Selection of Landmarks for Efficient Active Geolocation

Shinyoung Cho, Zachary Weinberg,

Arani Bhattacharya, Sophia Dai, Ramsha Rauf

    

Landmark Selection for Active Geolocation

- **Goal**: Enhance the efficiency of geolocation procedures by reducing the number of required landmarks

Landmark Selection for Active Geolocation

- **Goal**: Enhance the efficiency of geolocation procedures by reducing the number of required landmarks
- **Active geolocation**: Determine the location of an Internet host by measuring packet round-trip times (RTTs)

T Geolocation target

Landmark Selection for Active Geolocation

- **Goal**: Enhance the efficiency of geolocation procedures by reducing the number of required landmarks
- **Active geolocation**: Determine the location of an Internet host by measuring packet round-trip times (RTTs)

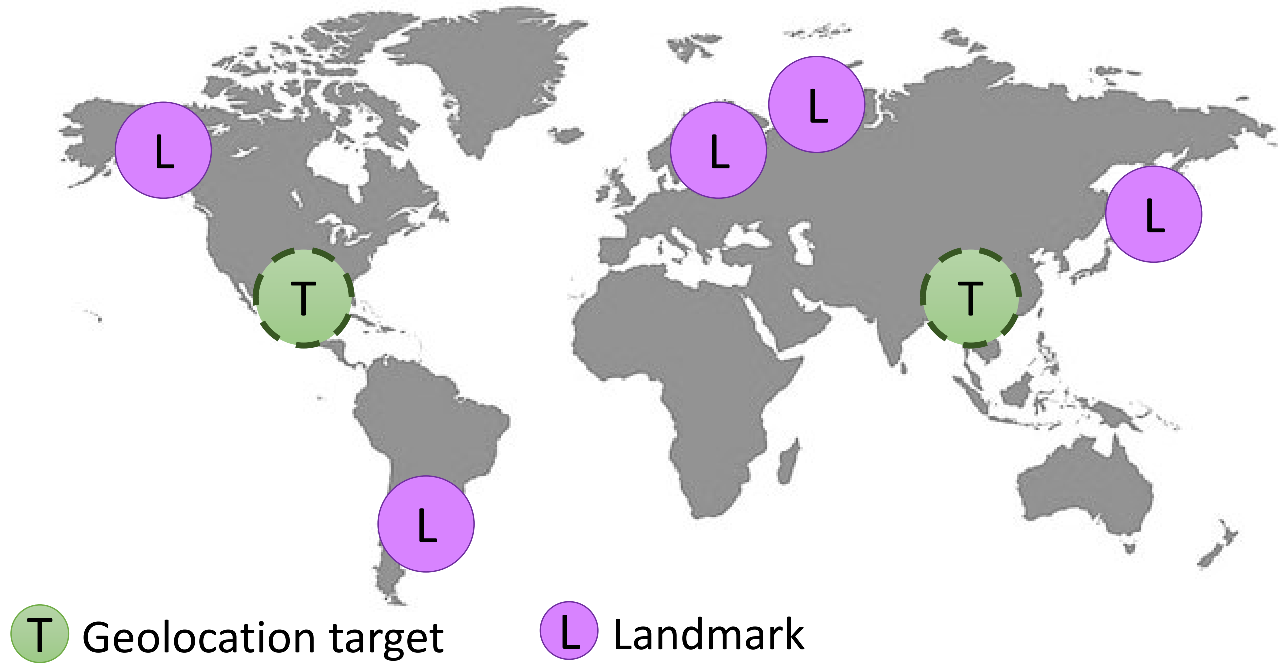

Landmark Selection for Active Geolocation

- **Goal**: Enhance the efficiency of geolocation procedures by reducing the number of required landmarks
- **Active geolocation**: Determine the location of an Internet host by measuring packet round-trip times (RTTs)

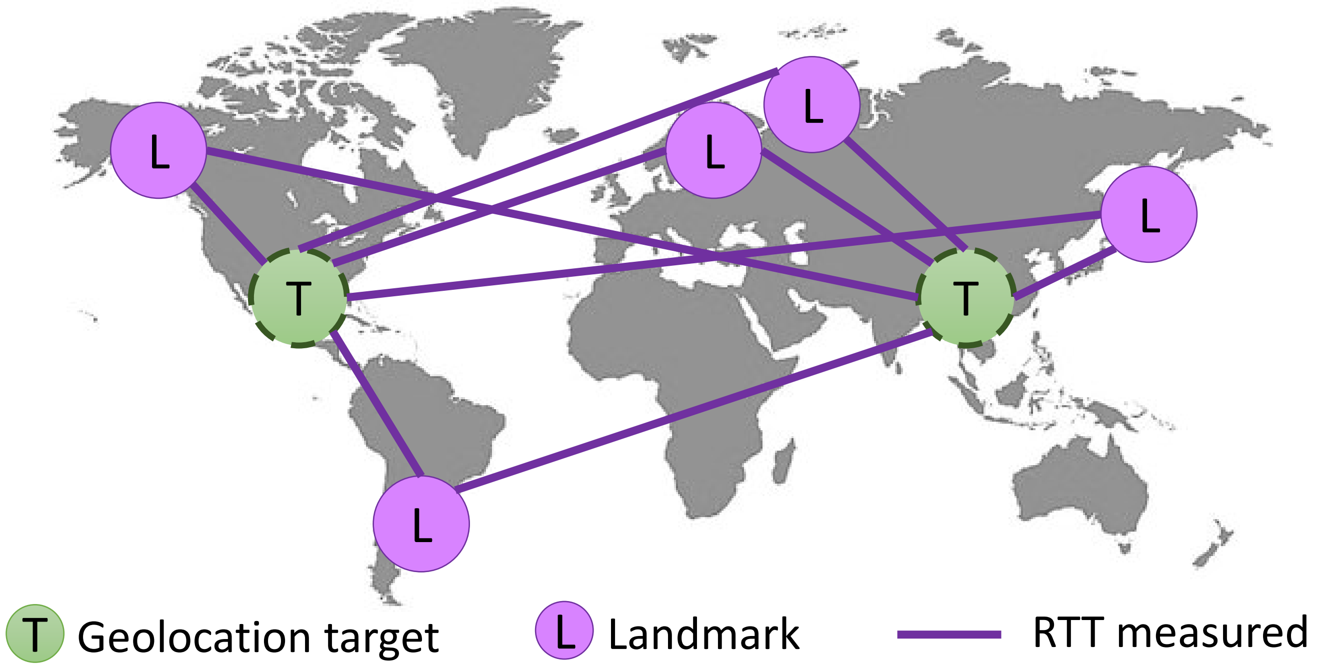

Active Geolocation Can Be Very Slow

- Especially, when **many targets** need to be located

Active Geolocation Can Be Very Slow

- Especially, when **many targets** need to be located
- **Reliable vantage points:** Crucial for accurate global Internet metrics

Active Geolocation Can Be Very Slow

- Especially, when **many targets** need to be located

- **Reliable vantage points:** Crucial for accurate global Internet metrics
- **Frequent moves**: Some vantage points are relocated often [Ramesh2022]
  - Requiring daily location checks

[Ramesh2022] R. Ramesh, L. Evdokimov, D. Xue, and R. Ensafi, "VPNalyzer: Systematic Investigation of the VPN Ecosystem," in *Network and Distributed System Security*, 2022, pp. 24–28.

Active Geolocation Can Be Very Slow

- Especially, when **many targets** need to be located

- **Reliable vantage points:** Crucial for accurate global Internet metrics
- **Frequent moves**: Some vantage points are relocated often [Ramesh2022]
  - Requiring daily location checks
- **Network traffic**: High volumes from "ping" packets could lead to network overload, resembling a DDoS attack [Hu2012]

[Ramesh2022] R. Ramesh, L. Evdokimov, D. Xue, and R. Ensafi, "VPNalyzer: Systematic Investigation of the VPN Ecosystem," in *Network and Distributed System Security*, 2022, pp. 24–28.
[Hu2012] Z. Hu, J. Heidemann, and Y. Pradkin, "Towards geolocation of millions of IP addresses," in Proceedings of the 2012 Internet Measurement Conference, 2012, p. 123–130.

Active Geolocation Can Be Very Slow

- Especially, when **many targets** need to be located

- **Reliable vantage points:** Crucial for accurate global Internet metrics
- **Frequent moves**: Some vantage points are relocated often [Ramesh2022]
  - Requiring daily location checks
- **Network traffic**: High volumes from "ping" packets could lead to network overload, resembling a DDoS attack [Hu2012]

- Necessary to optimize geolocation speed and minimize network impact

[Ramesh2022] R. Ramesh, L. Evdokimov, D. Xue, and R. Ensafi, "VPNalyzer: Systematic Investigation of the VPN Ecosystem," in *Network and Distributed System Security*, 2022, pp. 24–28.
[Hu2012] Z. Hu, J. Heidemann, and Y. Pradkin, "Towards geolocation of millions of IP addresses," in Proceedings of the 2012 Internet Measurement Conference, 2012, p. 123–130.

Our Focus

- **Our research focus**: Determine the smallest effective subset of landmarks for accurate geolocation of many worldwide targets

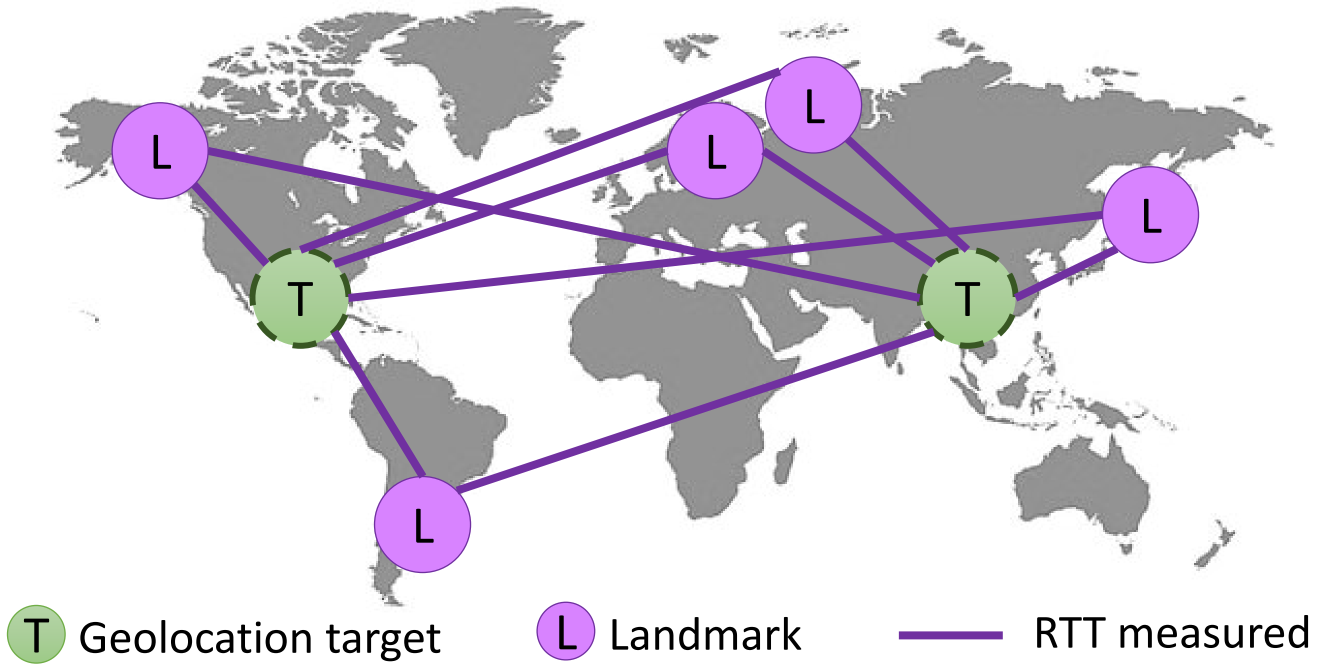

Our Focus

- **Our research focus**: Determine the smallest effective subset of landmarks for accurate geolocation of many worldwide targets

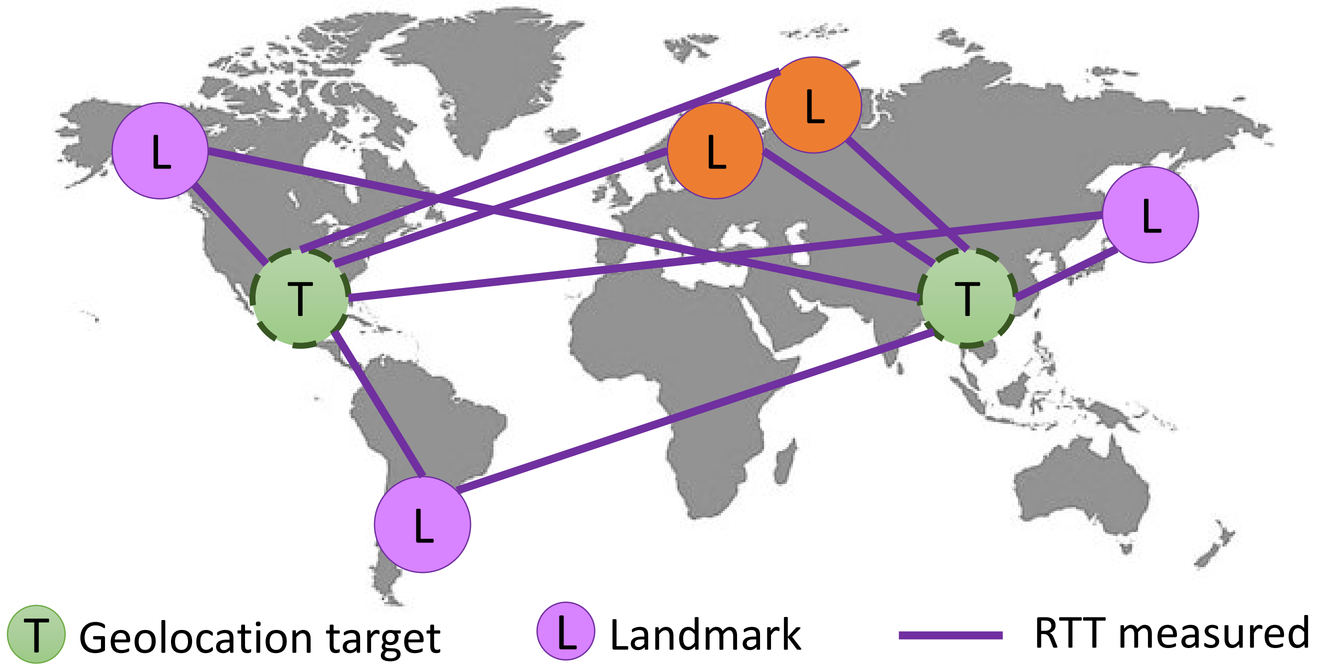

Our Focus

- **Our research focus**: Determine the smallest effective subset of landmarks for accurate geolocation of many worldwide targets

