METASCRITIC, REFRAMING TOPOLOGY DISCOVERY AS A RECOMMENDER SYSTEM

<u>LOQMAN SALAMATIAN</u>, KEVIN VERMEULEN, ITALO CUNHA, VASILIS GIOTSAS, ETHAN KATZ-BASSETT

COLUMBIA UNIVERSITY IN THE CITY OF NEW YORK

CNRS

CLOUDFLARE

UFMG UNIVERSIDADE FEDERAL DE MINAS GERAIS

Published at ACM IMC '24

Punchline:

Inferred 34x AS links than in measured Internet, with > 80% recall and precision across multiple datasets

METASCRITIC, REFRAMING TOPOLOGY DISCOVERY AS A RECOMMENDER SYSTEM

LOQMAN SALAMATIAN, KEVIN VERMEULEN, ITALO CUNHA, VASILIS GIOTSAS, ETHAN KATZ-BASSETT

COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK
cnrs
CLOUDFLARE
UFMG
UNIVERSIDADE FEDERAL DE MINAS GERAIS

Published at ACM IMC '24

THE POWER OF A MORE COMPLETE AS-TOPOLOGY

- **Researchers:** Build accurate models, assess network resilience, and analyze the impact of peering connectivity on Internet properties.

- **Policymakers:** Make informed decisions on Internet governance, infrastructure, and security.

- **Operators:** Gain insights for better network planning, traffic optimization, and troubleshooting.

SOME PART OF THE INTERNET TOPOLOGY IS HIDDEN.

Reason: ASes typically advertise their peering links only to their customers, so only vantage points in the customers can measure the links.

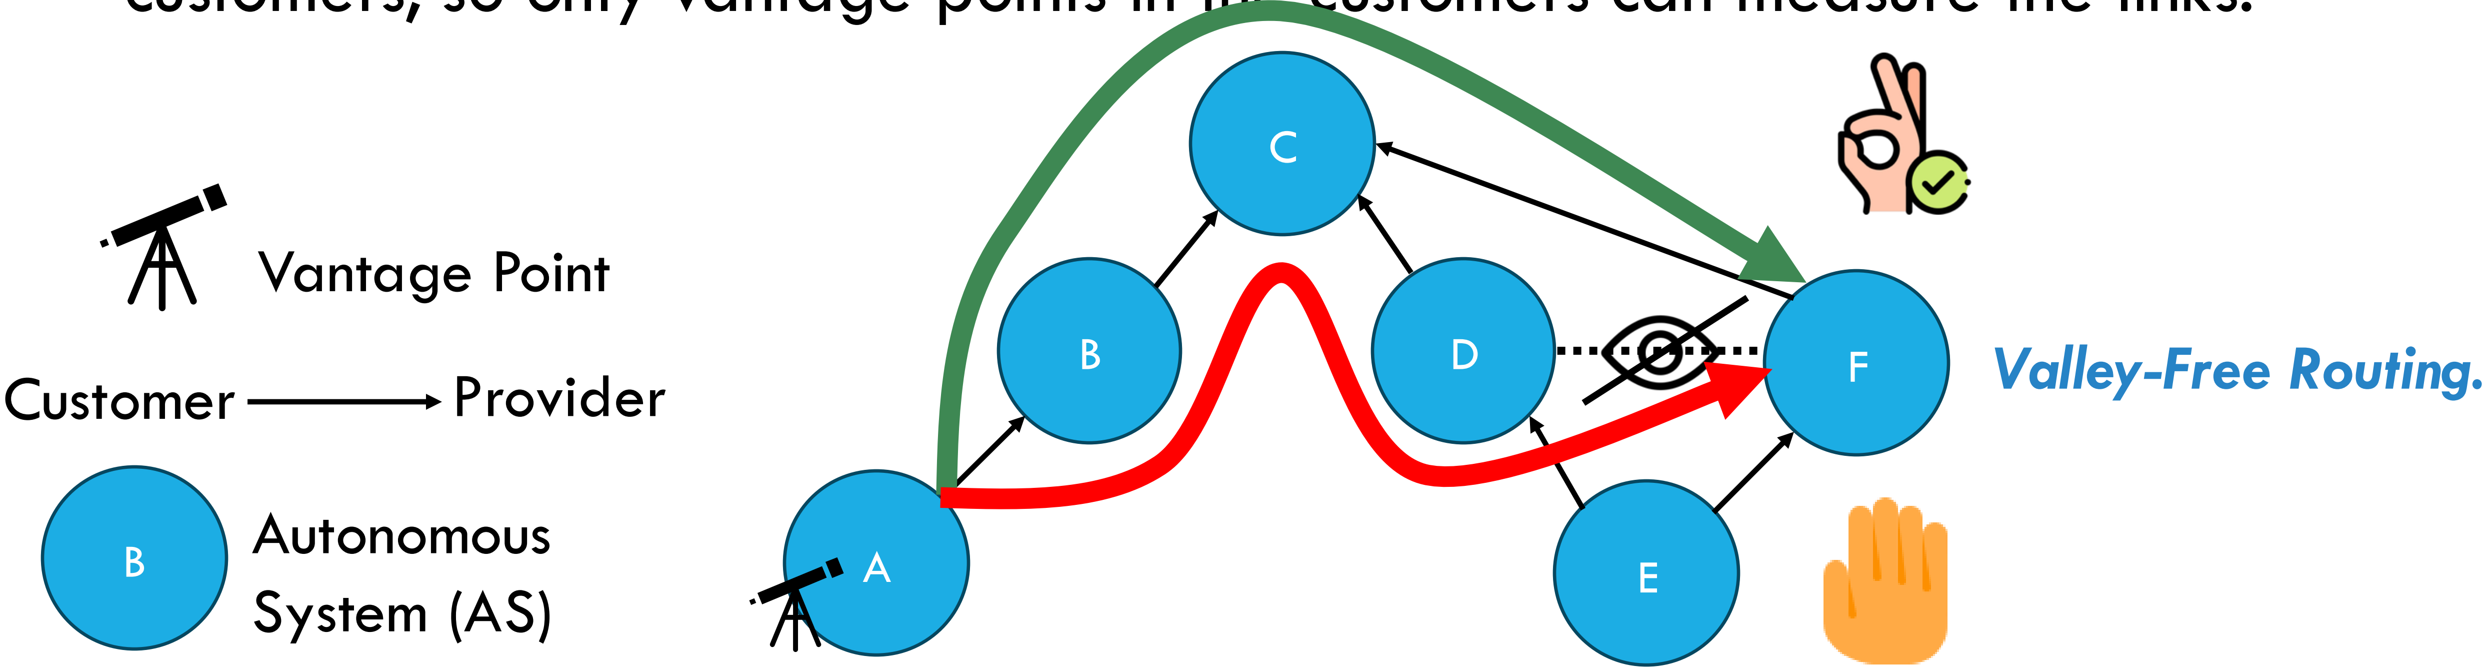

SOME PART OF THE INTERNET TOPOLOGY IS HIDDEN.

Reaso
custon ... s.

Based on available vantage points, only 6% of the potential peering links can be faithfully recovered!

V

Customer—

ee Routing.

B Autonomous System (AS)

DESPITE DECADES OF EFFORTS, OUR VISIBILITY HAS LIKELY DWINDLED.

For a single IXP: *Ager et al*. found nearly **50K** peering interconnections, more than the number observed by publicly available monitors [1].

[1] Anatomy of a Large European IXP – Ager et al. in ACM SIGCOMM 2012

DESPITE DECADES OF EFFORTS, OUR VISIBILITY HAS LIKELY DWINDLED.

For a single IXP: *Ager et al*. found nearly **50K** peering interconnections, more than the number observed by publicly available monitors [1].

For a single AS: *Arnold et al*. found that more than **90%** of the Google and Microsoft peering links were invisible from BGP feeds [2].

[1] Anatomy of a Large European IXP – Ager et al. in ACM SIGCOMM 2012

[2] Cloud Providers Connectivity – Arnold et al. in ACM IMC 2020

DESPITE DECADES OF EFFORTS, OUR VISIBILITY HAS LIKELY DWINDLED.

For a single IXP: *Ager et al.* found nearly **50K** peering interconnections, more than the number observed by publicly available monitors [1].

For a single AS: *Arnold et al.* found that more than **90%** of the Google and Microsoft peering links were invisible from BGP feeds [2].

Richer connectivity structure and **stagnating VPs [3]:**

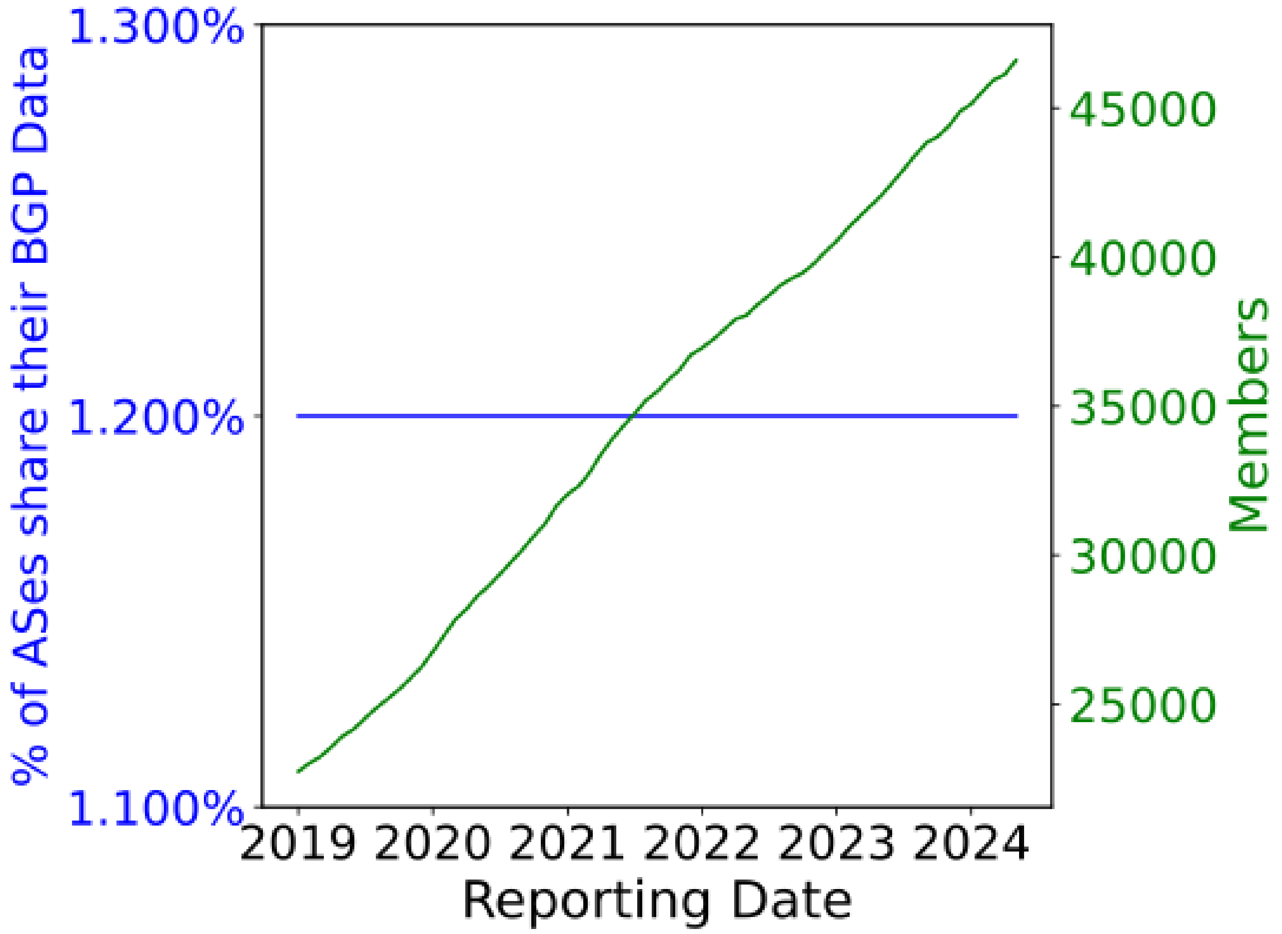

[1] Anatomy of a Large European IXP – Ager et al. in ACM SIGCOMM 2012

[2] Cloud Providers Connectivity – Arnold et al. in ACM IMC 2020

[3] The Next Generation of BGP Data Collection Platforms – Alfroy et al. in ACM SIGCOMM 2024

THE NEED FOR A FUNDAMENTAL SHIFT: INFERENTIAL APPROACHES TO THE RESCUE

Inferential approaches extend our limited coverage by using patterns in the visible topology to make educated guesses about the unseen parts.

Challenge: **Inferential techniques** introduce a new kind of uncertainty.

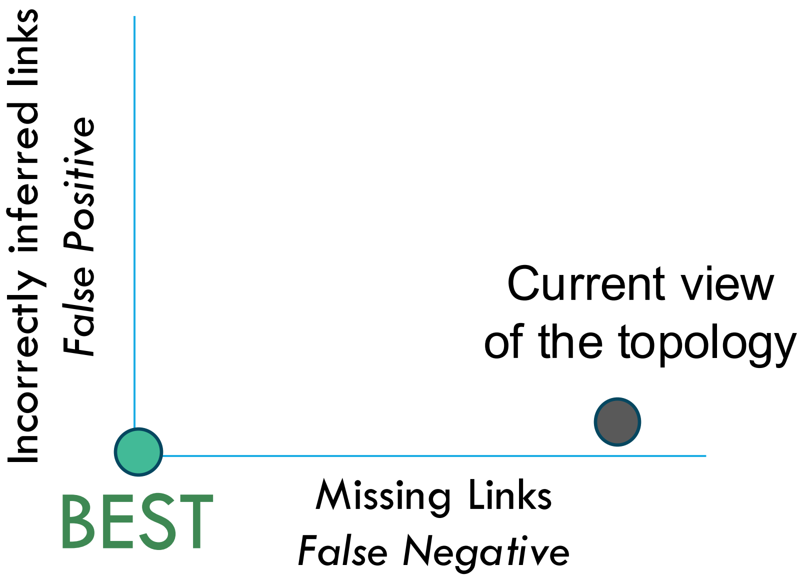

THE NEED FOR A FUNDAMENTAL SHIFT: INFERENTIAL APPROACHES TO THE RESCUE

Inferential approaches extend our limited coverage by using patterns in the visible topology to make educated guesses about the unseen parts.

Challenge: Inferential techniques introduce a new kind of uncertainty.

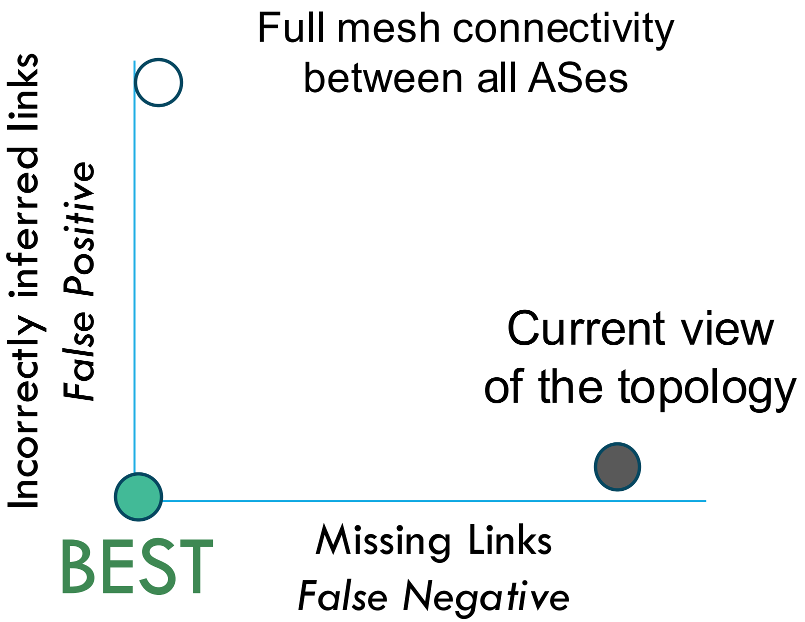

THE NEED FOR A FUNDAMENTAL SHIFT: INFERENTIAL APPROACHES TO THE RESCUE

The insights gained from a more complete picture of the topology can outweigh the inherent uncertainty of inferential methods.

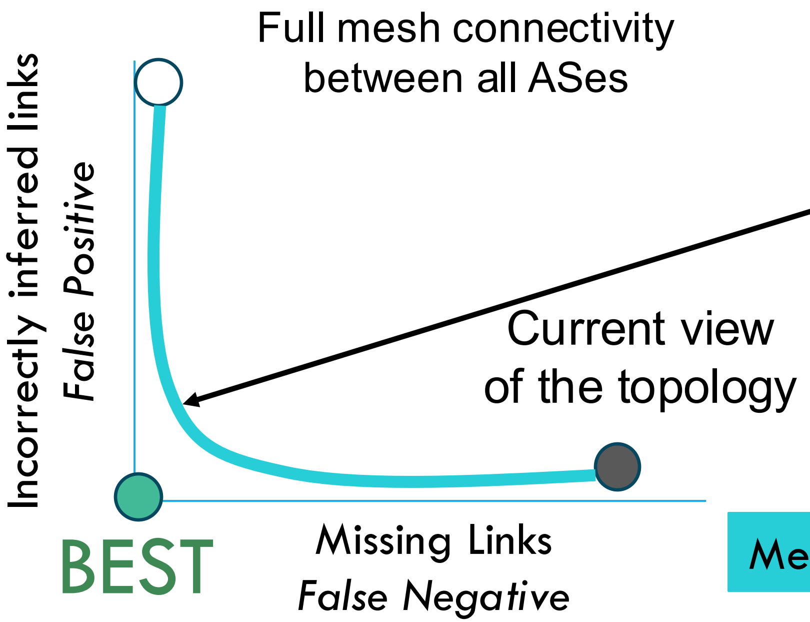

All of this trade-off space is unexplored. MetAScritic enables operating at various points on this tradeoff line.

Metro-AS critic (shortened to MetAScritic)

BUILDING APPLICATION-DEPENDENT VIEWS OF THE TOPOLOGY

The insights gained from a more complete picture of the topology can outweigh the inherent uncertainty of inferential methods.

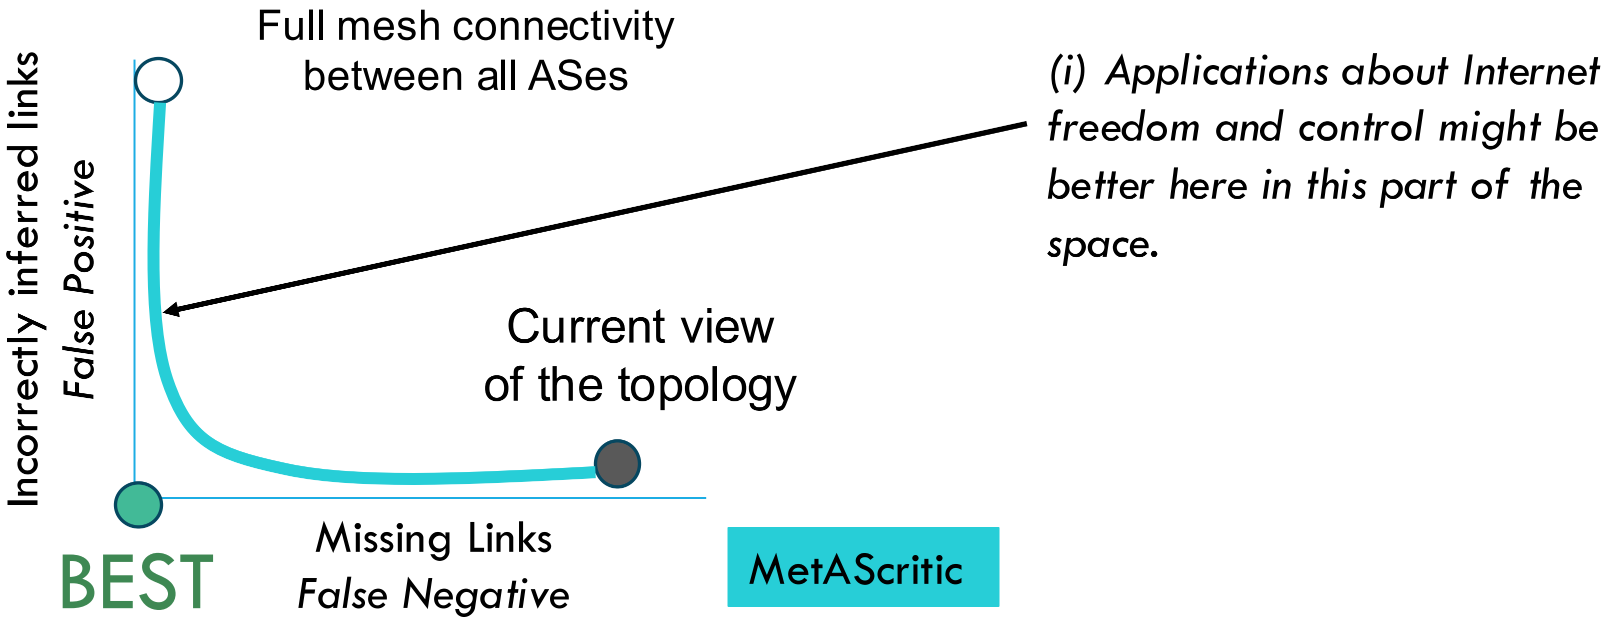

MetAScritic

BUILDING APPLICATION-DEPENDENT VIEWS OF THE TOPOLOGY

The insights gained from a more complete picture of the topology can outweigh the inherent uncertainty of inferential methods.

