METASCRITIC:
TOPOLOGY DISCOVERY AS A RECOMMENDER SYSTEM

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

Published at ACM IMC '24

Punchline:

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

METASCRITIC:
TOPOLOGY DISCOVERY AS A RECOMMENDER SYSTEM

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

Published at ACM IMC '24

DESPITE DECADES OF EFFORT, VISIBILITY HAS 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 EFFORT, VISIBILITY HAS 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 EFFORT, VISIBILITY HAS 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].

More and more Ases peering, but vantage points plateauing [3]:

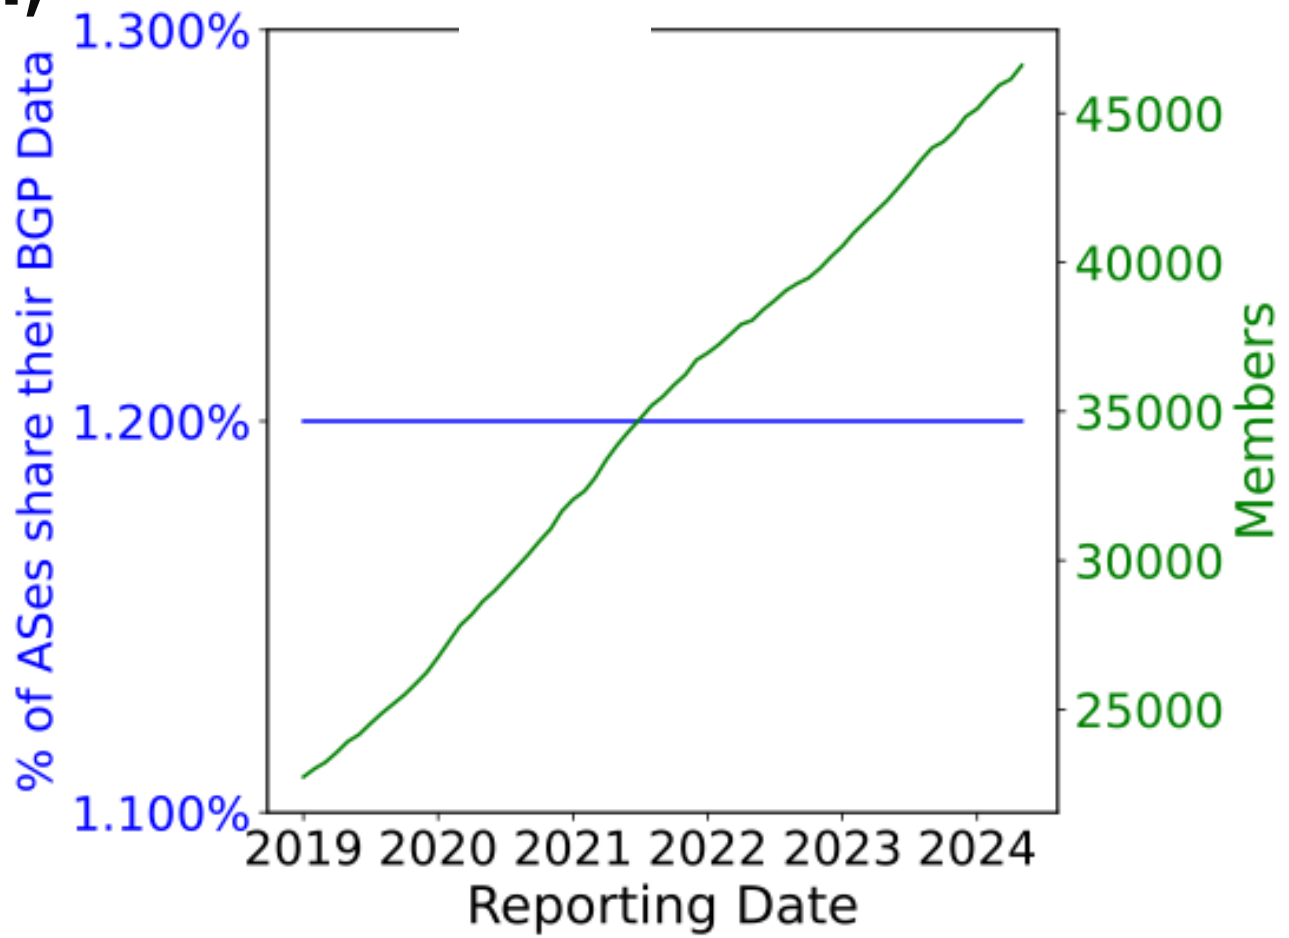

[1] Anatomy of a Large European IXP – Ager et al. in ACM SIGCOMM 2012

[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

WE NEED 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.

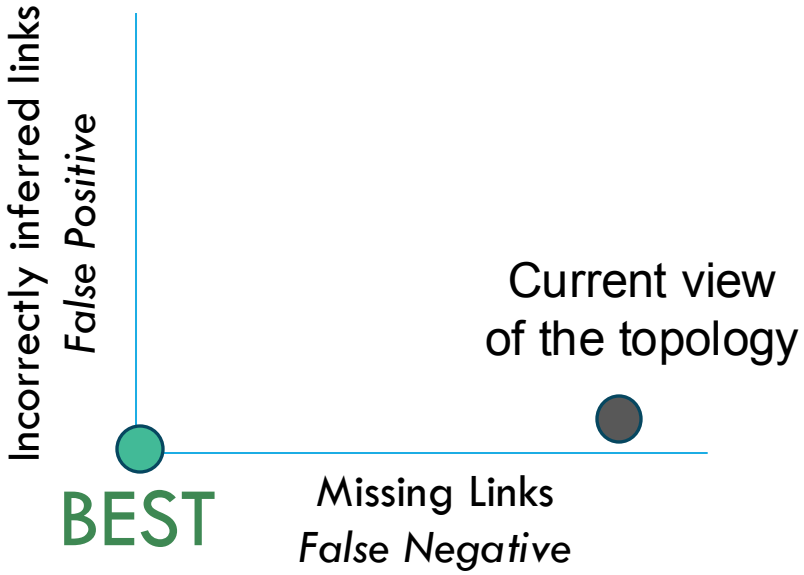

WE NEED A FUNDAMENTAL SHIFT: INFERENTIAL APPROACHES TO THE RESCUE

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.