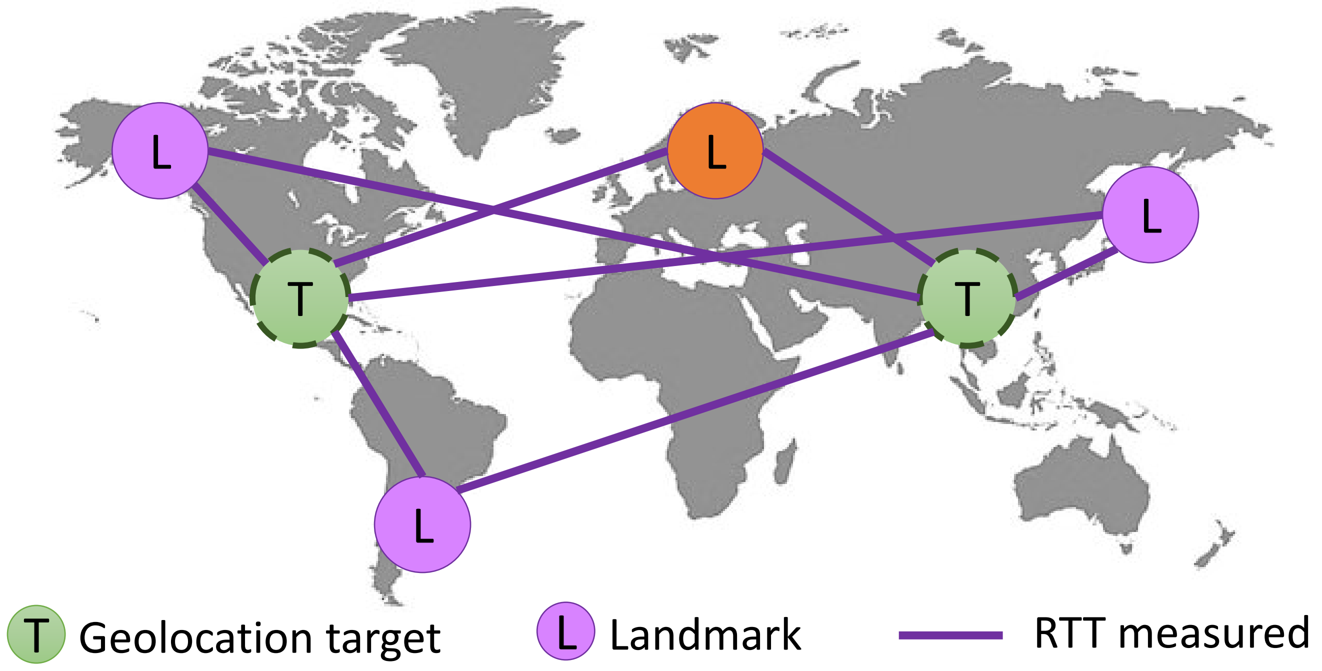

Our Focus

- **Our research focus**: Determine the smallest effective subset of landmarks for accurate geolocation of many worldwide targets
- **Prior work:** shows that active geolocation should use only landmarks near the target [Darwich2023]

[Darwich2023] O. Darwich, H. Rimlinger, M. Dreyfus, M. Gouel, and K. Vermeulen, "Replication: Towards a publicly available internet scale ip geolocation dataset," in Internet Measurement Conference, 2023, pp. 1–15

Our Focus

- **Our research focus**: Determine the smallest effective subset of landmarks for accurate geolocation of many worldwide targets
- **Prior work:** shows that active geolocation should use only landmarks near the target [Darwich2023]
- **Challenge in geographic uncertainty**: selecting nearby landmarks when the target's global location is unknown is challenging

[Darwich2023] O. Darwich, H. Rimlinger, M. Dreyfus, M. Gouel, and K. Vermeulen, "Replication: Towards a publicly available internet scale ip geolocation dataset," in Internet Measurement Conference, 2023, pp. 1–15

Our Focus

- **Our research focus**: Determine the smallest effective subset of landmarks for accurate geolocation of many worldwide targets
- **Prior work:** shows that active geolocation should use only landmarks near the target [Darwich2023]
- **Challenge in geographic uncertainty:** selecting nearby landmarks when the target's global location is unknown is challenging
- **Algorithm evaluation:** Assess various algorithms to select an optimal subset of landmarks from a larger pool

[Darwich2023] O. Darwich, H. Rimlinger, M. Dreyfus, M. Gouel, and K. Vermeulen, "Replication: Towards a publicly available internet scale ip geolocation dataset," in Internet Measurement Conference, 2023, pp. 1–15

Experimental Setup

- Our landmarks
- Our targets
- Active geolocation algorithm

Our Landmarks

- **Landmark requirements:** Large set of hosts, known locations, worldwide distribution, reliable and always available for pinging

Our Landmarks

- **Landmark requirements:** Large set of hosts, known locations, worldwide distribution, reliable and always available for pinging
- **Used 780 RIPE Atlas anchors:**
  - More stable and reliably operational than RIPE Atlas probes

Our Landmarks

- **Landmark requirements:** Large set of hosts, known locations, worldwide distribution, reliable and always available for pinging
- **Used 780 RIPE Atlas anchors:**
  - More stable and reliably operational than RIPE Atlas probes

RIPE NCC
RIPE Atlas

Anchors

Probes

Our Landmarks

- **Landmark requirements:** Large set of hosts, known locations, worldwide distribution, reliable and always available for pinging
- **Used 780 RIPE Atlas anchors:**
    - More stable and reliably operational than RIPE Atlas probes

Continent	# of countries	cities	landmarks
Asia	31	71	122
Europe	36	270	438
South America	8	20	28
Oceania	3	11	25
Africa	9	14	18
North America	9	95	149

Our Targets

- Utilized 559 commercial VPN endpoints

Our Targets

- Utilized 559 commercial VPN endpoints
  - Advantages of VPNs: Convenient for global network measurements

Our Targets

- Utilized 559 commercial VPN endpoints
  - Advantages of VPNs: Convenient for global network measurements
  - Caveats: VPNs may exaggerate coverage [Weinberg2018]

[Weinberg2018] Z. Weinberg, S. Cho, N. Christin, V. Sekar, and P. Gill, "How to Catch When Proxies Lie: Verifying the Physical Locations of Network Proxies with Active Geolocation," in *Internet Measurement Conference*, 2018, pp. 203–217.

Our Targets

- Utilized 559 commercial VPN endpoints
  - Advantages of VPNs: Convenient for global network measurements
  - Caveats: VPNs may exaggerate coverage [Weinberg2018]
  - Verification Needed: Confirm VPN location claims before research use

[Weinberg2018] Z. Weinberg, S. Cho, N. Christin, V. Sekar, and P. Gill, "How to Catch When Proxies Lie: Verifying the Physical Locations of Network Proxies with Active Geolocation," in *Internet Measurement Conference*, 2018, pp. 203–217.

Our Targets

- Utilized 559 commercial VPN endpoints
  - Advantages of VPNs: Convenient for global network measurements
  - Caveats: VPNs may exaggerate coverage [Weinberg2018]
  - Verification Needed: Confirm VPN location claims before research use

Continent	# of countries	targets
Asia	51	110
Europe	47	120
South America	13	27
Oceania	20	41
Africa	50	103
North America	34	176

[Weinberg2018] Z. Weinberg, S. Cho, N. Christin, V. Sekar, and P. Gill, "How to Catch When Proxies Lie: Verifying the Physical Locations of Network Proxies with Active Geolocation," in *Internet Measurement Conference*, 2018, pp. 203–217.

Active Geolocation Algorithm

- Determines if a target's claimed country is accurate

Active Geolocation Algorithm

- Determines if a target's claimed country is accurate
- **Method**: Assumes packets travel the shortest great-circle route to the nearest border of the claimed country, converting minimum RTT into travel speed

Active Geolocation Algorithm

- Determines if a target’s claimed country is accurate
- **Method**: Assumes packets travel the shortest great-circle route to the nearest border of the claimed country, converting minimum RTT into travel speed
- **Location verification:** If the calculated speed exceeds a predefined limit, the location claim is rejected

Active Geolocation Algorithm

- Determines if a target's claimed country is accurate
- **Method**: Assumes packets travel the shortest great-circle route to the nearest border of the claimed country, converting minimum RTT into travel speed
- **Location verification:** If the calculated speed exceeds a predefined limit, the location claim is rejected

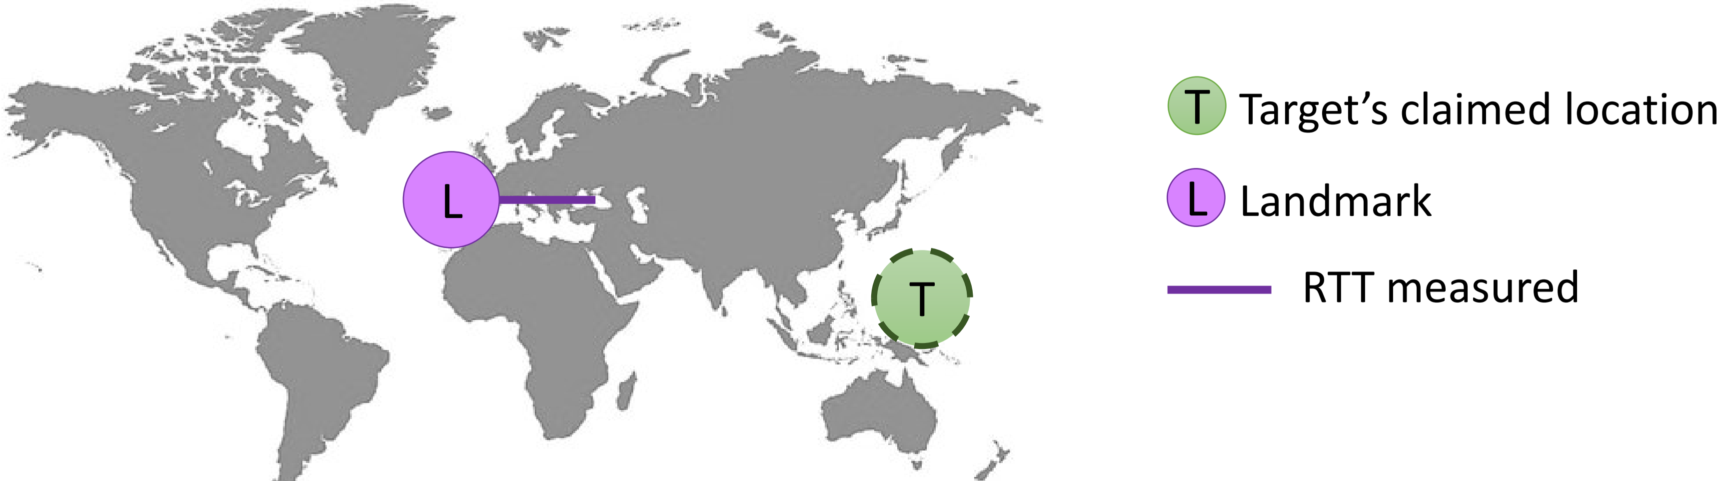

Active Geolocation Algorithm

- Determines if a target's claimed country is accurate
- **Method**: Assumes packets travel the shortest great-circle route to the nearest border of the claimed country, converting minimum RTT into travel speed
- **Location verification:** If the calculated speed exceeds a predefined limit, the location claim is rejected

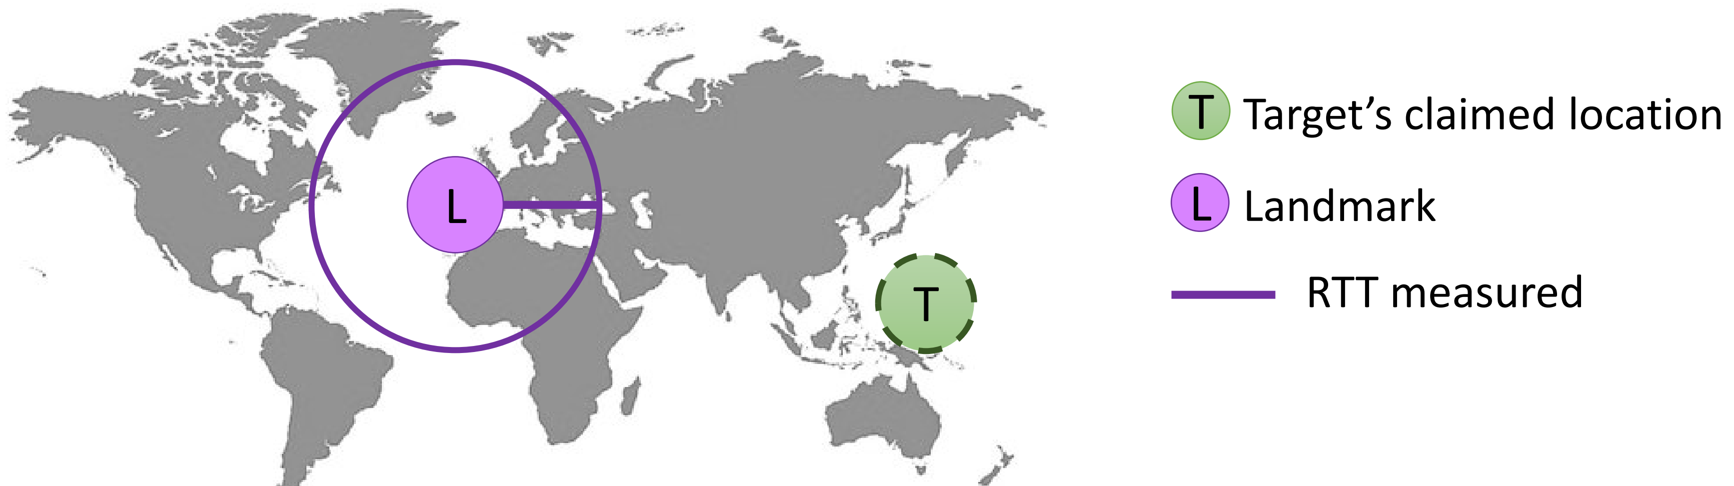

Overview

- Motivation and Research Focus
- Our Experimental Setup
- Our Landmark Selections
  - LS1: Random Selection
  - LS2: Clustering Selection
  - LS3: Greatest-Distance Selection
  - LS4: Hybrid Selection

Evaluation

We evaluated several algorithms by comparing their performance to the performance of the full set of landmarks

Evaluation

We evaluated several algorithms by comparing their performance to the performance of the full set of landmarks

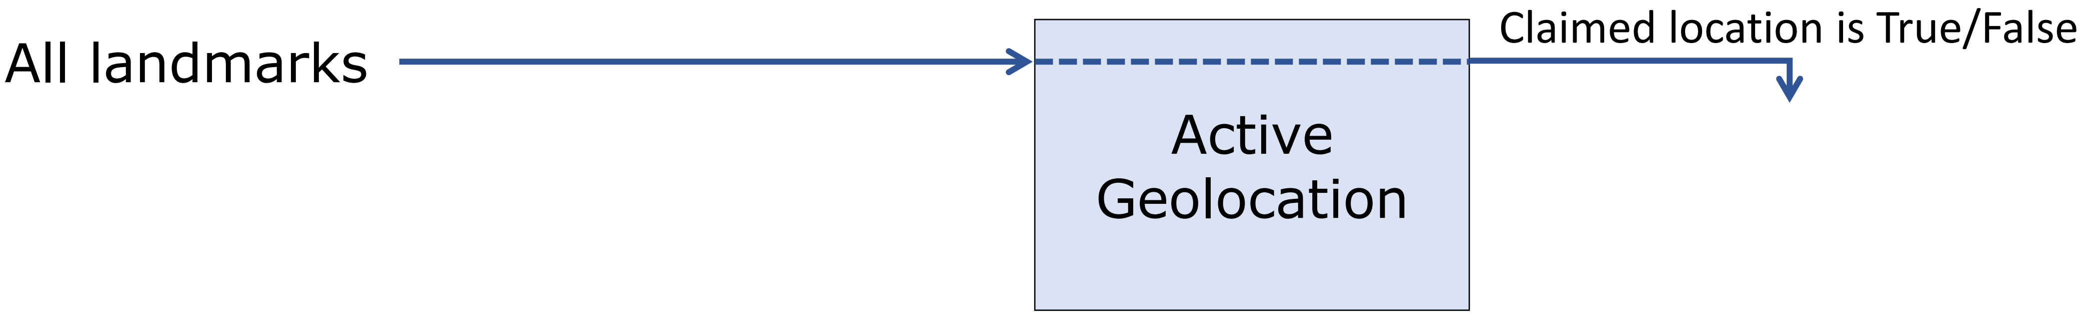

Evaluation

We evaluated several algorithms by comparing their performance to the performance of the full set of landmarks

