least significant word while maintaining the standard RFC syntax.

Additionally, many Inbound communities are effectively multi-label with respect to the unified taxonomy proposed by Bonaventure [8]. Among observed communities, this is exclusive to Inbound communities used as route tags to encode passive information. 523 such communities (5.22%) were collected; the majority of these encoded both Type of Peer and Geographic information about a particular route.

Finally, this research discovered that community definitions are not uniformly atomic. In the schema used by at least one AS, AS 209, communities have different semantics depending on the presence of other communities. For example, the community string 209:888 209:64520 209:64749 instructs AS 209 not to announce to any peers *except* Deutsche Telekom, while 209:888 alone is simply a route tag indicating a type of peer community [48]. This indicates that, at least in some cases, the entirety of the community string must be analyzed to diagnose the correct signaling semantic in use.

These complexities were not fully captured by the models developed in Chapter 4, and thus represent a limitation of our research. However, the majority of communities collected (93.36%) are clearly differentiable according to the Bonaventure taxonomy and—to the best of our knowledge—exhibit atomicity with respect to their signaling semantics.

Furthermore, our research demonstrates that an MLP-based One-vs-All model is able to predict eight subclasses of the Bonaventure taxonomy with 78.65% aggregate accuracy when presented with novel testing data, and is likewise capable of predicting the three primary classes (Inbound, Outbound, and Blackhole) of the Bonaventure taxonomy with 93.84% accuracy. A Random Forest-based multi-class model attains even higher aggregate accuracy: 90.64% accuracy when trained to recognize eight subclasses of the Bonaventure taxonomy and 97.80% when trained on the the primary three classes.

An unexpected result of this model’s iterative development process is that Blackhole communities, with an F_1 score of 0.76 in the best model, are the most difficult of the primary three classes for our model to effectively classify in novel data, even with relatively high support in testing data. This suggests that Blackhole communities are the most variable (i.e.,

have the fewest common features) between ASes. This must be considered when examining the estimated taxonomic distribution presented in Figure 4.8. Table 4.9 demonstrates that the majority of the model’s mispredictions for this class are, in fact, Inbound communities.

Using this model, we were able to estimate the taxonomic distribution of communities in three vantages at RIPE during two separate capture intervals: between 20:00 and 20:35 on August 10, 2020 and between 4:00 and 4:35 on September 1, 2020. Our model predicts that, among communities seen during training (approximately 15% of all communities), the majority (between 96.71 and 98.43%) of communities collected over this interval were Inbound communities. A much smaller fraction (between 1.08% and 1.24%) of these communities were predicted to be Outbound communities; this may indicate that the intervals studied were periods of relative stability in the global routing tables or that ASes are simply more likely to use Outbound communities sparingly given the associated computational cost with path convergence and exploration. An unexpectedly-high percentage (between 0.49% and 2.04%) of historical communities were classified as Blackhole communities. It is unclear if this is due to a spike in Blackhole traffic (i.e., in response to a DoS attack(s)) during the interval studied, or if the model is simply misclassifying Inbound communities given its relatively lower F_1 score with respect to Blackhole communities. This pattern appeared to hold when the model was applied to the remaining 85% of communities observed at each vantage, with between 94.94% and 94.98% of communities predicted to be Inbound communities. However, this is necessarily a less confident result.

Additionally, the length of community strings over this interval was found to be relatively stable. Among UPDATES utilizing the optional community attribute, approximately 85% of community strings seen during both capture intervals included ten or fewer communities. However, 1.13% of these community strings contained thirty or more communities; between 0.04% and 0.05% of communities contained more than 50. The longest among these contained as many as 386 communities.

To obtain further intuition over the instantaneous distribution of communities in the wild, communities in recent BGP data were additionally fed to our model to create stacked area plots measuring the proportion of community services seen at a given moment in time for a vantage. This analysis suggests that the majority of communities seen at any given moment at a vantage are used to tag ingress and the economic relationship between ASes.

This is a logical result of the constraints of inter-domain routing with respect to economic policy among ASes. Finally, spikes in predicted Blackholing communities during this time may correspond to an RTBH event during one of our capture intervals; this may partially explain the observed increase in the proportion of Blackholing communities between capture intervals.

Unfortunately, this model cannot fully account for the usage of communities over this time interval; in particular, it is unclear whether such spikes in community string length, the instantaneous distribution of predicted community semantics, and other observations represent anomalies or regular routing events. This is an essential question left for future work.