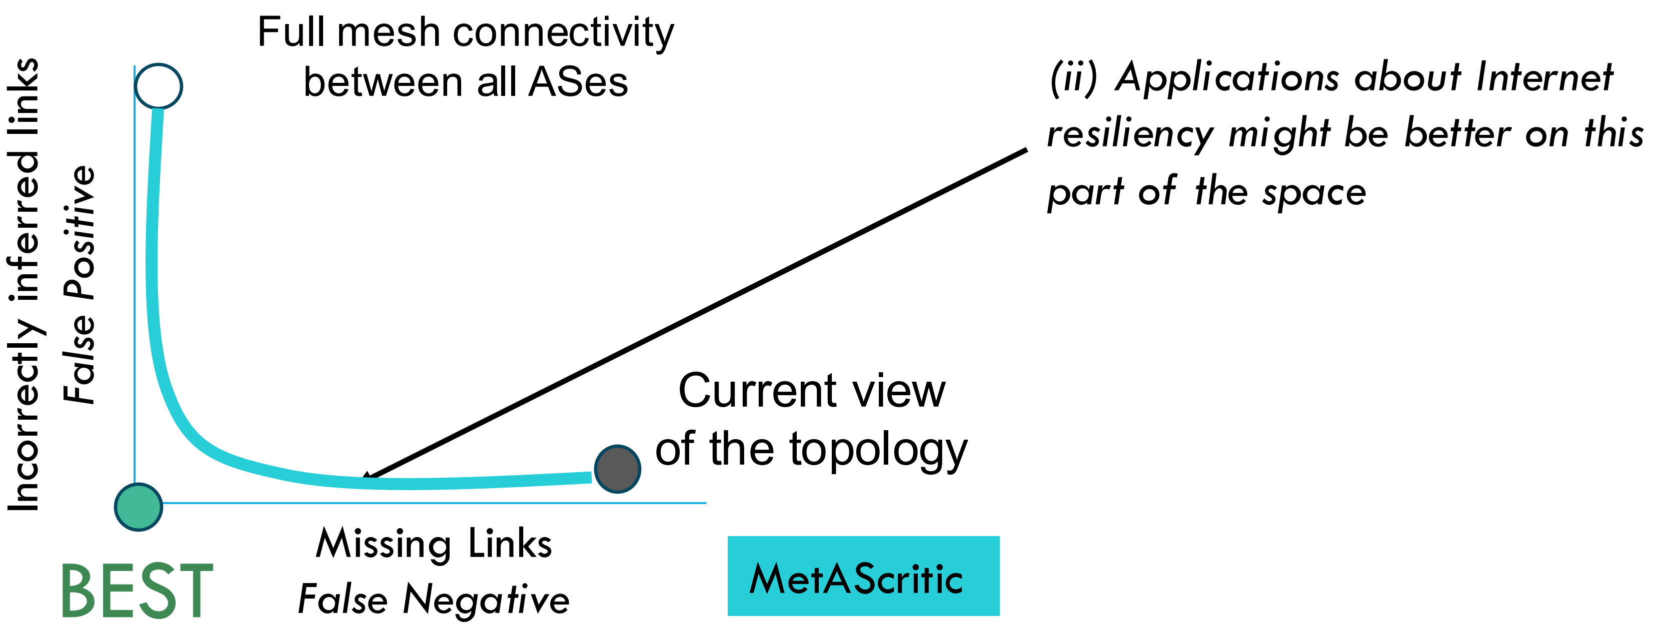

OBSERVED PEERING REVEALS MANY VALUABLE INSIGHTS ABOUT THE UNSEEN ONES.

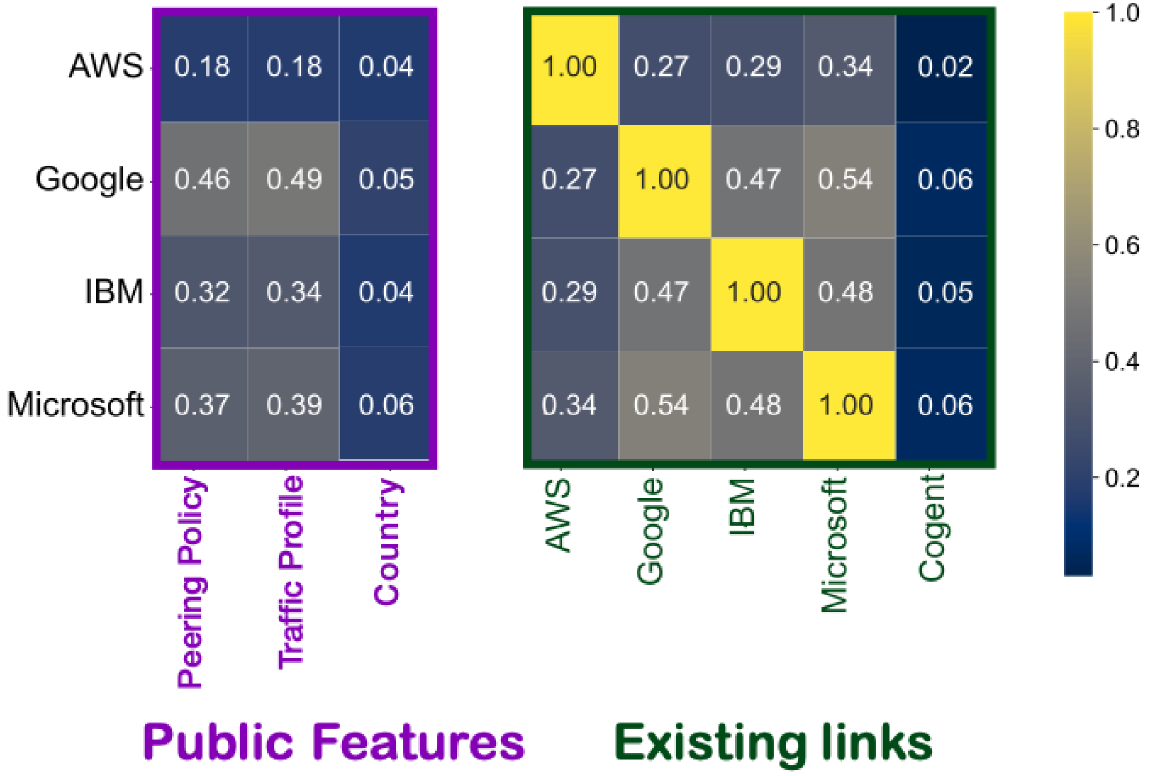

Setup:

We measure peering links for (i) AWS, (ii) Google, (iii) IBM, and (iv) Microsoft.

We analyze all available features to identify the best predictor for the existence of a peering link with a given cloud provider for a given AS.

OBSERVED PEERING REVEALS MANY VALUABLE INSIGHTS ABOUT THE UNSEEN ONES.

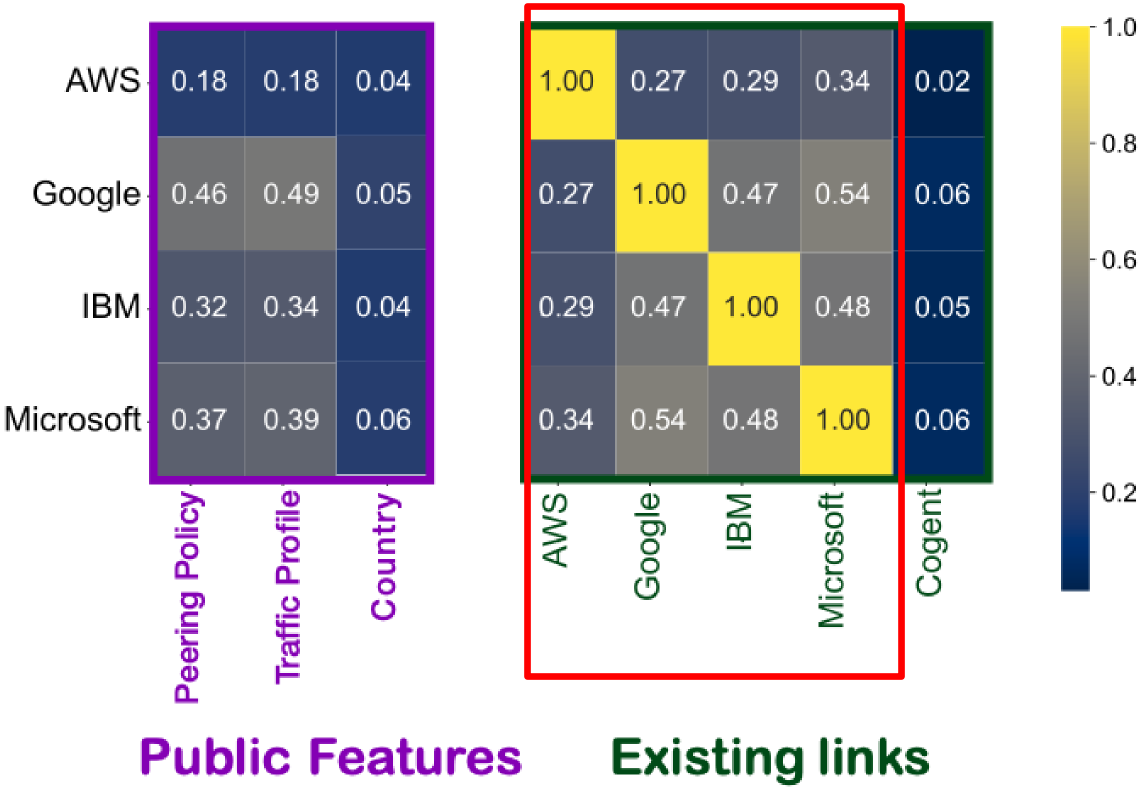

Setup:

We measure peering links for (i) AWS, (ii) Google, (iii) IBM, and (iv) Microsoft.

We analyze all available features to identify the best predictor for the existence of a peering link with a given cloud provider for a given AS.

Result:

ASes that (i) already peer with one cloud provider, are much more likely to establish peering links with other cloud providers.

OBSERVED PEERING REVEALS MANY VALUABLE INSIGHTS ABOUT THE UNSEEN ONES.

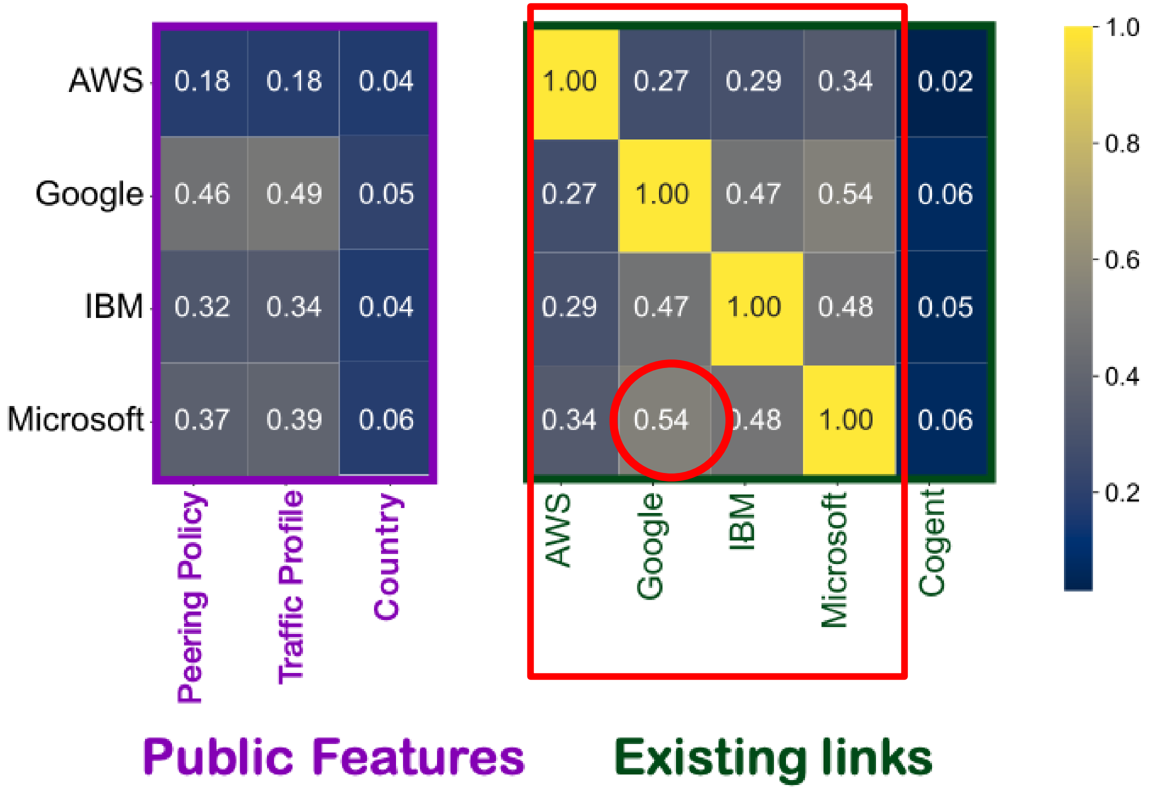

Setup:

We measure peering links for (i) AWS, (ii) Google, (iii) IBM, and (iv) Microsoft.

We analyze all available features to identify the best predictor for the existence of a peering link with a given cloud provider for a given AS.

Result:

ASes that (i) already peer with one cloud provider, are much more likely to establish peering links with other cloud providers.

OBSERVED PEERING REVEALS MANY VALUABLE INSIGHTS ABOUT THE UNSEEN ONES.

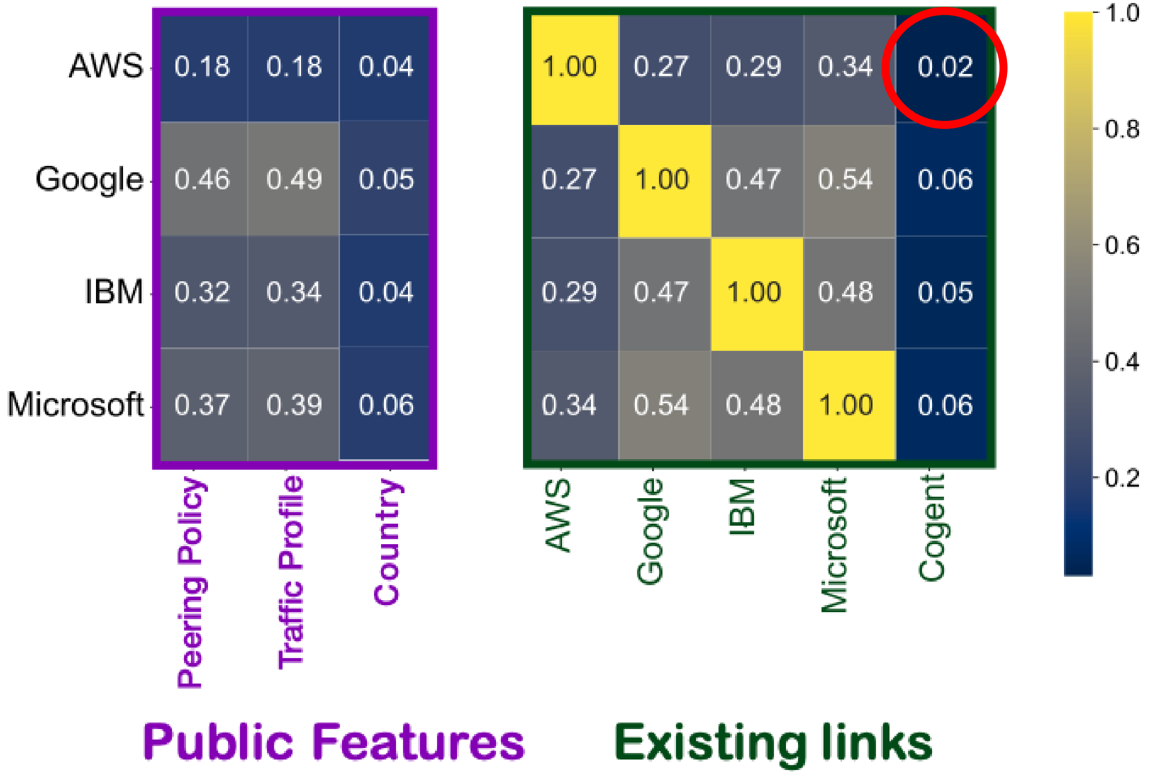

Setup:

We measure peering links for (i) AWS, (ii) Google, (iii) IBM, and (iv) Microsoft.

We analyze all available features to identify the best predictor for the existence of a peering link with a given cloud provider for a given AS.

Result:

ASes that (i) already peer with one cloud provider, are much more likely to establish peering links with other cloud providers.

OBSERVED PEERING REVEALS MANY VALUABLE INSIGHTS ABOUT THE UNSEEN ONES.

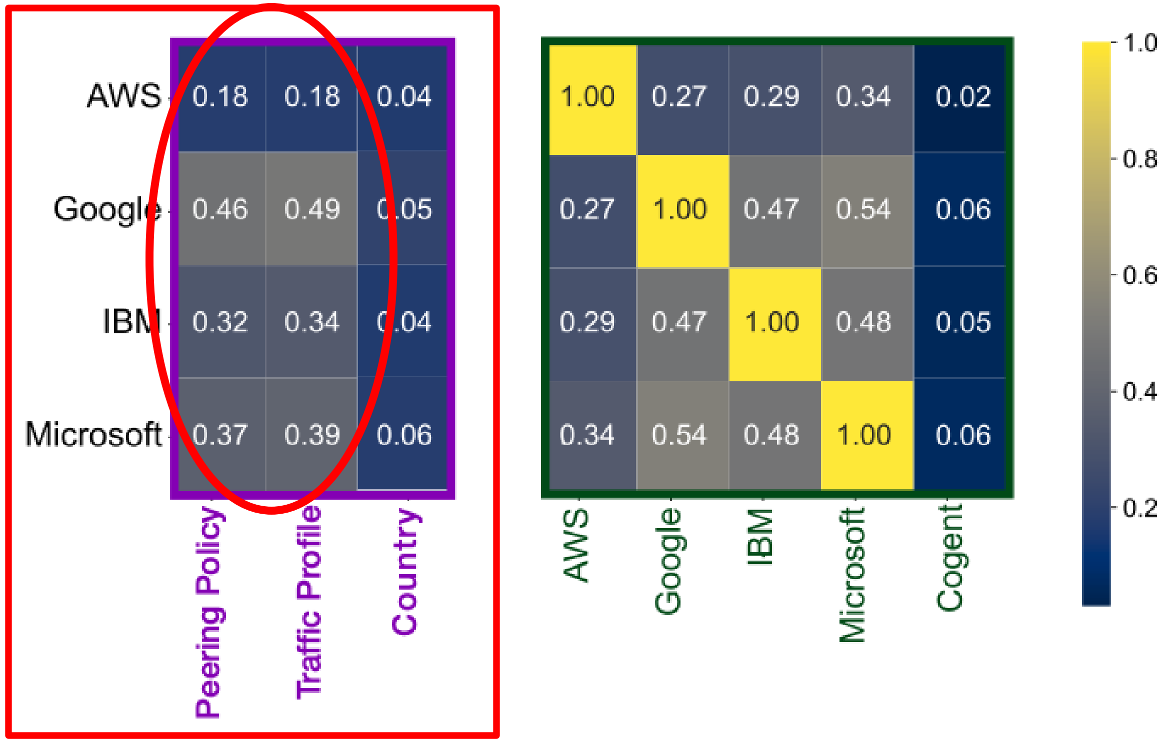

Public Features **Existing links**

Setup:

We measure peering links for (i) AWS, (ii) Google, (iii) IBM, and (iv) Microsoft.

We analyze all available features to identify the best predictor for the existence of a peering link with a given cloud provider for a given AS.

Result:

ASes that (i) already peer with one cloud provider, (ii) with selective peering policy or (iii) heavy inbound traffic profile are much more likely to establish peering links with other cloud providers.

OUR SOLUTION – METASCRITIC INSPIRED BY RECOMMENDER SYSTEMS

Key Idea: ASes with aligned peering strategies in a metropolitan area (metro)—driven by factors like **infrastructure, traffic profiles, business models, geopolitics,** and **history**—are likely to share similar peers.

OUR SOLUTION – METASCRITIC INSPIRED BY RECOMMENDER SYSTEMS

Treating AS connectivity as a recommendation system:

- Tinder or Netflix predict whether a user will like another user/movie based on user characteristics and interaction history.
- metAScritic leverages **AS features** and **known peering links** to infer missing connections.

UNDERSTANDING RECOMMENDATION IN THE CONTEXT OF TINDER.

Tinder analogy:

Intrinsic Properties:
Age: 32 years
Height: 1,75 m
Profession: Magician
Likes: Gandalf
Gender: Male

Intrinsic Properties:
Age: 29 years
Height: 1,55 m
Profession: Scientist
Likes: MetAScritic
Gender: Female

Existing Behavior:
Likes people who love the Lord of the Rings.

Existing Behavior:
Dislikes people who are into magic.

UNDERSTANDING RECOMMENDATION IN THE CONTEXT OF TINDER.

