The RIPE Atlas Measurements Patterns Analysis tool (by the AI4NetMon project)

Pavlos Sermpezis

AI4NetMon project (https://ai4netmon.csd.auth.gr/)

Data & Web Science Laboratory (Datalab), https://datalab.csd.auth.gr/
Computer Science Dept., Aristotle University of Thessaloniki

AI4NetMon project: dataset, code, API, Web app

- AI4NetMon project https://ai4netmon.csd.auth.gr/
  - You can find all the information about the project!

RIPE NCC Community Projects Fund 2022

- **RIPE Atlas Measurements Patterns Analysis tool** https://ai4netmon.csd.auth.gr/atlas-app/

- **Web app** https://app-ai4netmon.csd.auth.gr/
- **API** https://ai4netmon.csd.auth.gr/api/
- **Code & Data @ GitHub** https://github.com/sermpezis/ai4netmon/
- **Paper @ TMA'23** https://arxiv.org/abs/2307.09958

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis

The tool… **in a nutshell**

- ***AI4NetMon's RIPE Atlas Measurements Patterns Analysis tool***
  … or just ***"The tool"*** (for brevity in this presentation)
- Goal: Analyze patterns in a set of measurements
- Bias & generic patterns
- Complementary to RIPE Atlas UI

The tool… vs. RIPE Atlas UI

- **RIPE Atlas UI** presents (very nicely!) the **result** of **one measurement**

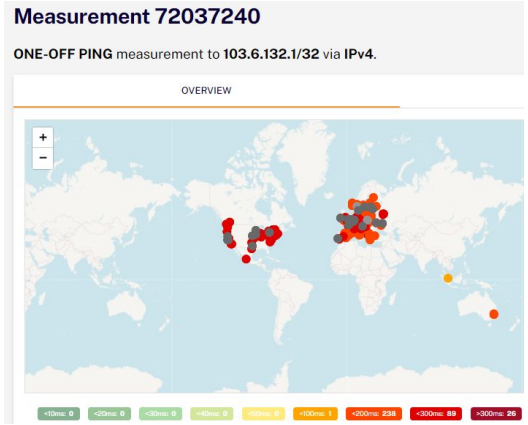

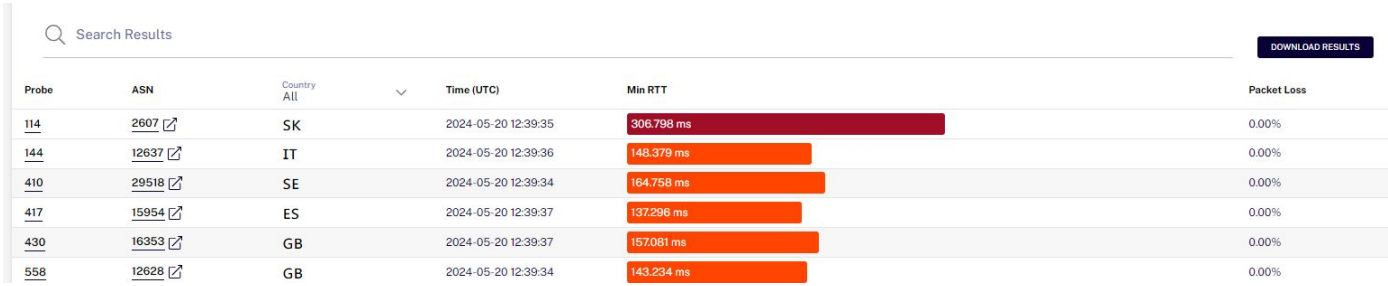

- **The AI4NetMon's tool** presents data about the **set of probes (not the results)** in **a set of measurements**

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis
AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

The tool… **in a nutshell**

- **The user** provides a set of RIPE Atlas measurement IDs
  - e.g., “1018338, 1004340, 1017820, 1019139, 1017005”

- **The tool (in the backend)** retrieves information about the measurements (RIPE Atlas API, AI4NetMon API) and performs analyses

- **The tool (through its UI)** presents a set of visualizations with the characteristics and patterns of measurements

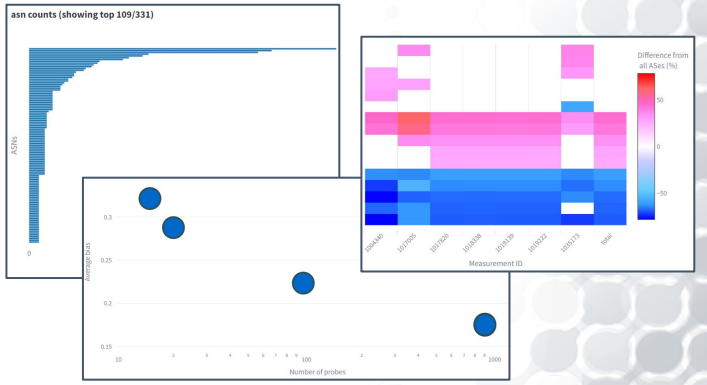

The tool… **in a nutshell**

- General statistics of measurements
  - Top most frequent ASNs and probes
  - Number of probes and ASNs used per measurement
- Bias analysis of measurements
  - Bias comparison: Measurements (average) vs random sample of probes
  - Number of probes vs Avg Bias per measurement
  - Bias Causes
  - CDF of Bias per Bias dimension
  - Data details