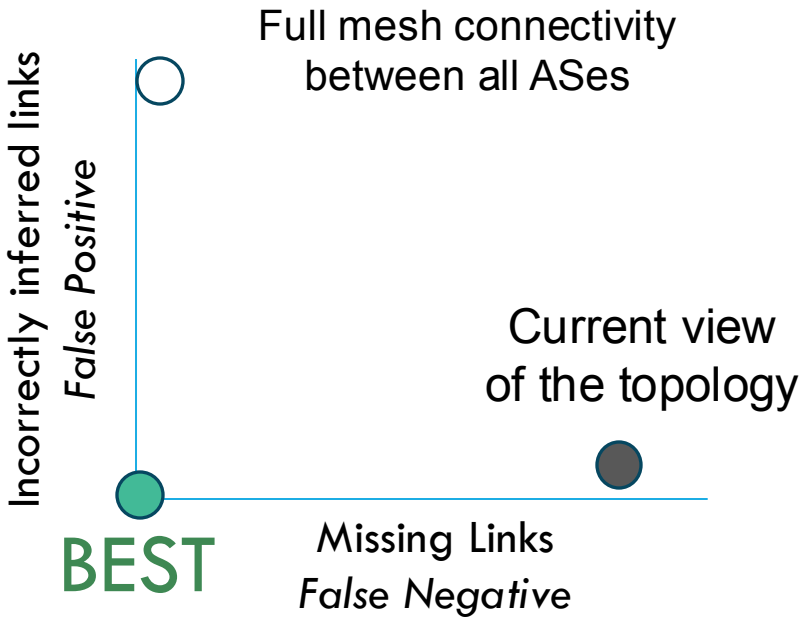

WE NEED A FUNDAMENTAL SHIFT: INFERENTIAL APPROACHES TO THE RESCUE

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

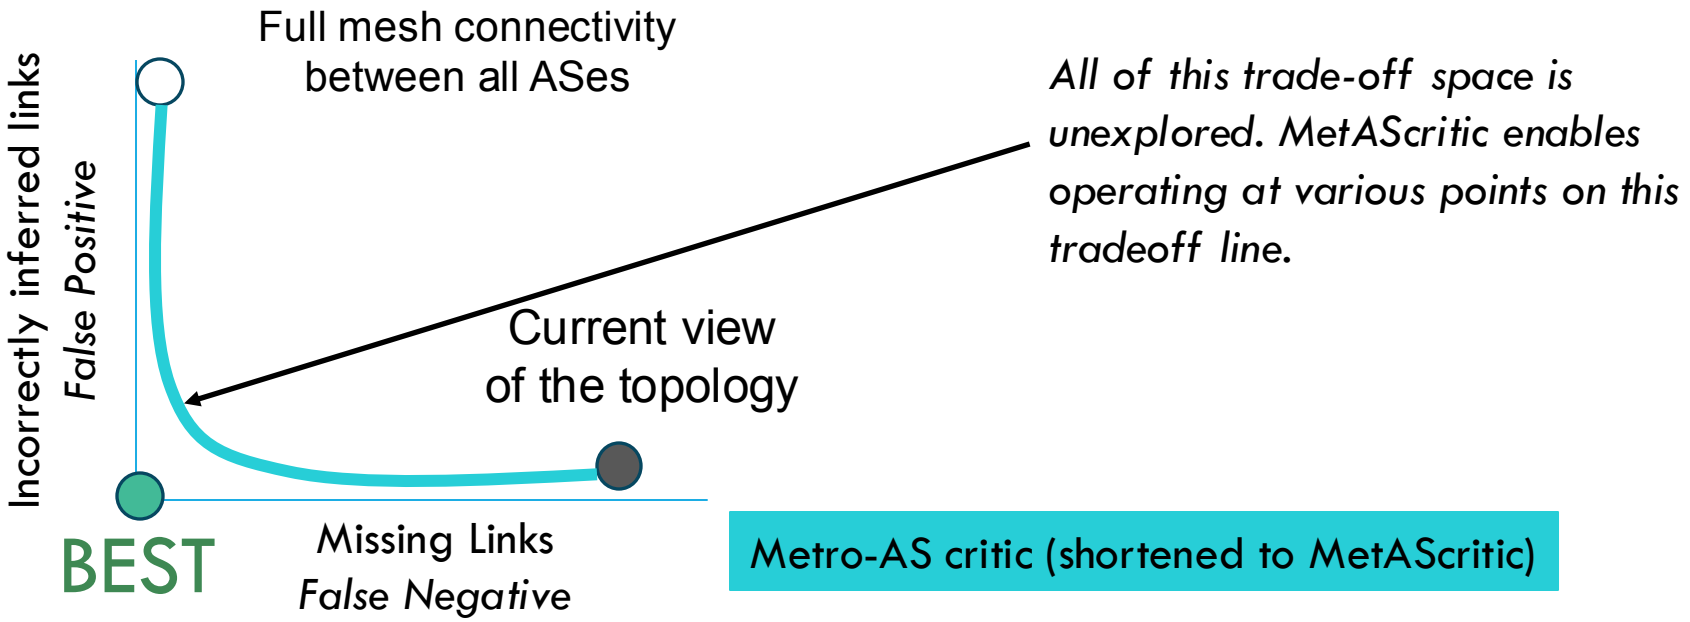

OUR SOLUTION: METASCRITIC, INSPIRED BY RECOMMENDER SYSTEMS

Key Idea: ASes with similar peering strategies—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.
- Similarly, metAScritic uses **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

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

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

Existing Behavior:
Dislikes people who are into magic.

UNDERSTANDING RECOMMENDATION IN THE CONTEXT OF TINDER.

Tinder analogy:

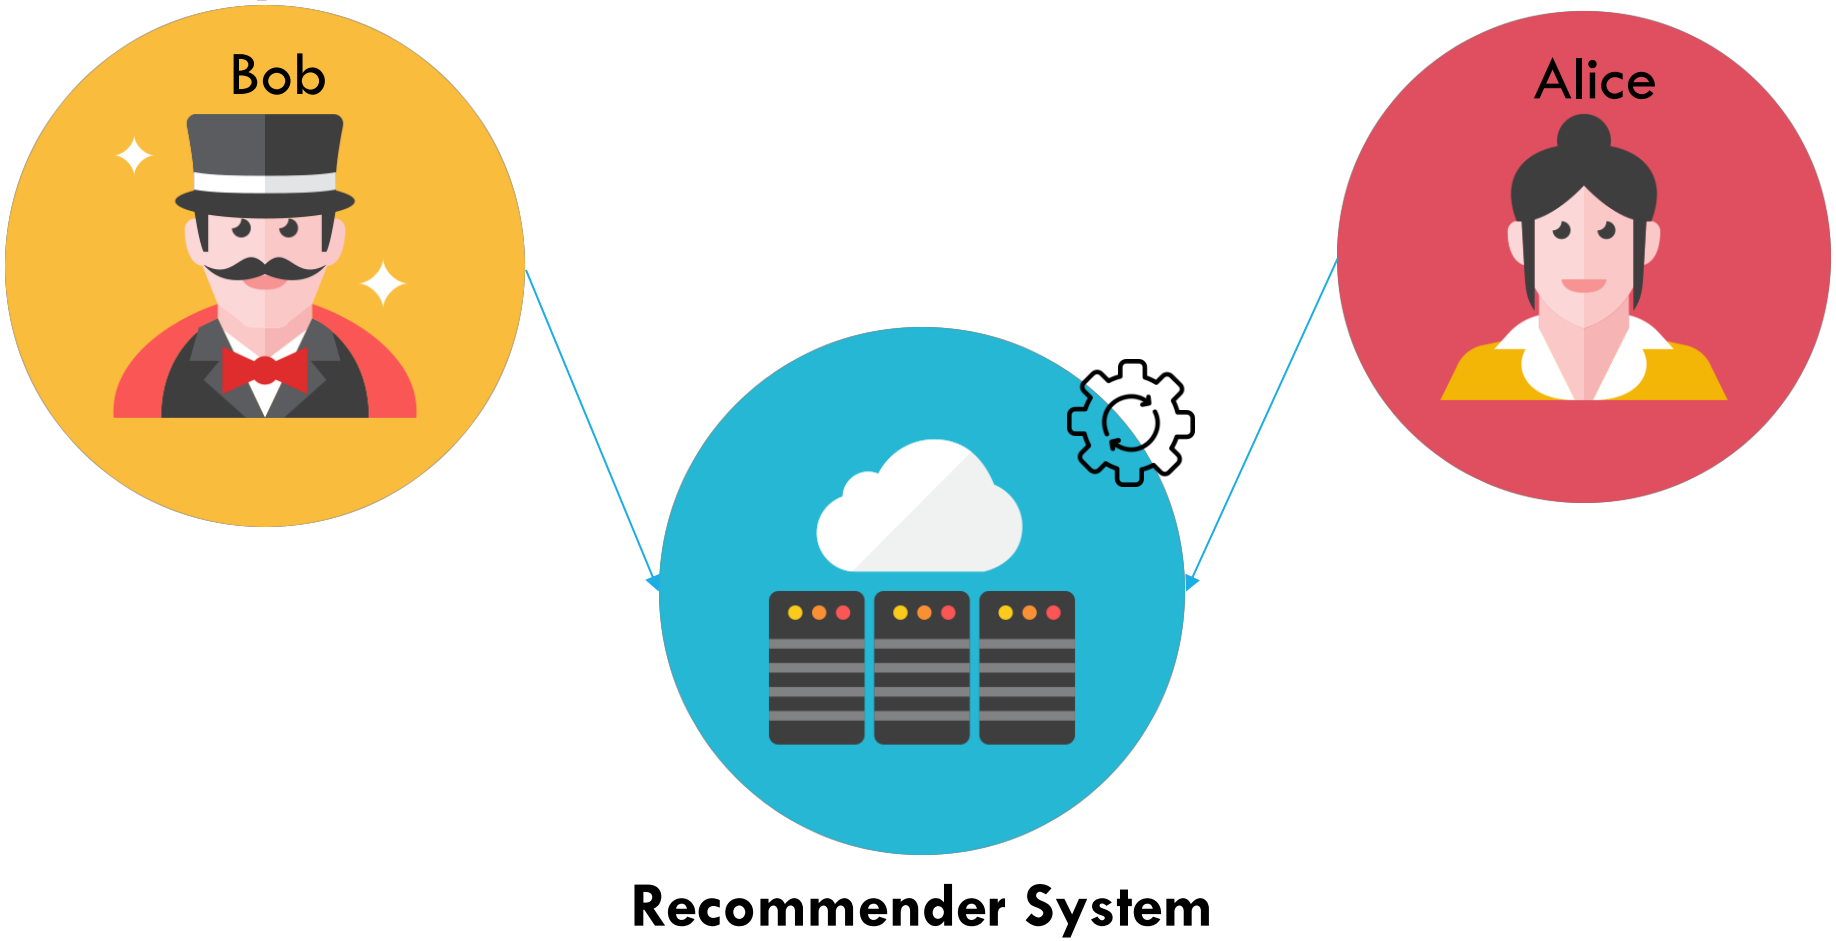

UNDERSTANDING RECOMMENDATION IN THE CONTEXT OF TINDER.

Tinder analogy:

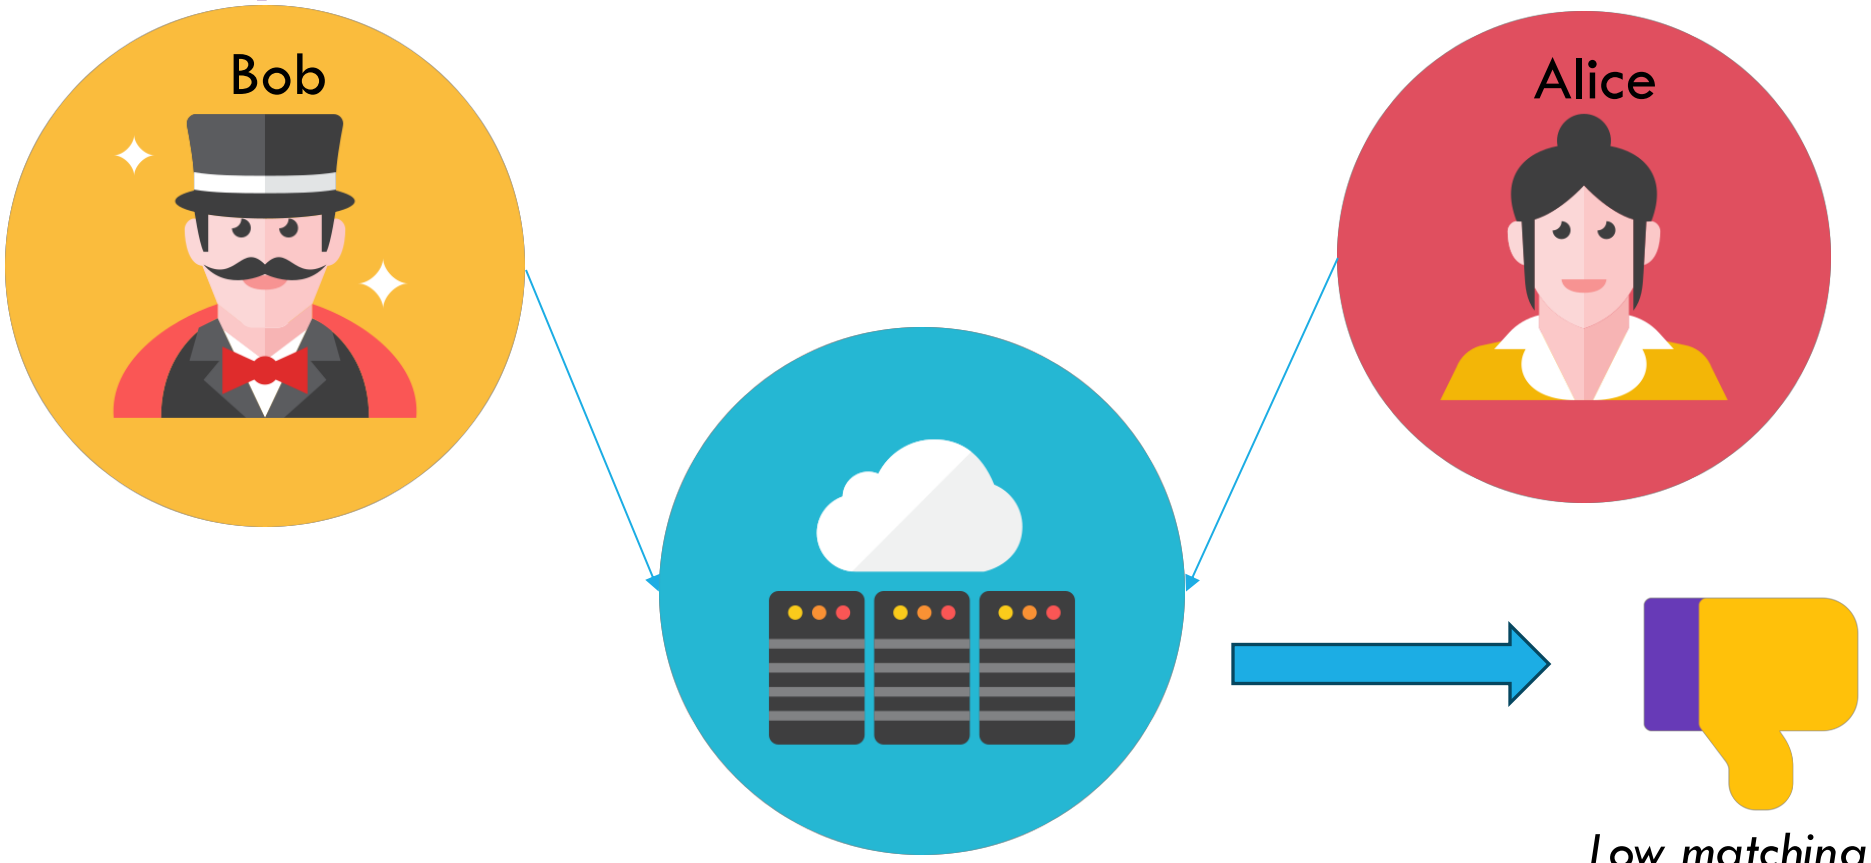

METASCRITIC DOES THE SAME, BUT WITH PEERING CONNECTIVITY.

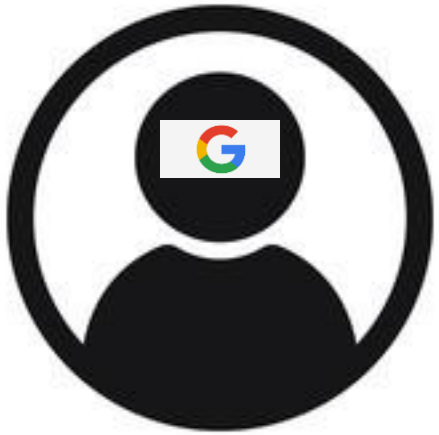

Intrinsic Properties:

Peering Policy: Open

Traffic Profile: Heavily Outbound

Number of Eyeballs: 1M

Customer Cone Size: 23

…

Existing Behavior:

Is peering with large access networks.

Is peering with ASes that peer with other Cloud Providers and CDNs.

Intrinsic Properties:

Peering Policy: Selective

Traffic Profile: Heavily Inbound

Number of Eyeballs: 42M

Customer Cone Size: 2372

…

Existing Behavior:

Is peering with Cloud Providers.

Is unlikely to peer with Open ASes.

METASCRITIC DOES THE SAME, BUT WITH PEERING CONNECTIVITY.

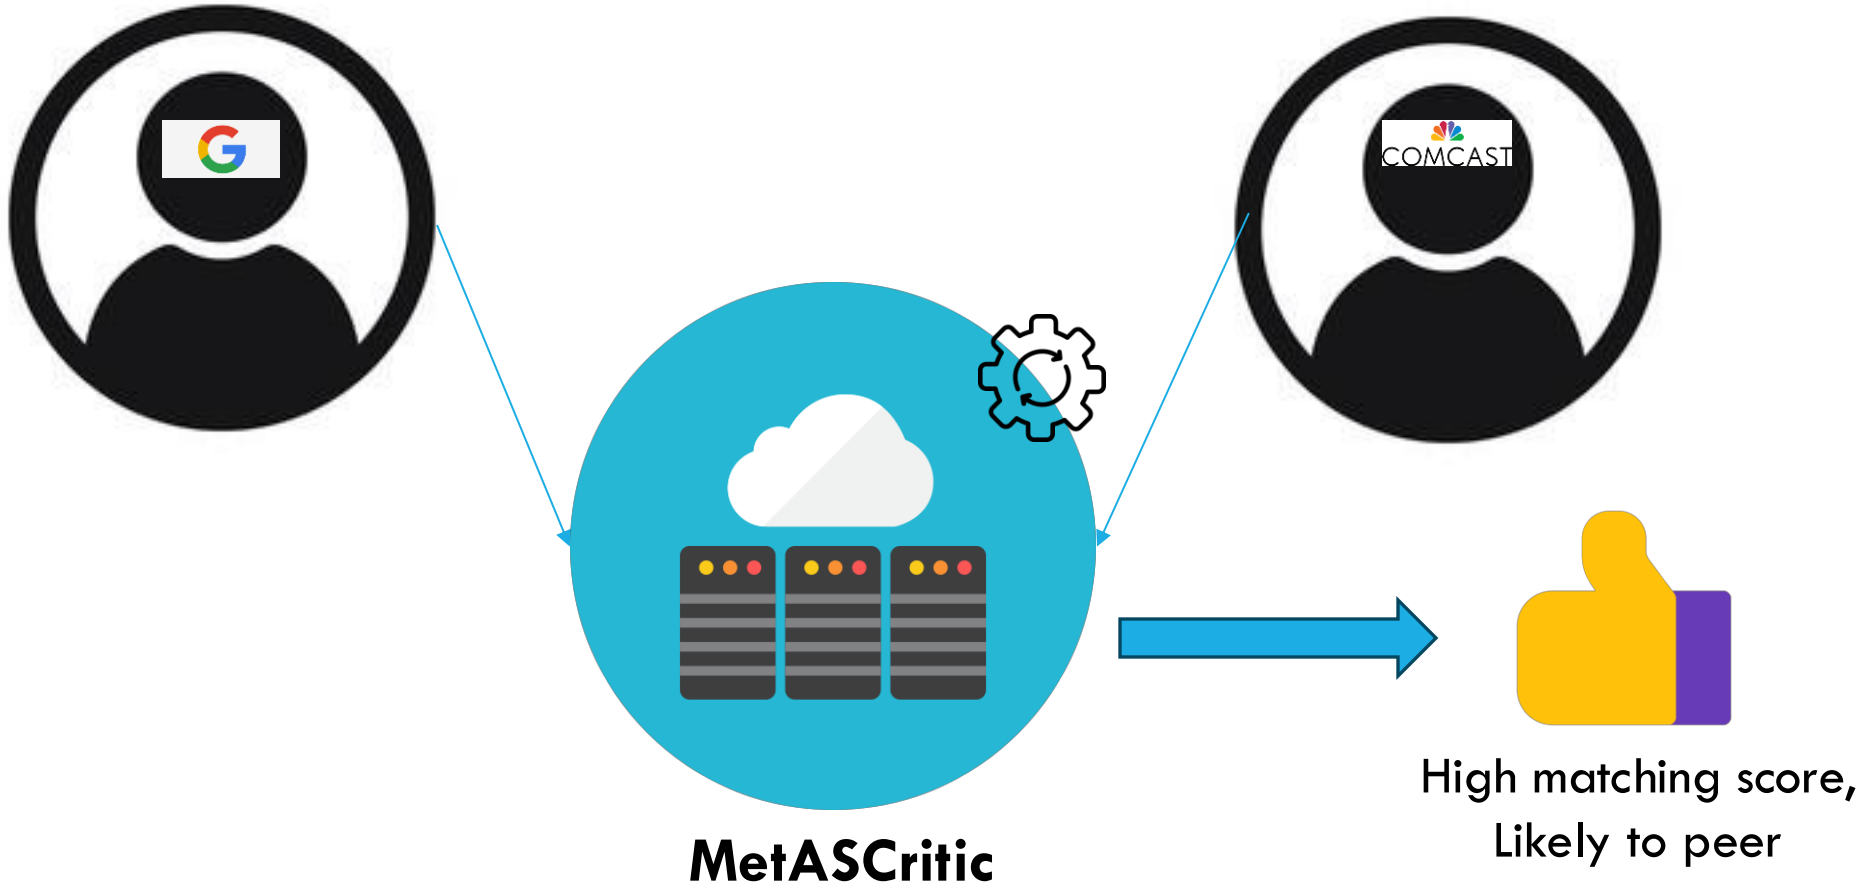

METHODOLOGY: HOW DOES METASCRITIC WORK?