Tinder analogy:

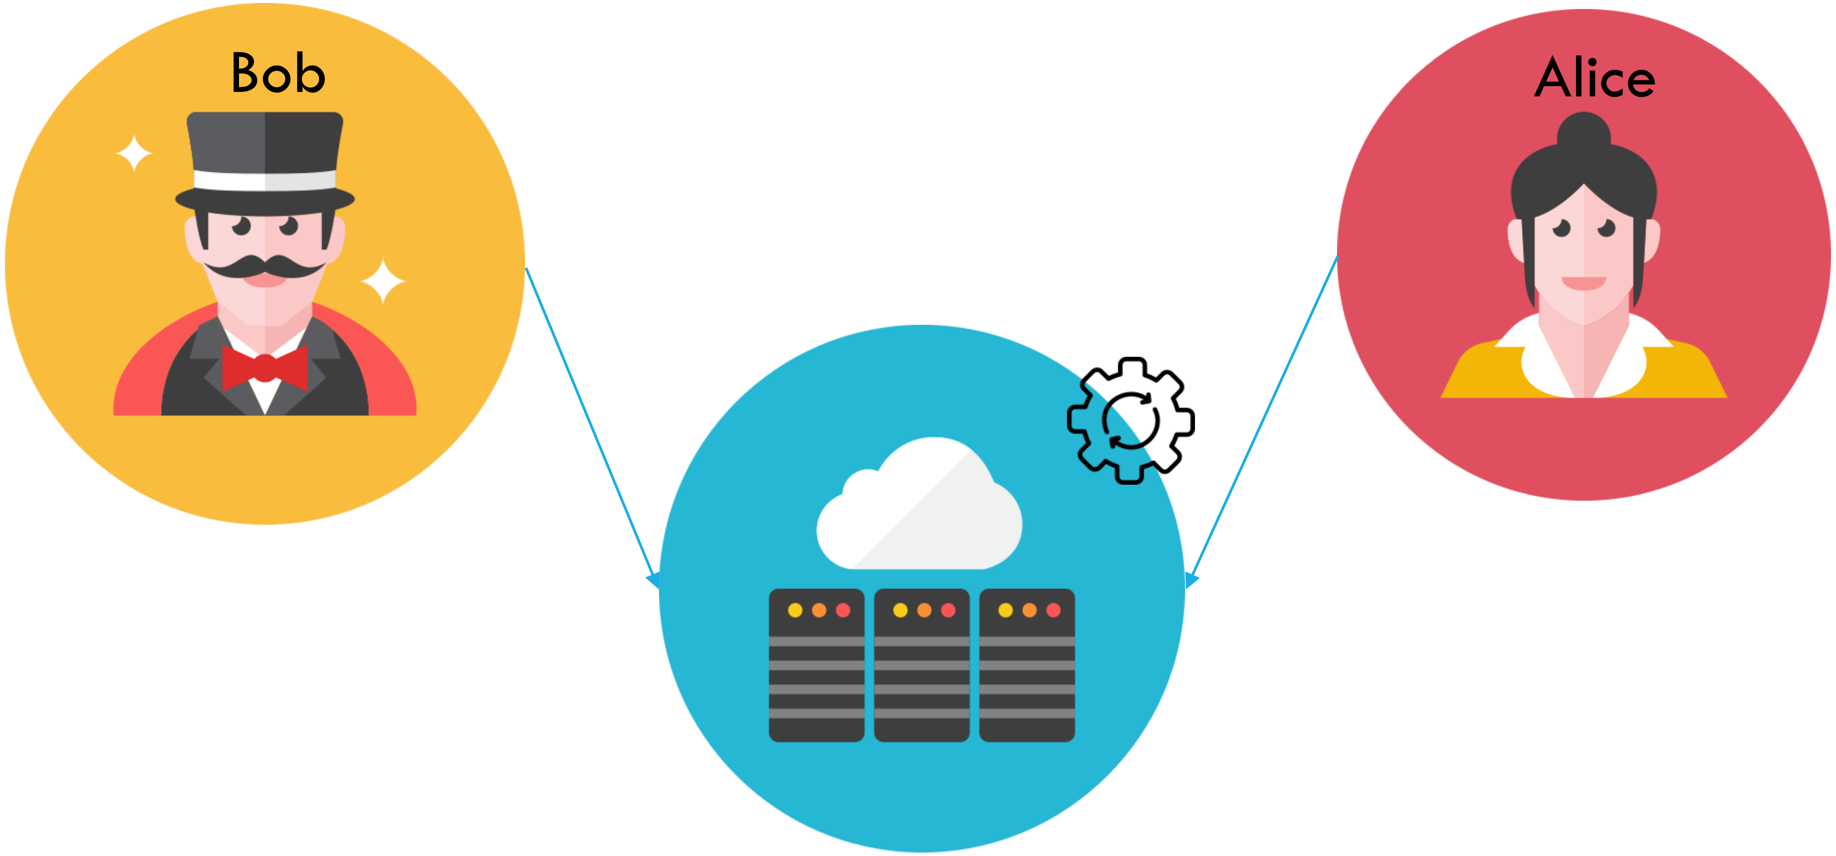

Recommender System

UNDERSTANDING RECOMMENDATION IN THE CONTEXT OF TINDER.

Tinder analogy:

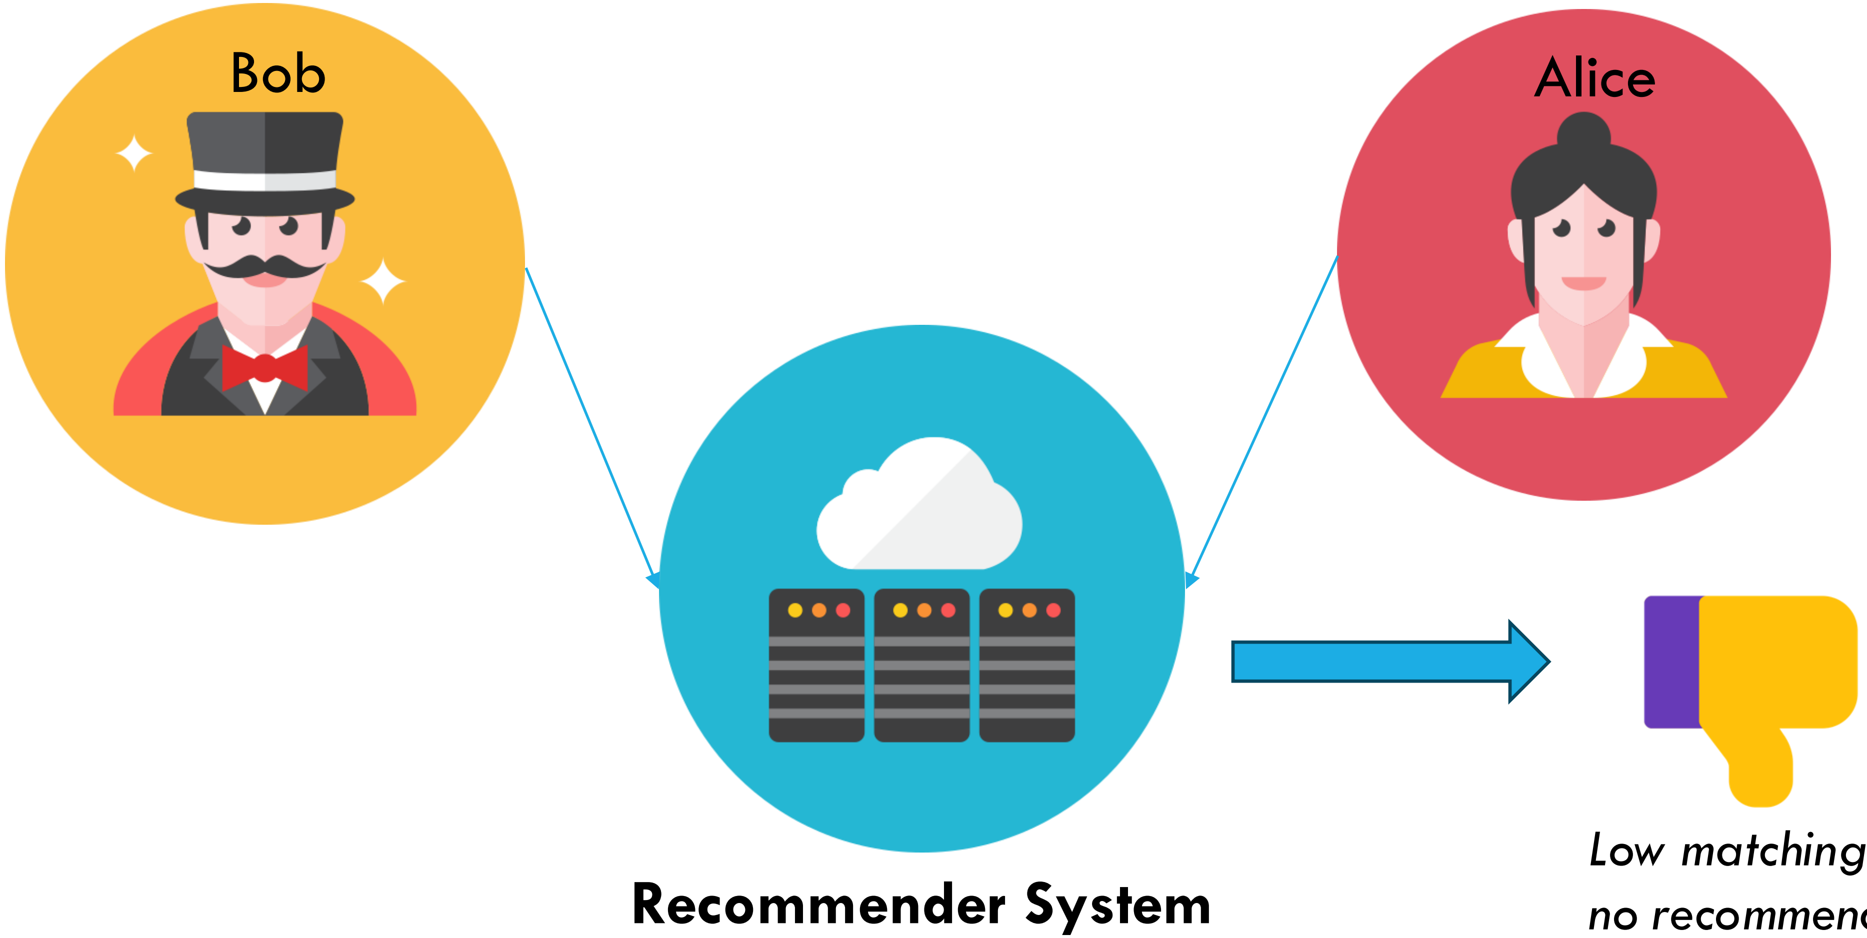

Low matching score, no recommendation.

METASCRITIC DOES THE SAME BUT WITH PEERING CONNECTIVITY INSTEAD.

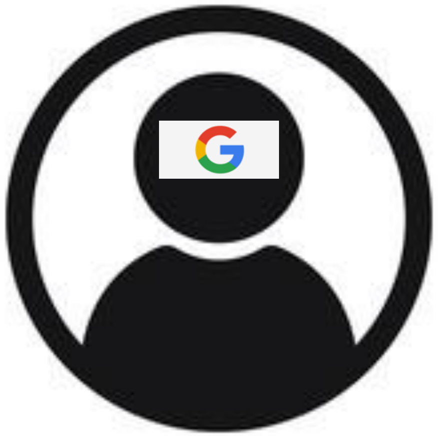

Intrinsic Properties:
Peering Policy: Open
Traffic Profile: Heavily Outbound
Number of Eyeballs: 1M
Customer Cone Size: 23
…

Intrinsic Properties:
Peering Policy: Selective
Traffic Profile: Heavily Inbound
Number of Eyeballs: 42M
Customer Cone Size: 2372
…

Existing Behavior:
Is peering with large access networks.
Is peering with ASes that peer with other Cloud Providers and CDNs.

Existing Behavior:
Is peering with Cloud Providers.
Is unlikely to peer with Open ASes.

METASCRITIC DOES THE SAME BUT WITH PEERING CONNECTIVITY INSTEAD.

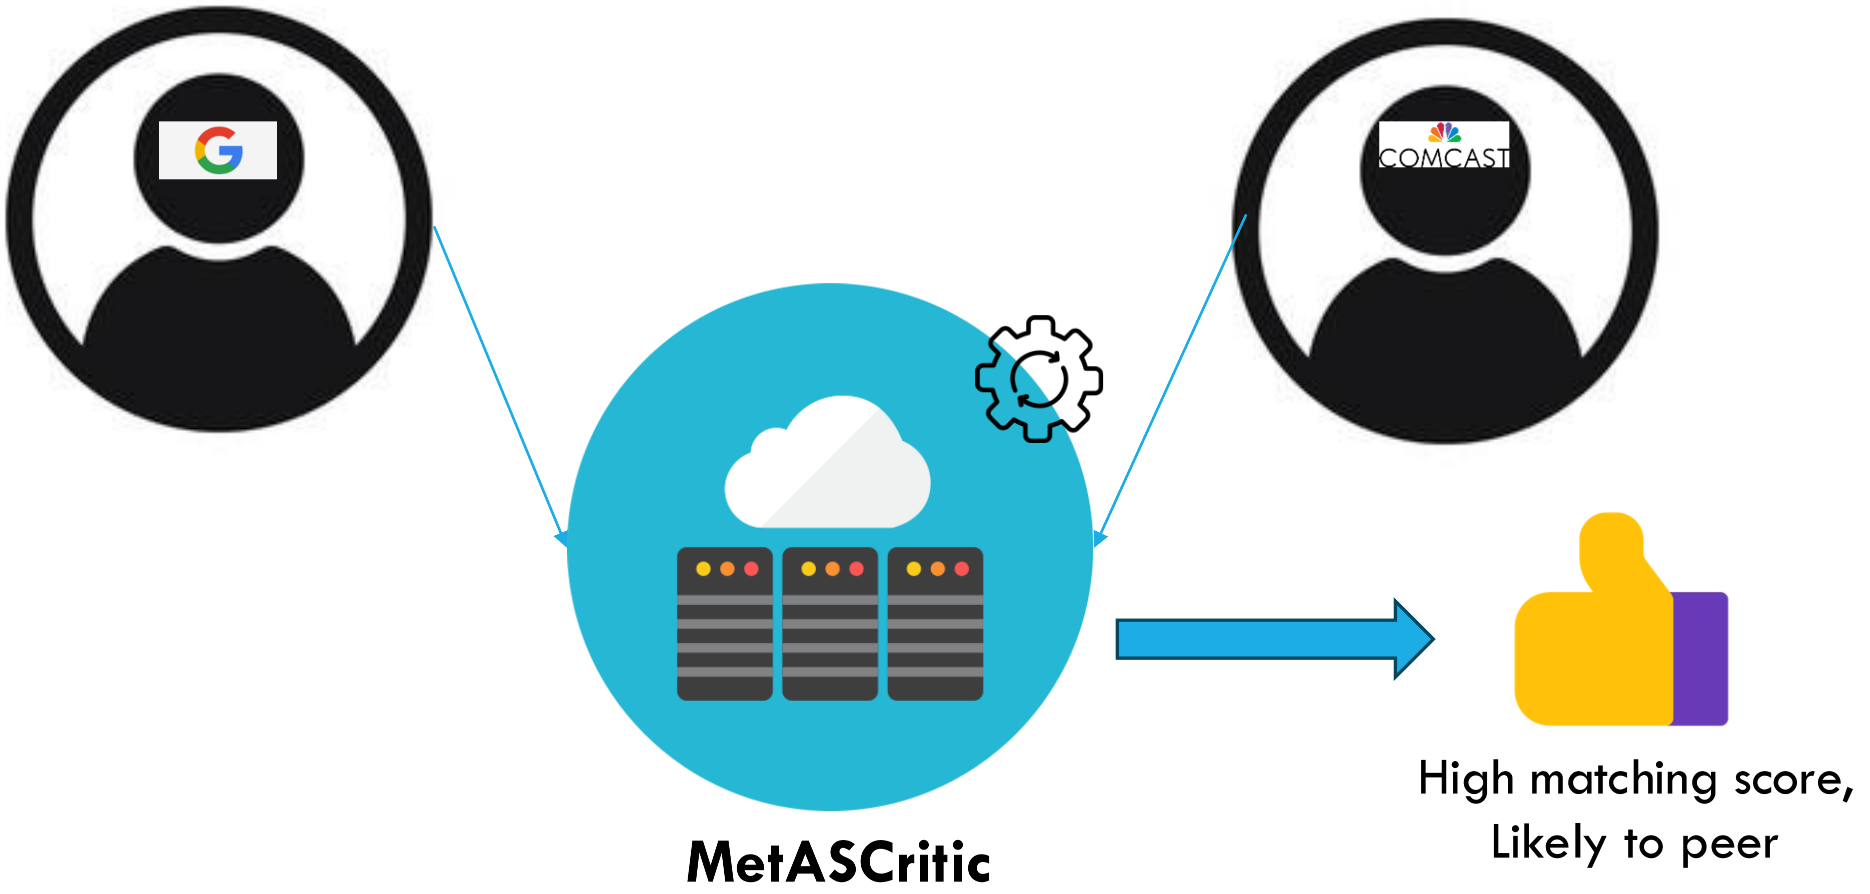

USING RECOMMENDER SYSTEMS FOR INFERENCES.

Recommender systems typically suggest potential connections, but for AS peering, these suggestions often transition into accurate predictions.

Why?

1. **Rational decision-making based on optimization considerations.**
2. **Dedicated staff tasked with peering decision.**
3. **Limited options of peering partners.**

METHODOLOGY: HOW DOES METASCRITIC WORK?

METASCRITIC FOCUSES ON A METRO

- **Step 1: Pick a metro of interest**

Facilities in Madrid

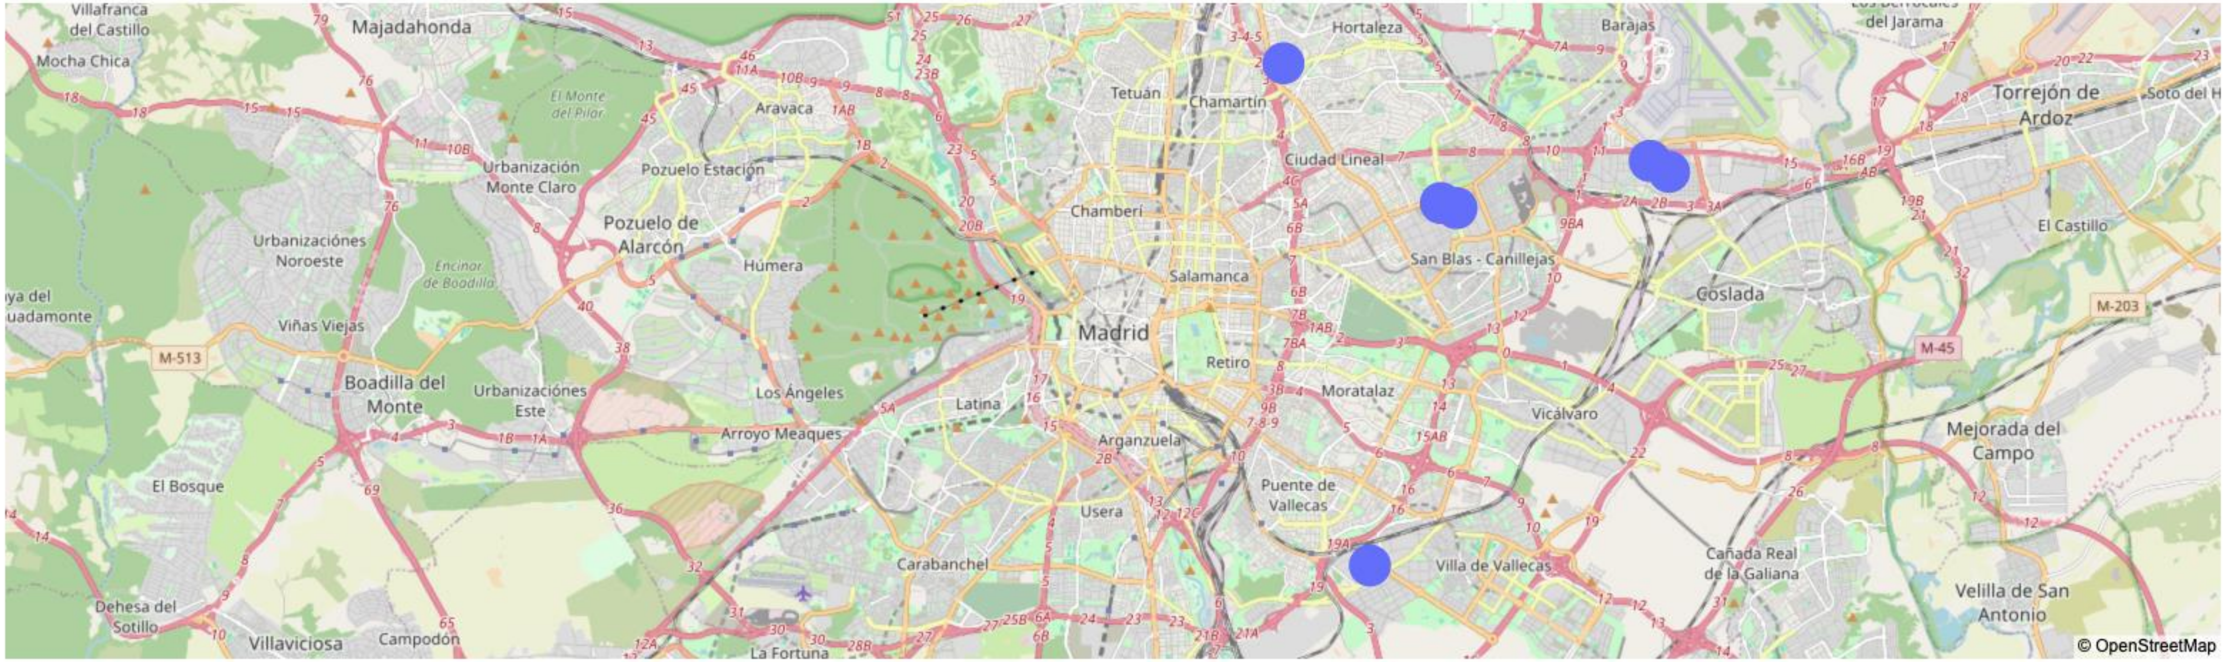

METASCRITIC IDENTIFIES ASES OF INTEREST

- **Step 2: Identify the ASes that have a presence in metro:**

In the facilities - Equinix Madrid

Peer Name	ASN
Adamo Telecom Iberia S.A	35699
Aire Networks del Mediterraneo	29119
Akamai Technologies	20940
Algeria Telecom	36947
Altecom	16030
Amazon IVS	46489
Amazon.com	16509
Angola Cables	37468
Arelion (Twelve99)	1299

In the IXPs – DE-CIX Madrid

Peer Name… / IPv4	ASN / IPv6	Speed / Port L…	Policy…
A J SHERIFF ELECTRICAL	202087	10G	Open
185.1.192.27	2001:7f8:a0:0:3:1567:0:1		
Abaclouda	56987	10G	Open
185.1.192.32	2001:7f8:a0::de9b:0:1		
ADAM	15699	2G	Open
185.1.192.19	2001:7f8:a0::3d53:0:1		
Adamo Telecom Iberia S.A	35699	100G	Open
185.1.192.1	2001:7f8:a0::8b73:0:1		
Adamo Telecom Iberia S.A	35699	100G	Open
185.1.192.218	2001:7f8:a0::8b73:0:2		

METASCRITIC COLLECTS DATA AVAILABLE ABOUT THE ASES

- **Step 3: Data Collection**

Gather **AS features** (e.g., peering policy, traffic profile, customer cone size) and **connectivity from existing measurements** (e.g., RIPE Atlas, Arkipelago).

METASCRITIC BULDS FROM INTRINSIC PROPERTIES ON TOP OF EXTERNAL BEHAVIOR.