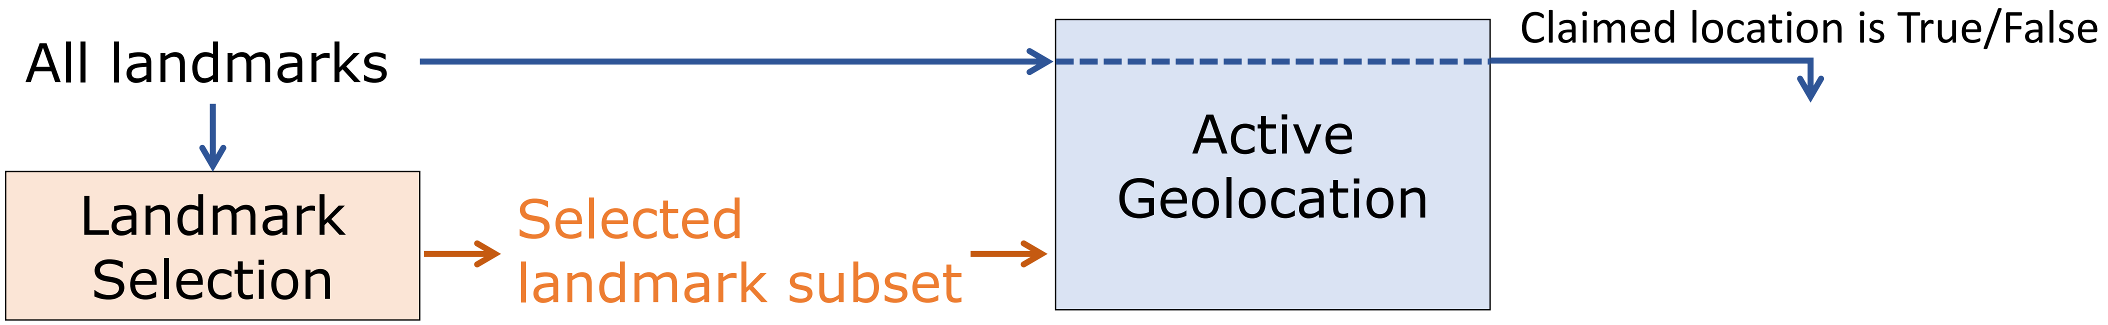

Evaluation

We evaluated several algorithms by comparing their performance to the performance of the full set of landmarks

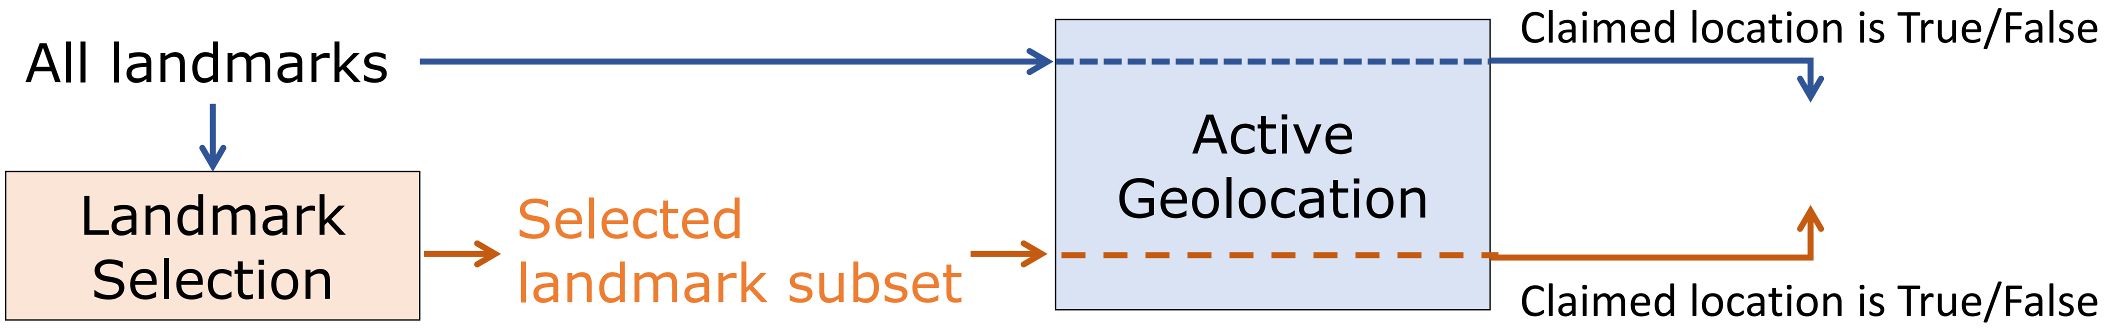

Evaluation

We evaluated several algorithms by comparing their performance to the performance of the full set of landmarks

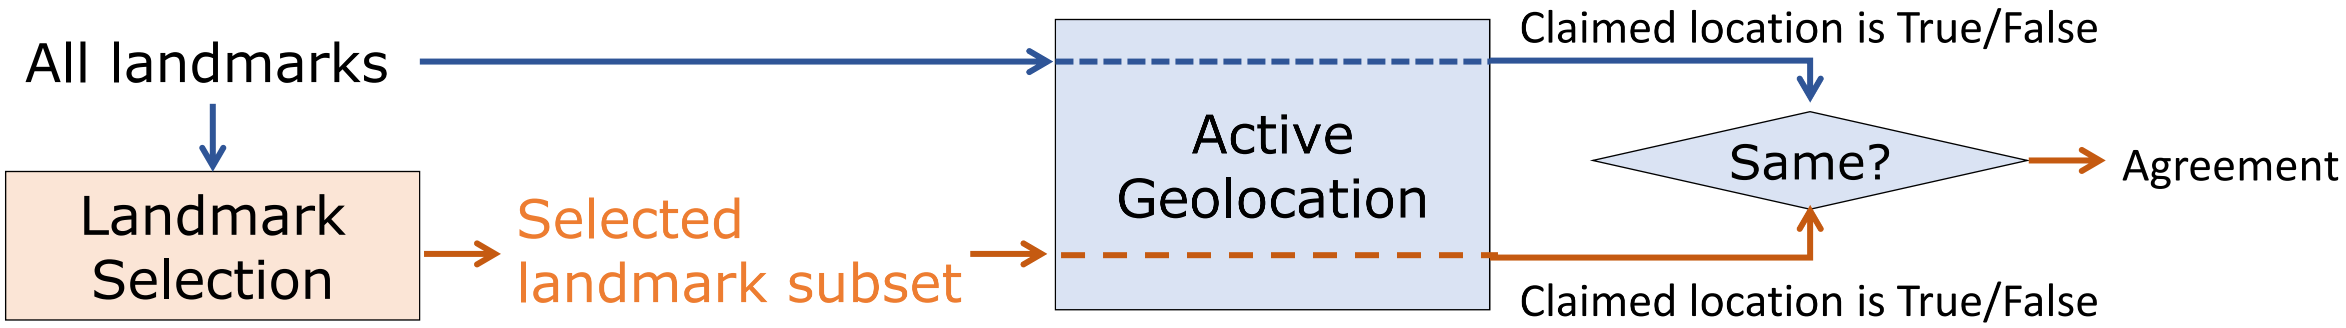

LS1: Random Selection

LS1: Random Selection

- In large scale-free graphs, small random samples often represent the complete graph as well or better than structured samples [Leskovec06]

[Leskovec06] J. Leskovec and C. Faloutsos, “Sampling from Large Graphs,” in *Knowledge Discovery and Data Mining*, 2006, pp. 631–636. DOI:10.1145/ 1150402.1150479

LS1: Random Selection

- In large scale-free graphs, small random samples often represent the complete graph as well or better than structured samples [Leskovec06]
- **Expectation**: A small random landmark subset can estimate the number of landmarks needed for accurate geolocation

[Leskovec06] J. Leskovec and C. Faloutsos, "Sampling from Large Graphs," in *Knowledge Discovery and Data Mining*, 2006, pp. 631–636. DOI:10.1145/ 1150402.1150479

LS1: Random Selection

- In large scale-free graphs, small random samples often represent the complete graph as well or better than structured samples [Leskovec06]
- **Expectation**: A small random landmark subset can estimate the number of landmarks needed for accurate geolocation

All landmarks

Select subsets of the landmark pool **uniformly at random** (without replacement); 1,000 subsets for each size (1 to 780)

Selected landmark subsets

[Leskovec06] J. Leskovec and C. Faloutsos, "Sampling from Large Graphs," in *Knowledge Discovery and Data Mining*, 2006, pp. 631–636. DOI:10.1145/ 1150402.1150479