The tool… **in action**

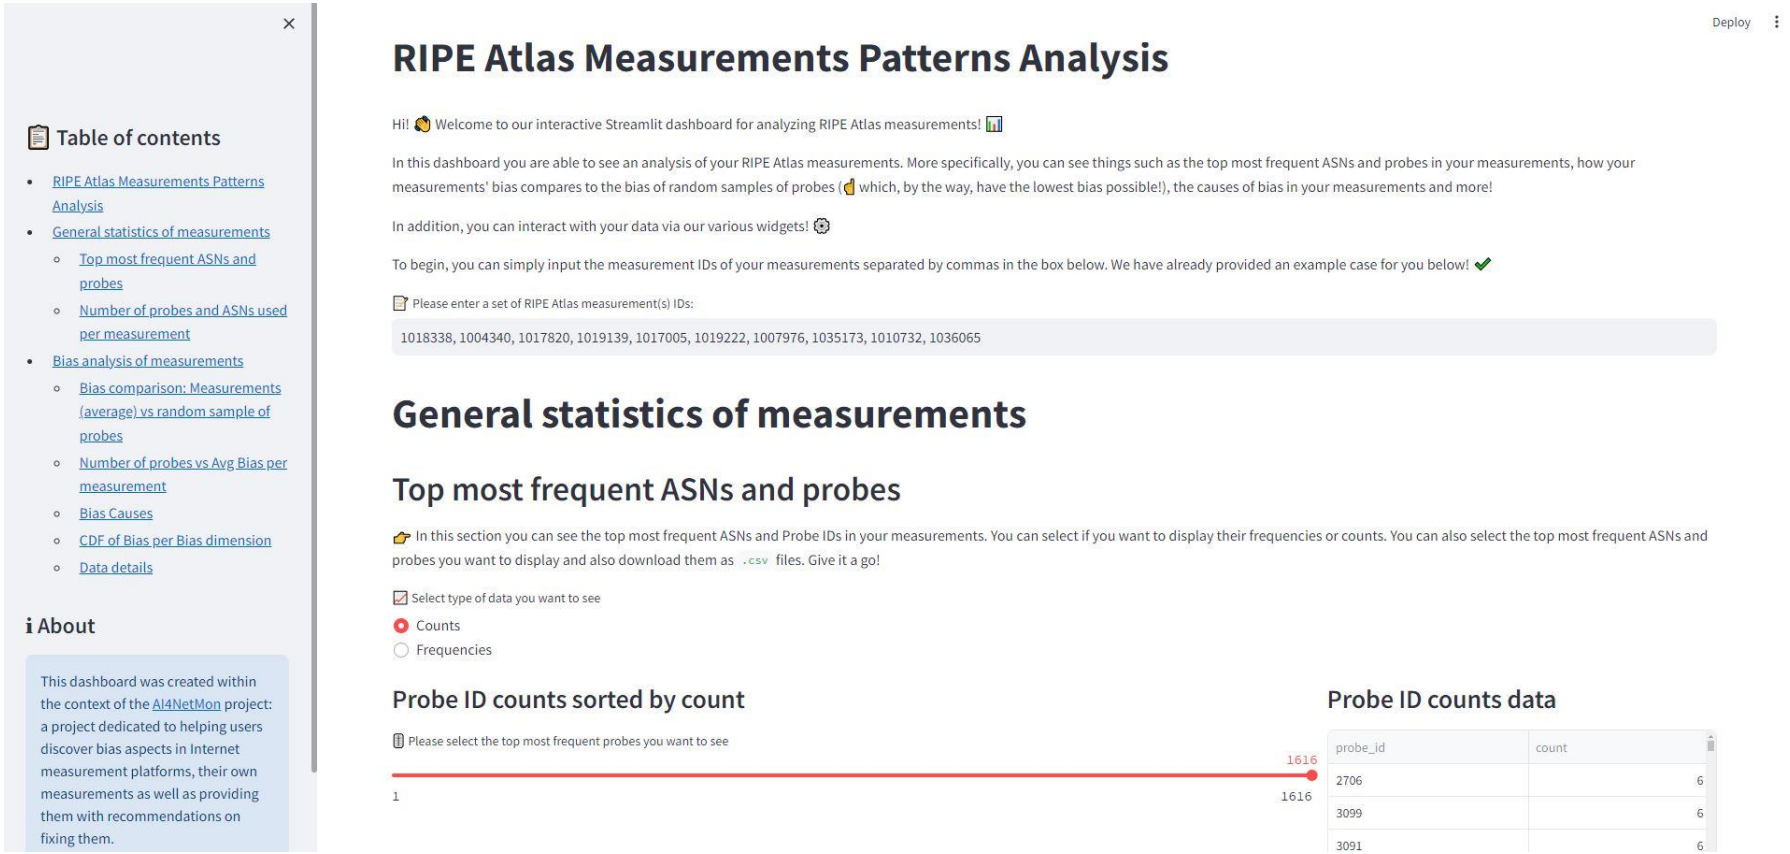

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis
AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