METASCRITIC COMBINES OBSERVED LINKS WITH KNOWN PROPERTIES.

Measurements (RIPE Atlas, Ark)

Known Properties

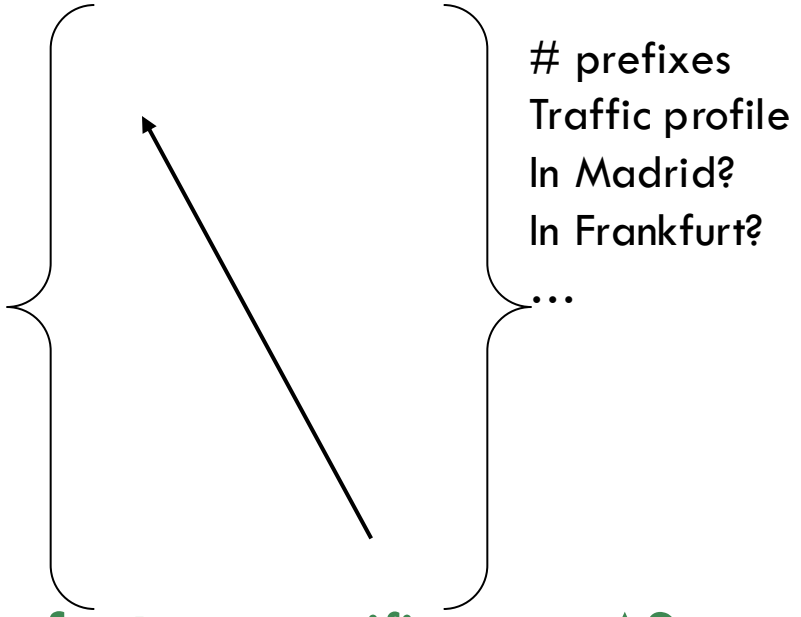

value of a **feature** specific to an AS,
(e.g., the number of prefixes)

METASCRITIC COMBINES OBSERVED LINKS WITH KNOWN PROPERTIES.

Measurements

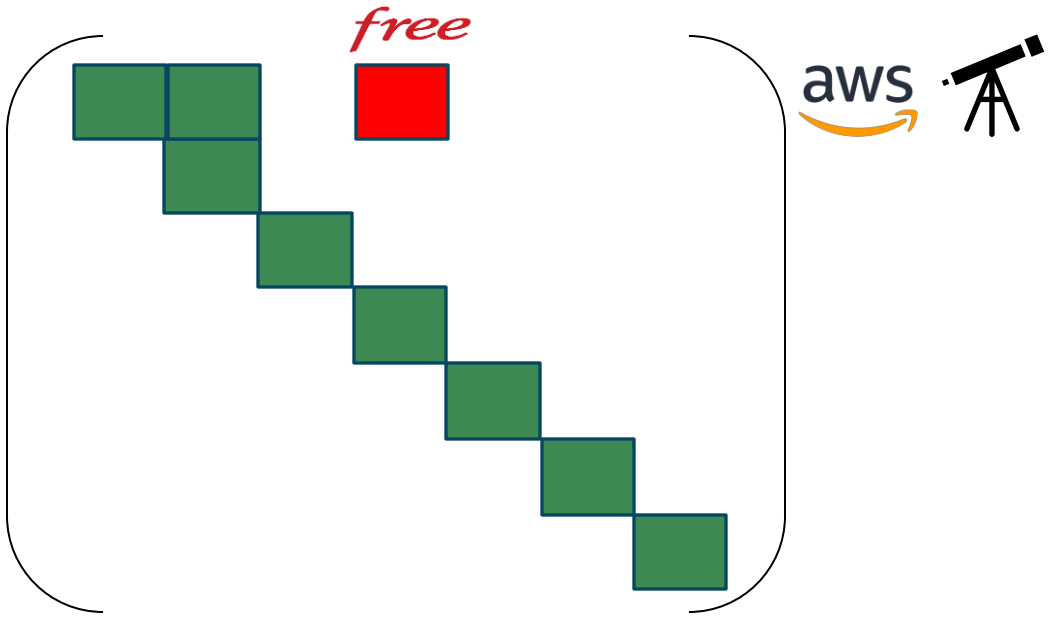

aws

Known Properties

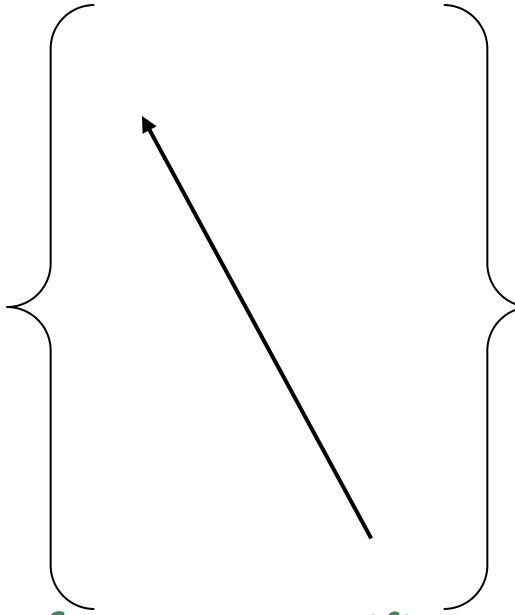

prefixes
Traffic profile
In Madrid?
In Frankfurt?
…

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

METASCRITIC COMBINES OBSERVED LINKS WITH KNOWN PROPERTIES.

Measurements

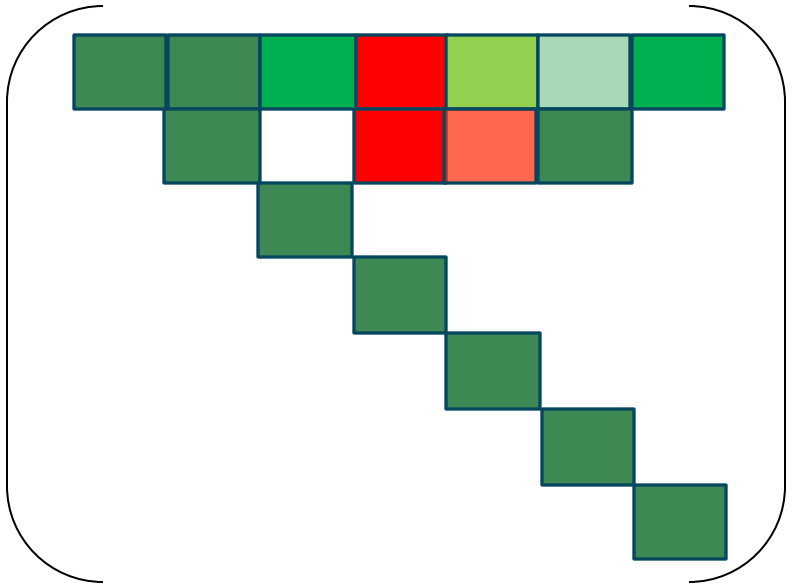

+

Known Properties

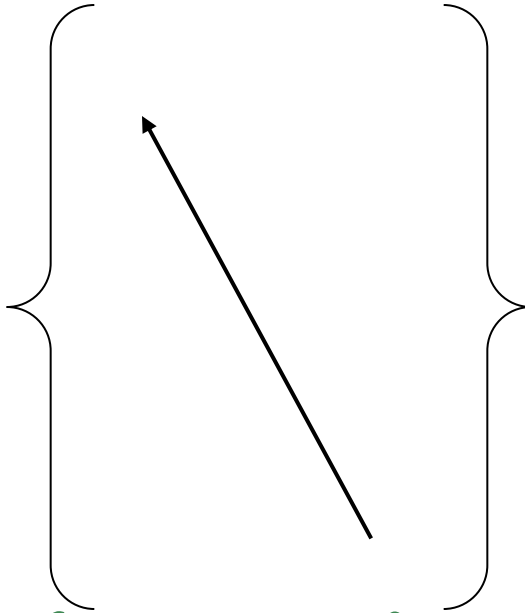

value of a **feature** specific to an AS,
(e.g., the number of prefixes)

COMPLETING THE MATRIX.