LS1: Random Selection Result

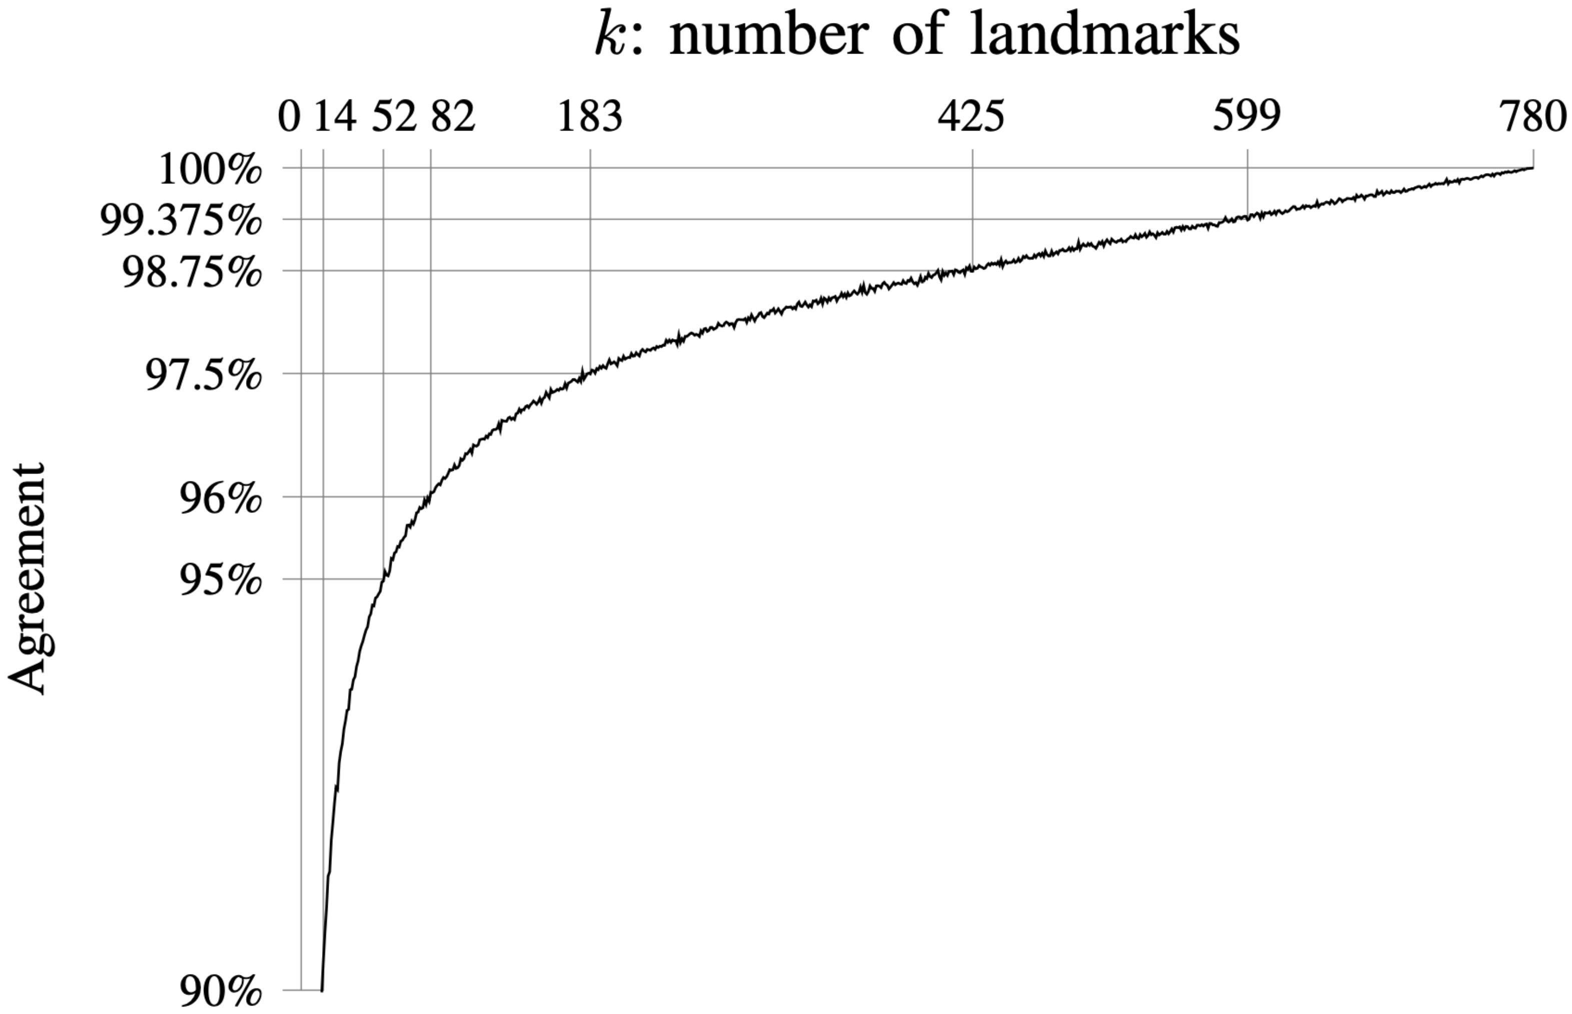

LS1: Random Selection Result

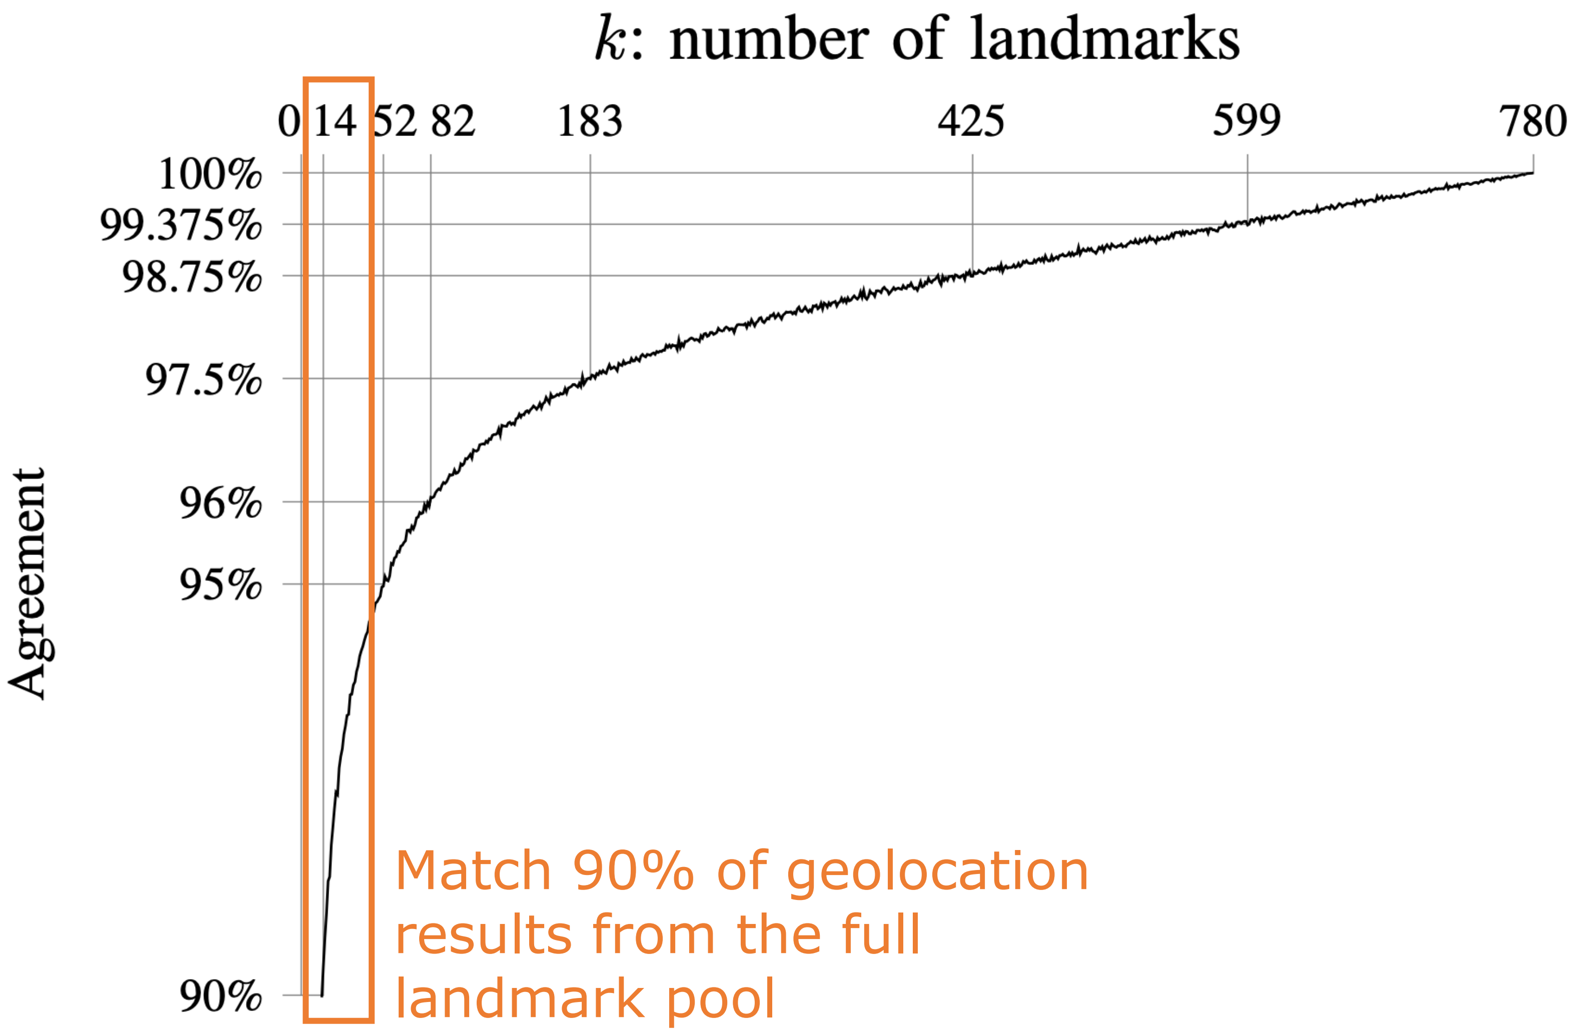

LS1: Random Selection Result

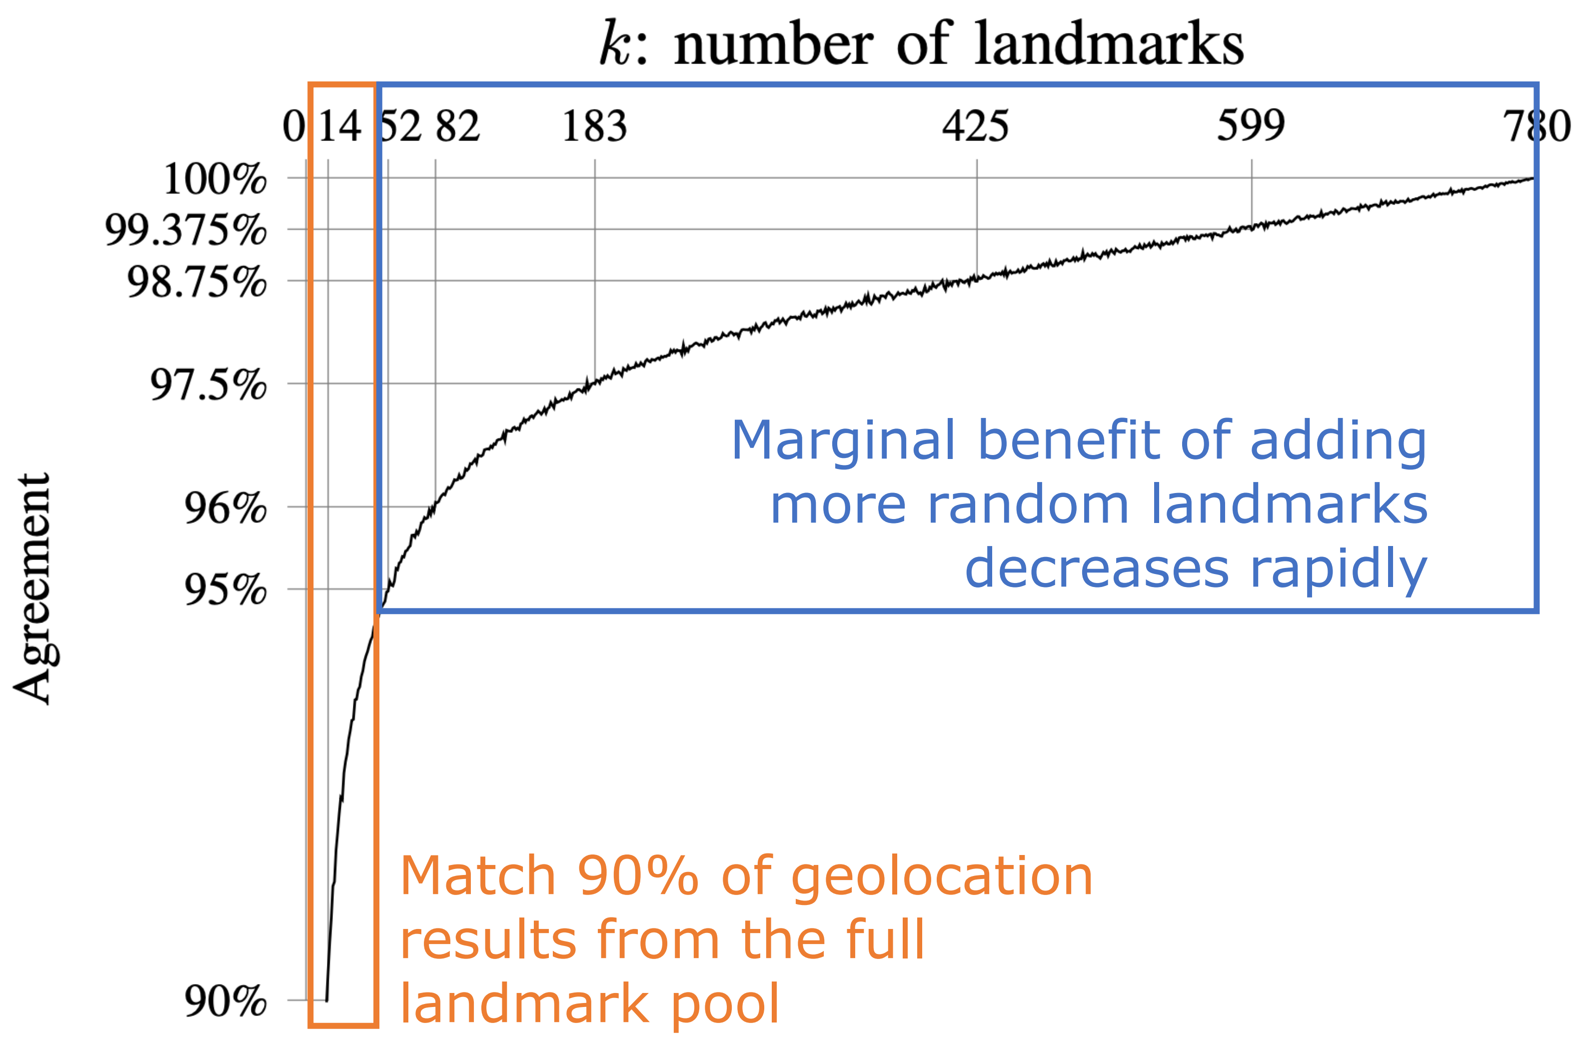

LS1: Random Selection Result

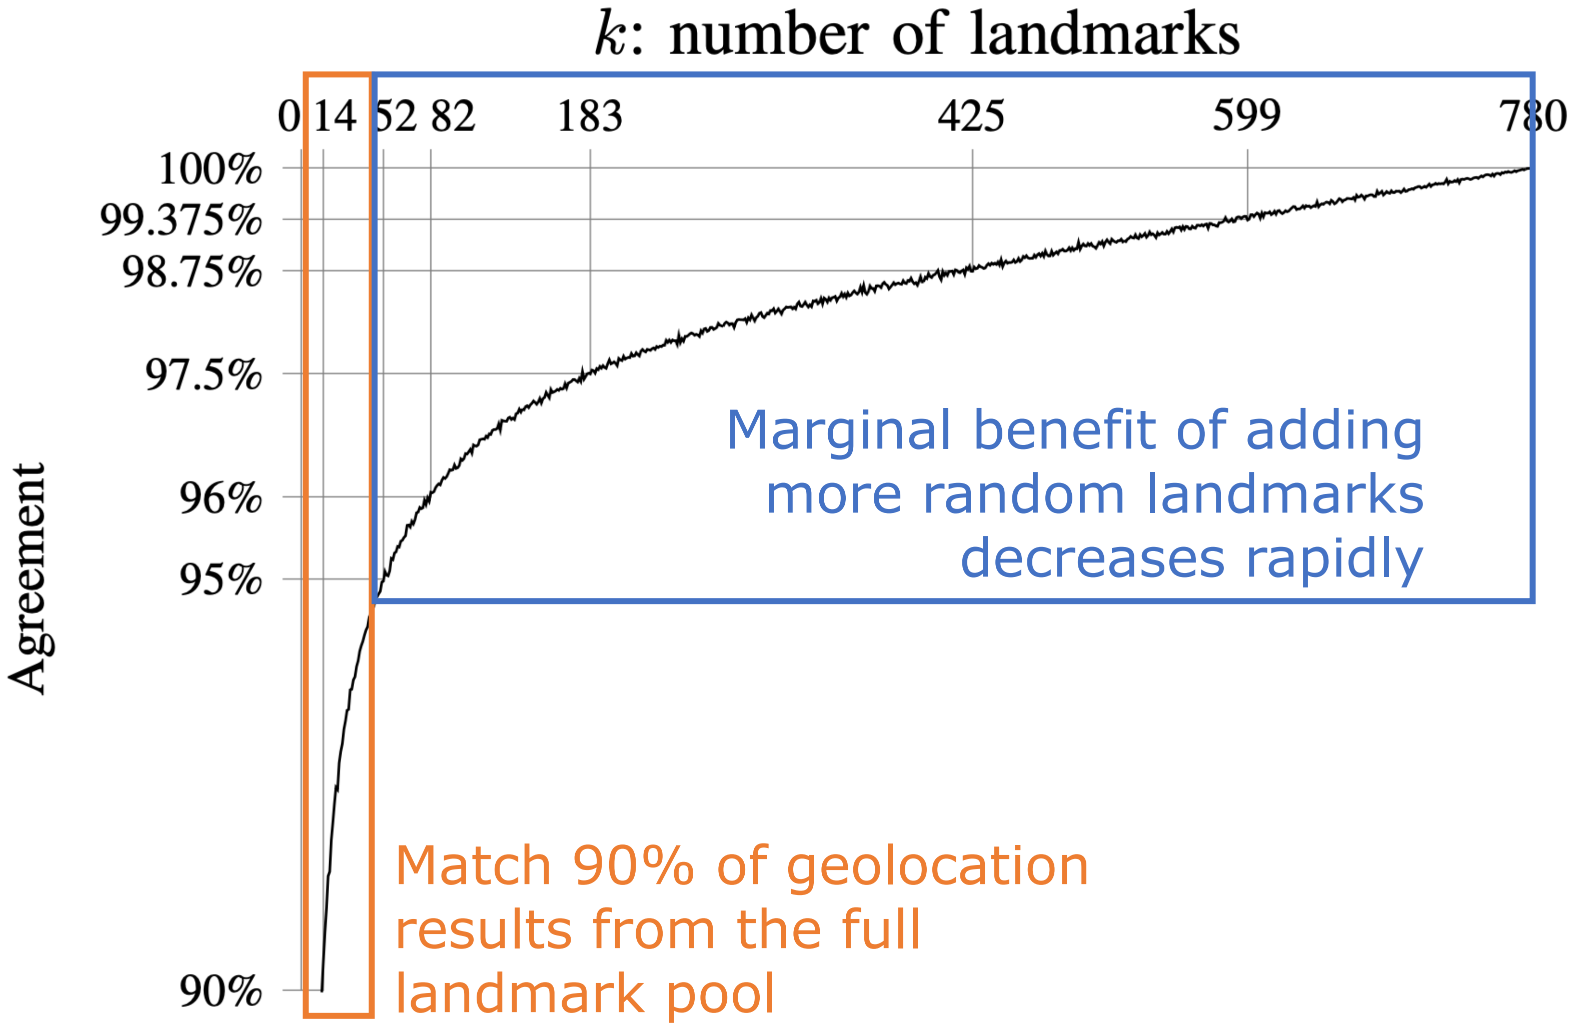

Full Agreement: Reached only when all landmarks are in use

Overview

- Motivation and Research Focus
- Our Experimental Setup
- Our Landmark Selections
    - LS1: Random Selection
    - LS2: Clustering Selection
    - LS3: Greatest-Distance Selection
    - LS4: Hybrid Selection

Overview

- Motivation and Research Focus
- Our Experimental Setup
- Our Landmark Selections
  - LS1: Random Selection
  - LS2: Clustering Selection
  - LS3: Greatest-Distance Selection
  - LS4: Hybrid Selection

Overview

- Motivation and Research Focus
- Our Experimental Setup
- Our Landmark Selections
  - LS1: Random Selection
  - LS2: Clustering Selection
  - LS3: Greatest-Distance Selection
  - LS4: Hybrid Selection

Diversity Metrics

Optimal Selection of a landmark that **maximize diversity metrics**

LS2: Clustering Selection

- Four types of cluster

Type	# clusters	Mean agreement vs. random
ASes	534	$99.28\% > 99.20\%$
Cities	481	$99.62 > 98.99$
Countries	96	$93.88 \ll 96.32$
Continents	6	$83.08 < 84.05$

LS2: Clustering Selection

- Four types of cluster

	Type	# clusters	Mean agreement vs. random
Internet topology	**ASes**	534	99.28% > 99.20%
Geographical distribution	**Cities**	481	99.62 > 98.99
	Countries	96	93.88 ≪ 96.32
	Continents	6	83.08 < 84.05

LS2: Clustering Selection

- Four types of cluster

Compare subset sizes (=number of clusters) to randomly selected subsets of the same size

	Type	# clusters	Mean agreement vs. random
Internet topology	ASes	534	99.28% $>$ 99.20%
Geographical distribution	Cities	481	99.62 $>$ 98.99
	Countries	96	93.88 $\ll$ 96.32
	Continents	6	83.08 $<$ 84.05

LS2: Clustering Selection

- Four types of cluster

Compare subset sizes (=number of clusters) to randomly selected subsets of the same size

	Type	# clusters	Mean agreement vs. random
Internet topology	ASes	534	$99.28\% > 99.20\%$
Geographical distribution	Cities	481	$99.62 > 98.99$
	Countries	96	$93.88 \ll 96.32$
	Continents	6	$83.08 < 84.05$

LS2: Clustering Selection

- Four types of cluster

Compare subset sizes (=number of clusters) to randomly selected subsets of the same size

	Type	# clusters	Mean agreement vs. random	
Internet topology	ASes	534	99.28% $>$ 99.20%	The need for fine-grained diversity
Geographical distribution	Cities	481	99.62 $>$ 98.99	
	Countries	96	93.88 $\ll$ 96.32	
	Continents	6	83.08 $<$ 84.05	

LS2: Clustering Selection Result

- **Extended to all sizes:** Landmarks are randomly selected, aiming for equal contribution from each cluster where possible