5.2 Future Work

This thesis seeks to serve as a basis for future investigations into anomalous routing events, particularly community-based threats. For example, identification of communities not intended for a given AS may aid in the creation of an intelligent filter capable of offering protection against unintended (or malicious) community-initiated consequences, such as traffic blackholing. Such work should prioritize expanding the database of defined communities to improve the generalizability of our model and elucidate further structure—if such structure indeed exists—in the community space among ASes from which communities have not yet been collected. This process may be aided by further work in web-scraping, and perhaps even in natural language parsing. While our model is able to predict classifications for communities defined by ASes from which definitions were not collected, additional ground-truth data will help to improve accuracy, particularly in poorly-represented subclasses such as Inbound: AS, and perhaps reveal unexpected additional complications such as are outlined in this chapter and in Section 3.1. These complexities should be addressed in future research for the most accurate results.

Future work may also explore a performance analysis on a per-AS basis; that is, are there some ASes which our model is able to classify with greater accuracy? Besides the vantages studied in this thesis, future work may also explore whether taxonomic distribution of communities is heavily influenced by the vantage at which UPDATES are collected. For example, does the type of AS affect this distribution significantly? Does the time of day at a given AS

affect it significantly? Do known anomalous events such as prefix hijacking affect it significantly? In general, future work should expand the application of this model to additional vantages over multiple capture intervals. Finally, can we more strictly estimate prediction confidence? That is, using estimation statistics, can we quantify the uncertainty for a given prediction outcome, given the difficulties in understanding the underlying distribution of community semantics?

Additionally, while extended and large communities have not yet seen the same level of adoption as their 32-bit predecessors [15], future work may expand the scope of this work to examine semantic structure and predictability to these communities. Given that extended and large communities are becoming increasingly prevalent, this will likely be increasingly important to understand the full implications of community usage generally, particularly BGP anomalies which incorporate them.

Anomaly detection, however, is perhaps the most important subject for future work in this area. The greatest challenge in regard to detection of anomalies is the lack of ground-truth knowledge about community usage in the wild. This research can only partially answer this question, insofar as it applies to community semantics and their distribution over time as seen from a particular vantage. While some research has examined features by which to classify these; e.g., longitudinal persistence and prevalence of communities in such strings [57], discovering signatures for attacks and other anomalous behavior is a more challenging problem.

Furthermore, detection of anomalies in community usage over time will require a model architecture capable of responding to changes over time in multi-dimensional data. Future work might involve collection of community strings over time from a specific vantage, including semantic predictions and computations of metrics identified in related work, for use as input to a model capable of classifying complex, multi-dimensional tensors.

One neural network architecture capable of extracting meaningful details from time-series data is a Temporal Convolutional Network (TCN). This architecture applies temporal convolution to an input tensor, such that the model learns to predict subsequent values in the time series; the residuals between the predictor values and the real values in test data can then be calculated and fitted to a multivariate Gaussian distribution. Prediction errors can be used as an indicator for the probability of individual points in the time series being anomalous. Un-

fortunately, while He et al. demonstrated that such a model works well for anomaly detection in time-series data with regular patterns, such as an Electrocardiogram (EKG) [58], it may be the case that changes in the global routing table are too frequent and long-lasting for a TCN to be effective alone in classifying changes in the community string. Furthermore, the communities in use will necessarily reflect very vantage-specific factors, such as economic relationships, geographic location for route ingress, etc. between ASes, and may change over time as prefixes are withdrawn or first announced, further complicating the issue.

Thus, it will likely be necessary to build upon this architecture. For example, a TCN is capable of acting as a feature extraction module in a Fully Convolutional Network (FCN). In fact, Karim et al. demonstrated that an FCN augmented with an LSTM RNN can achieve better than state-of-the-art performance in several applied time-series classification problems. Such models can learn temporal dependencies in sequences, and with the addition of an attention mechanism can learn such dependencies even over longer intervals [50]. We offer this as a suggestion, though many architectural approaches are possible and should be explored in future research for best results.

Building a full anomaly detection engine is a daunting task, but this research has demonstrated that machine learning methods can be applied successfully to BGP community semantics. Demystifying these semantics gives AS operators and researchers more tools to understand community usage in their networks; moreover, prediction of community semantics is likely to be an integral component of any efforts to identify and deter BGP anomalies, particularly community-based threats.