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

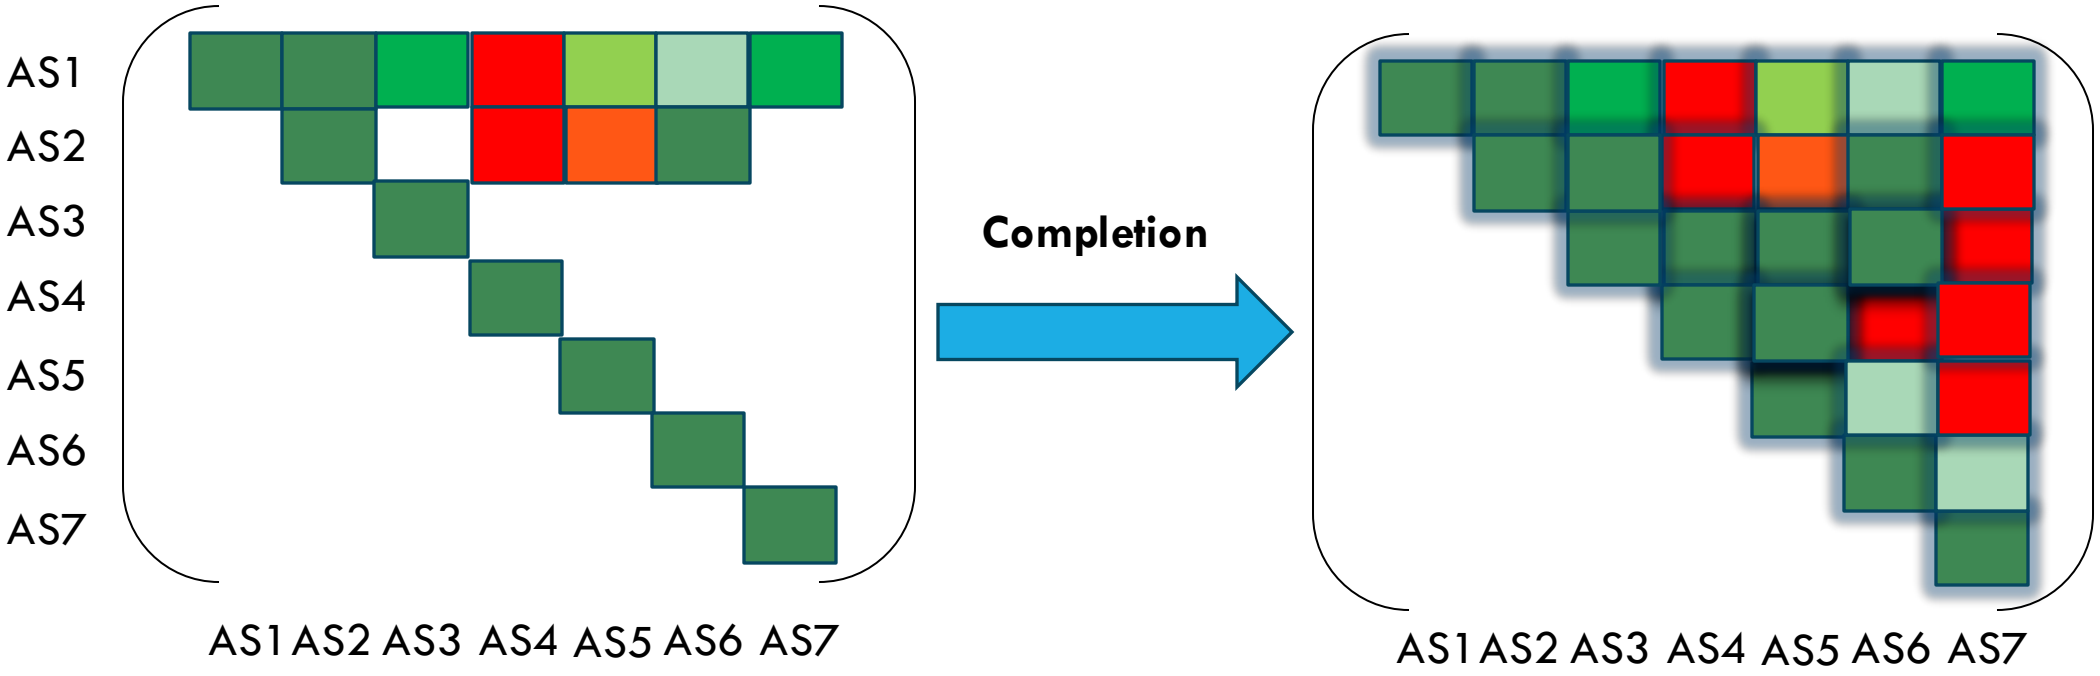

Existing Connectivity

Predicted Connectivity

COMPLETING THE MATRIX.

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

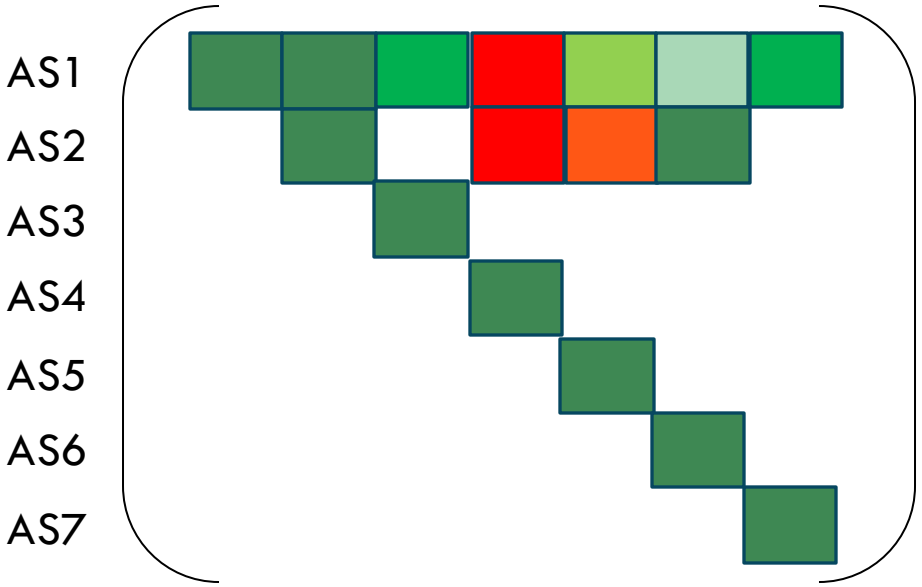

Existing Connectivity

Completion

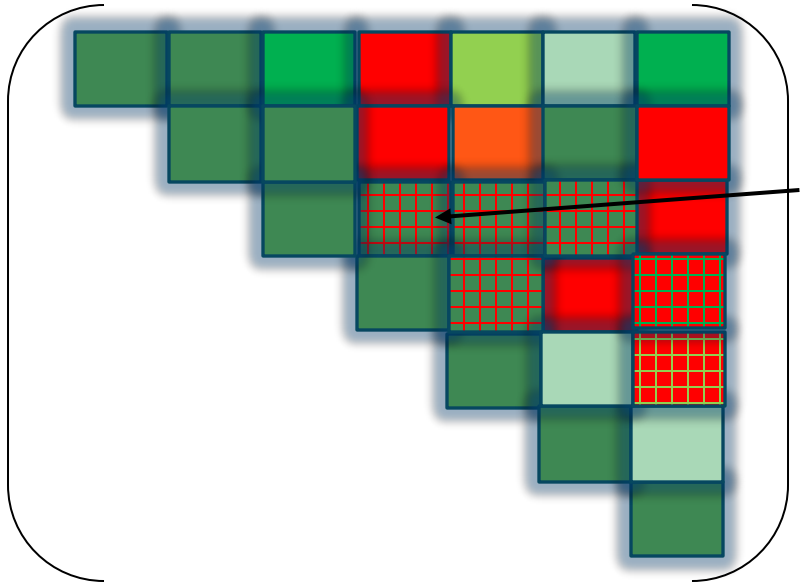

Predicted Connectivity

Incorrect inference of a link that does not exist (i.e., false positive)

False Negative

False Positive

NAVIGATING THE SPACE OF POSSIBLE TOPOLOGIES BY TRADING OFF FALSE POSITIVES AND NEGATIVES.