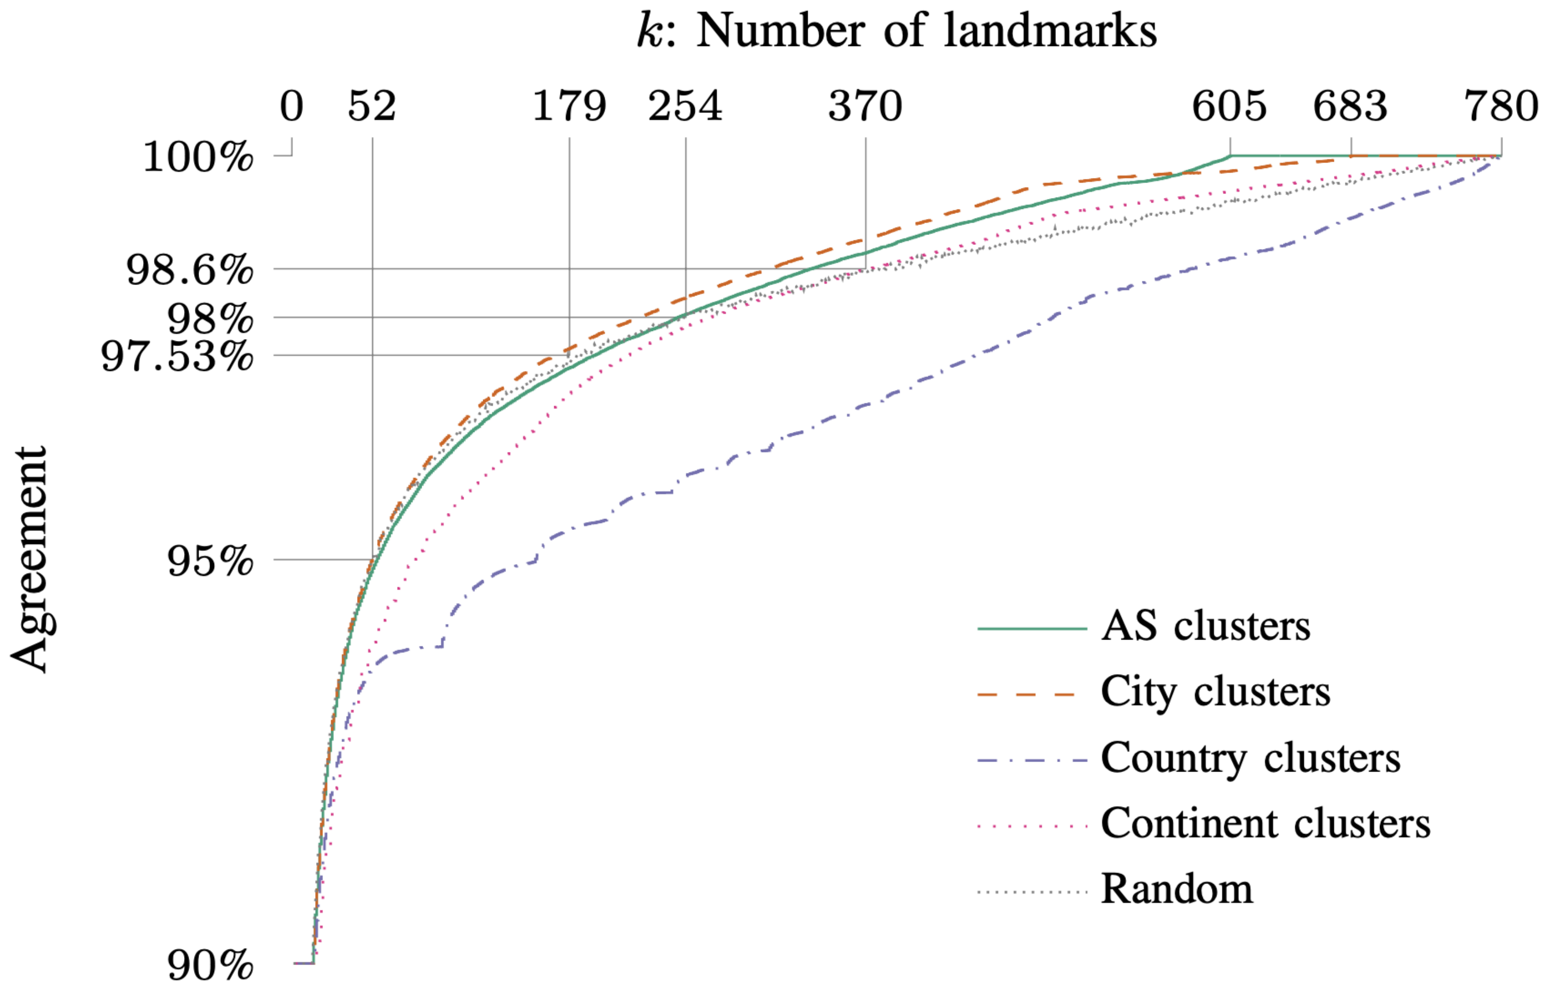

LS2: Clustering Selection Result

- **Extended to all sizes:** Landmarks are randomly selected, aiming for equal contribution from each cluster where possible

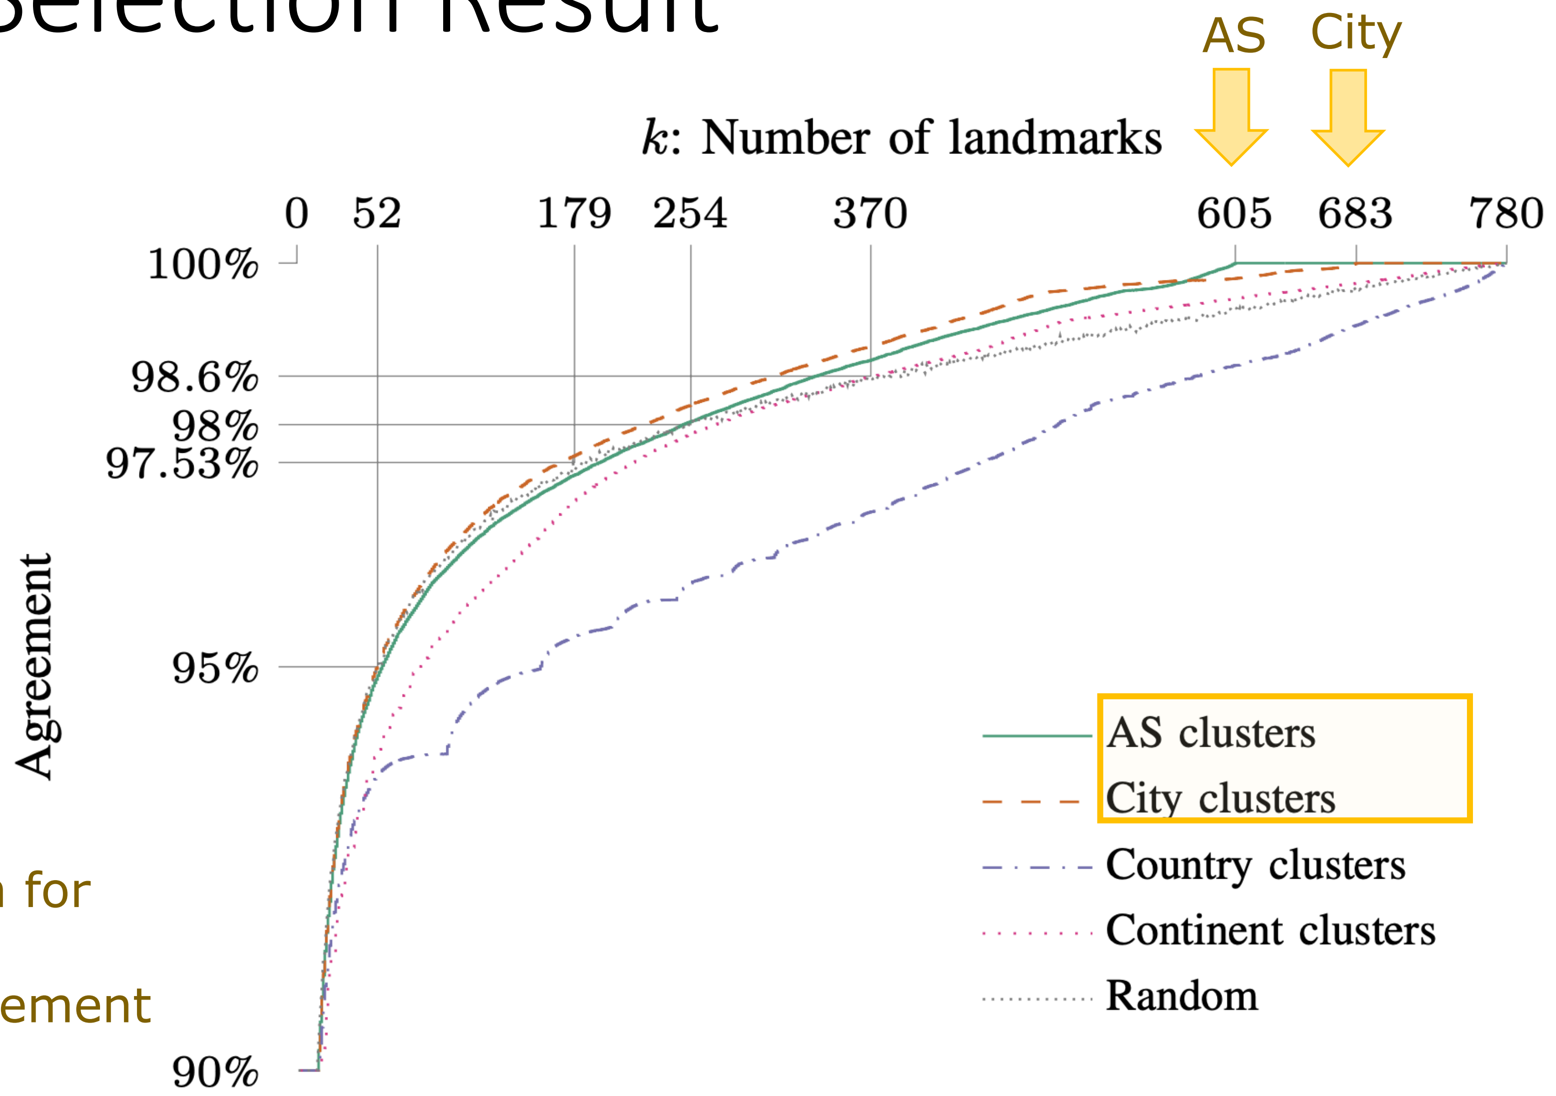

City and AS Clusters:

- Outperform random selection for most sizes
- Achieve perfect (100%) agreement without full pool use

LS2: Clustering Selection Result

- **Extended to all sizes:** Landmarks are randomly selected, aiming for equal contribution from each cluster where possible

City and AS Clusters:

- Outperform random selection for most sizes
- Achieve perfect (100%) agreement without full pool use

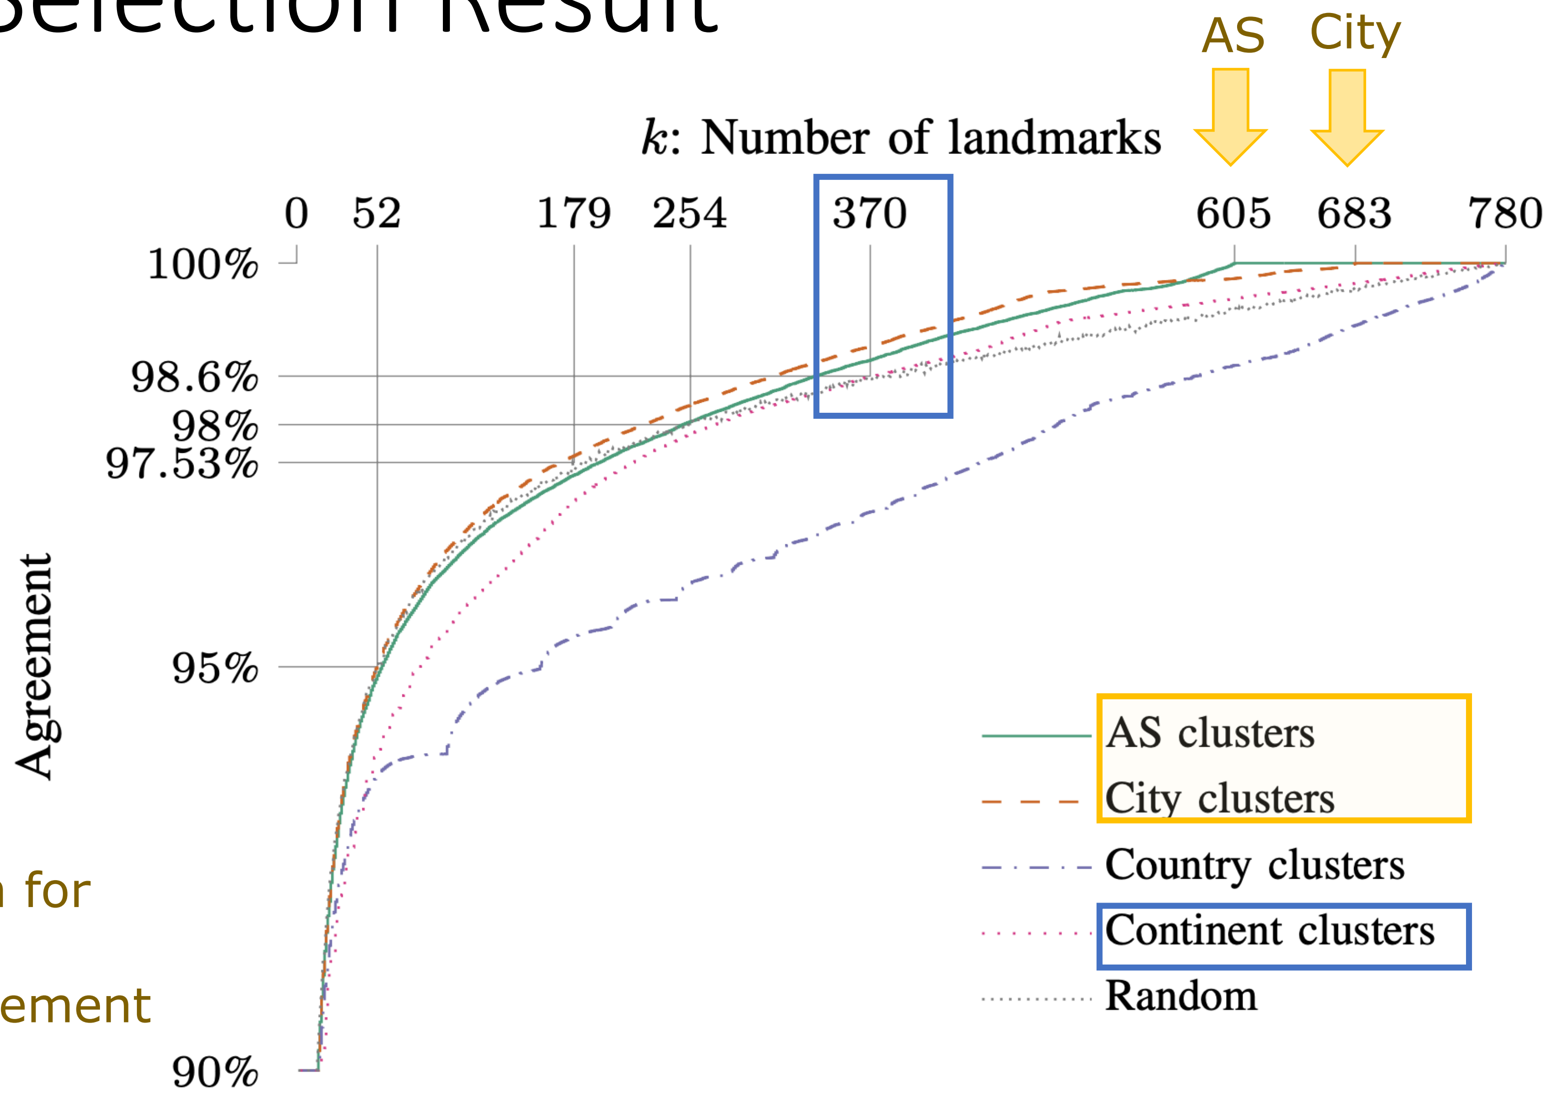

LS2: Clustering Selection Result

- **Extended to all sizes:** Landmarks are randomly selected, aiming for equal contribution from each cluster where possible

City and AS Clusters:

- Outperform random selection for most sizes
- Achieve perfect (100%) agreement without full pool use