List of References

- [1] S. Hares, Y. Rekhter, and T. Li, “A Border Gateway Protocol 4 (BGP-4),” Internet Engineering Task Force, Tech. Rep., Jan. 2006. Available: <https://tools.ietf.org/html/rfc4271>
- [2] P. Faratin, D. D. Clark, S. Bauer, W. Lehr, P. W. Gilmore, and A. Berger, “The growing complexity of Internet interconnection,” *Communications & Strategies*, no. 72, p. 51, 2008.
- [3] W. Shao, F. Devienne, L. Iannone, and J.-L. Rougier, “On the use of BGP communities for fine-grained inbound traffic engineering,” *arXiv:1511.08336 [cs]*, Nov. 2015, arXiv: 1511.08336. Available: <http://arxiv.org/abs/1511.08336>
- [4] R. Chandra, P. Traina, and T. Li, “BGP communities attribute,” Internet Engineering Task Force, Tech. Rep., Aug. 1996. Available: <https://www.rfc-editor.org/info/rfc1997>
- [5] F. Streibelt, F. Lichtblau, R. Beverly, A. Feldmann, C. Pelsser, G. Smaragdakis, and R. Bush, “Bgp communities: Even more worms in the routing can,” in *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. New York, NY, USA: ACM, 2018, pp. 279–292. Available: <http://doi.acm.org/10.1145/3278532.3278557>
- [6] “BGP community guides,” One Step Consulting, Inc. Accessed: Aug. 20, 2020. [Online]. Available: <https://onestep.net/communities/>
- [7] “Routing,” Nippon Telegraph and Telephone, Tech. Rep., Mar. 2015. Available: <https://www.gin.ntt.net/support-center/policies-procedures/routing/>
- [8] B. Donnet and O. Bonaventure, “On BGP communities,” *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 55–59, Mar. 2008. Available: <http://doi.acm.org/10.1145/1355734.1355743>
- [9] F. Streibelt, “BGP communities: A weapon for the Internet (Part 1),” RIPE Labs, Tech. Rep., Mar. 2019. Accessed: Aug. 10, 2020. [Online]. Available: https://labs.ripe.net/Members/florian_streibelt/bgp-communities-a-weapon-for-the-internet-part-1
- [10] R. Bush, S. Bayraktar, R. Bonica, and J. Borkenhagen, “Policy behavior for well-known BGP communities,” Internet Engineering Task Force, Tech. Rep., Aug. 2019. Available: <https://tools.ietf.org/html/rfc8642>

- [11] O. Nordström and C. Dovrolis, “Beware of BGP attacks,” *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 2, pp. 1–8, Apr. 2004. Available: <https://dl.acm.org/doi/10.1145/997150.997152>
- [12] Q. Li, J. Liu, Y.-C. Hu, M. Xu, and J. Wu, “BGP with BGPsec: Attacks and countermeasures,” *IEEE Network*, vol. 33, no. 4, pp. 194–200, July 2019, conference Name: IEEE Network.
- [13] R. Mahajan, D. Wetherall, and T. Anderson, “Understanding BGP misconfiguration,” in *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM ’02)*. New York, NY, USA: Association for Computing Machinery, Aug. 2002, pp. 3–16. Available: <https://doi.org/10.1145/633025.633027>
- [14] G. Doering, J. Durand, and I. Pepelnjak, “BGP operations and security,” Internet Engineering Task Force, Tech. Rep., Feb. 2015. Available: <https://tools.ietf.org/html/rfc7454>
- [15] E. Aben, “BGP large communities uptake: An update,” RIPE Labs, Tech. Rep., Mar. 2019. Accessed: Aug. 12, 2020. [Online]. Available: <https://labs.ripe.net/Members/emileaben/bgp-large-communities-uptake-an-update>
- [16] “AS rank: A ranking of the largest autonomous systems (AS) in the Internet.” Accessed: Aug. 11, 2020. [Online]. Available: <https://asrank.caida.org/>
- [17] J. Honig, D. Katz, M. Mathis, Y. Rekhter, and J. Yu, “RFC 1164: Application of the Border Gateway Protocol in the Internet,” Internet Engineering Task Force, Tech. Rep., June 1990. Available: <https://datatracker.ietf.org/doc/rfc1164/>
- [18] S. E. Deering, “Internet Protocol, Version 6 (IPv6) specification,” Internet Engineering Task Force, Tech. Rep., Dec. 1998. Available: <https://tools.ietf.org/html/rfc2460>
- [19] R. Hinden, “Internet Protocol, Version 6 (IPv6) specification,” Internet Engineering Task Force, Tech. Rep., July 2017. Available: <https://tools.ietf.org/html/rfc8200>
- [20] V. Fuller and T. Li, “Classless inter-domain routing (CIDR): The internet address assignment and aggregation plan,” Internet Engineering Task Force, Tech. Rep., Aug. 2006. Available: <https://tools.ietf.org/html/rfc4632>
- [21] “Practical usage of the blackhole community,” Noction, Tech. Rep., Aug. 2017. Available: <https://www.noction.com/blog/bgp-blackhole-community>
- [22] “BGP timers detecting dead neighbors and BFD,” Noction, Tech. Rep., Jan. 2016. Available: <https://www.noction.com/blog/bgp-timers>