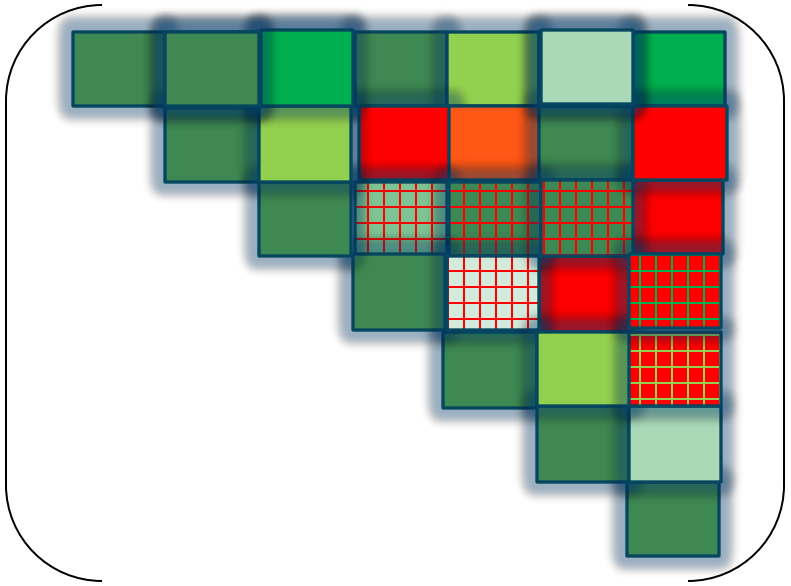

All green links

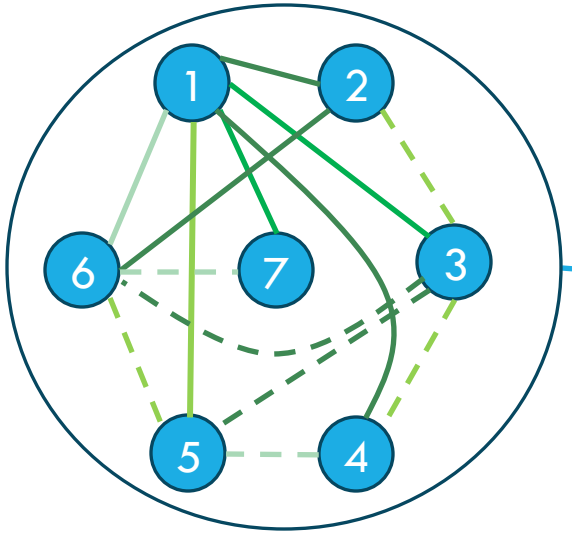

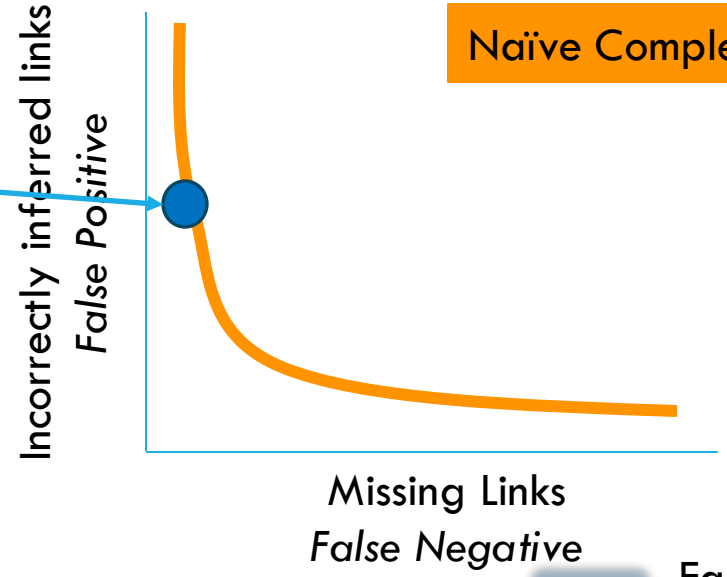

Measured Link

Inferred Link

Naïve Completion

False Negative

False Positive

NAVIGATING THE SPACE OF POSSIBLE TOPOLOGIES BY TRADING OFF FALSE POSITIVES AND NEGATIVES.

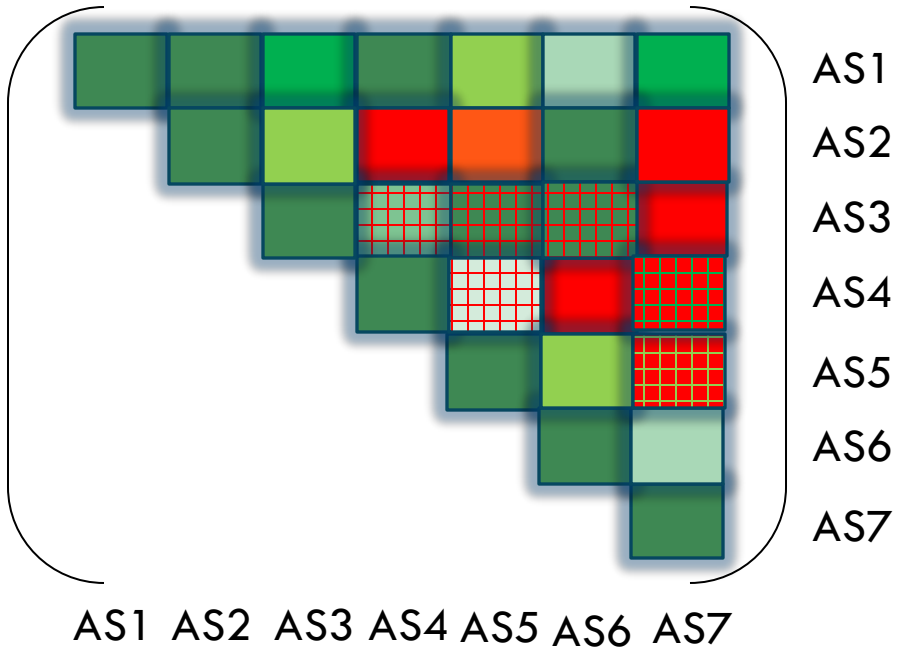

Only high-confidence links

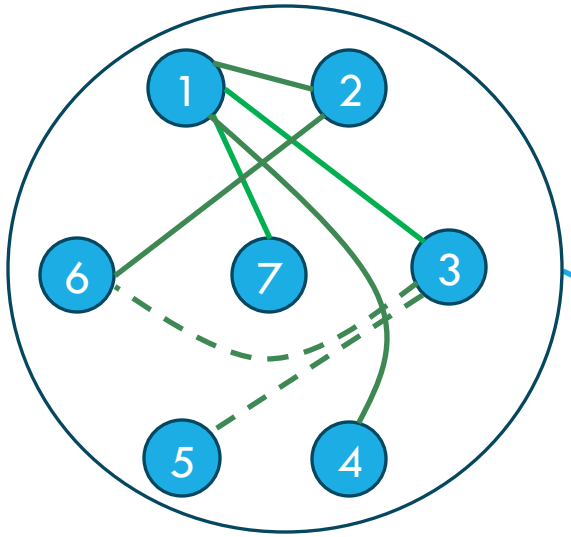

Measured Link

Inferred Link

Naïve Completion

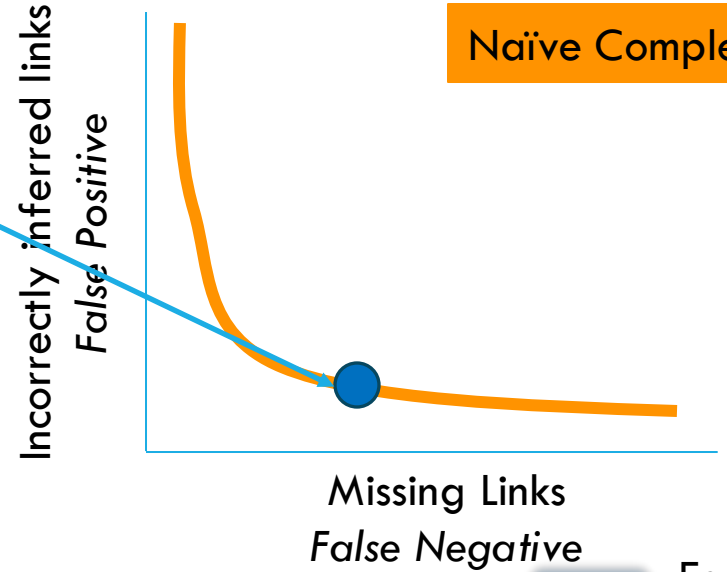

False Negative

False Positive

CHALLENGE: IMBALANCED MEASUREMENTS RESULT IN IMBALANCED INFERENCES.

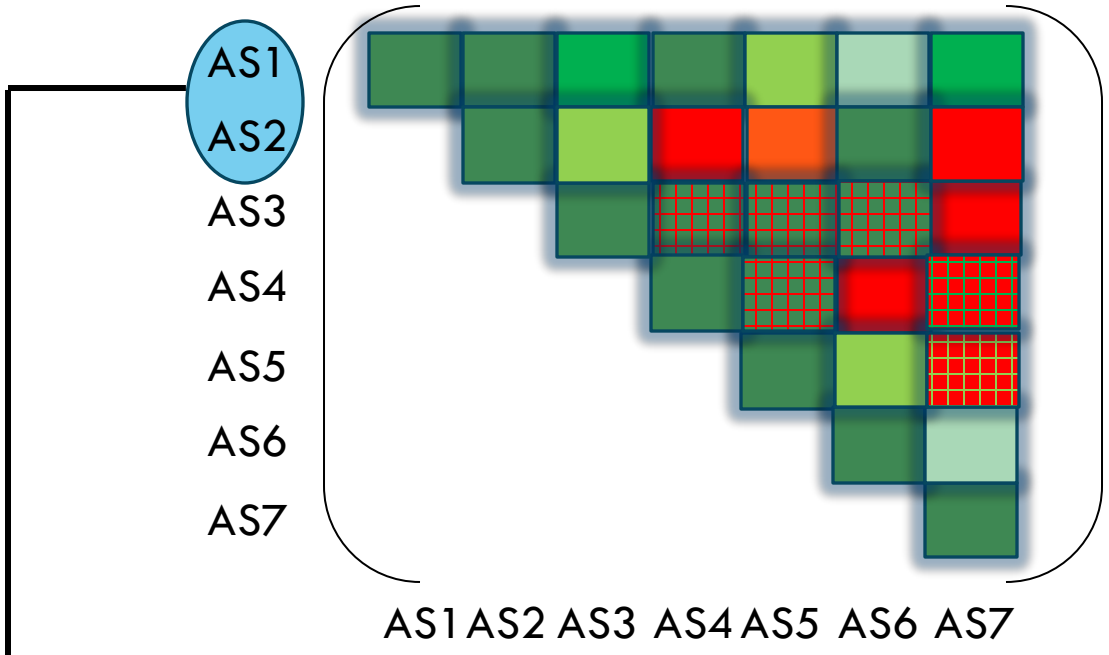

0 False Positive
0 False Negative

Great visibility for Ases that host a VP, so great predictive power

False Negative

False Positive

CHALLENGE: IMBALANCED MEASUREMENTS RESULT IN IMBALANCED INFERENCES.