The tool… **in action**

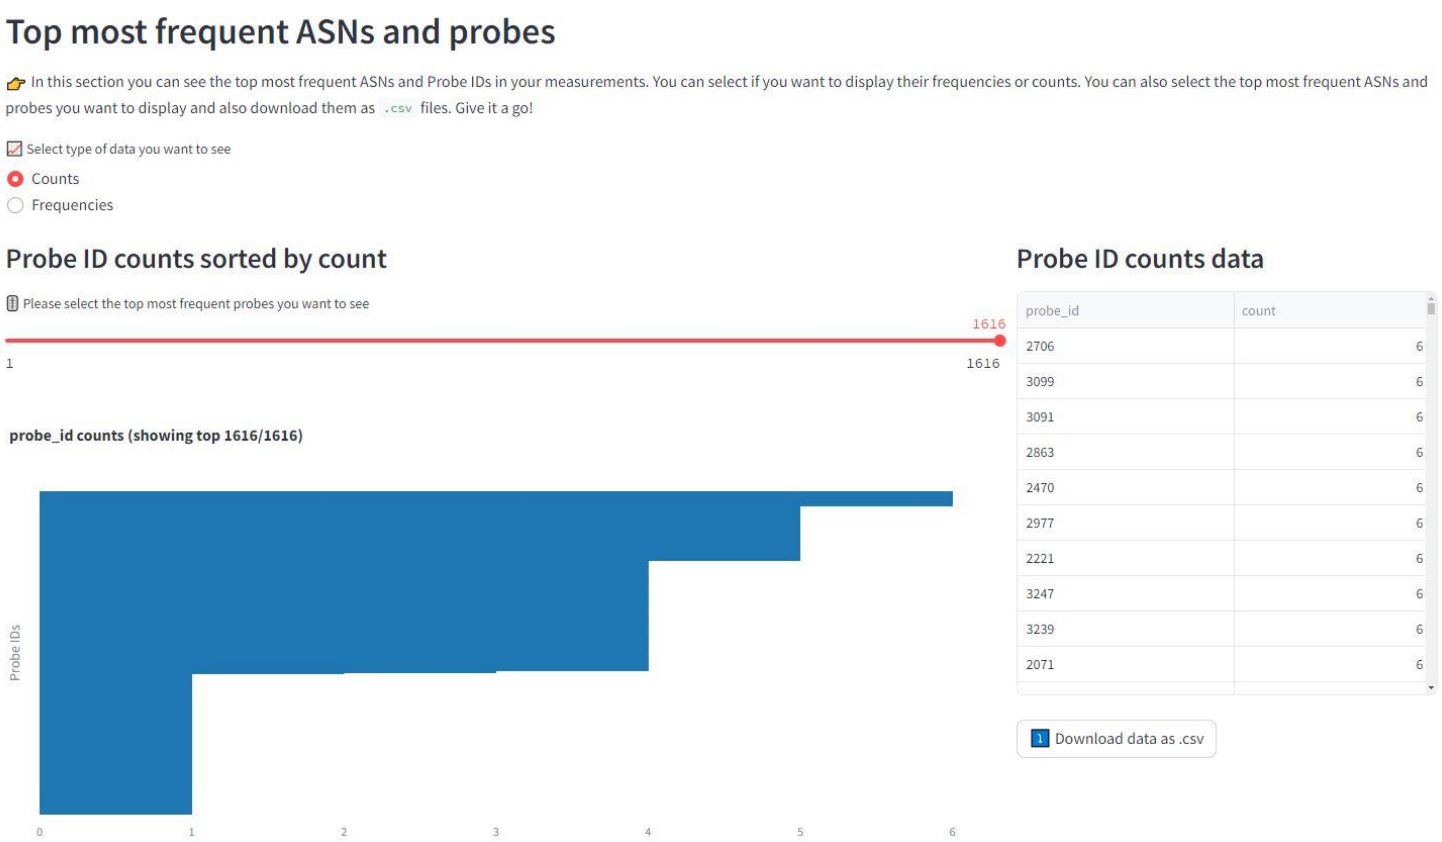

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis

AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

The tool… **in action**

ASN counts sorted by count

Please select the top most frequent ASNs you want to see

28

1 642

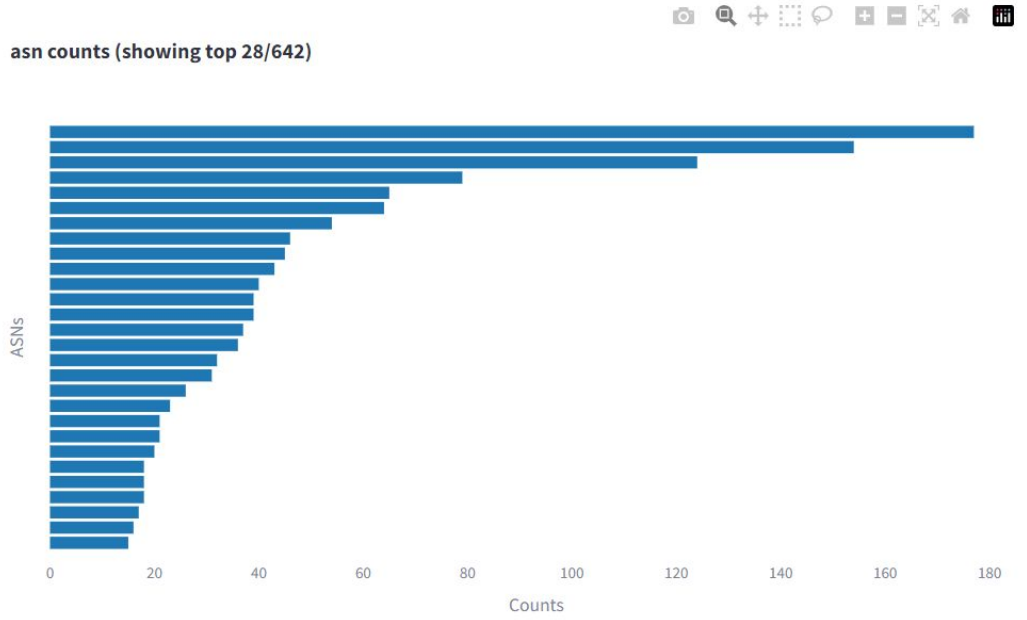

ASN counts data

asn	count
7922	177
3320	154
9	124
12322	79
3333	65
20712	64
3265	54
3215	46
6830	45
13030	43