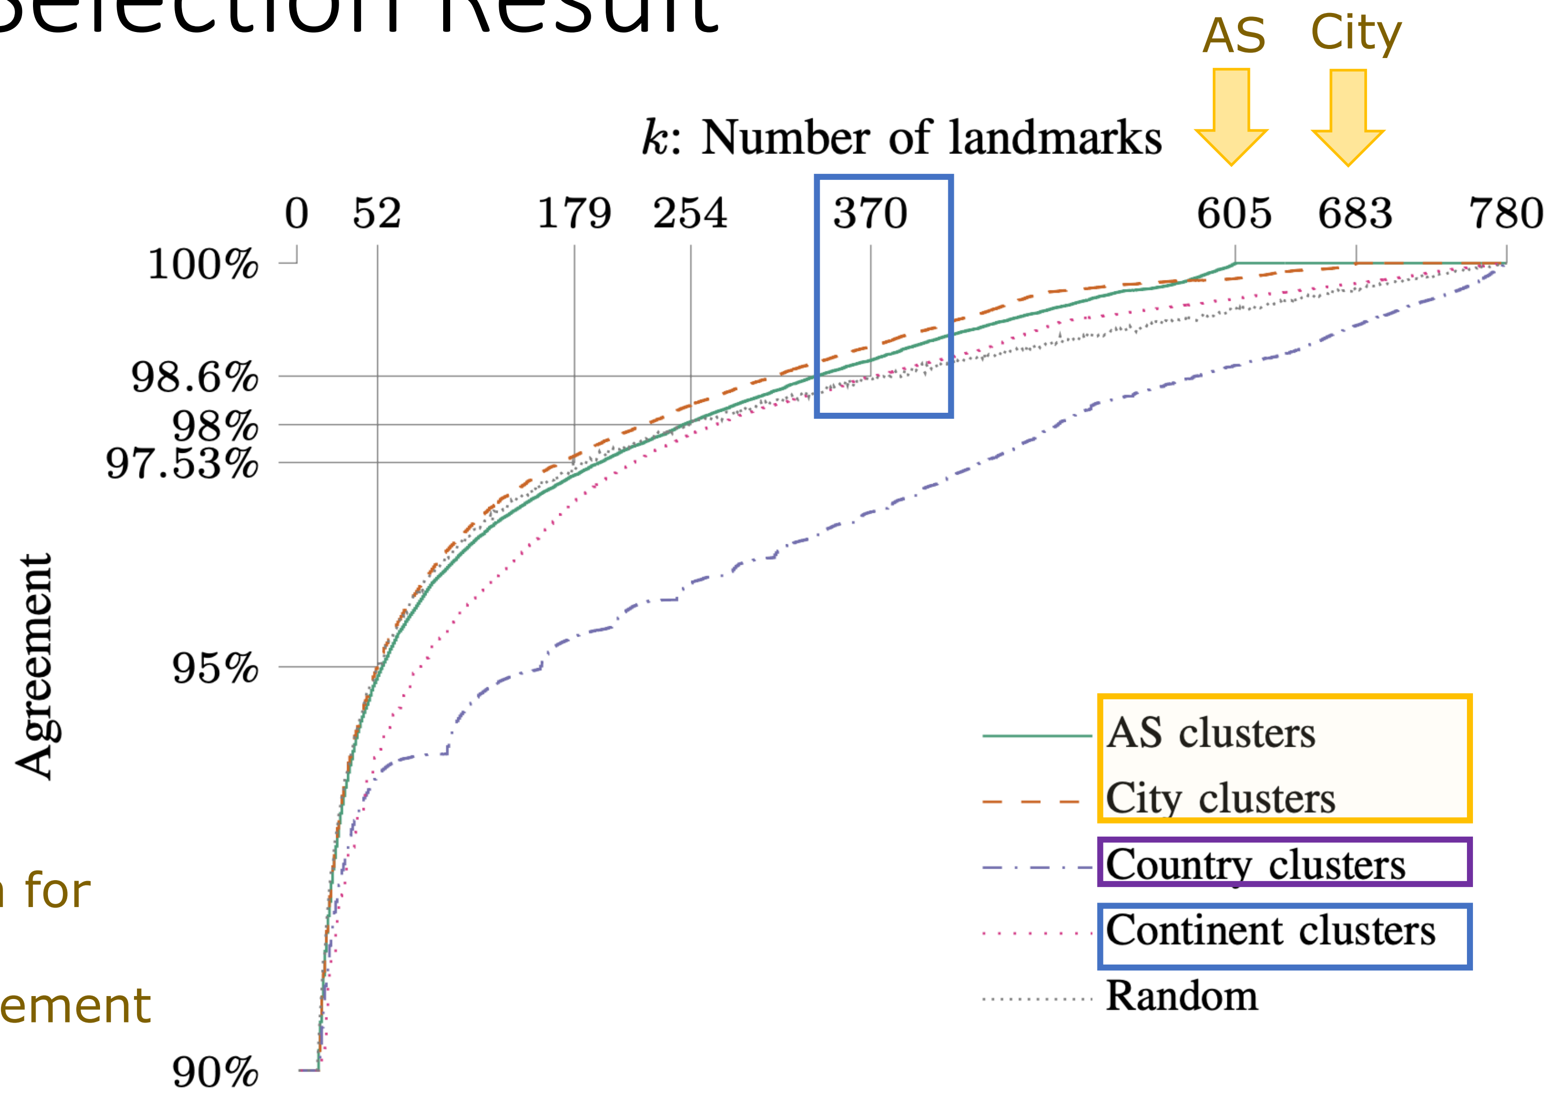

LS3: Greatest-Distance Selection

- Geographic distances between landmarks
- Minimum RTT between landmarks

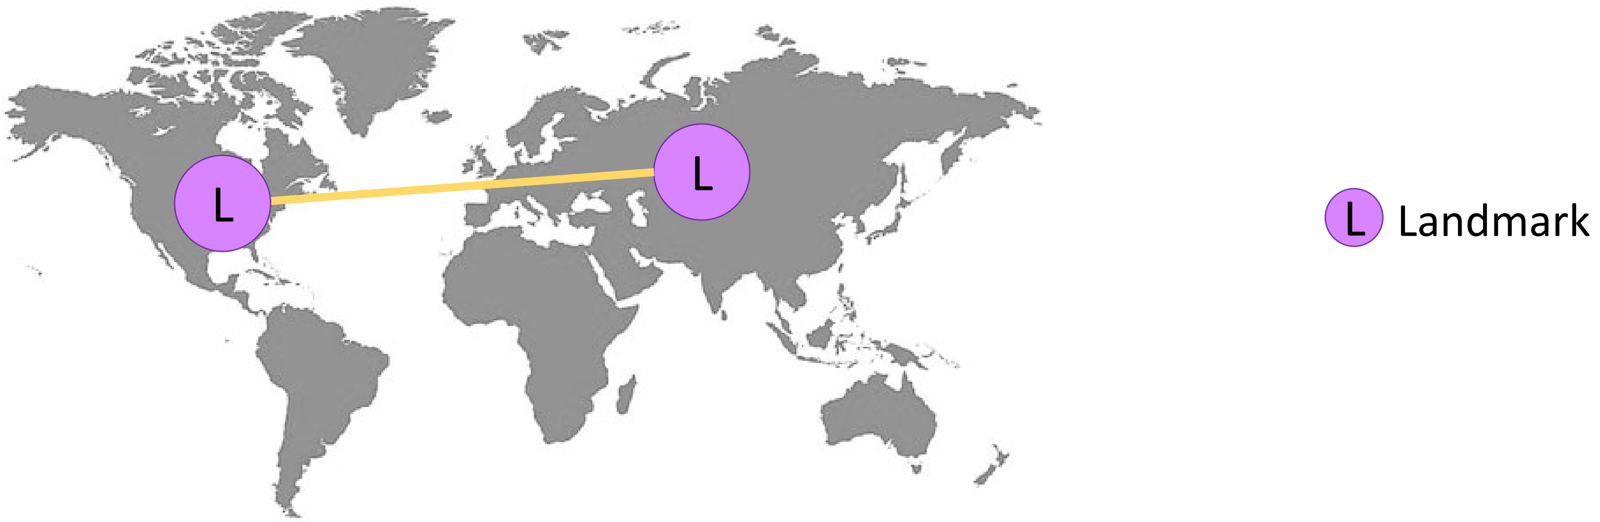

LS3: Greatest-Distance Selection

- Geographic distances between landmarks
- Minimum RTT between landmarks
  - *Anchors continuously measure and upload RTT data between each other, eliminating the need for additional measurements*

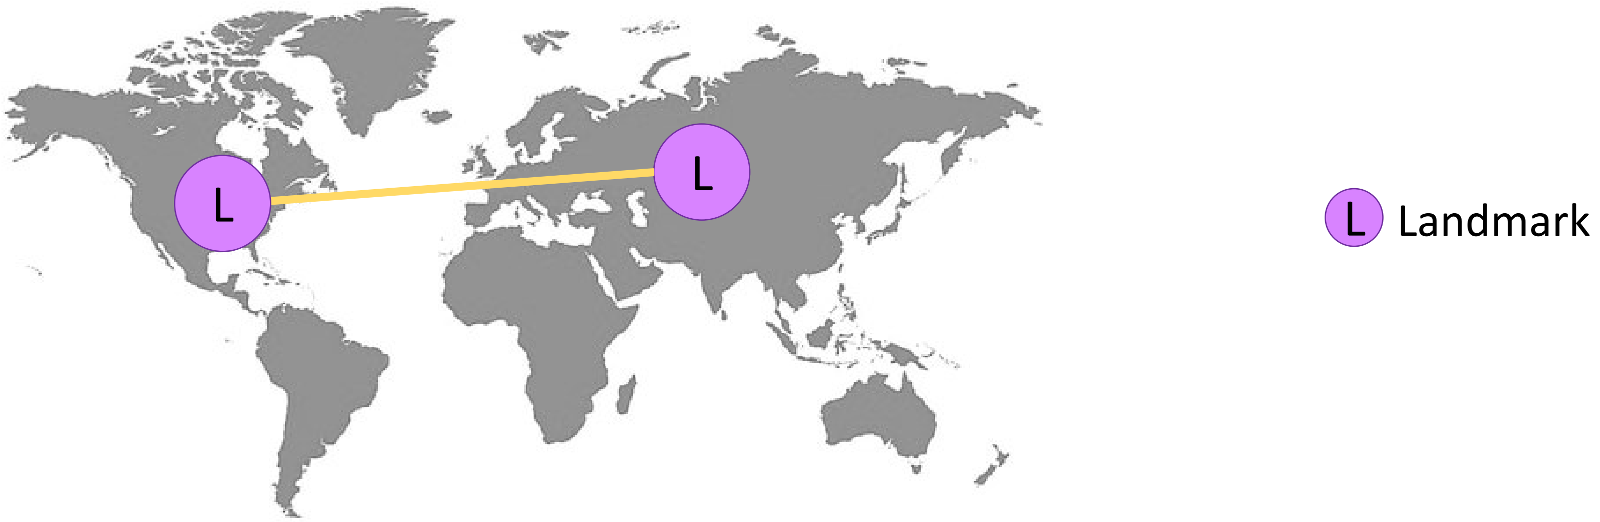

LS3: Greatest-Distance Selection (cont’d)

Selection with greedy algorithm for maximum spanning trees

LS3: Greatest-Distance Selection (cont'd)

Selection with greedy algorithm for maximum spanning trees

- Initial selection: starts with a landmark
  - Within the target's claimed location, or …

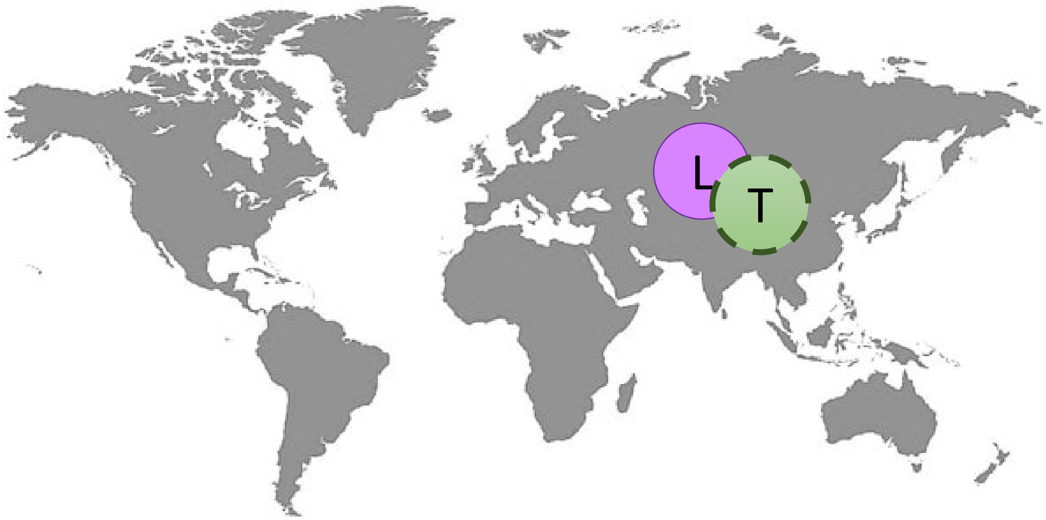

L Landmark

T Target's claimed location

LS3: Greatest-Distance Selection (cont’d)

Selection with greedy algorithm for maximum spanning trees

- Initial selection: starts with a landmark
  - Within the target's claimed location, or …
  - That maximizes the distance matric

LS3: Greatest-Distance Selection (cont'd)

Selection with greedy algorithm for maximum spanning trees

- Initial selection: starts with a landmark
  - Within the target's claimed location, or …
  - That maximizes the distance matric
- Continuation:
  - Selects landmarks maximizing diversity until desired subset size is reached