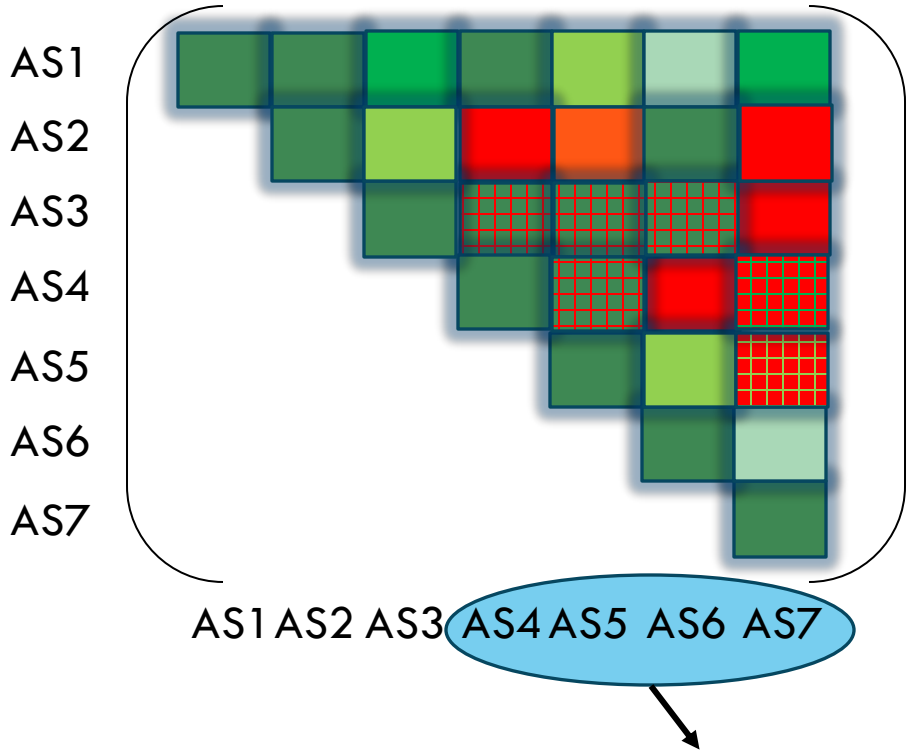

4 False Positives
2 False Negatives

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

False Negative

False Positive

THE IMPORTANCE OF DEBIASING THE DATASET.

From our collections of inferred and geolocated links:

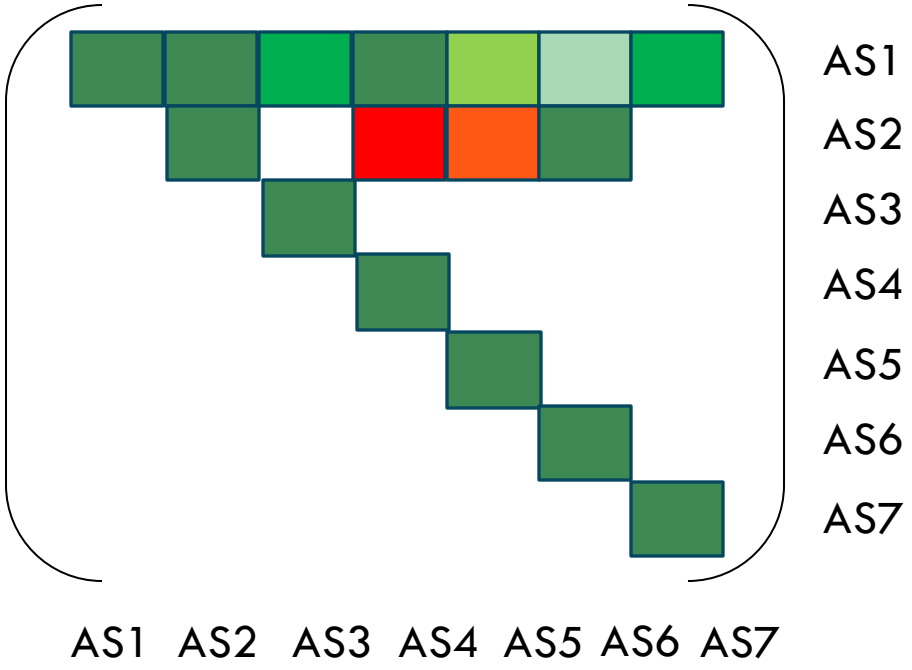

Problem: The public datasets are heavily skewed toward ASes that host vantage points.

APPROACH: ISSUE TRACEROUTES TO IMPROVE THE LEARNING PROCESS

Idea # 1: Use theoretical foundation to establish how many entries must be known per AS (for accurate matrix completion).

Idea # 1: Identify unknown links that are likely to be the most informative.

Idea # 2: Model how likely each possible traceroute is to uncover presence (or absence) of a link.

BY INCORPORATING A FEW MEASUREMENTS, METASCRITIC IMPROVES ITS COMPLETION.

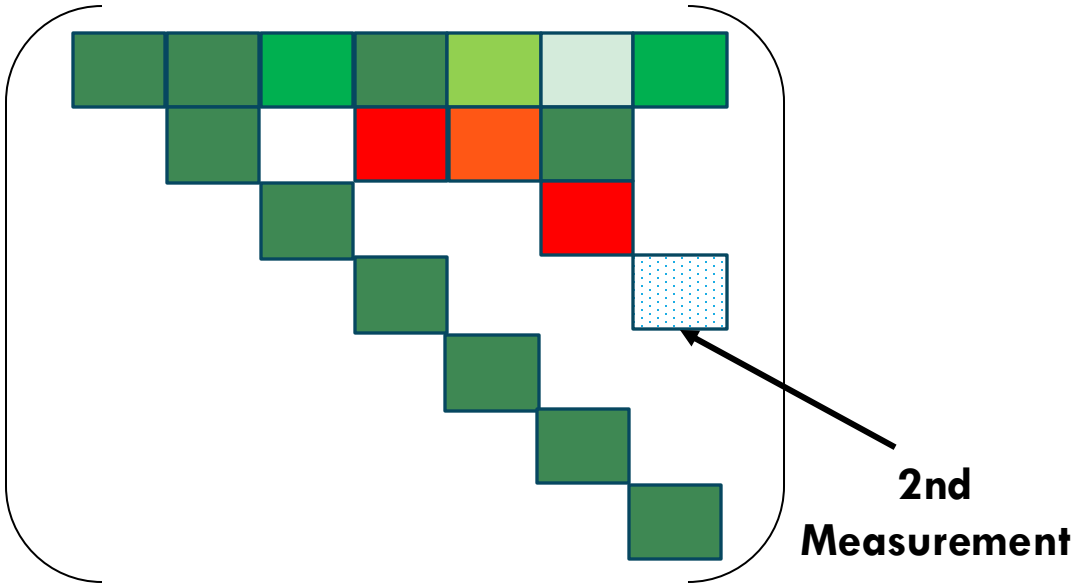

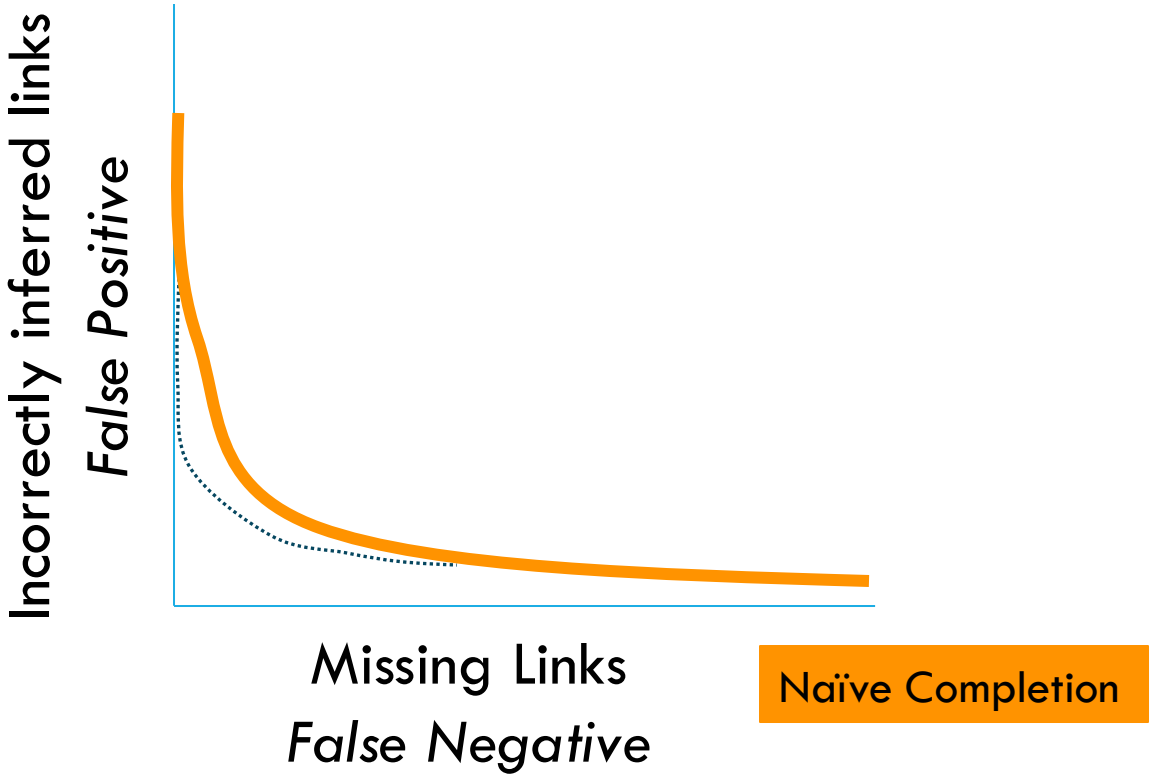

BY INCORPORATING A FEW MEASUREMENTS, METASCRITIC IMPROVES ITS COMPLETION.