Download data as .csv

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis

AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

The tool… **in action**

Bias analysis of measurements

Bias comparison: Measurements (average) vs random sample of probes

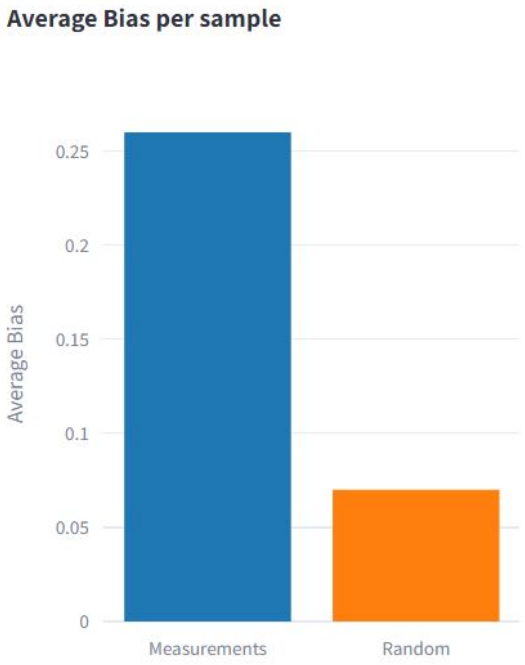

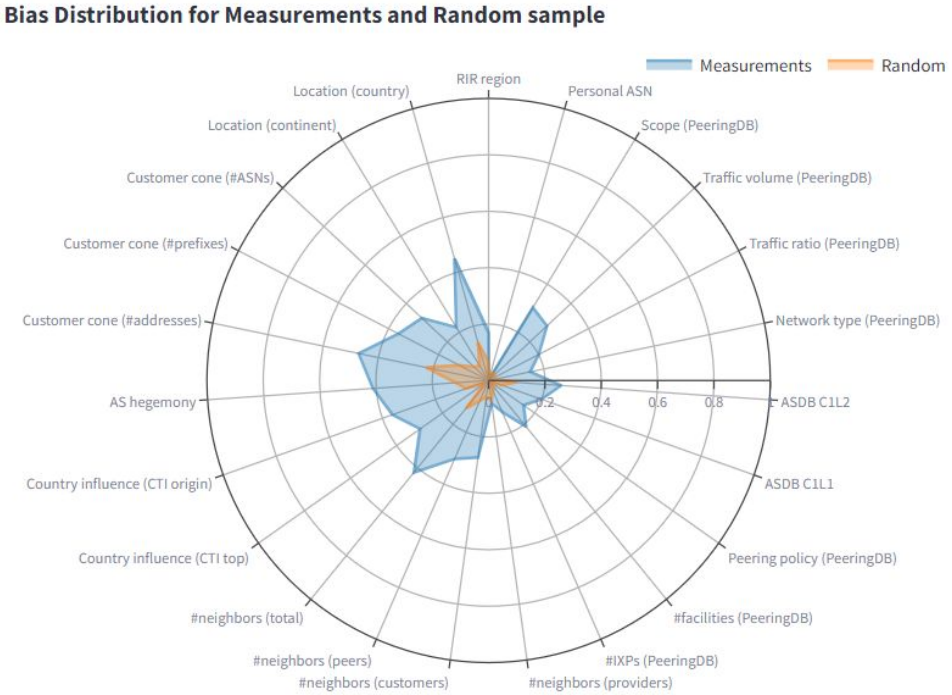

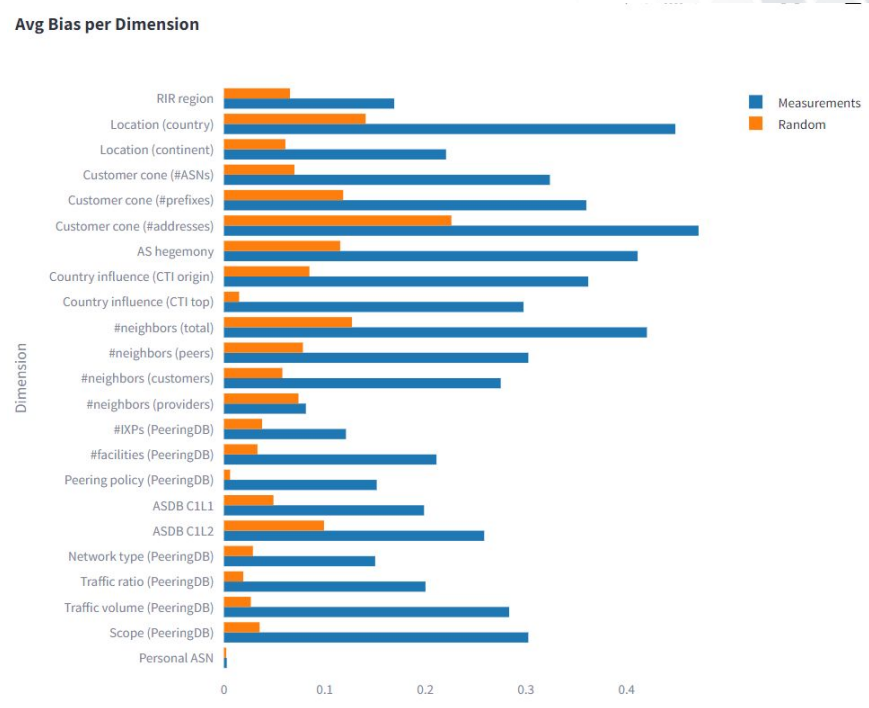

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis
AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

The tool… **in action**

Number of probes vs Avg Bias per measurement

👉 Below you can see a scatter plot, showing you the average bias per measurement vs its number of probes. Each dot in the plot represents one of your measurements.