LS3: Greatest-Distance Selection Result

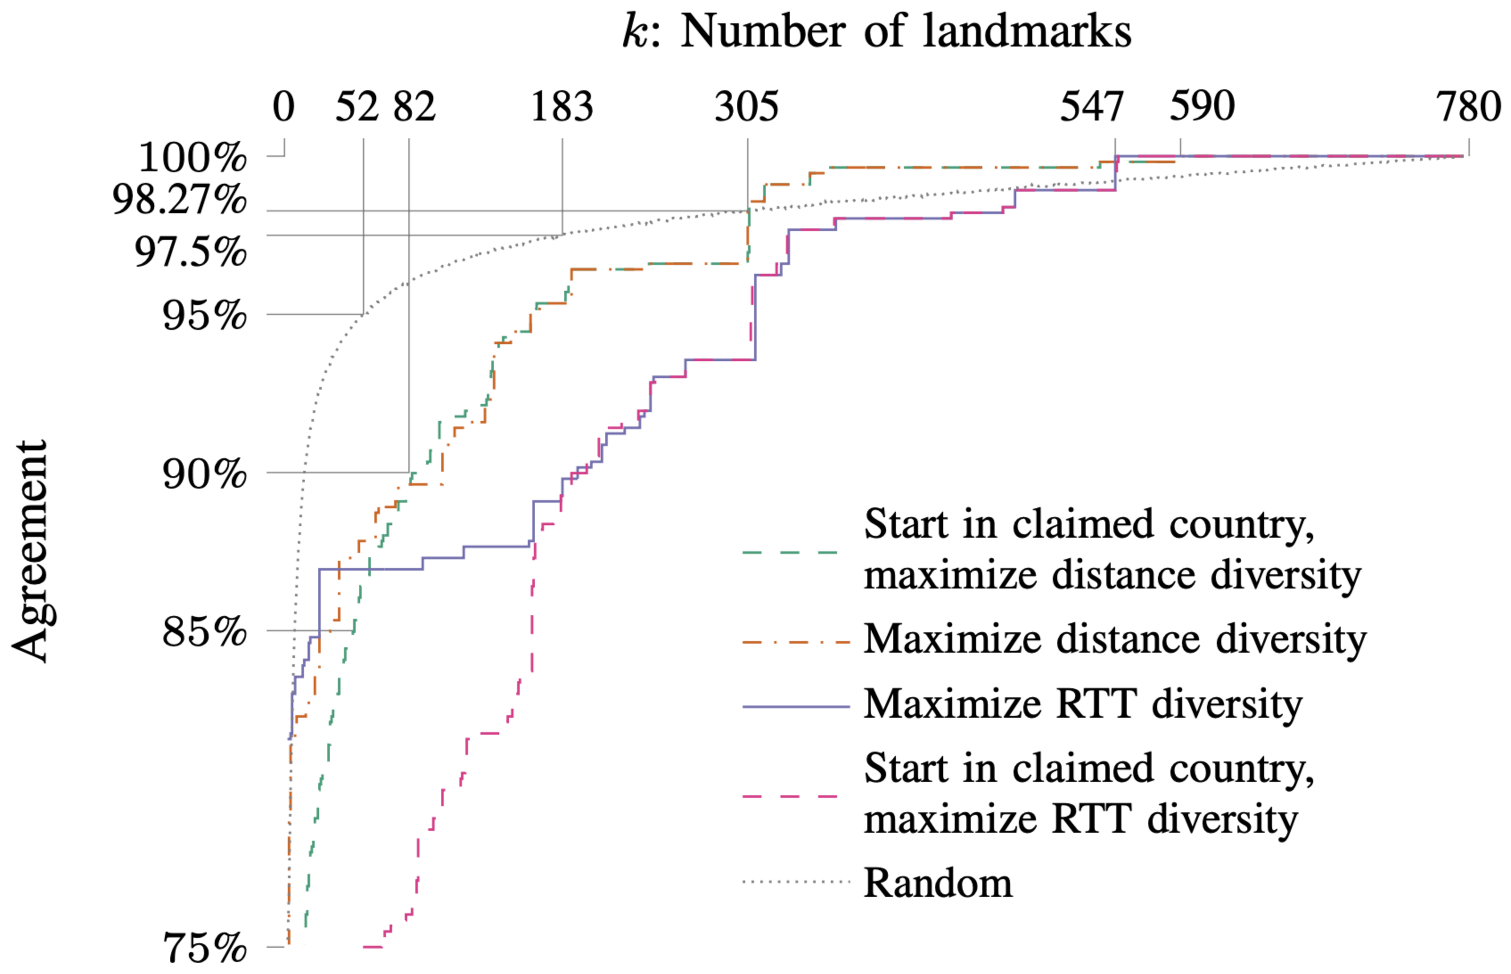

LS3: Greatest-Distance Selection Result

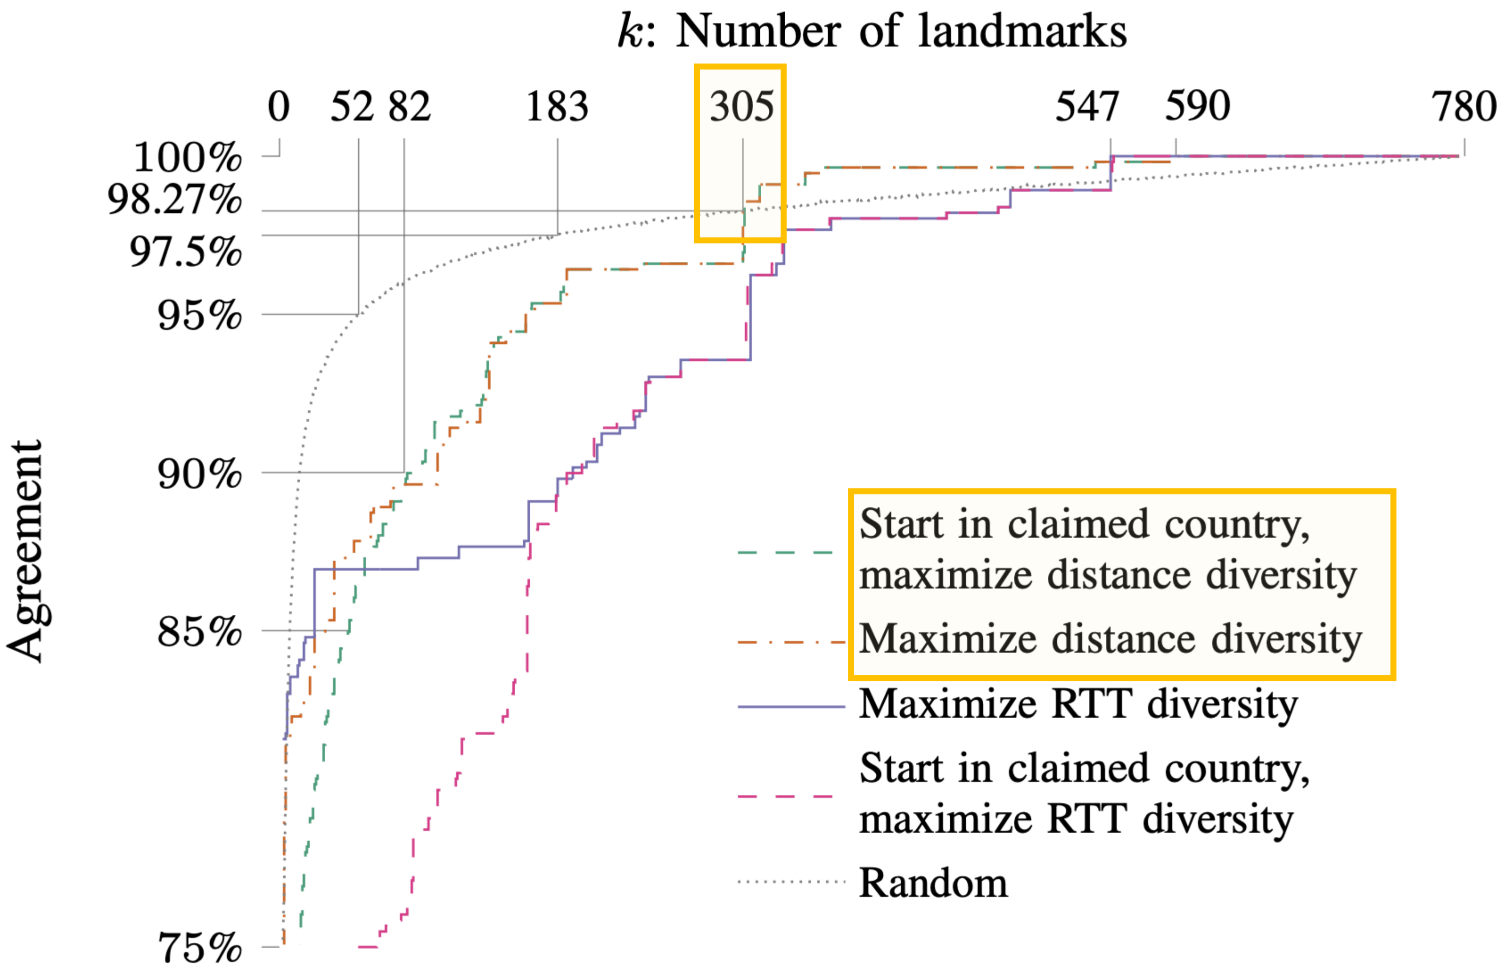

Outperforms random selection when 305+ landmarks are used **(39% of the pool)**

LS3: Greatest-Distance Selection Result

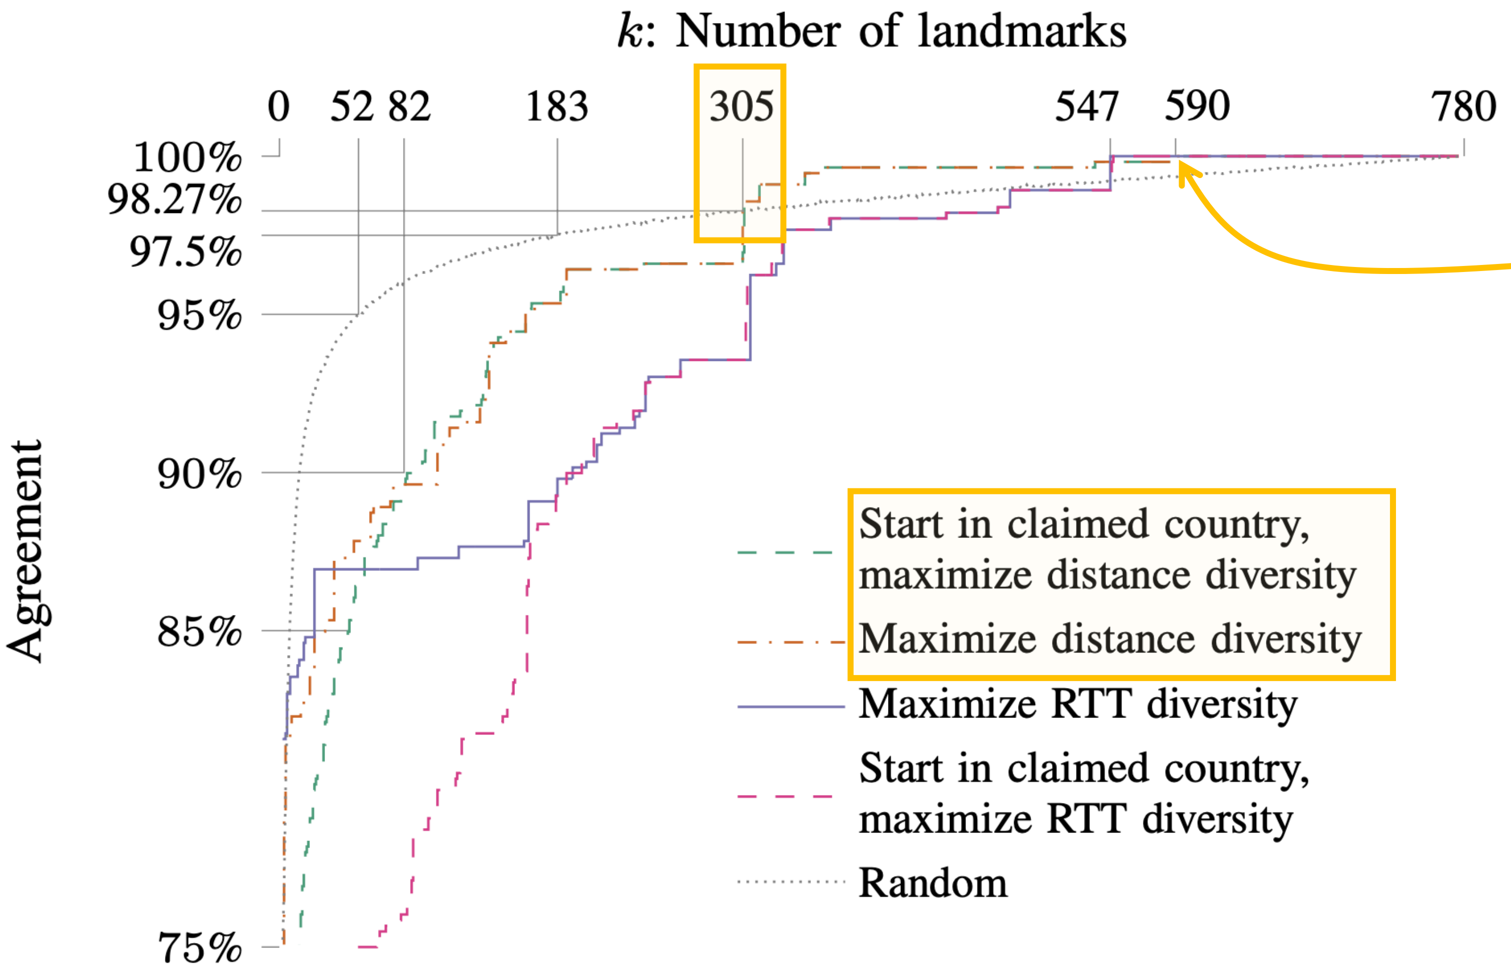

Outperforms random selection when 305+ landmarks are used **(39% of the pool)**

Achieves perfect agreement with 590+ landmarks **(75.6% of the pool)**

LS3: Greatest-Distance Selection Result

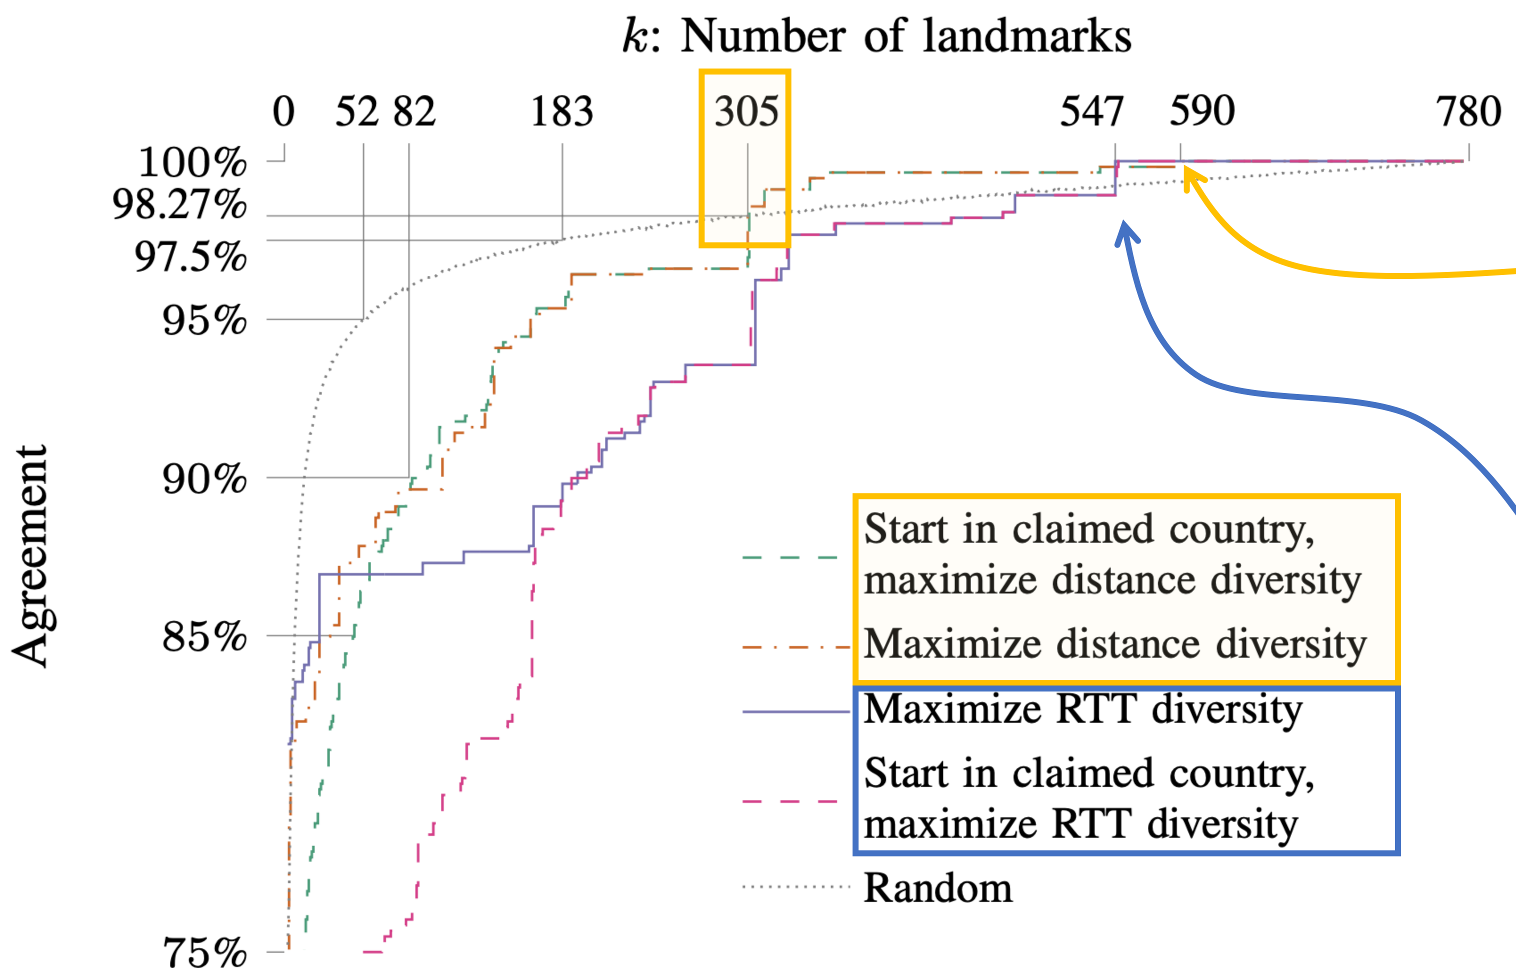

Outperforms random selection when 305+ landmarks are used **(39% of the pool)**

Achieves perfect agreement with 590+ landmarks **(75.6% of the pool)**

Outperforms random and geographic distance with 547+ landmarks (70% of the pool)