- **Step 4: Data Representation**

Existing Behavior

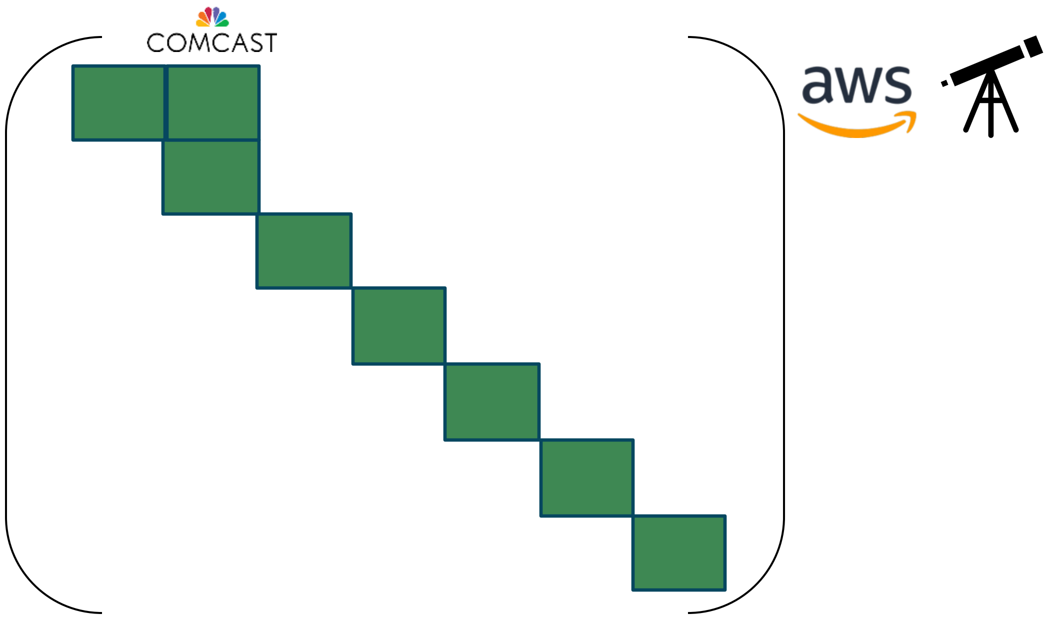

Intrinsic Properties

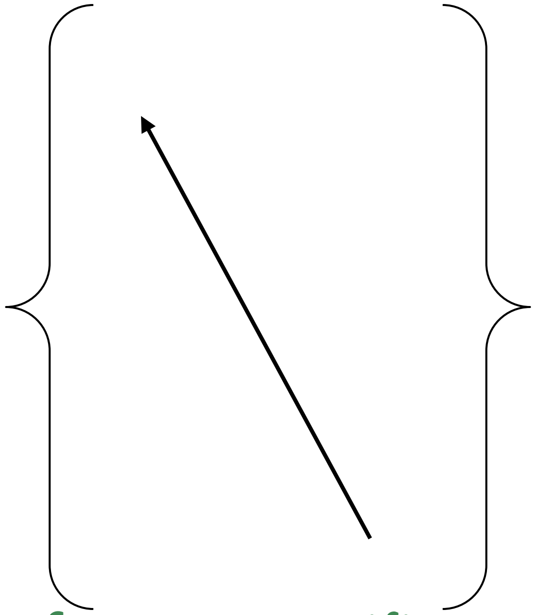

value of a **feature** specific to an AS, (e.g., the number of prefixes)

METASCRITIC BULDS FROM INTRINSIC PROPERTIES ON TOP OF EXTERNAL BEHAVIOR.

- **Step 4: Data Representation**

Existing Behavior **Intrinsic Properties**

value of a **feature** specific to an AS, (e.g., the number of prefixes)

METASCRITIC BULDS FROM INTRINSIC PROPERTIES ON TOP OF EXTERNAL BEHAVIOR.

- **Step 4: Data Representation**

Existing Behavior

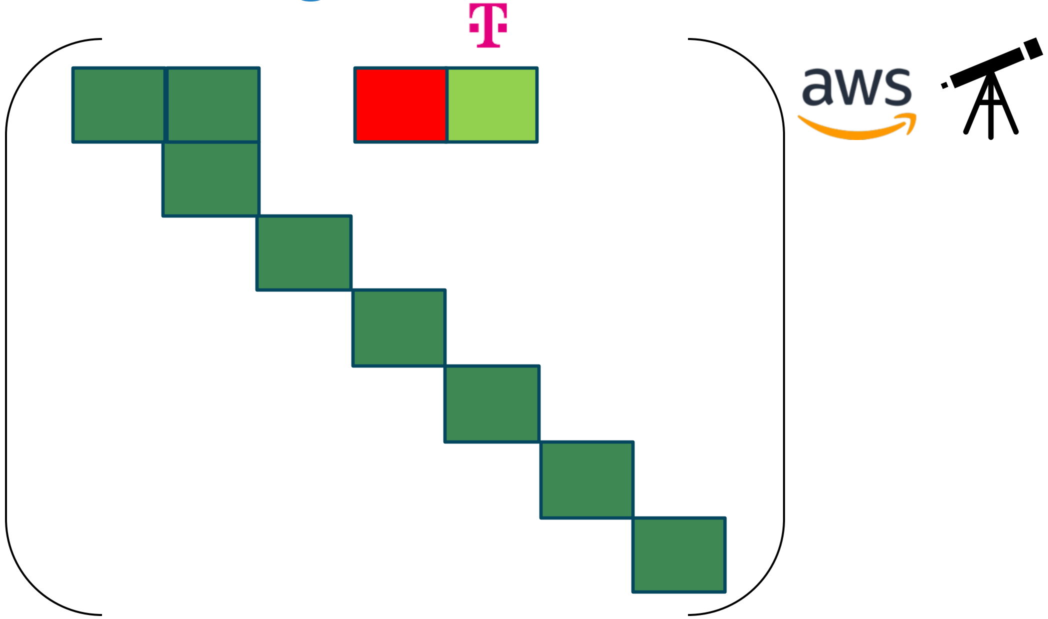

Intrinsic Properties

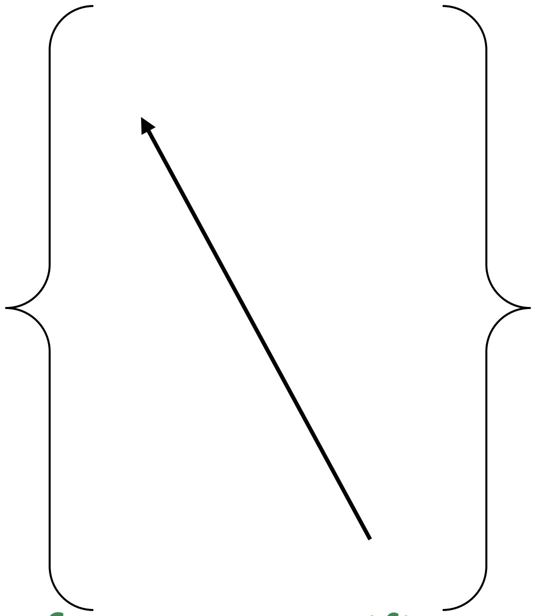

value of a **feature** specific to an AS,
(e.g., the number of prefixes)

METASCRITIC BULDS FROM INTRINSIC PROPERTIES ON TOP OF EXTERNAL BEHAVIOR.

- **Step 4: Data Representation**

Existing Behavior

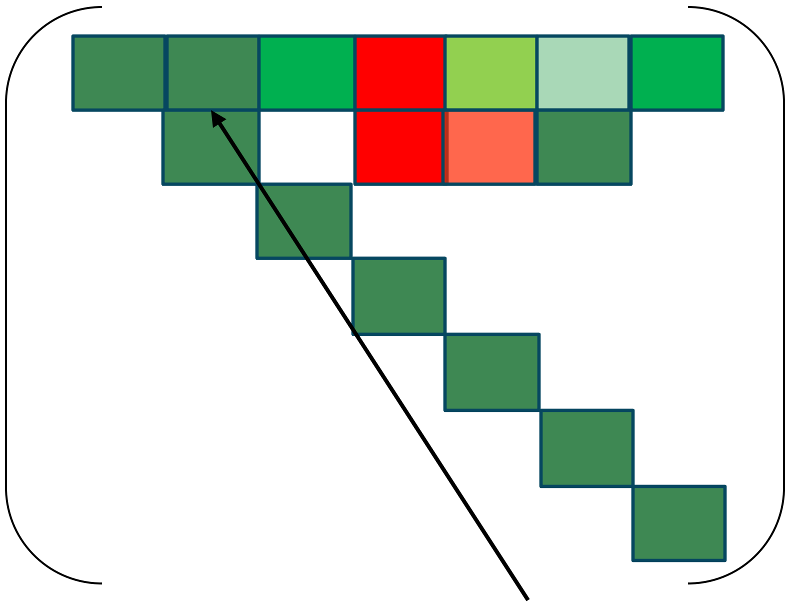

presence (>0), **non-presence** (<0) of a peering link between AS k and AS l. **Absolute value (color intensity)** reflects the confidence that the link exists.

Intrinsic Properties

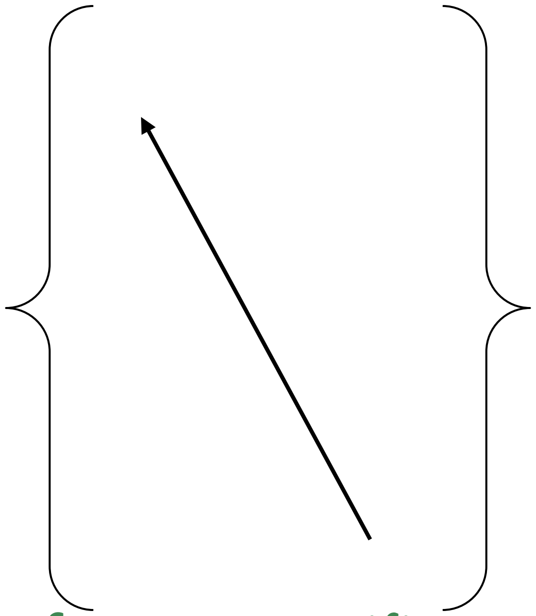

value of a **feature** specific to an AS, (e.g., the number of prefixes)

(NAIVELY) COMPLETING THE MATRIX.

We can complete the missing entries of the existing connectivity matrix.

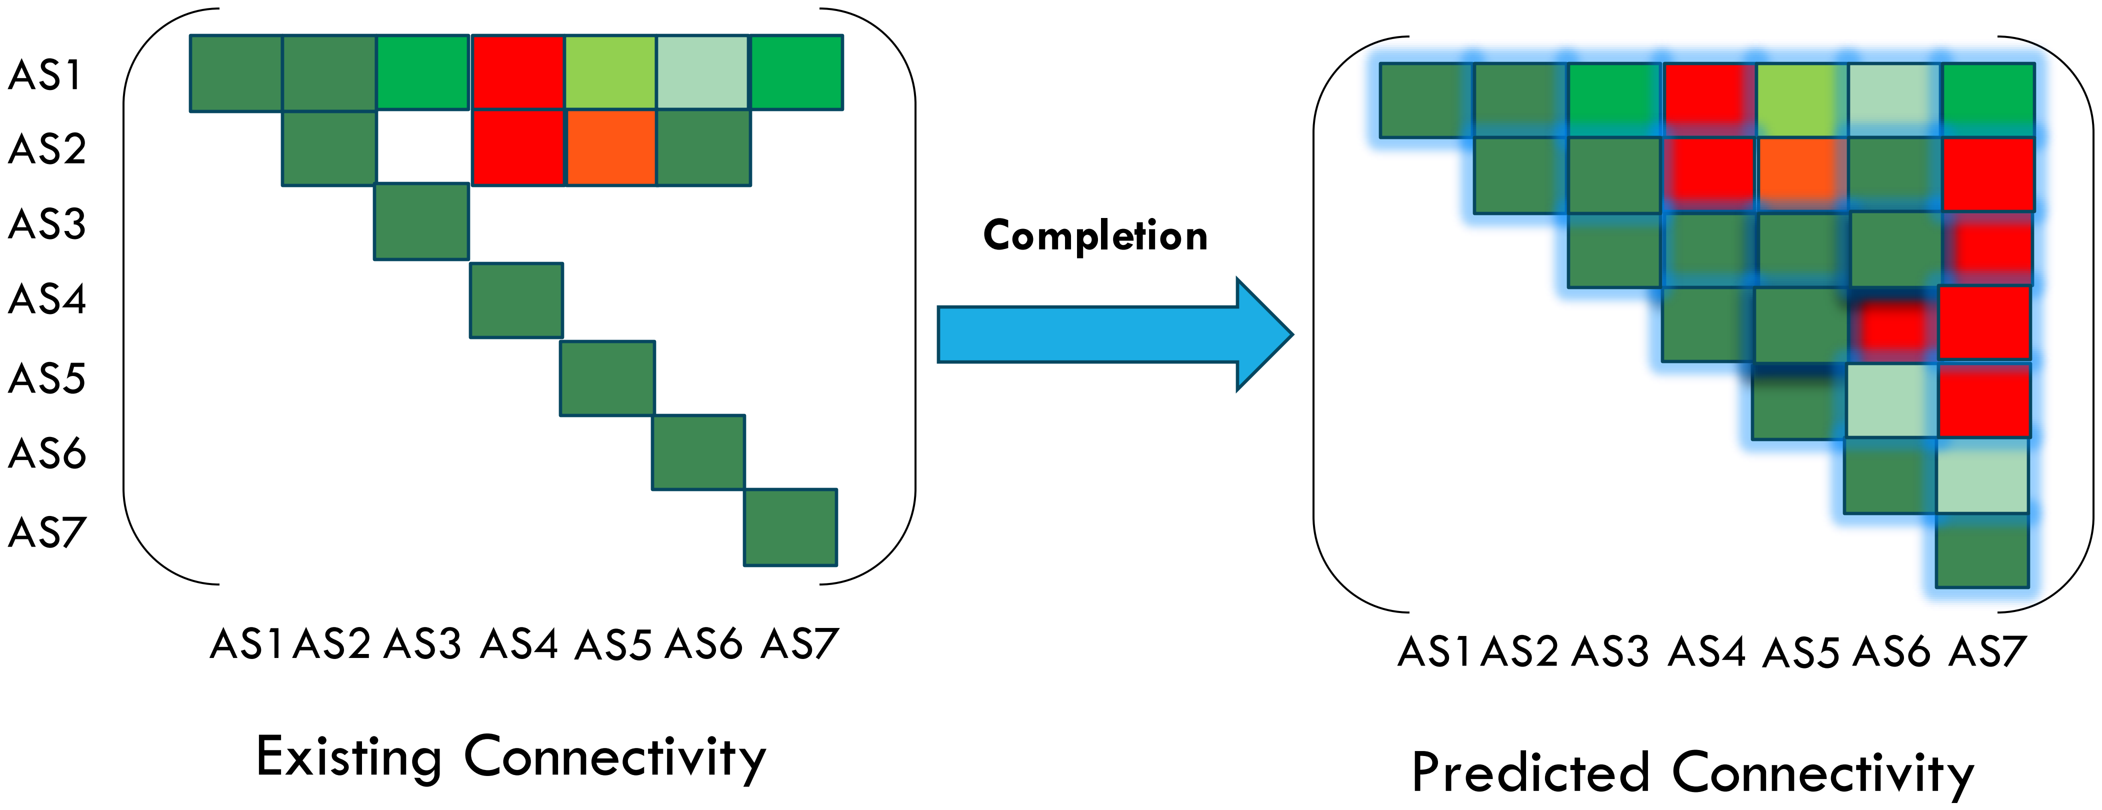

(NAIVELY) COMPLETING THE MATRIX.

We can complete the missing entries of the existing connectivity matrix.

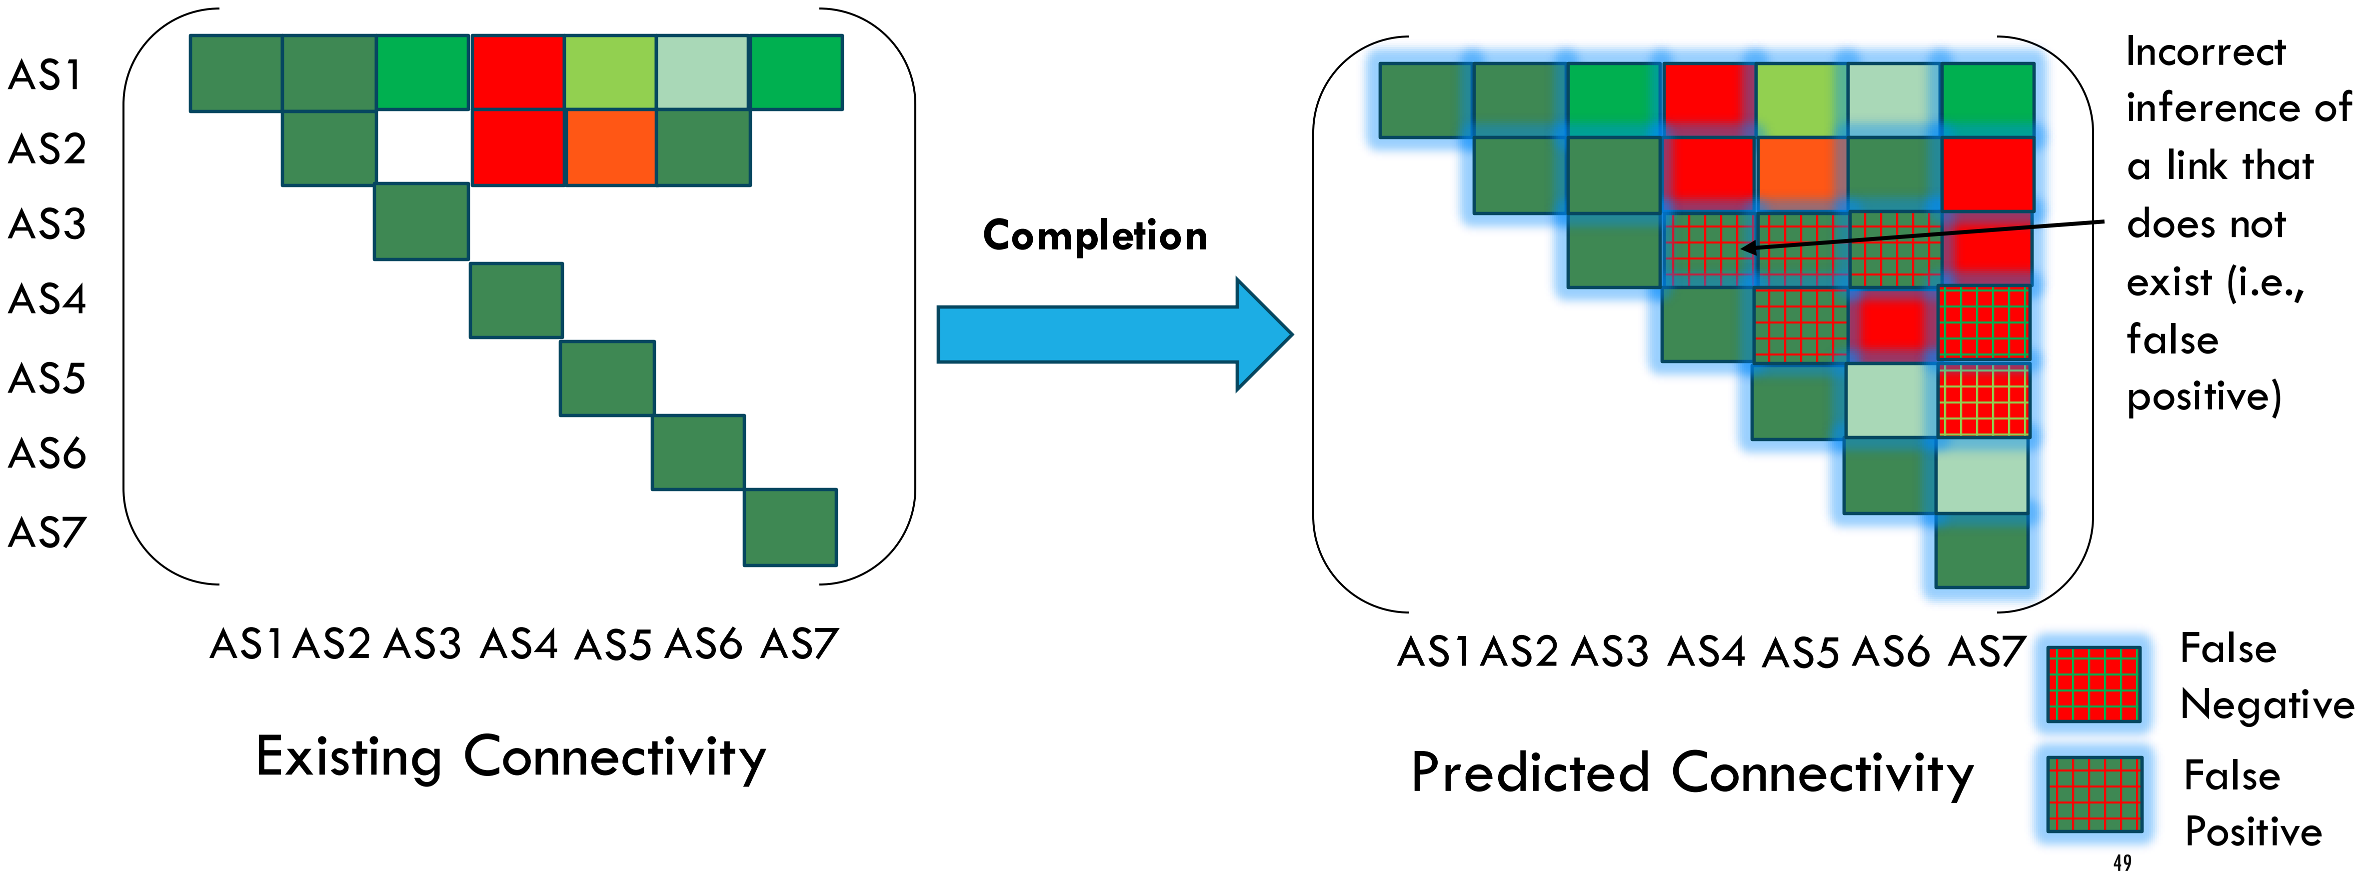

NAVIGATING THE SPACE OF AVAILABLE TOPOLOGY BY TRADING-OFF FALSE POSITIVE AND NEGATIVE.