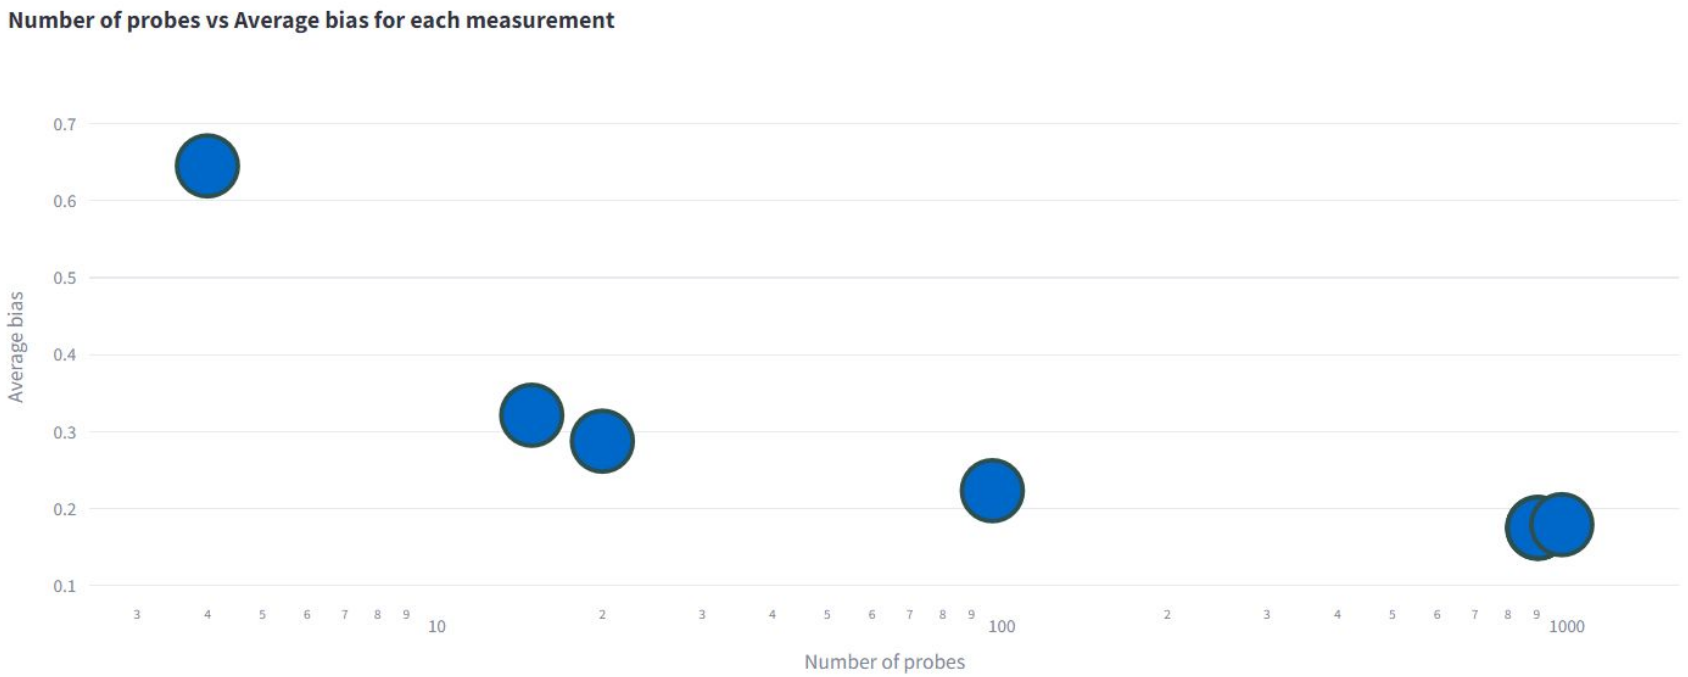

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis
AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