- [23] C. Villamizar, R. Govindan, and R. Chandra, “BGP route flap damping,” Internet Engineering Task Force, Tech. Rep., Nov. 1998. Available: <https://tools.ietf.org/html/rfc2439>
- [24] “BGP path hunting,” Noction, Tech. Rep., Mar. 2016. Available: <https://www.noction.com/blog/bgp-path-hunting>
- [25] R. Fernando, J. Uttaro, E. Rosen, and P. Mohapatra, “The accumulated IGP metric attribute for BGP,” Internet Engineering Task Force, Tech. Rep., Aug. 2014. Available: <https://tools.ietf.org/html/rfc7311>
- [26] “BGP-based inter-domain traffic engineering,” Noction, Tech. Rep., Aug. 2013. Available: <https://www.noction.com/blog/bgp-based-inter-domain-traffic-engineering>
- [27] L. Gao and J. Rexford, “Stable internet routing without global coordination,” *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 681–692, Dec. 2001. Available: <https://doi.org/10.1109/90.974523>
- [28] “How does BGP select the best routing path,” Noction, Tech. Rep., Jan. 2013. Available: <https://www.noction.com/blog/bgp-best-path-selection-algorithm>
- [29] E. Gregori, A. Improta, L. Lenzi, L. Rossi, and L. Sani, “BGP and inter-AS economic relationships,” in *NETWORKING 2011*, vol. 6641, D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, J. Domingo-Pascual, P. Manzoni, S. Palazzo, A. Pont, and C. Scoglio, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 54–67. Available: http://link.springer.com/10.1007/978-3-642-20798-3_5
- [30] L. Gao, “On inferring autonomous system relationships in the Internet,” *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, Dec. 2001.
- [31] Y. Rekhter and P. Gross, “Application of the Border Gateway Protocol in the Internet,” Internet Engineering Task Force, Tech. Rep., Mar. 1995. Available: <https://tools.ietf.org/html/rfc1772>
- [32] A. Mitseva, A. Panchenko, and T. Engel, “The state of affairs in BGP security: A survey of attacks and defenses,” *Computer Communications*, vol. 124, pp. 45–60, June 2018. Available: <http://www.sciencedirect.com/science/article/pii/S014036641731068X>
- [33] “What is RPKI?” RIPE Network Coordination Centre, Tech. Rep., May 2019. Available: <https://www.ripe.net/manage-ips-and-asns/resource-management/certification/what-is-rpki>

- [34] F. Baker, “Internet routing with MANRS,” Internet Society, Tech. Rep., June 2018. Available: <https://www.manrs.org/wp-content/uploads/2018/11/Internet-Routing-with-MANRS.pdf>
- [35] “Border Gateway Protocol (BGP) well-known communities,” Internet Assigned Number Authority, Tech. Rep., Sep. 2019. Available: <https://www.iana.org/assignments/bgp-well-known-communities/bgp-well-known-communities.xhtml>
- [36] J. Scudder, P. Mohapatra, D. Smith, J. Uttaro, and R. Raszuk, “BGP AC-CEPT_OWN community attribute,” Internet Engineering Task Force, Tech. Rep., Aug. 2015. Available: <https://tools.ietf.org/html/rfc7611>
- [37] B. Decraene, K. Patel, C. Pelsser, C. Filsfils, and P. Francois, “Graceful BGP session shutdown,” Internet Engineering Task Force, Tech. Rep., Mar. 2018. Available: <https://tools.ietf.org/html/rfc8326>
- [38] G. Huston, “NOPEER community for Border Gateway Protocol (BGP) route scope control,” Internet Engineering Task Force, Tech. Rep., Apr. 2004. Available: <https://tools.ietf.org/html/rfc3765>
- [39] C. Dietzel, G. Doering, T. King, G. Hankins, and J. Snijders, “Blackhole community,” Internet Engineering Task Force, Tech. Rep., Oct. 2016. Available: <https://tools.ietf.org/html/rfc7999>
- [40] S. Sangli, D. Tappan, and Y. Rekhter, “RFC 4360 - BGP extended communities attribute,” Internet Engineering Task Force, Tech. Rep., Feb. 2006. Available: <https://tools.ietf.org/html/rfc4360>
- [41] F. Cady, *The Data Science Handbook*. New York, United States: John Wiley & Sons, Incorporated, 2017. Available: <http://ebookcentral.proquest.com/lib/ebook-nps/detail.action?docID=4790656>
- [42] L. Van Efferen and A. M. Ali-Eldin, “A multi-layer perceptron approach for flow-based anomaly detection,” in *2017 International Symposium on Networks, Computers and Communications (ISNCC)*, May 2017, pp. 1–6.
- [43] C. Strobl, J. Malley, and G. Tutz, “An introduction to recursive partitioning: Rationale, application, and characteristics of classification and regression trees, bagging, and random forests,” *Psychological Methods*, vol. 14, no. 4, pp. 323–348, 2009.
- [44] R. Blagus and L. Lusa, “SMOTE for high-dimensional class-imbalanced data,” *BMC Bioinformatics*, vol. 14, no. 1, p. 106, Mar. 2013. Available: <https://doi.org/10.1186/1471-2105-14-106>