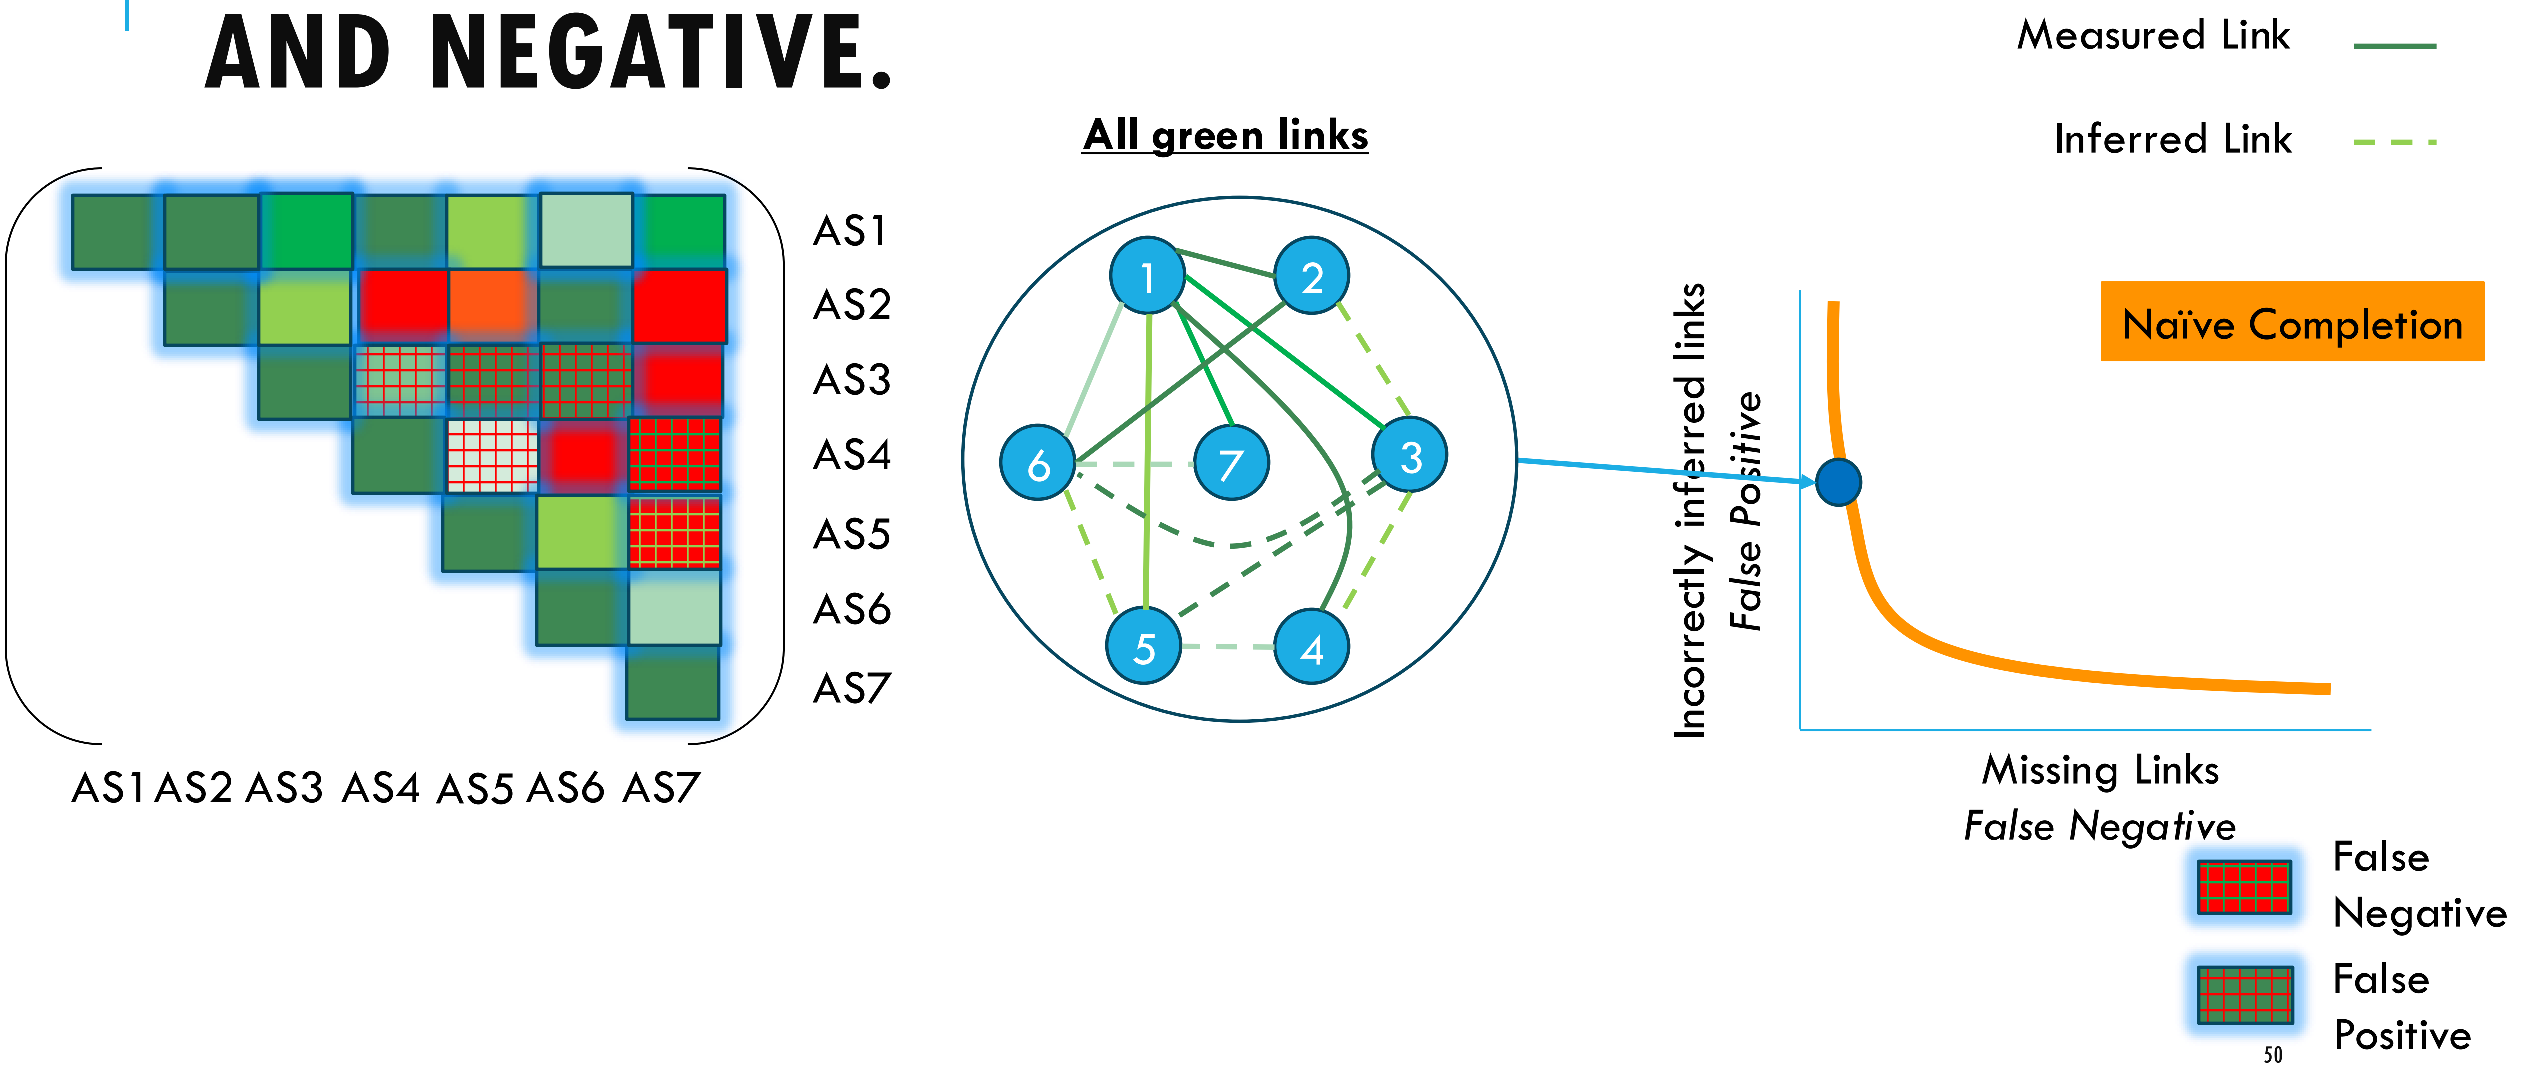

NAVIGATING THE SPACE OF AVAILABLE TOPOLOGY BY TRADING-OFF FALSE POSITIVE AND NEGATIVE.

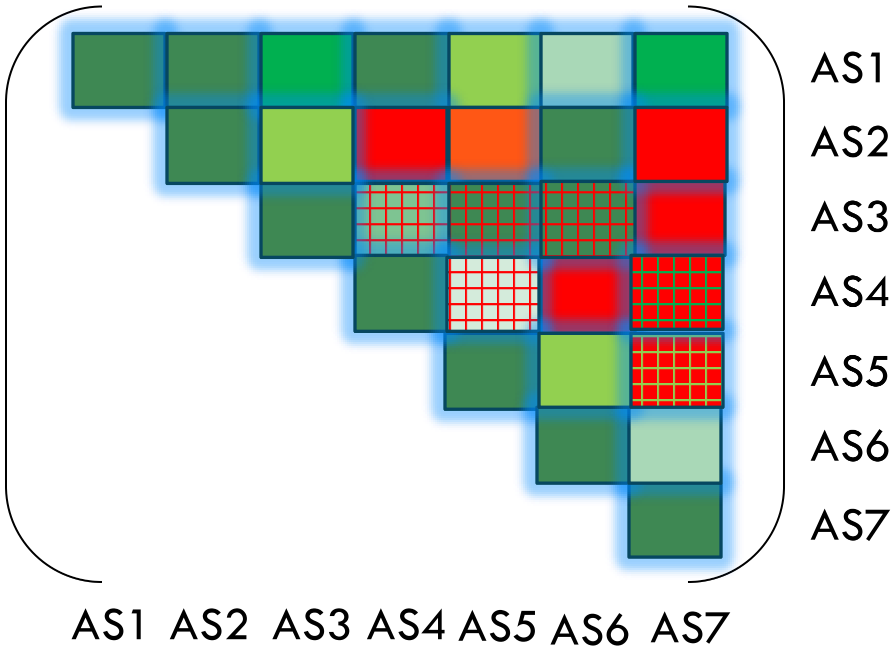

Only high-confidence links

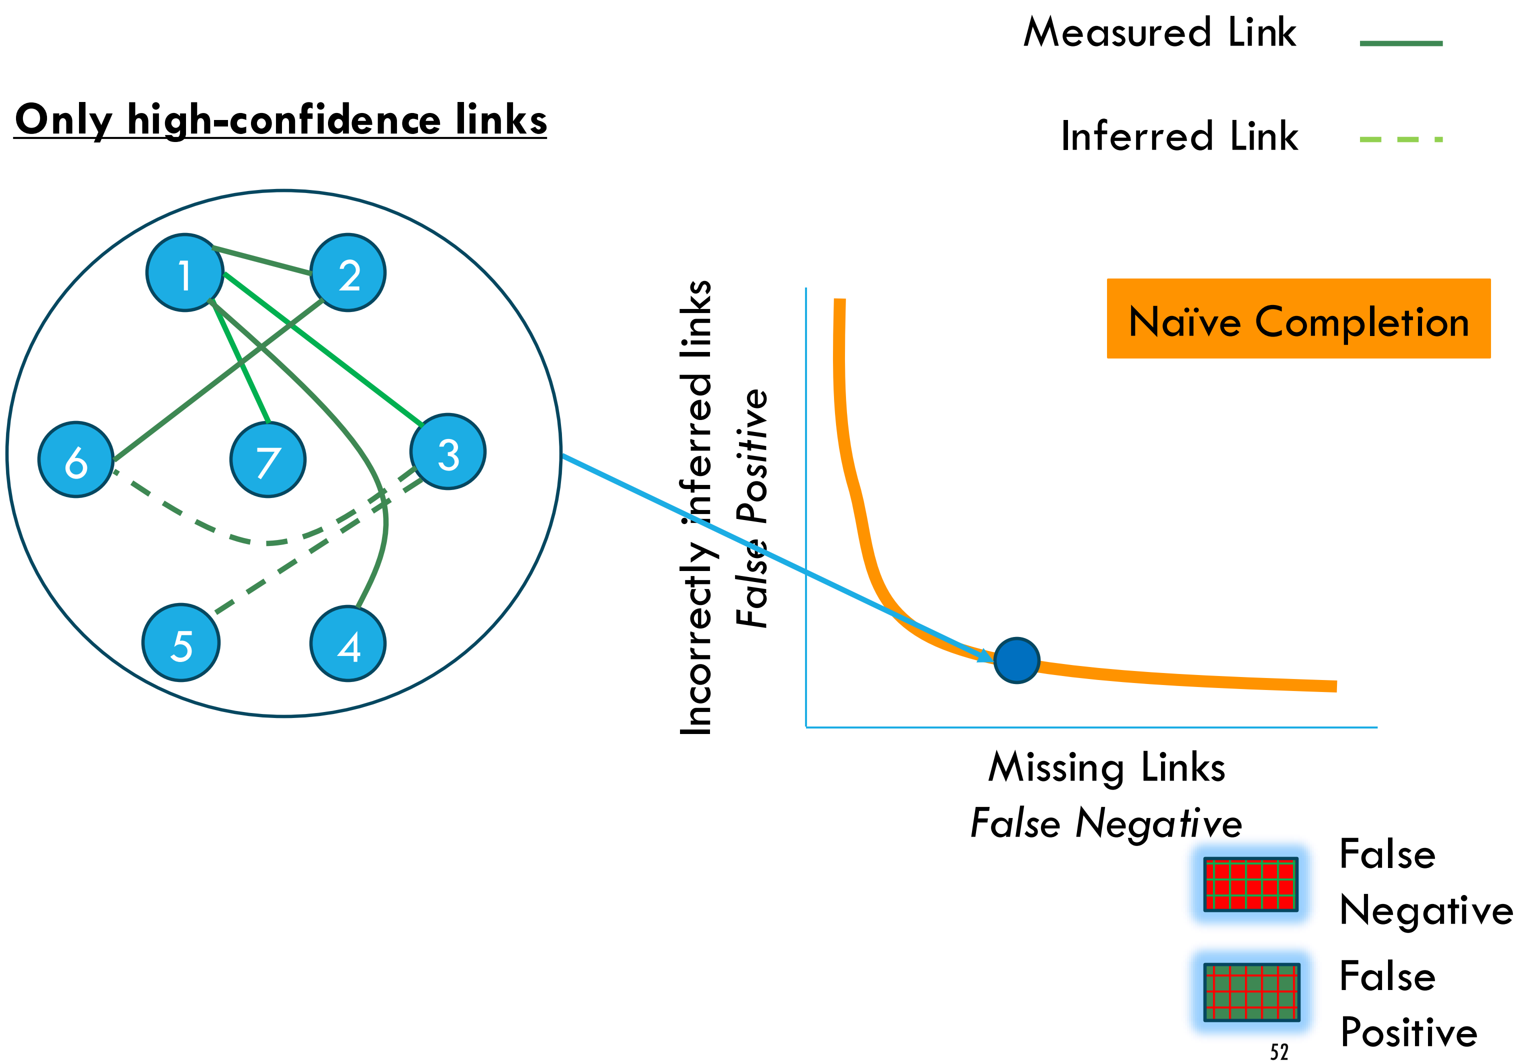

IMBALANCED DISTRIBUTION OF ENTRIES RESULTS IN IMBALANCED PERFORMANCE.

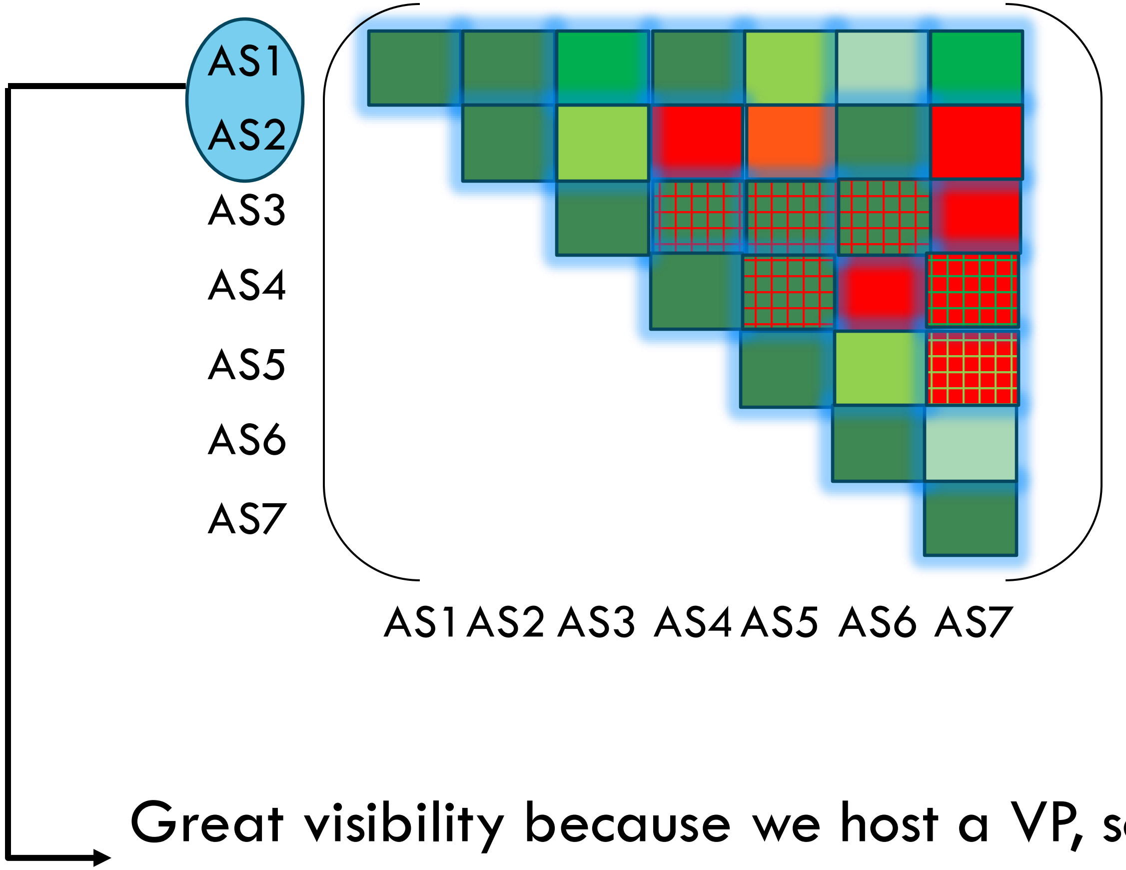

0 False Positive
0 False Negative

Great visibility because we host a VP, so great predictive power

False Negative

False Positive

IMBALANCED DISTRIBUTION OF ENTRIES RESULTS IN IMBALANCED PERFORMANCE.

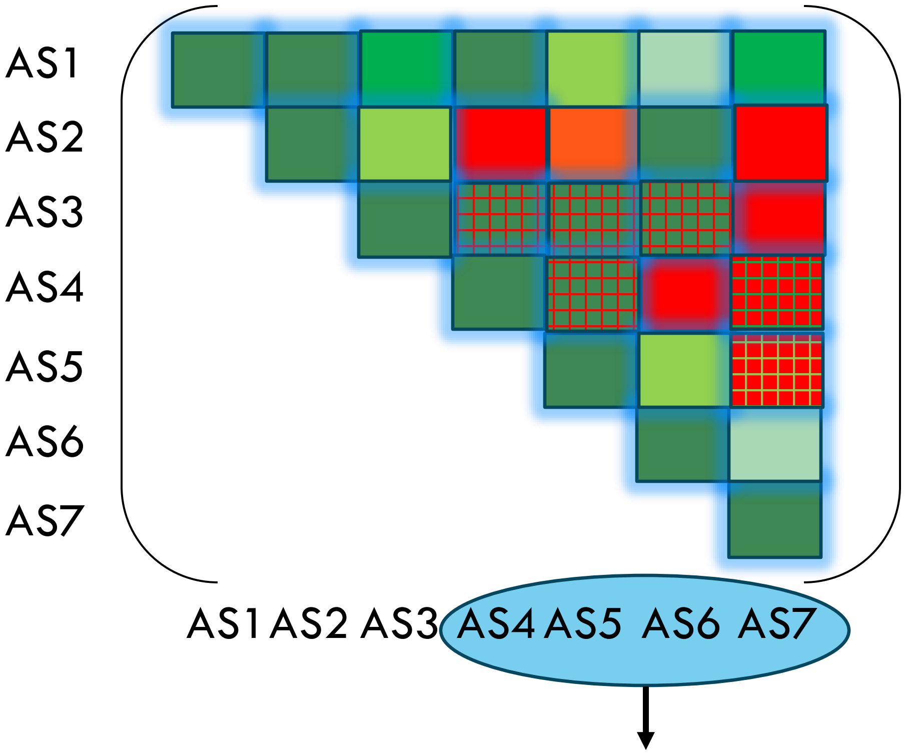

4 False Positives
2 False Negatives

Low visibility, no VP or poorly positioned VP so little predictive power

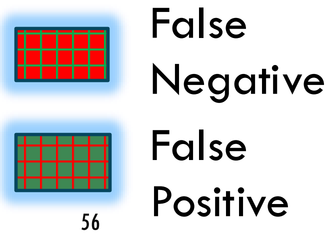

THE IMPORTANCE OF DEBIASING THE DATASET.

From our collections of inferred and geolocated links:

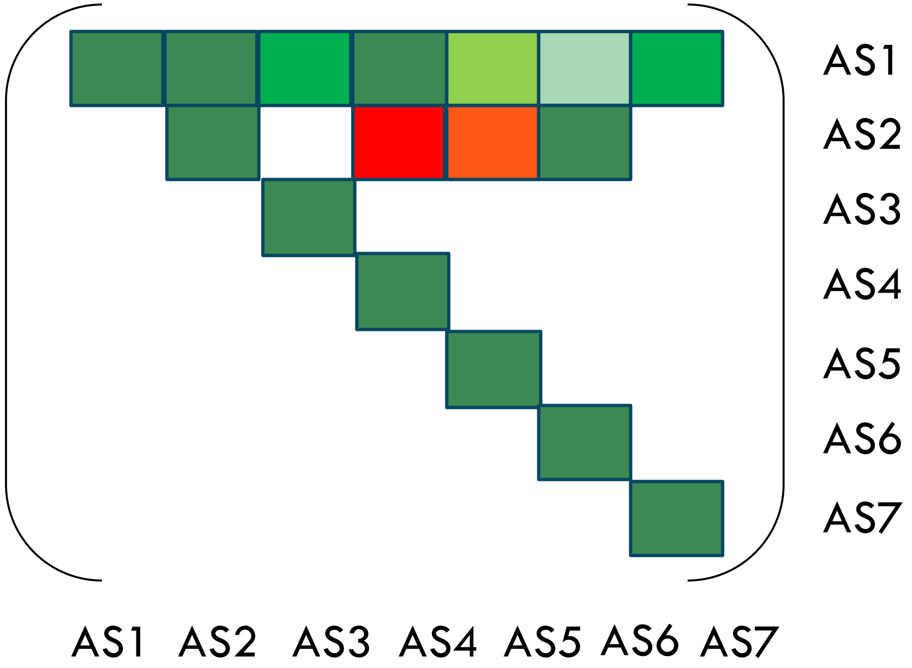

Problem: The public datasets are heavily skewed toward ASes that host vantage points.