The tool… **in action**

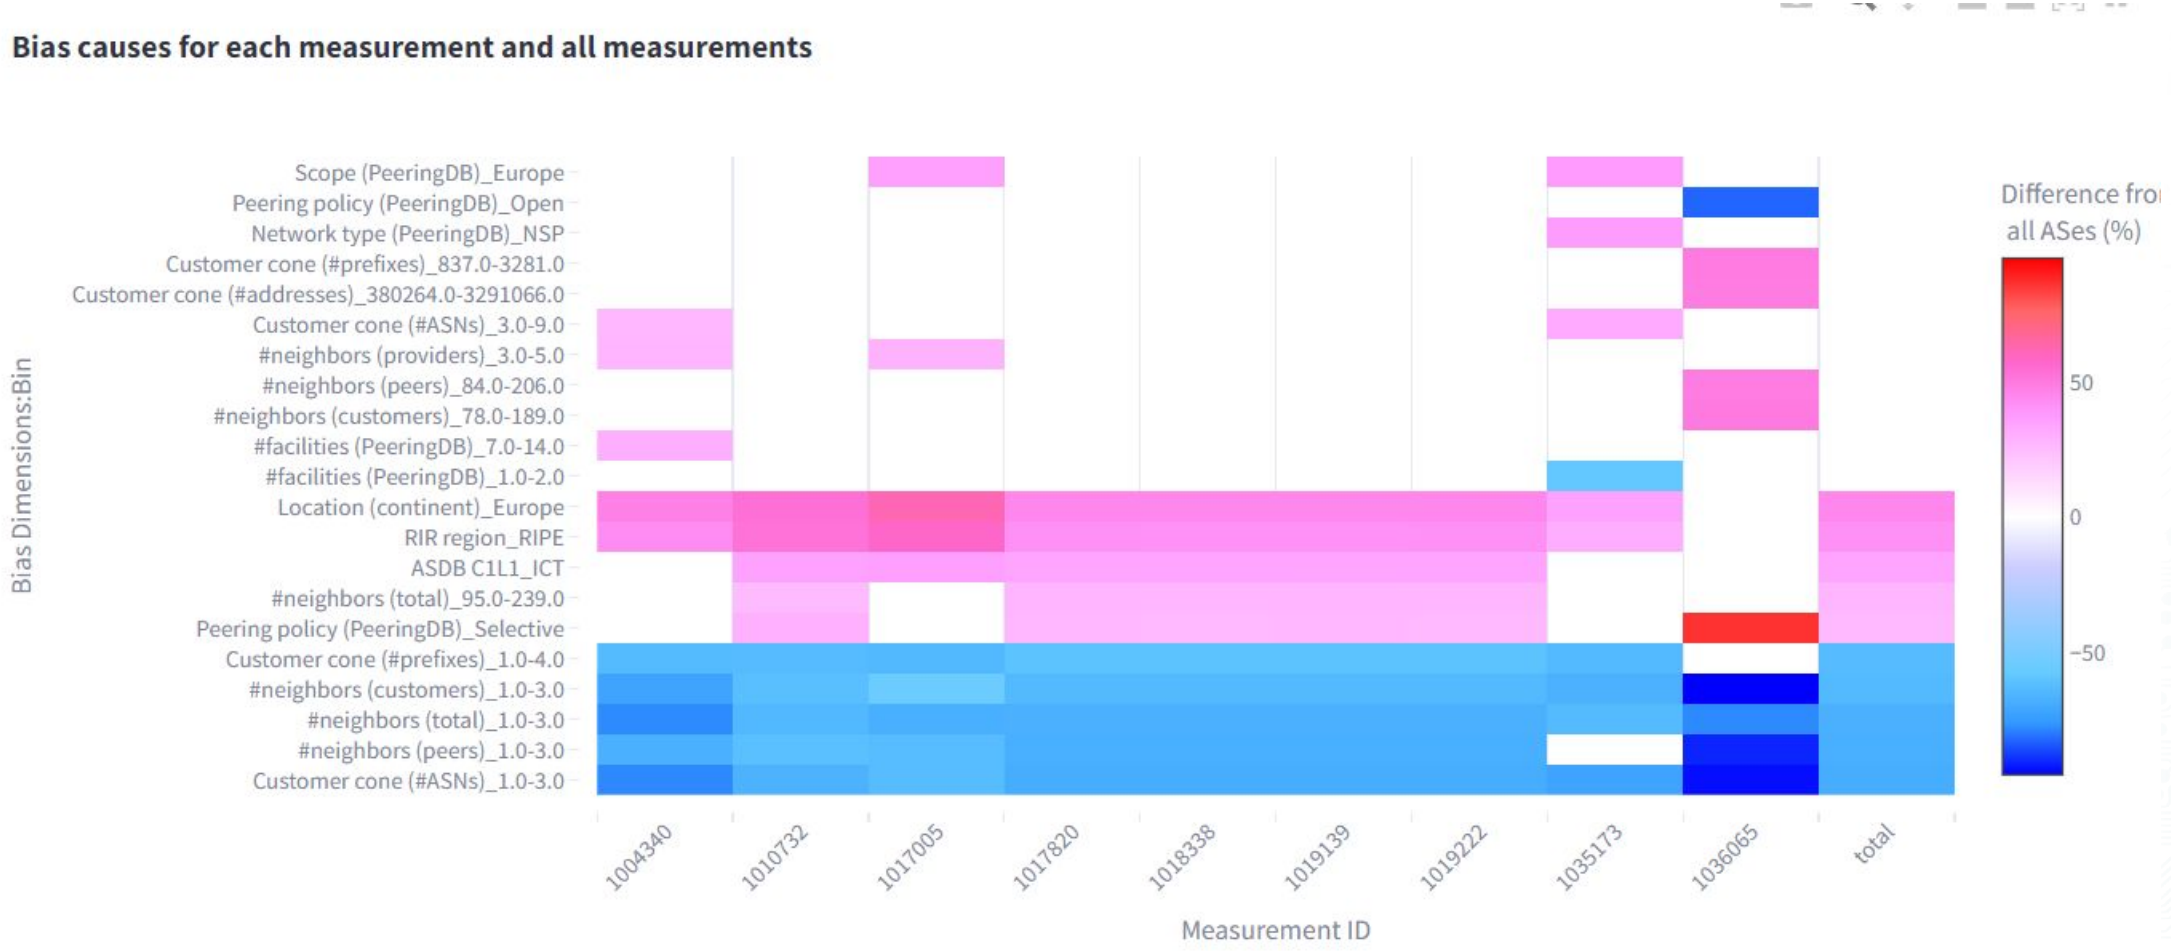

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis

AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

Summarizing…

- Contribution: An online tool to explore characteristics and patterns in RIPE Atlas measurements
  - AI4Netmon website https://ai4netmon.csd.auth.gr/
  - Online tool https://ai4netmon.csd.auth.gr/atlas-app/

- Who & Why to use it? → network operators & researchers
  - easily get deeper view your own measurements (reporting, recommendations, bias insights, etc.)
  - identify patterns in measurements of others (e.g., learn best practices).
  - analyze usage patterns in RIPE Atlas measurements at scale

- We want your feedback & ideas for extra functionality!

Our questions for you … give us your feedback!

- **Reporting** for your measurements?
  - would it be useful for you?
  - what info would you like to see?
  - how to provide results (e.g.,html, pdf, json, csv)?
- **User Interface**
  - what other visualizations to include?
  - feedback on current visualizations
- **Best practices** analysis
  - How to select past measurements to analyze?
    e.g., all measurements from country X to CDN networks?
  - What info / patterns to look at? (e.g., number of probes)

https://forms.gle/aTR8WZh8kYia8zfX6

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis

BACKUP SLIDES

Biases of Internet Measurement Platforms are known

Example 1 - Location bias:
RIPE Atlas has more probes in Europe

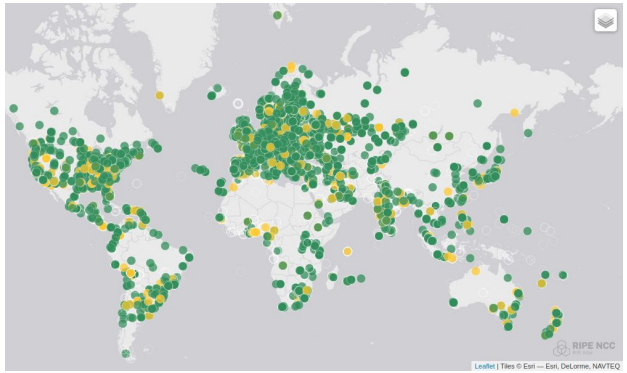

RIPE Atlas probes

https://atlas.ripe.net/results/maps/network-coverage/

Example 2 - Network type bias:
Peers of **RIPE RIS** and **RouteViews** do not equally represent all network types