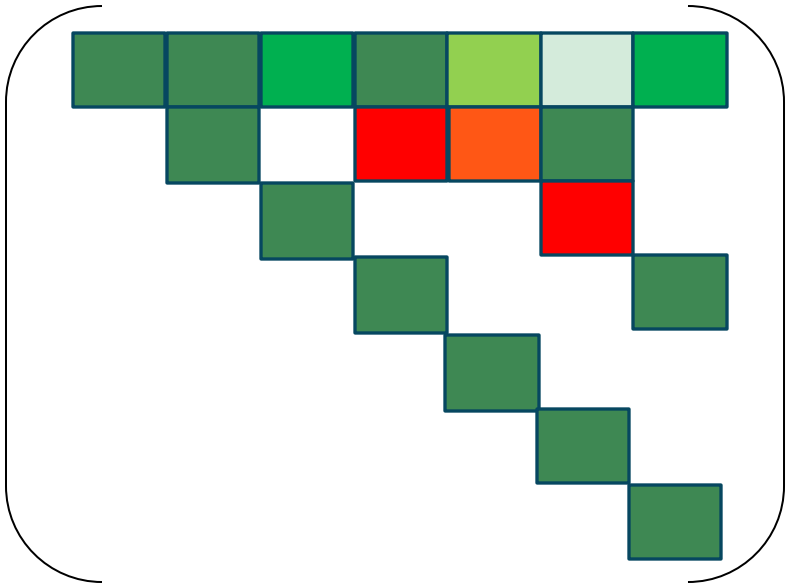

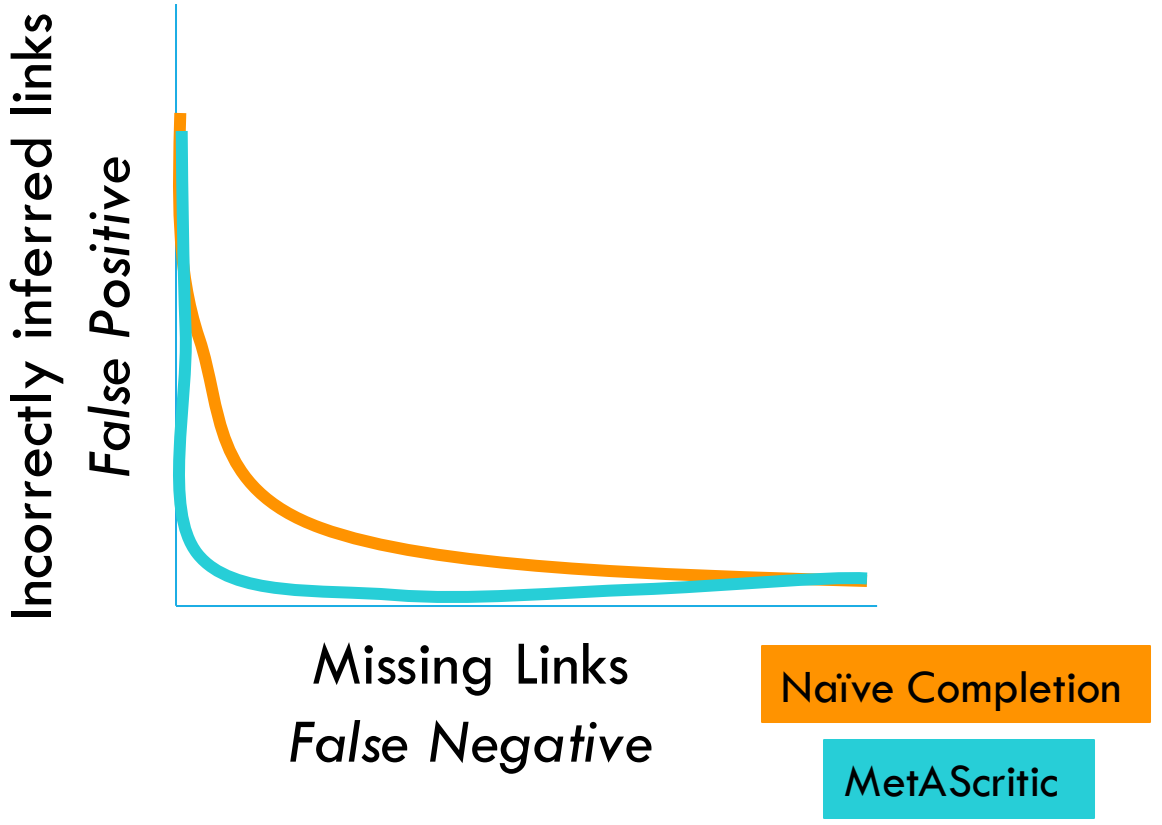

RESULTS

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 Communities	N.A.	0.9 – 1

Many more in the paper!

EXAMPLE USE CASES

More complete topology for simulations leads to more accurate results

- e.g., predicting the impact of a BGP hijack (see metAScritic IMC 2024)

Guide vantage point placement (e.g., for Ark or GILL BGP collector)

- Target parts of the Internet with many predicted but unobserved links?
- Target parts of the Internet with low confidence in predictions?

Rather than treating a measured topology as "the truth", analysis based on uncertainty

- (Measurements have missing links AND false links:
  bdrmapIT reported 1-9% error rate, my cloud interconnectivity paper reported 11-15% false links)
- Sweep thresholds to bound analysis (e.g., for AS hegemony and transit influence)
- Enables probabilistic reasoning (how likely is a link to exist, and what is its impact?)

A FINAL THOUGHT:

Applying machine learning for topology discovery is feasible and can help use cases.

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!

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.

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

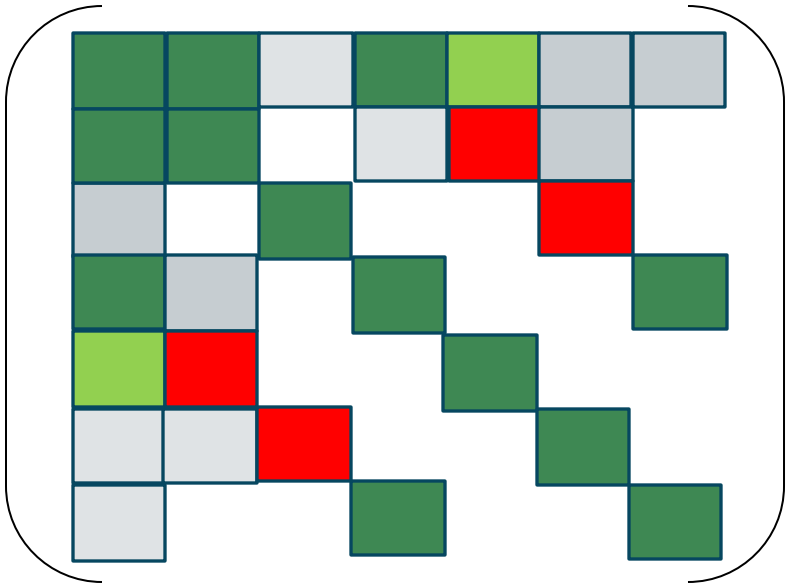

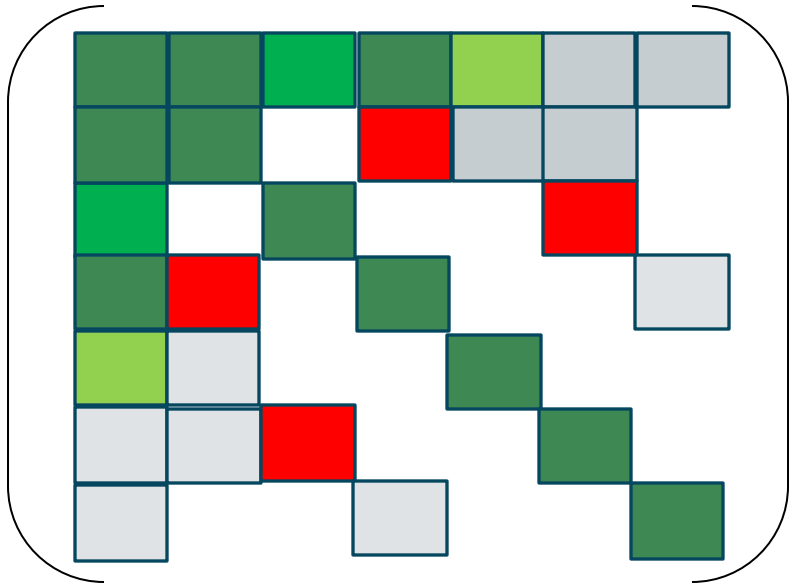

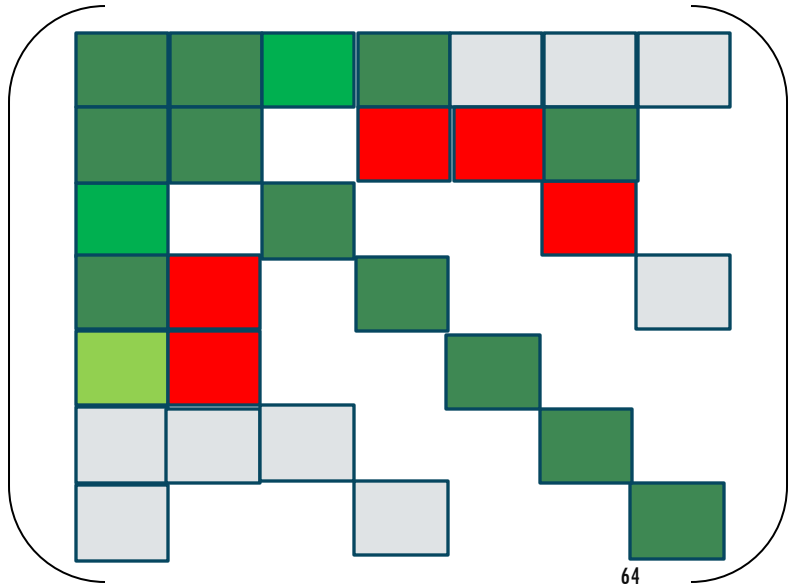

HIGH PRECISION WITH HIGH RECALL