MEASURING TO IMPROVE THE LEARNING PROCESS

Idea # 1: Identify the measurements to issue that are likely to be the most informative.

Use targeted traceroutes to efficiently gather information about the existence or non-existence of peering links.

Focus the traceroutes toward links of less observed ASes.

MEASURING WITH THE INTENT OF IMPROVING THE LEARNING PROCESS

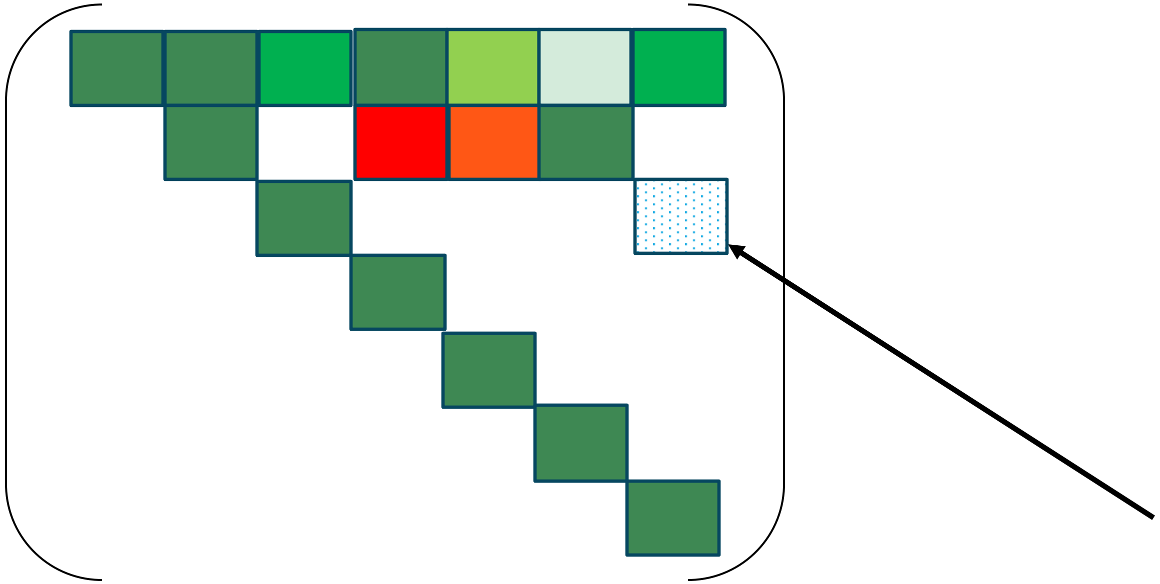

Measuring this link would help most with visibility, as both AS3 and AS7 have very few measurements.

MEASURING WITH THE INTENT OF IMPROVING THE LEARNING PROCESS

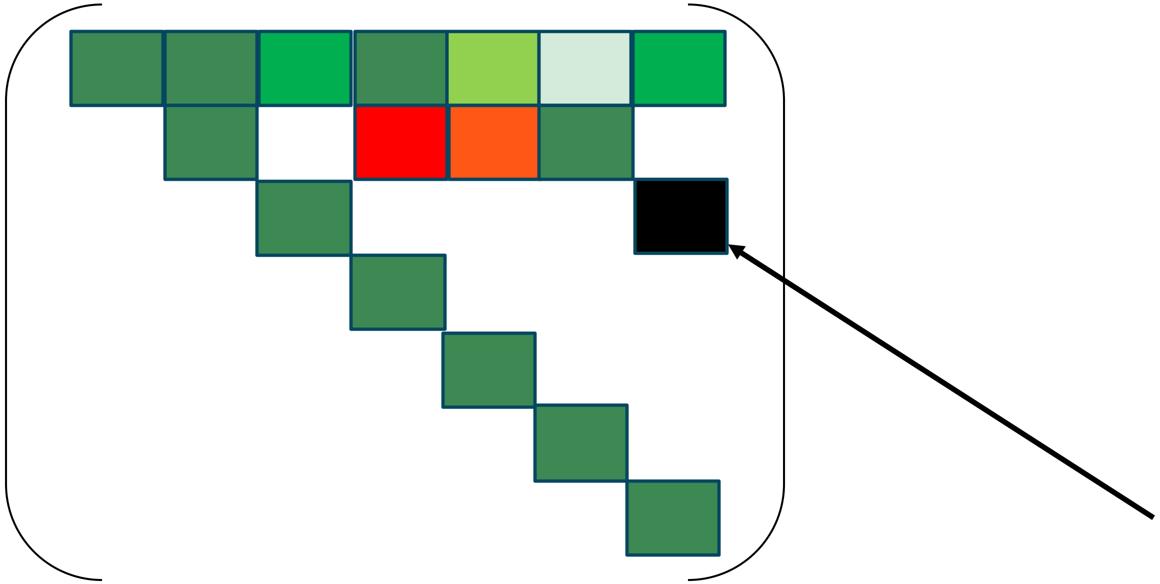

Problem: There might be no measurement that can discover the connectivity between AS3 and AS7 (because of valley-free routing).

MEASURING WITH THE INTENT OF IMPROVING THE LEARNING PROCESS

Idea # 2: Identify how likely it is to find a measurement that can infer the (non-)presence of a link.

WE COMBINE THESE TWO INSIGHTS INTO AN ALGORITHM THAT RUNS MEASUREMENT WITH THE INTENT OF IMPROVING THE FINAL COMPLETION*.

*Details and theoretical underpinning can be found in the paper.

OPTIMAL SET OF PEOPLE TO RECOMMEND TO UNDERSTAND ONE'S PREFERENCE

Tinder analogy:

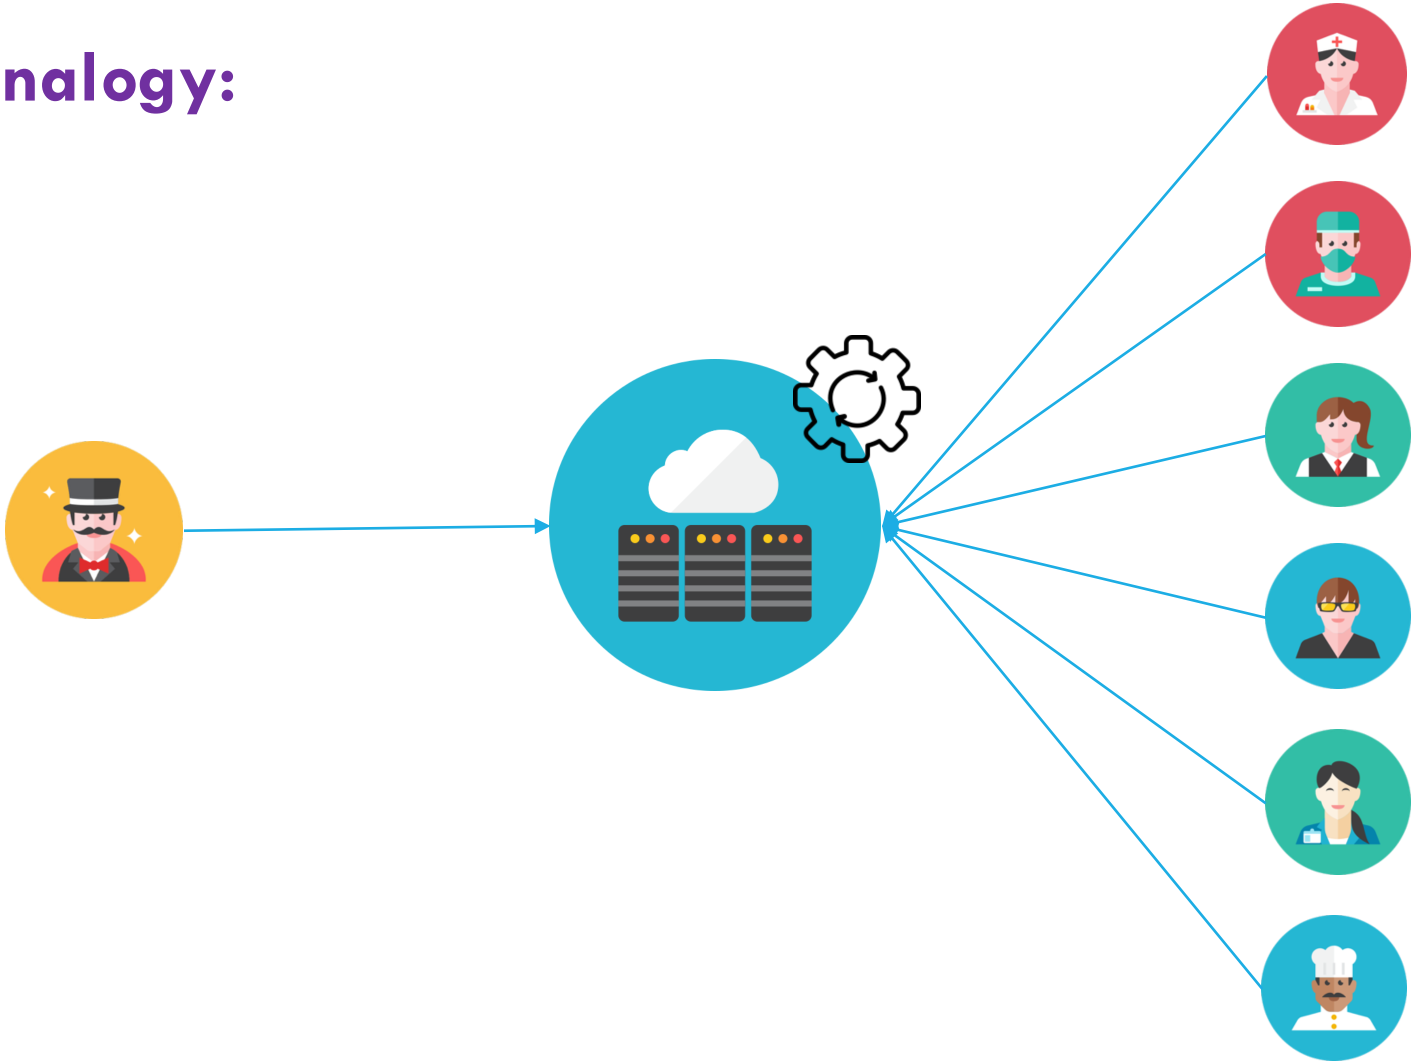

Intuition: Gathering ratings from diverse individuals with unique profiles is essential to well understand a magician's preferences.

ADDING SMART MEASUREMENTS RESULTS IN STRETCHING OUR SPACE OF POSSIBLE INFERENCES.

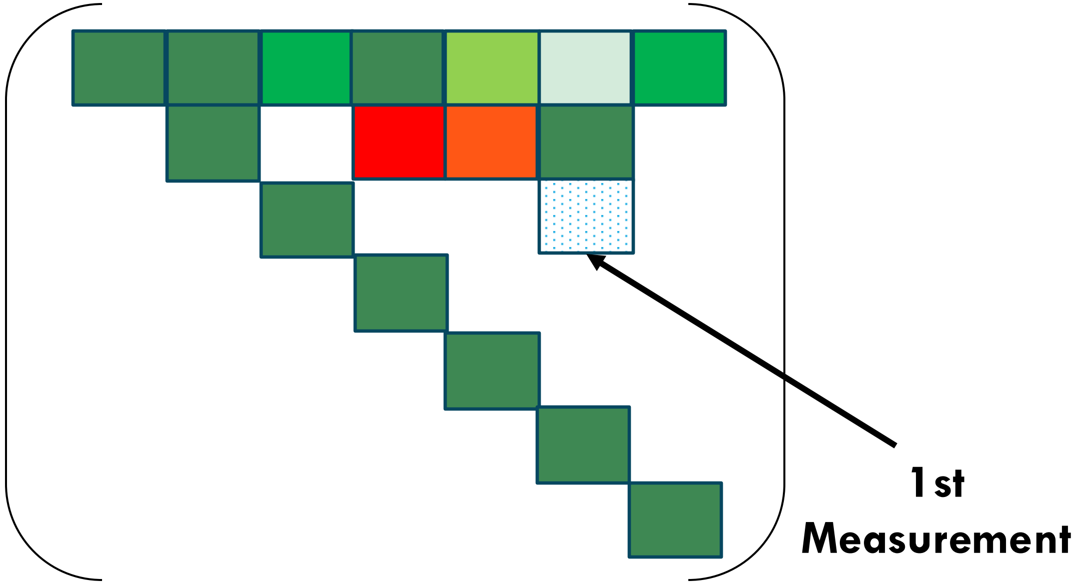

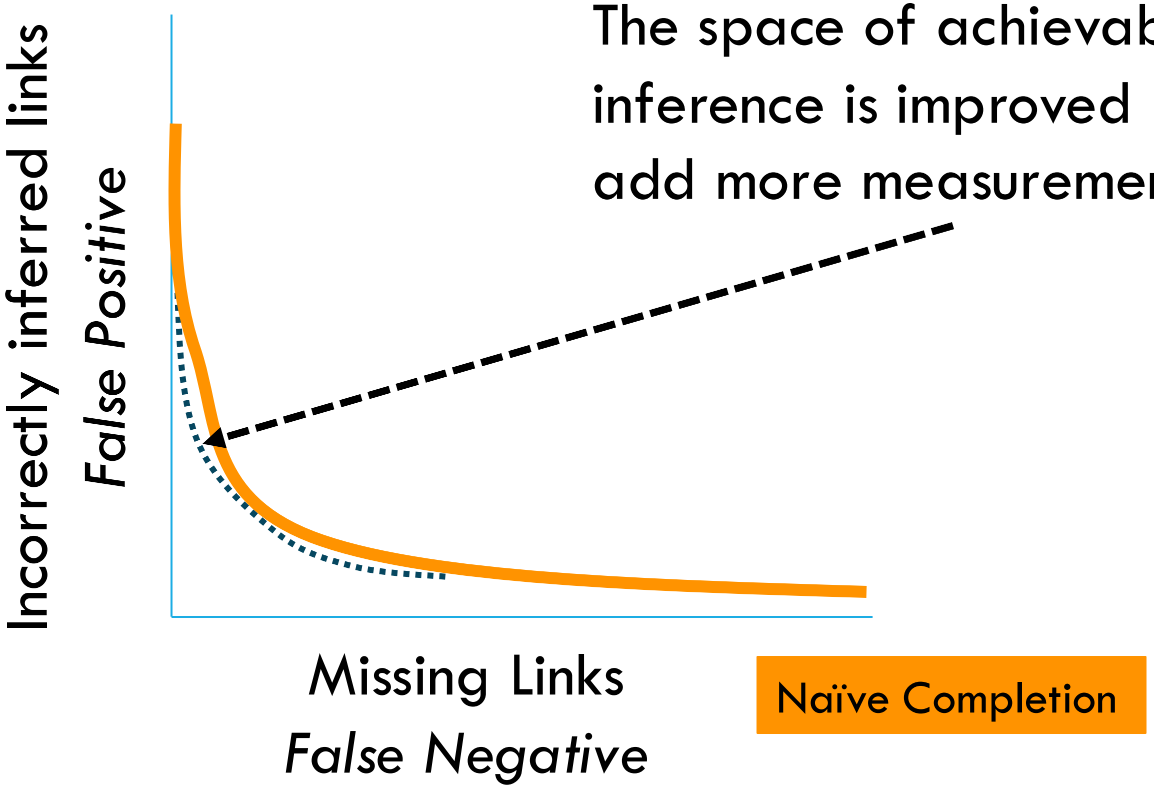

The space of achievable inference is improved as we add more measurements.

ADDING SMART MEASUREMENTS RESULTS IN STRETCHING OUR SPACE OF POSSIBLE INFERENCES.

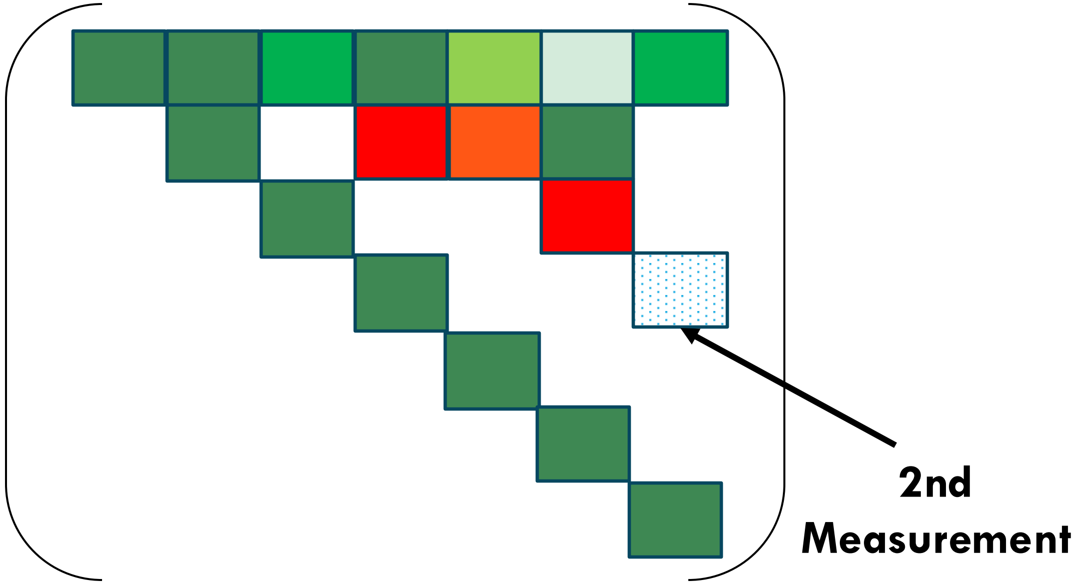

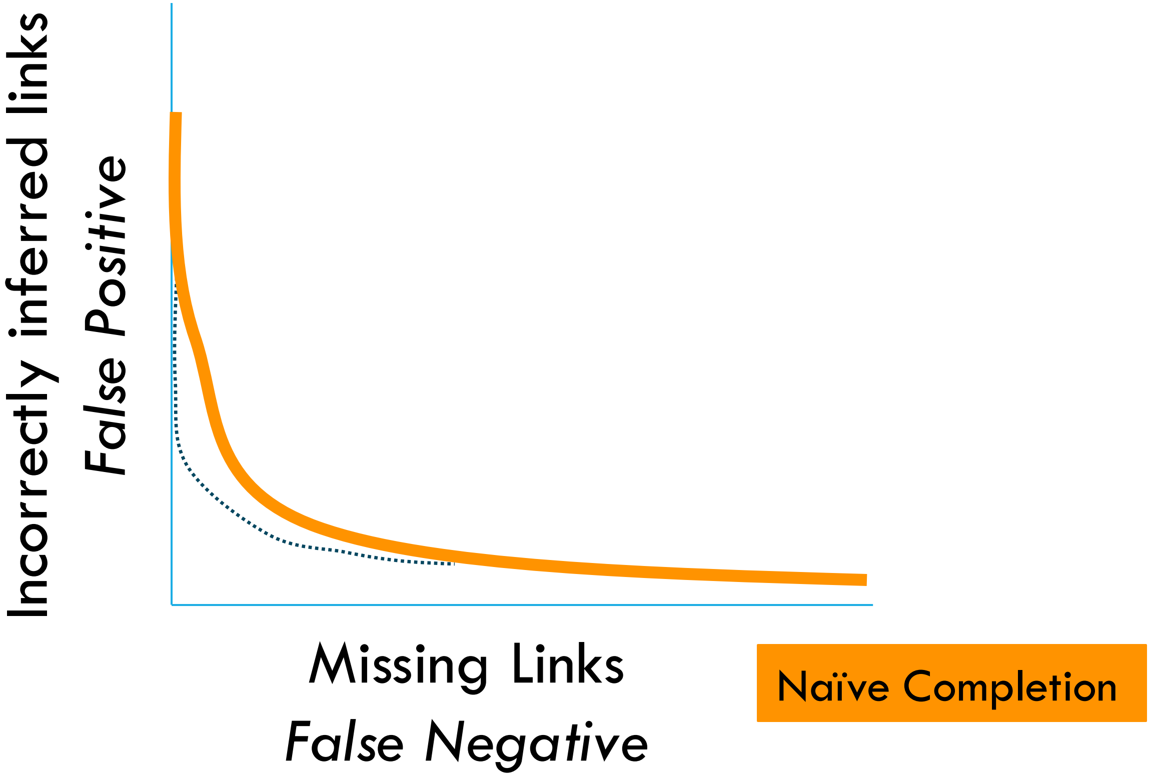

ADDING SMART MEASUREMENTS RESULTS IN STRETCHING OUR SPACE OF POSSIBLE INFERENCES.