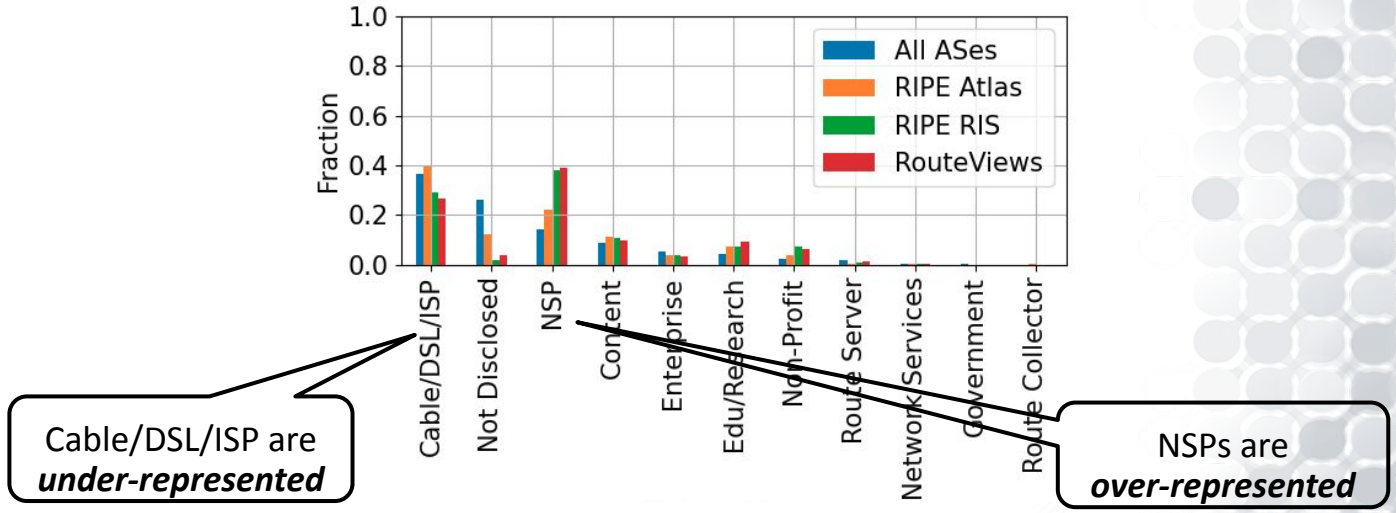

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis
AI4NetMon: The RIPE Atlas Measurements Patterns Analysis tool https://ai4netmon.csd.auth.gr/atlas-app/ - RIPE88

Quantifying bias

- Many dimensions of bias
  - *location, network size, topology, IXP connectivity, network type, etc.*
- Bias score per dimension
  - **Bias score**: a value between 0 (low bias) and 1 (high bias)
- Radar plot of bias
  - each radius ⟶ a bias dimension
  - colored lines/areas ⟶ bias score
  - high bias ⟶ far from center