- [45] Y. E. Kurniawati, A. E. Permanasari, and S. Fauziati, “Adaptive synthetic-nominal (ADASYN-N) and adaptive synthetic-KNN (ADASYN-KNN) for multiclass imbalance learning on laboratory test data,” in *2018 4th International Conference on Science and Technology (ICST)*, Aug. 2018, pp. 1–6.
- [46] UCSD - Center For Applied Internet Data Analysis, “CAIDA UCSD BGP community dictionary dataset - <01/01/2018 and 04/19/2018>,” 2018. Accessed: Mar. 20, 2020. [Online]. Available: <https://www.caida.org/data/bgp-communities/>
- [47] “ECIX’s new route server RTT communities,” European Commercial Internet Exchange, Tech. Rep., July 2017. Available: <https://www.ecix.net/about-us/news/ecixs-new-route-server-rtt-communities.html>
- [48] “AS-AS209,” Hurricane Electric Internet Services, Tech. Rep., Aug. 2020. Available: <https://bgp.he.net/irr/as-set/AS-AS209>
- [49] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-Learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [50] F. Karim, S. Majumdar, H. Darabi, and S. Chen, “LSTM fully convolutional networks for time series classification,” *IEEE Access*, vol. 6, pp. 1662–1669, 2018, conference Name: IEEE Access.
- [51] W. Koehrsen, “Beyond accuracy: Precision and recall,” Mar. 2018. Accessed: Aug. 14, 2020. [Online]. Available: <https://towardsdatascience.com/beyond-accuracy-precision-and-recall-3da06bea9f6c>
- [52] SciKit-Learn Developers, “Plotting precision-recall.” Accessed: Aug. 14, 2020. [On-line]. Available: https://scikit-learn.org/stable/auto_examples/model_selection/plot_precision_recall.html
- [53] RIPE Network Coordination Centre, “RIS raw data.” Accessed: Aug. 14, 2020. [Online]. Available: <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data>
- [54] “Bgpview,” Security Trails. Accessed: Sep. 8, 2020. [Online]. Available: <http://bgpview.io/>
- [55] T. Fawcett, “ROC graphs: Notes and practical considerations for researchers,” *Machine Learning*, vol. 31, pp. 1–38, Jan. 2004. Available: https://www.researchgate.net/publication/284043217_ROC_Graphs_Notes_and_Practical_Considerations_for_Researchers

- [56] T. Saito and M. Rehmsmeier, “The Precision-Recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets,” *PLoS ONE*, vol. 10, no. 3, Mar. 2015. Available: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4349800/>
- [57] A. Hardt, “Characterizing BGP community irregularities toward an anomaly detection engine,” Master’s thesis, Dept. Computer Science, Naval Postgraduate School, Monterey, California, Dec. 2019. Available: <http://hdl.handle.net/10945/64179>
- [58] Y. He and J. Zhao, “Temporal convolutional networks for anomaly detection in time series,” *Journal of Physics: Conference Series*, vol. 1213, p. 042050, June 2019, publisher: IOP Publishing. Available: <https://doi.org/10.1088/1742-6596/1213/4/042050>

Initial Distribution List

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California