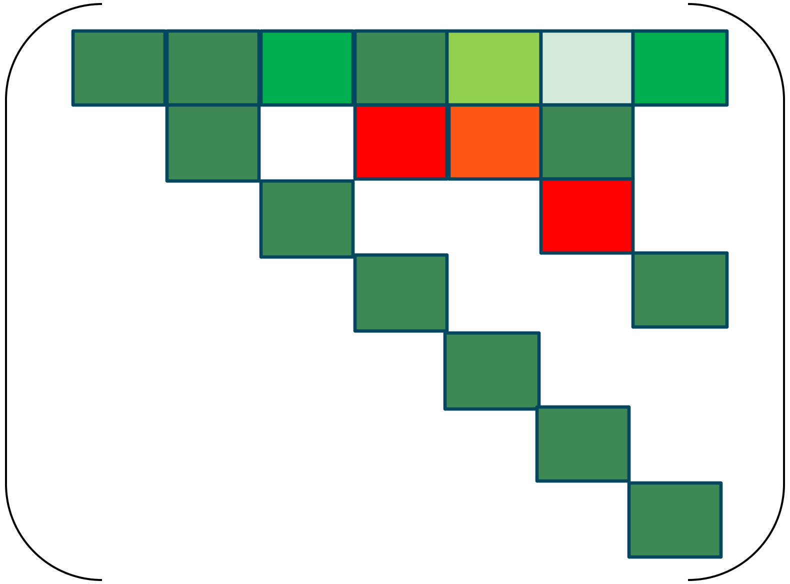

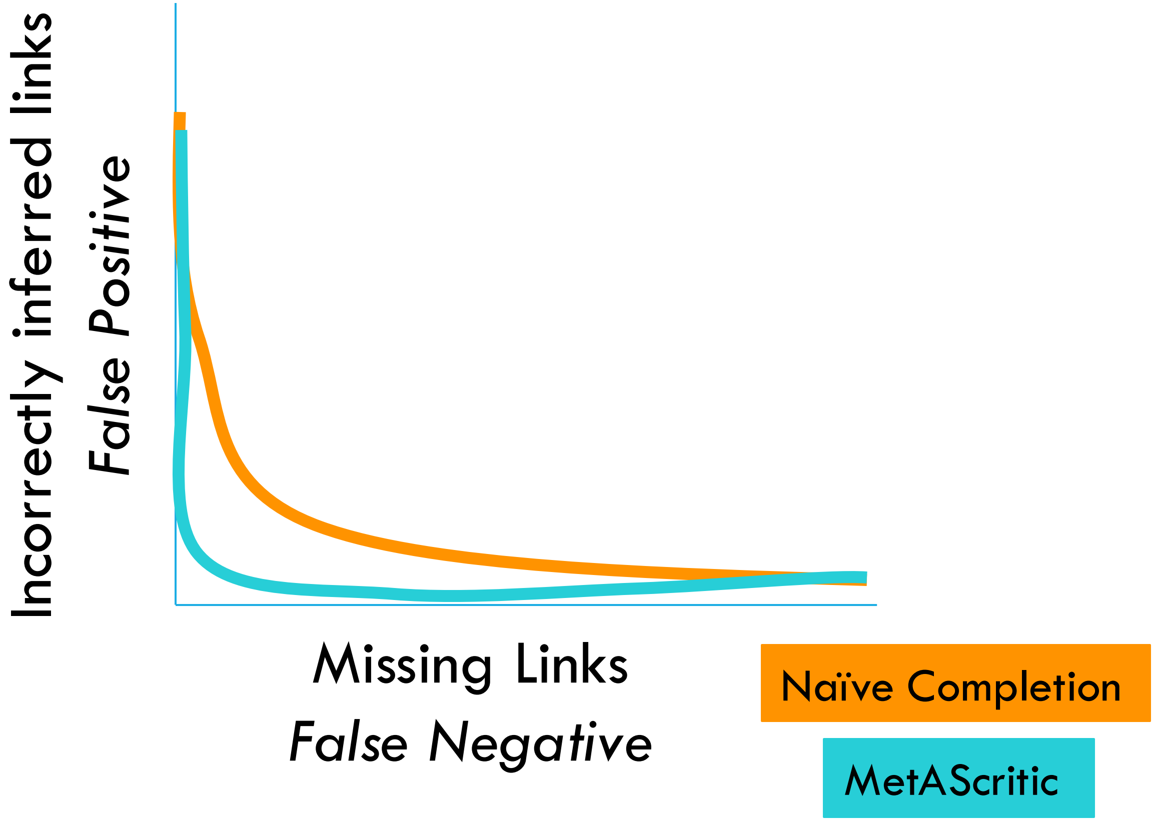

BY INCORPORATING A FEW MEASUREMENTS, METASCRITIC OUTPERFORMED A NAIVE COMPLETION.

The final completion is much better now!

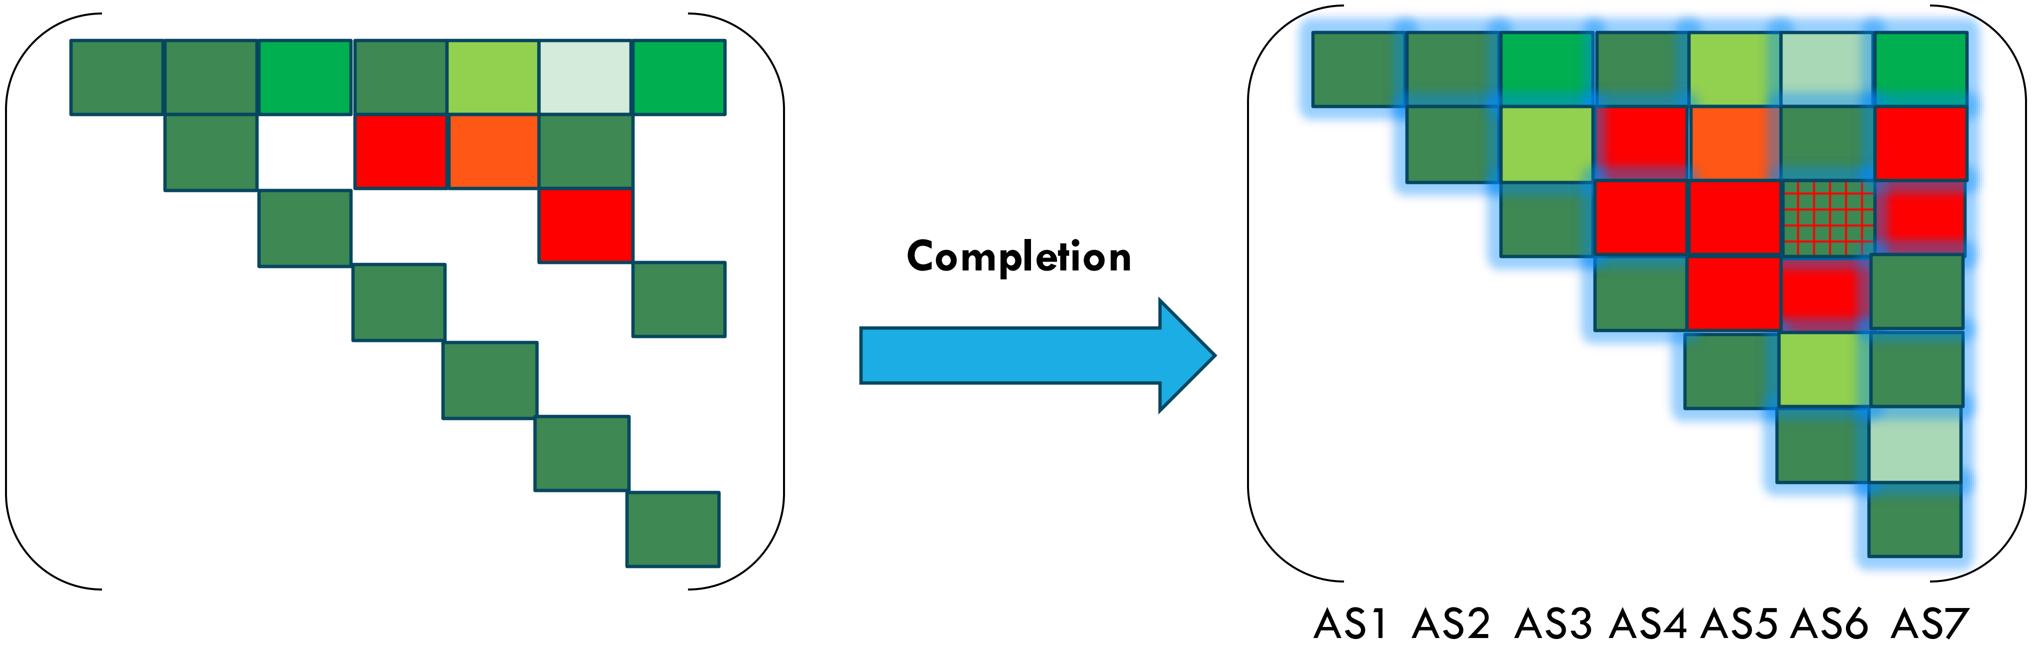

RESULTS

MetAScritic: 86K edges measured + 368K edges inferred with high confidence = **454K** edges across 6 evaluated metros

Public BGP data: 13K edges observed

34× increase compared to current visibility!

Dataset	Precision	Recall
Stratified Split	0.84 – 0.96	0.82 – 0.94
Ground-Truth (Vultr, Google, Looking Glasses)	0.78 – 0.95	0.84 – 0.97
BGP Community	N.A.	0.9 – 1

Many more in the paper!

A FINAL THOUGHT:

Applying machine learning for topology discovery is feasible and can help us improve many of our applications.

Our solution: MetASCritic, a recommender system for AS topology discovery.

Results: More than 34x increase in links compared to current visibility with an average 0.87 F1-score!

CONCLUSION AND LESSON LEARNED

1. Applying machine learning for topology discovery is feasible.

2. For machine learning to be more widely used, we need explainability on our inferences.

BACK-UP SLIDES

AN EXTENSIVE GROUND-TRUTH COLLECTED.

We study the effect of different splits on the accuracy:

Random	Stratified	Completely Left-Out
Randomly remove entries of the matrix	Remove the same fraction of entries from each row	Remove all the entries of a given row

We collect data from several sources:

Ground Truth	IXP Connectivity Matrix	BGP Community	Extensive Measurements	IP Aliases	iGDB
Cloud provider and Looking Glasses AS interconnections observed in the metros	IXP connectivity (both bilateral and through route-servers) matrices as ground truth to validate peering inferences.	BGP community geographic tags to infer AS interconnections at specific metros	AS links were observed from extensive measurement campaigns in a few metros.	IP alias to identify multiple IP in different networks belonging to the same router.	AS links from BGP that can be pinpointed to the specific locations.

IDENTIFYING NON-EXISTENCE USING VALLEY-FREE ROUTING

Valley-free routing:

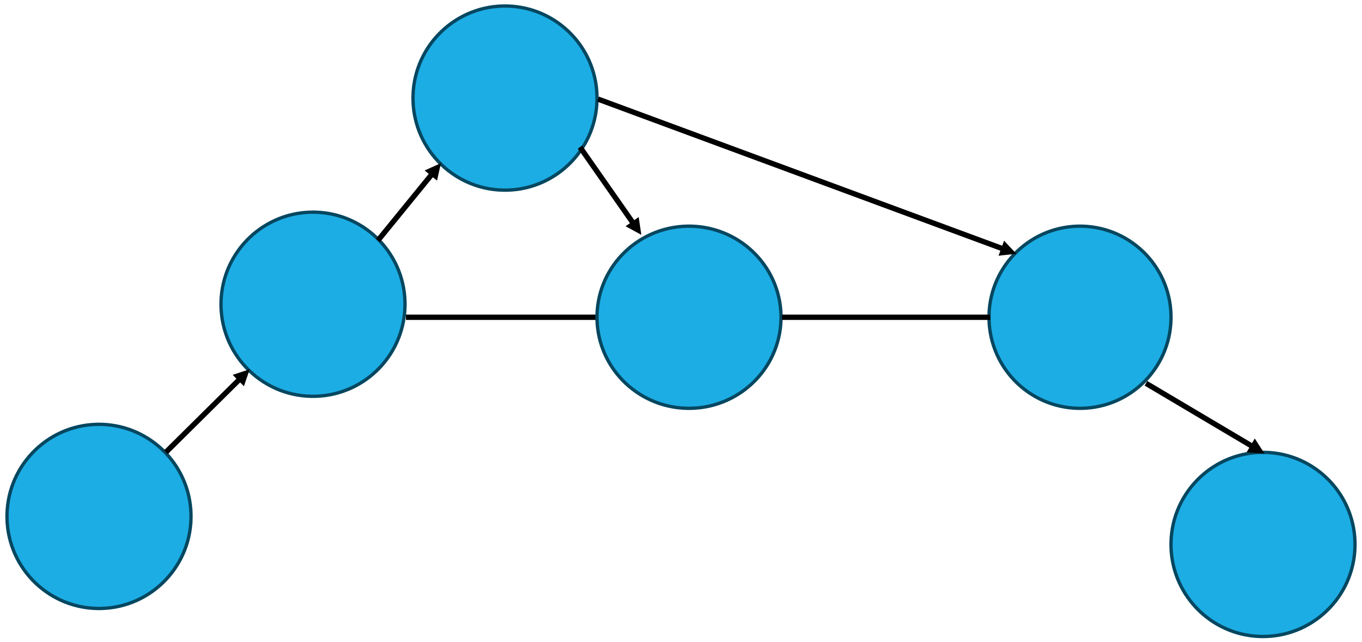

DIFFERENT SPLITS TO VERIFY FOR DIFFERENT PROPERTIES

We study the effect of different splits on the accuracy:

Random	Stratified (classical scenario)	Completely Left-Out (no VPs)
Randomly remove entries of the matrix	Remove the same fraction of entries from each row	Remove all the entries of a given row

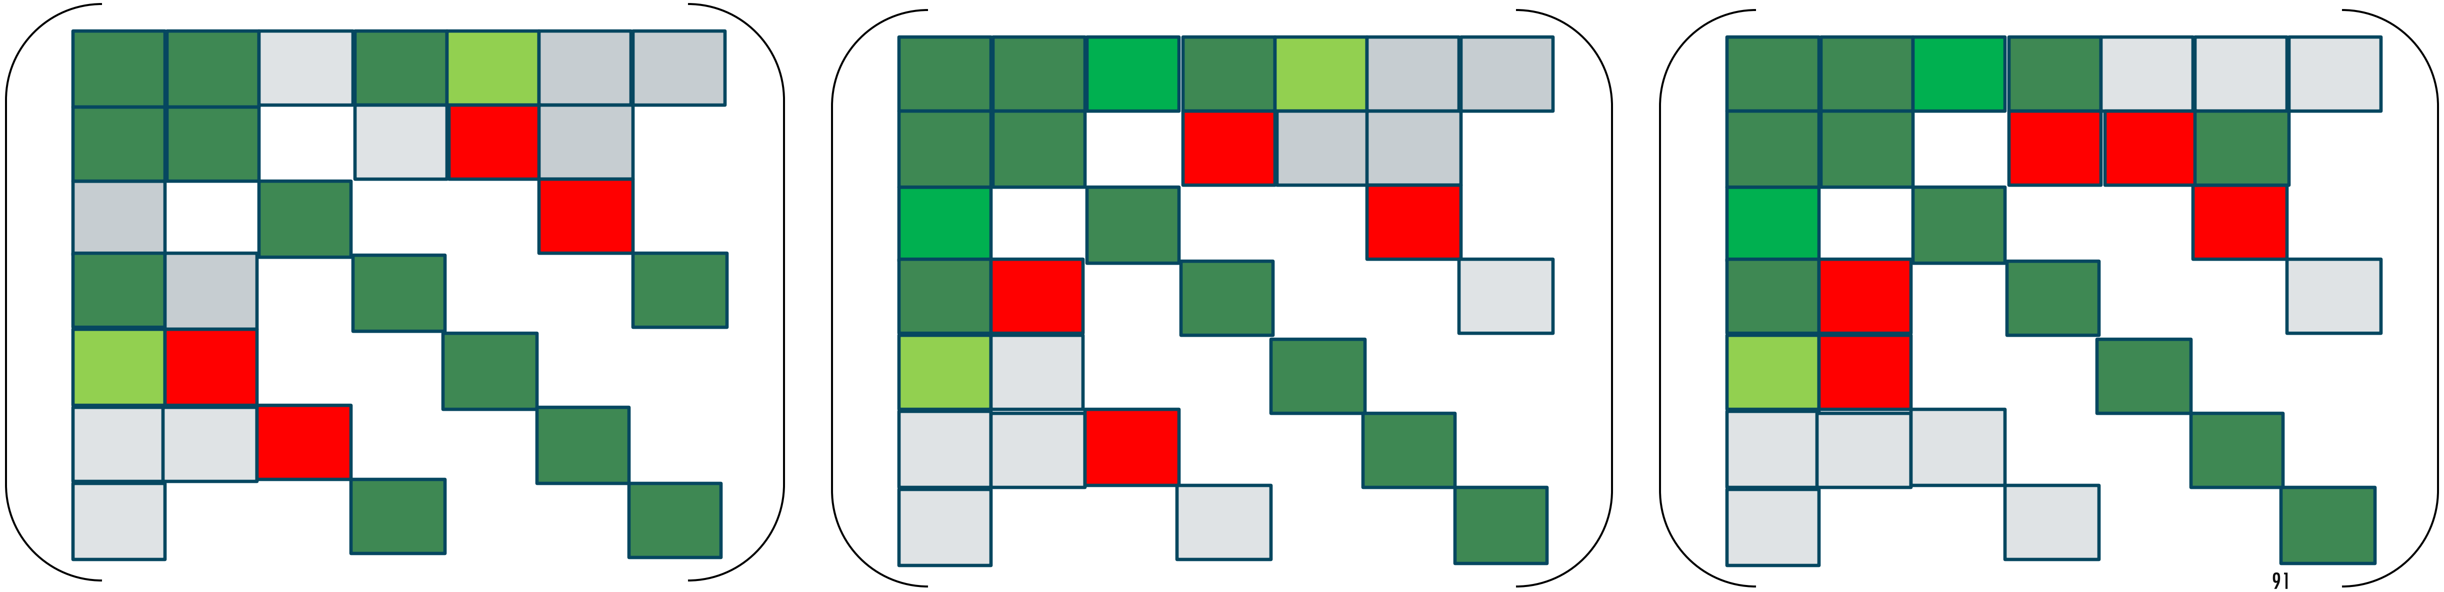

HIGH PRECISION WITH HIGH RECALL

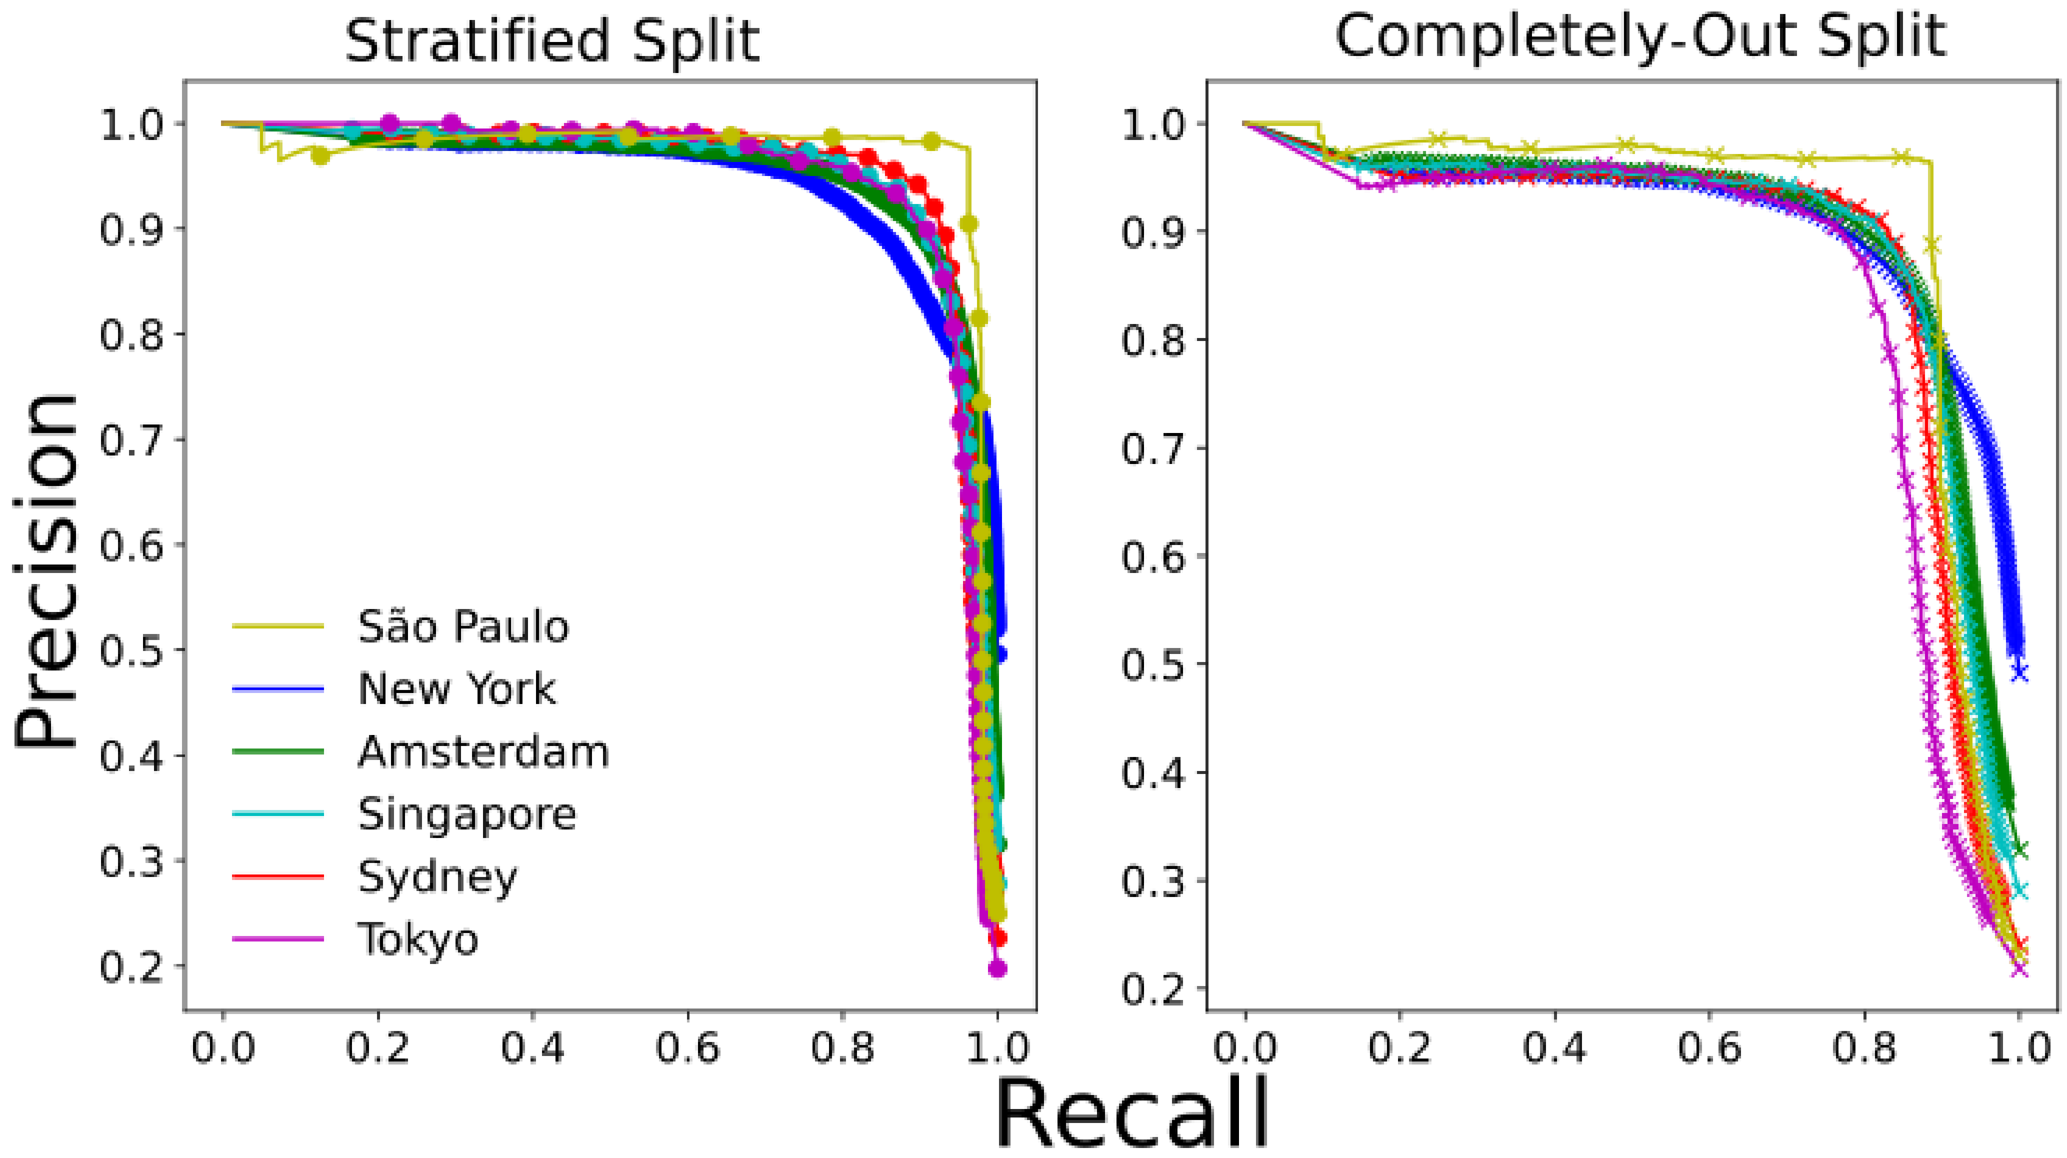

IDENTIFYING NON-EXISTENCE USING VALLEY-FREE ROUTING

Valley-free routing:

It's a valley-free compliant path because it only crosses at most 1 peering link and goes up and down.