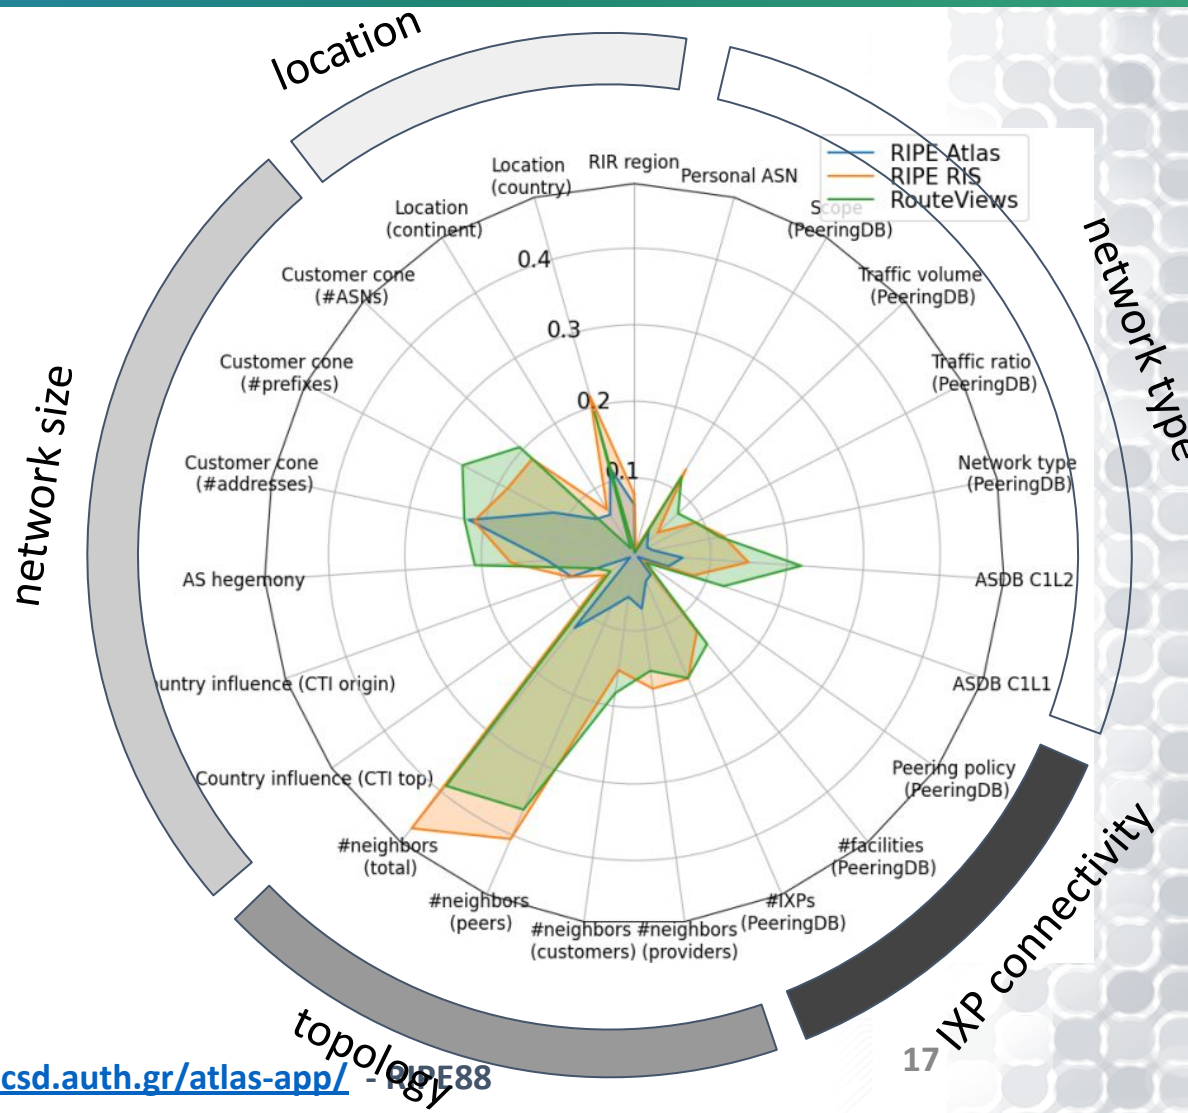

Use case: analyzing random sets of Atlas measurements

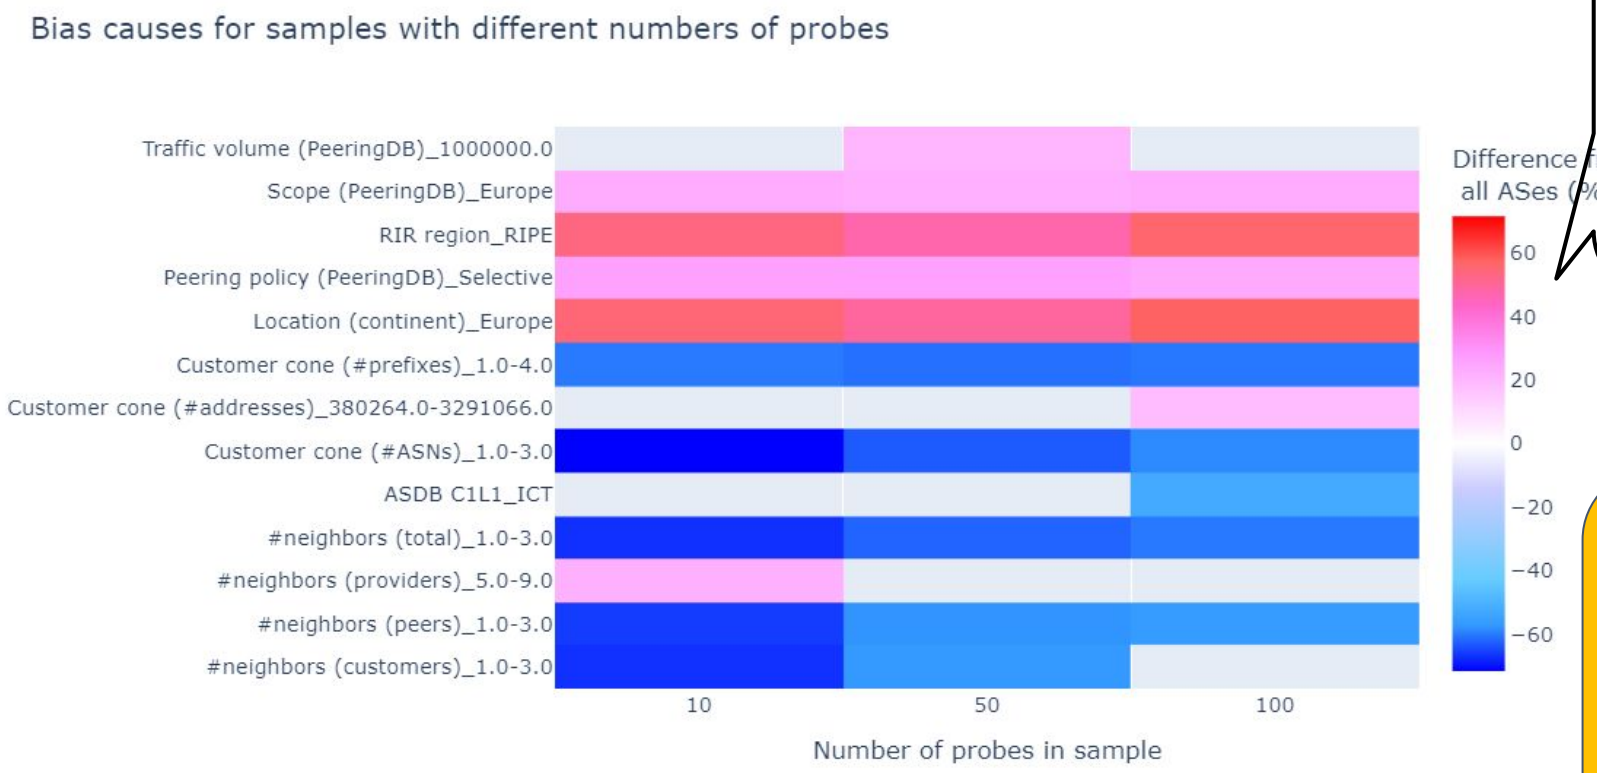

Users tend to *prefer*:

1. Medium/well-connected ASes with mostly providers as neighbours.
2. Medium/larger ASes.
3. European ASes.
4. ASes with Selective Peering Policy

Bias reduction recommendations

- Employ random sampling for the selection of their measurements' probes.
- Use more probes for their measurements.
- Use probes that lie in smaller, less connected ASes with Open Peering Policies.

Pavlos Sermpezis sermpezis@csd.auth.gr @PavlosSermpezis