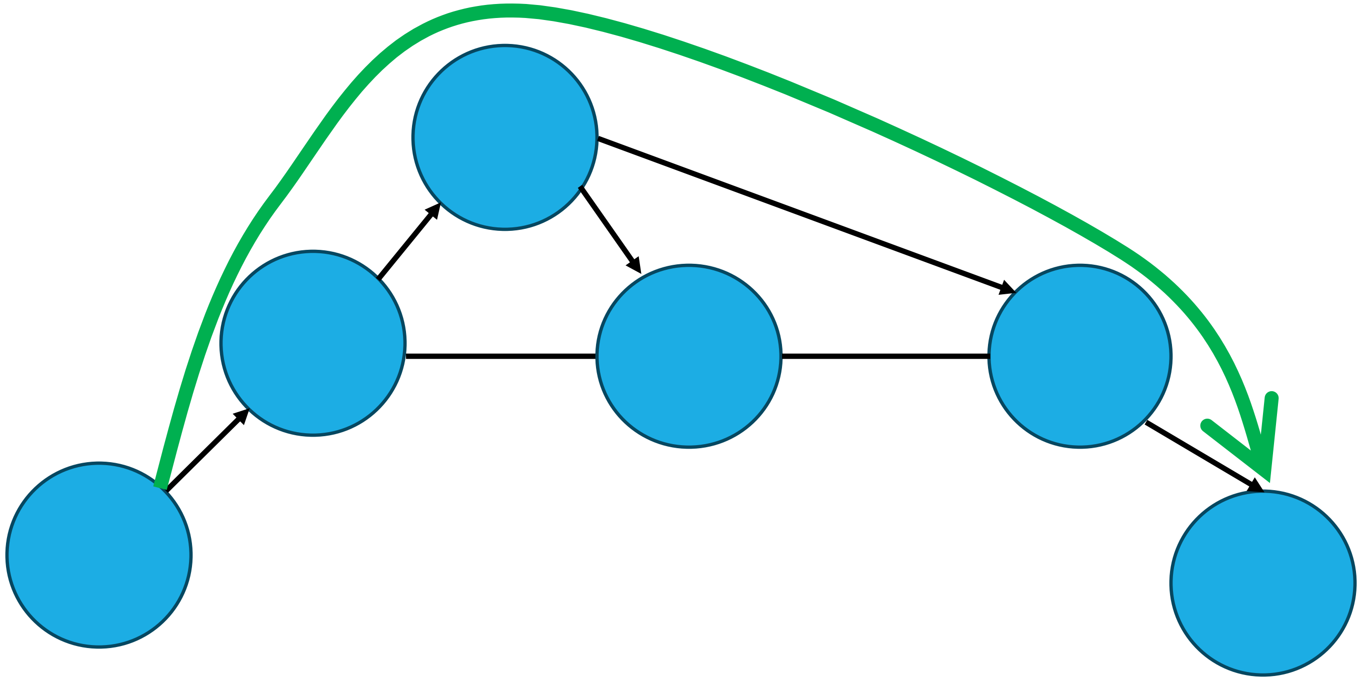

IDENTIFYING NON-EXISTENCE USING VALLEY-FREE ROUTING

Valley-free routing:

This is NOT a valley-free compliant path because it crosses 2 peering links in a row!

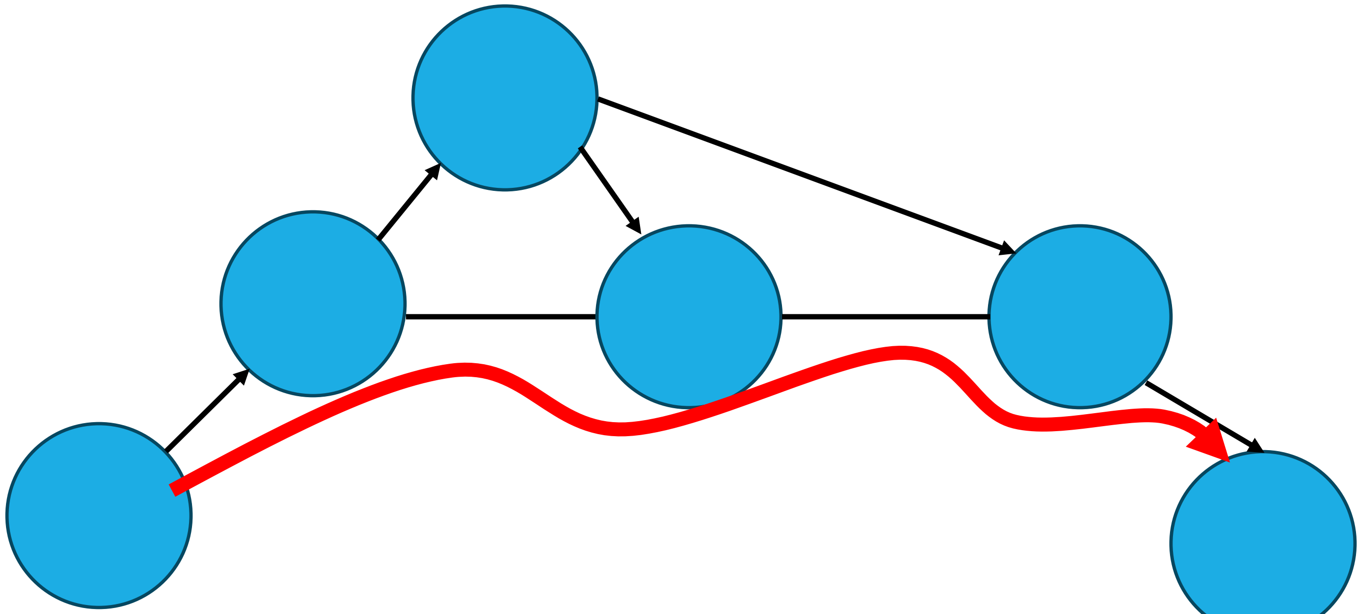

IDENTIFYING NON-EXISTENCE USING VALLEY-FREE ROUTING

Key observation: Valley-free routing often holds at a given geographic granularity.

An AS is said to have **consistent routing** at a given granularity if it consistently prefers peering links over transit to reach the rest of the Internet.

Example: Columbia University is consistent at all granularities because its AS will always prefer peering links over transit providers.

CONSISTENT ROUTING TO IDENTIFY WHERE WE CAN MAKE INFERENCE

Definition: For a consistent AS k, we can infer that a peering link with AS l does not exist (*non-existence*) if AS k routes traffic to a provider for packets destined to AS l.

This results in negative inferences (equivalent to a negative rating to a movie in Netflix or a left swipe on Tinder)

IMPROVING USABILITY WITH SHAPLEY VALUE

Challenge: Unlike errors in traceroute-inferred links, which have identifiable causes like unresponsive routers or IP address anomalies, inaccuracies in metAScritic's inferential approach are not as easily explained.

Solution: Shapley values quantify the contribution of each feature (e.g., geographic overlap, AS type, previous connections) to individual link predictions.

By calculating how the inclusion or exclusion of a feature affects the predicted rating, Shapley values reveal which factors are most influential in driving a particular inference.

METASCRITIC: HOW DOES IT WORK?

From our collections of inferred and geolocated links to ratings:

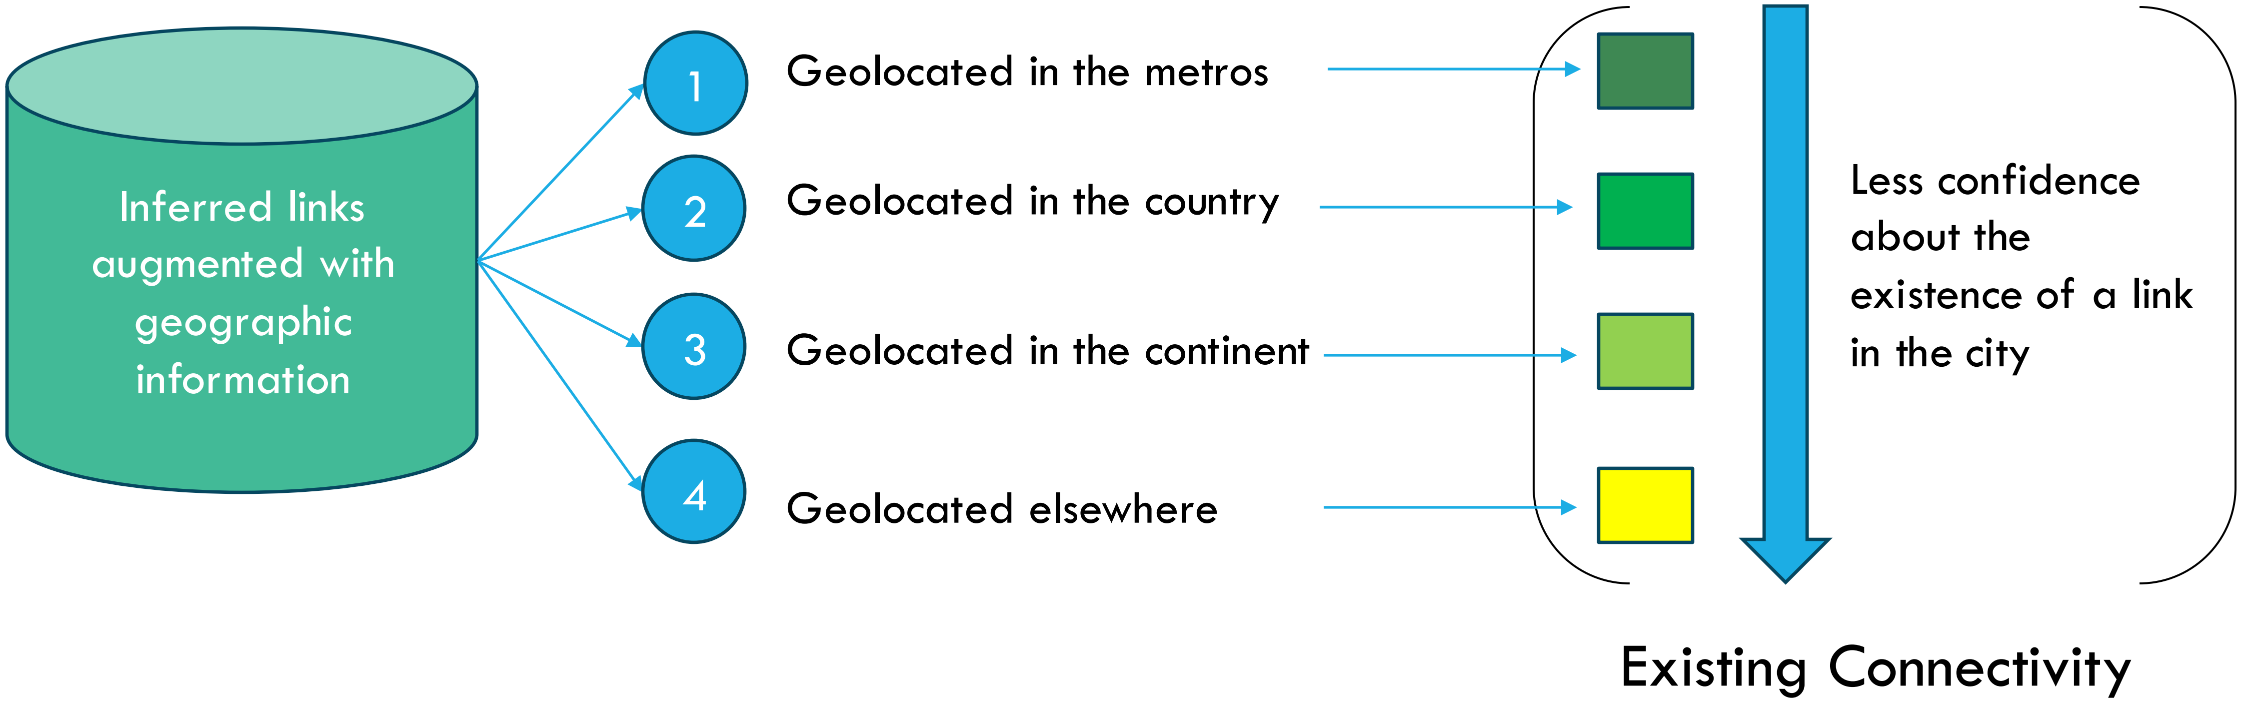

IMPROVING USABILITY WITH SHAPLEY VALUE

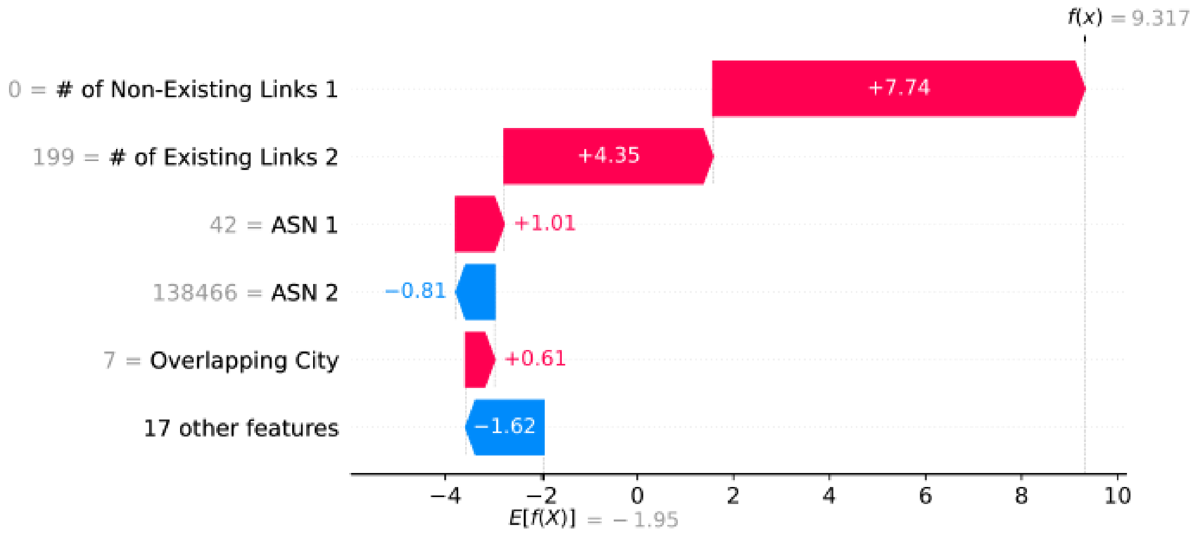

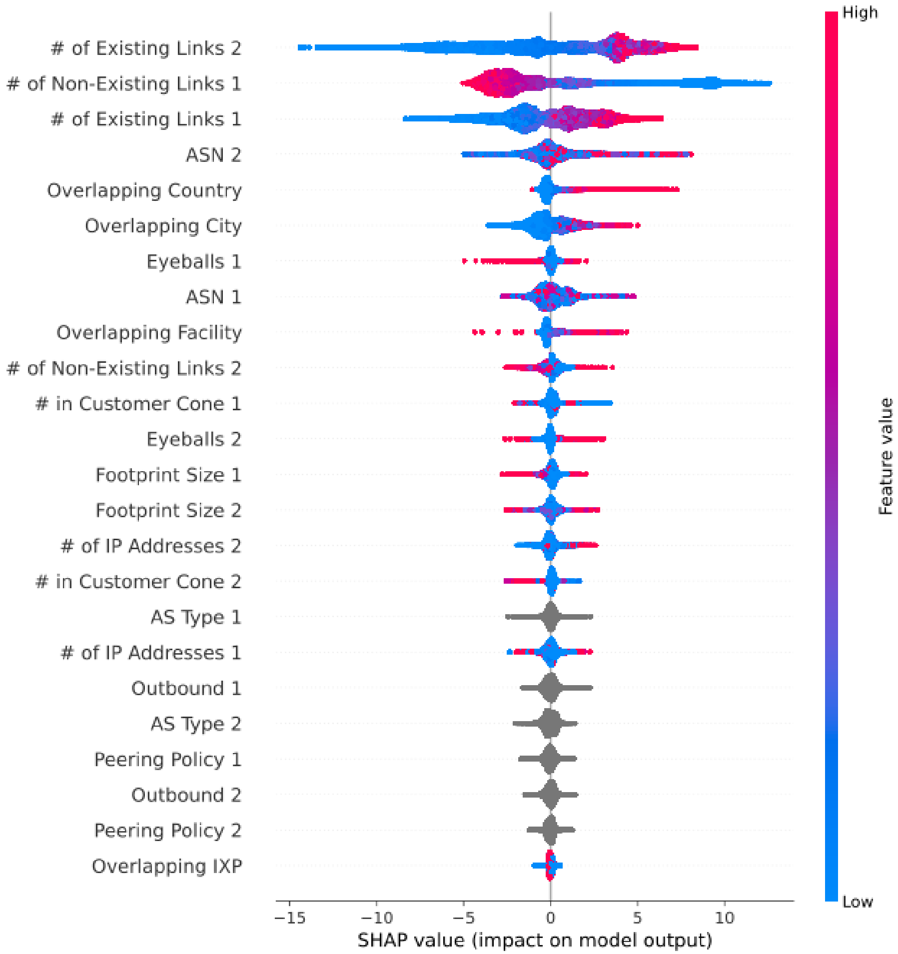

EVALUATION: CAN METASCRITIC SELECT THE RIGHT ENTRIES TO MEASURE TO IMPROVE ACCURACY?

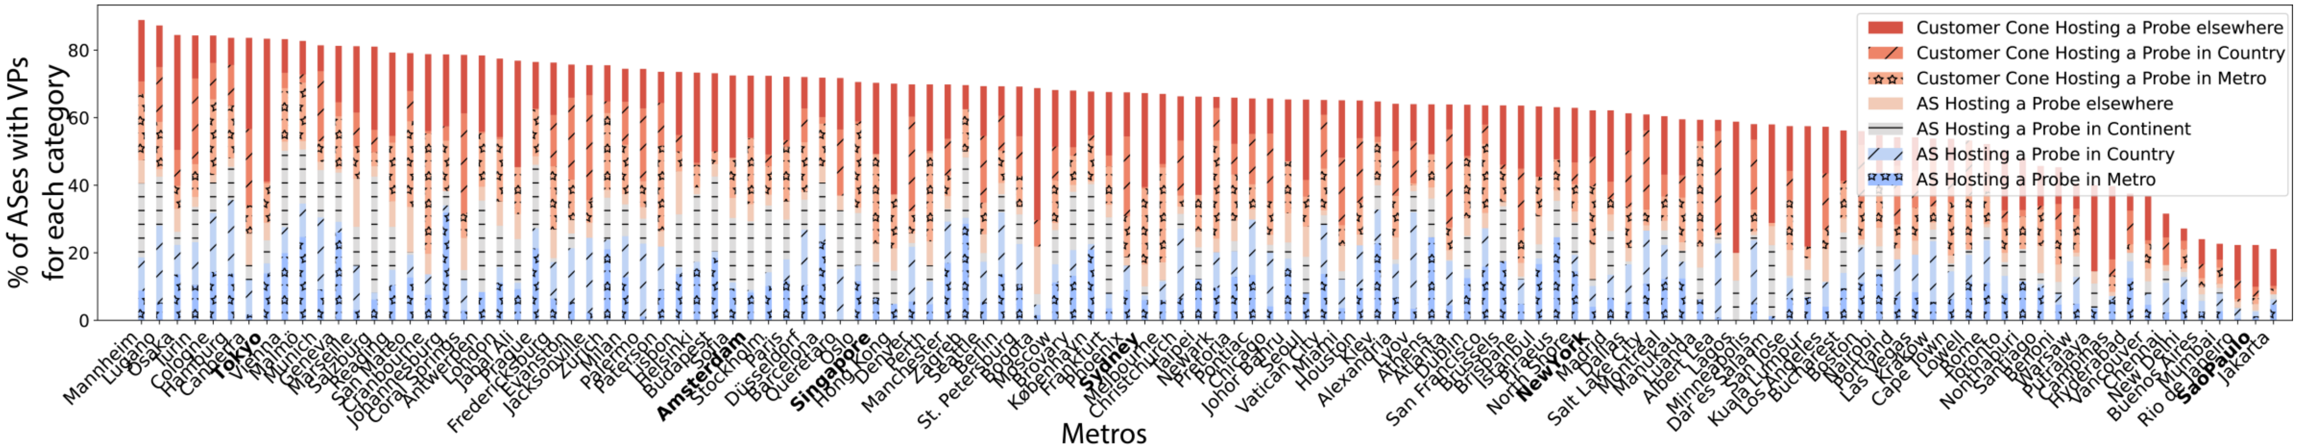

HOW WELL DOES METASCRITIC PERFORMS COMPARED TO ANOTHER ALGORITHM?