LS3: Greatest-Distance Selection Result

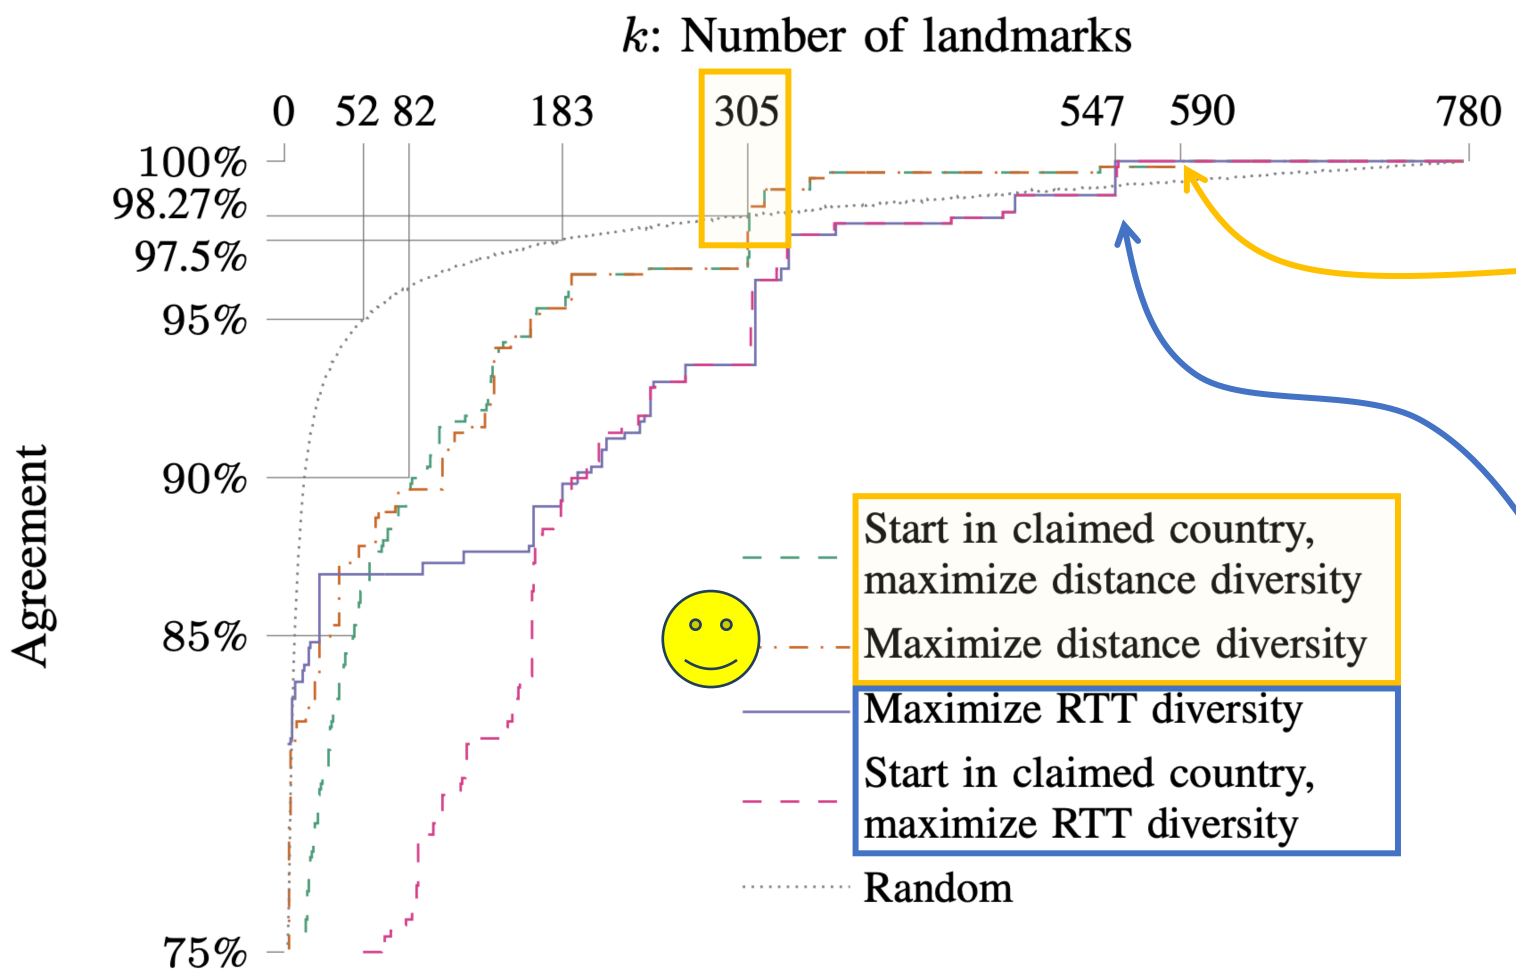

Outperforms random selection when 305+ landmarks are used **(39% of the pool)**

Achieves perfect agreement with 590+ landmarks **(75.6% of the pool)**

Outperforms random and geographic distance with 547+ landmarks (70% of the pool)

LS4: Hybrid Selection

- Observations from LS1–LS3
  - **Small Subsets**:

    Random and clustering selections **>** Geographic distance maximization
  - **Large Subsets:**

    Random and clustering selections **<** Geographic distance maximization

LS4: Hybrid Selection

- Observations from LS1–LS3
  - **Small Subsets**:

    Random and clustering selections **>** Geographic distance maximization
  - **Large Subsets:**

    Random and clustering selections **<** Geographic distance maximization

- Hybrid approaches may yield better results than any single method
  - Hybrid 1: Clustering and great distance
  - Hybrid 2: Random, then Hybrid 1

LS4: Hybrid 1: Clustering and Great Distance

LS4: Hybrid 1: Clustering and Great Distance

- **Initial Focus:** Prioritize cluster diversity over geographic distance

LS4: Hybrid 1: Clustering and Great Distance

- **Initial Focus:** Prioritize cluster diversity over geographic distance
- **Filling Gaps:** Select next landmark from unrepresented clusters if any are missing

LS4: Hybrid 1: Clustering and Great Distance

- **Initial Focus:** Prioritize cluster diversity over geographic distance
- **Filling Gaps:** Select next landmark from unrepresented clusters if any are missing
- **Subsequent Focus:** Once all clusters are represented, shift to purely geographic distance maximization

LS4: Hybrid 1: Clustering and Great Distance Result

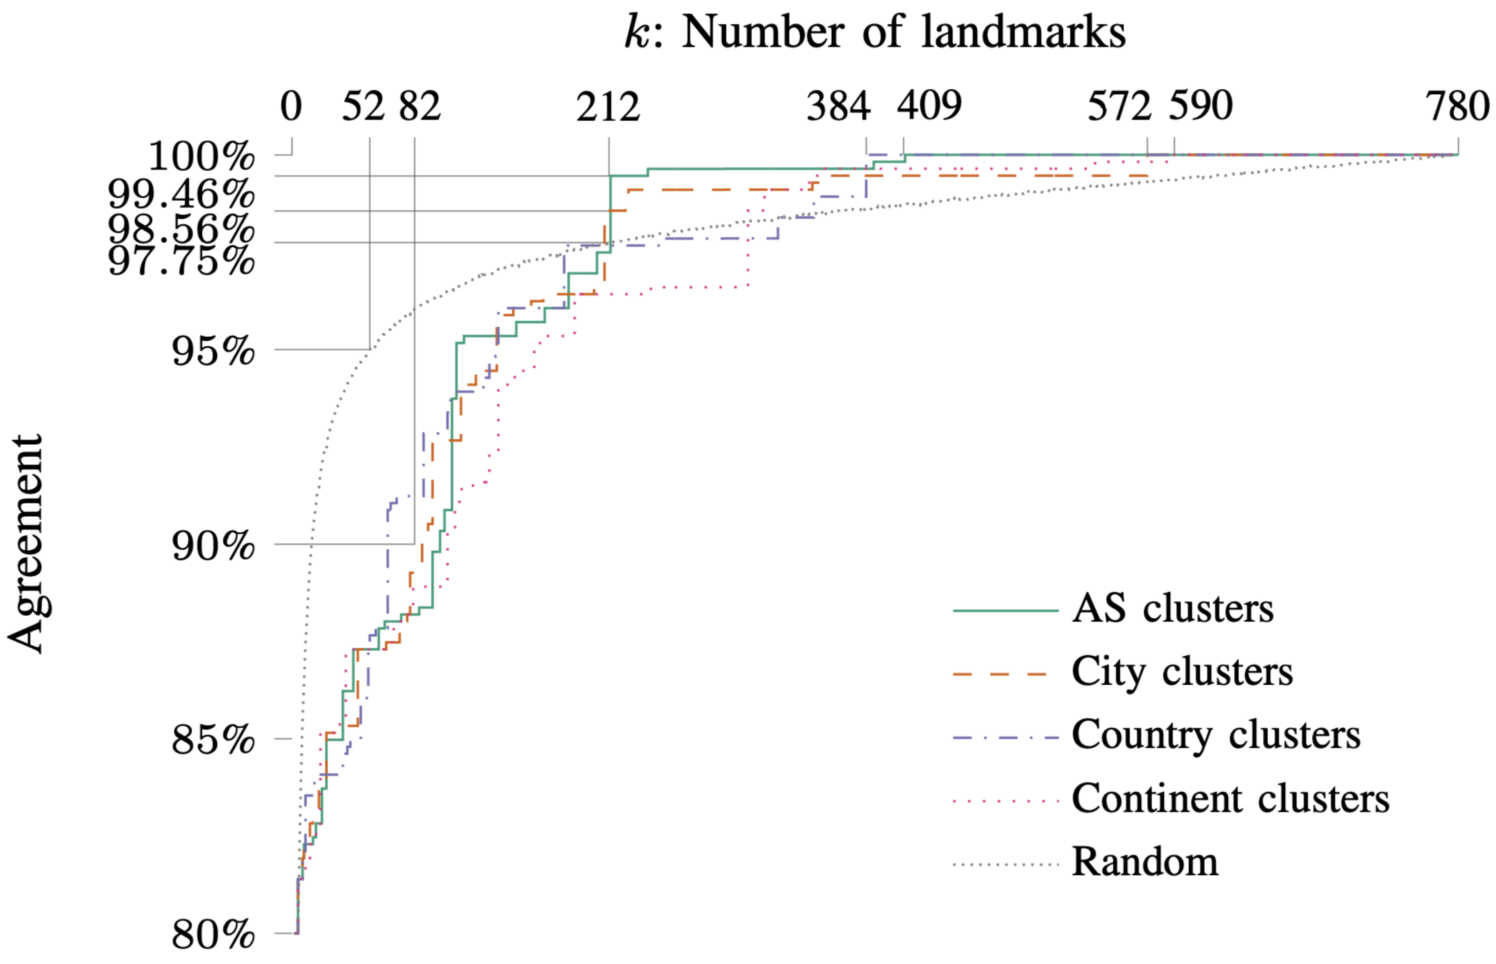

LS4: Hybrid 1: Clustering and Great Distance Result

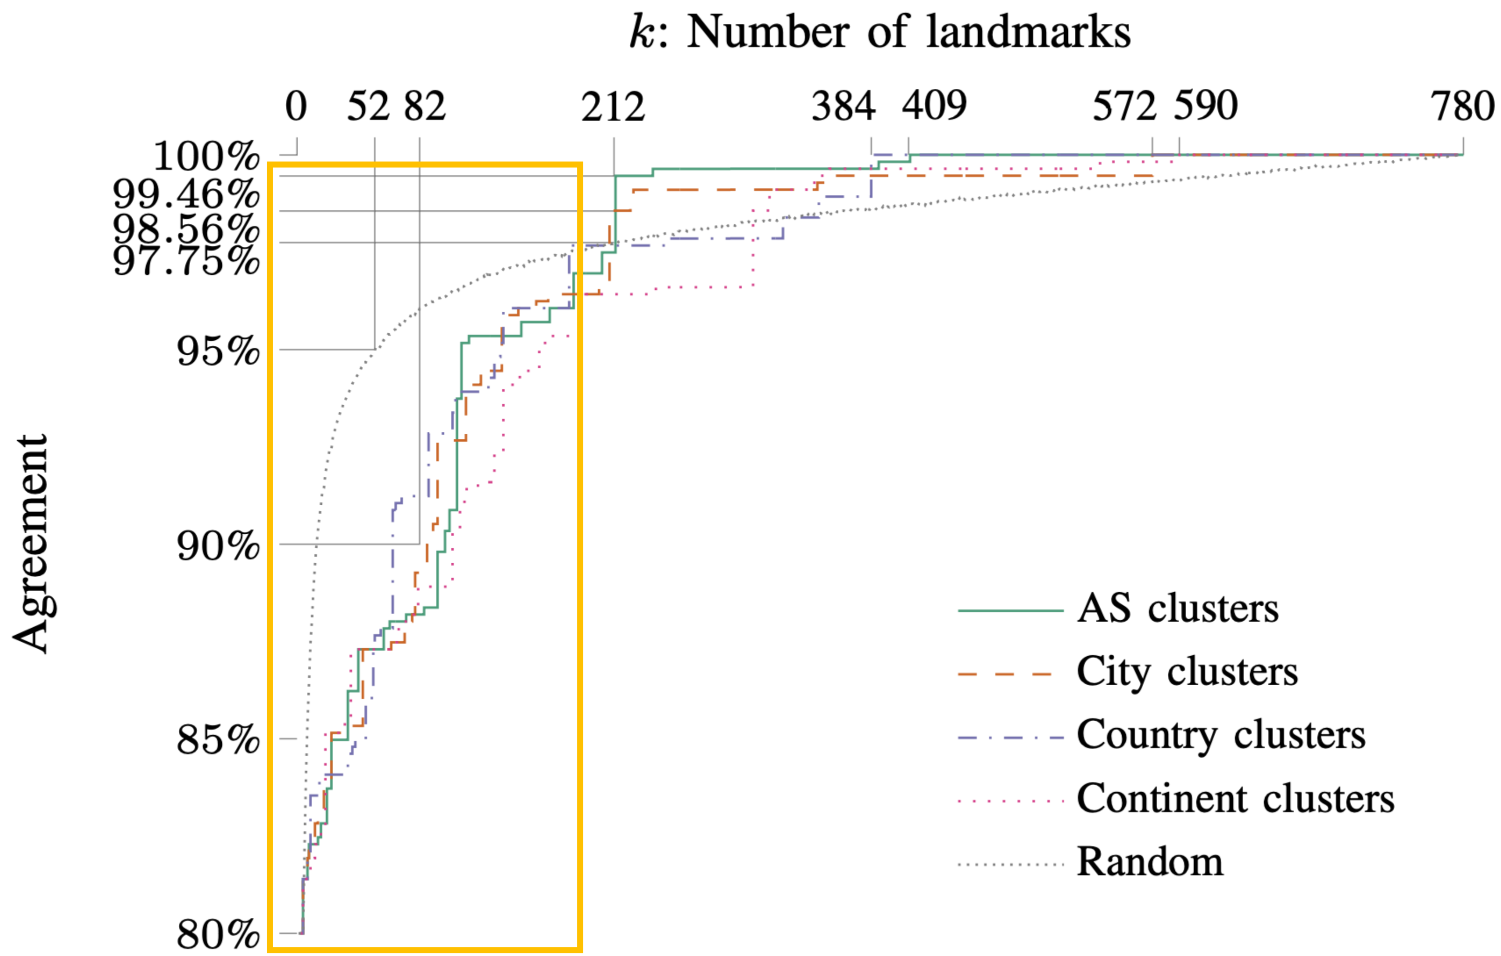

Worse than random at small subset sizes

LS4: Hybrid 1: Clustering and Great Distance Result

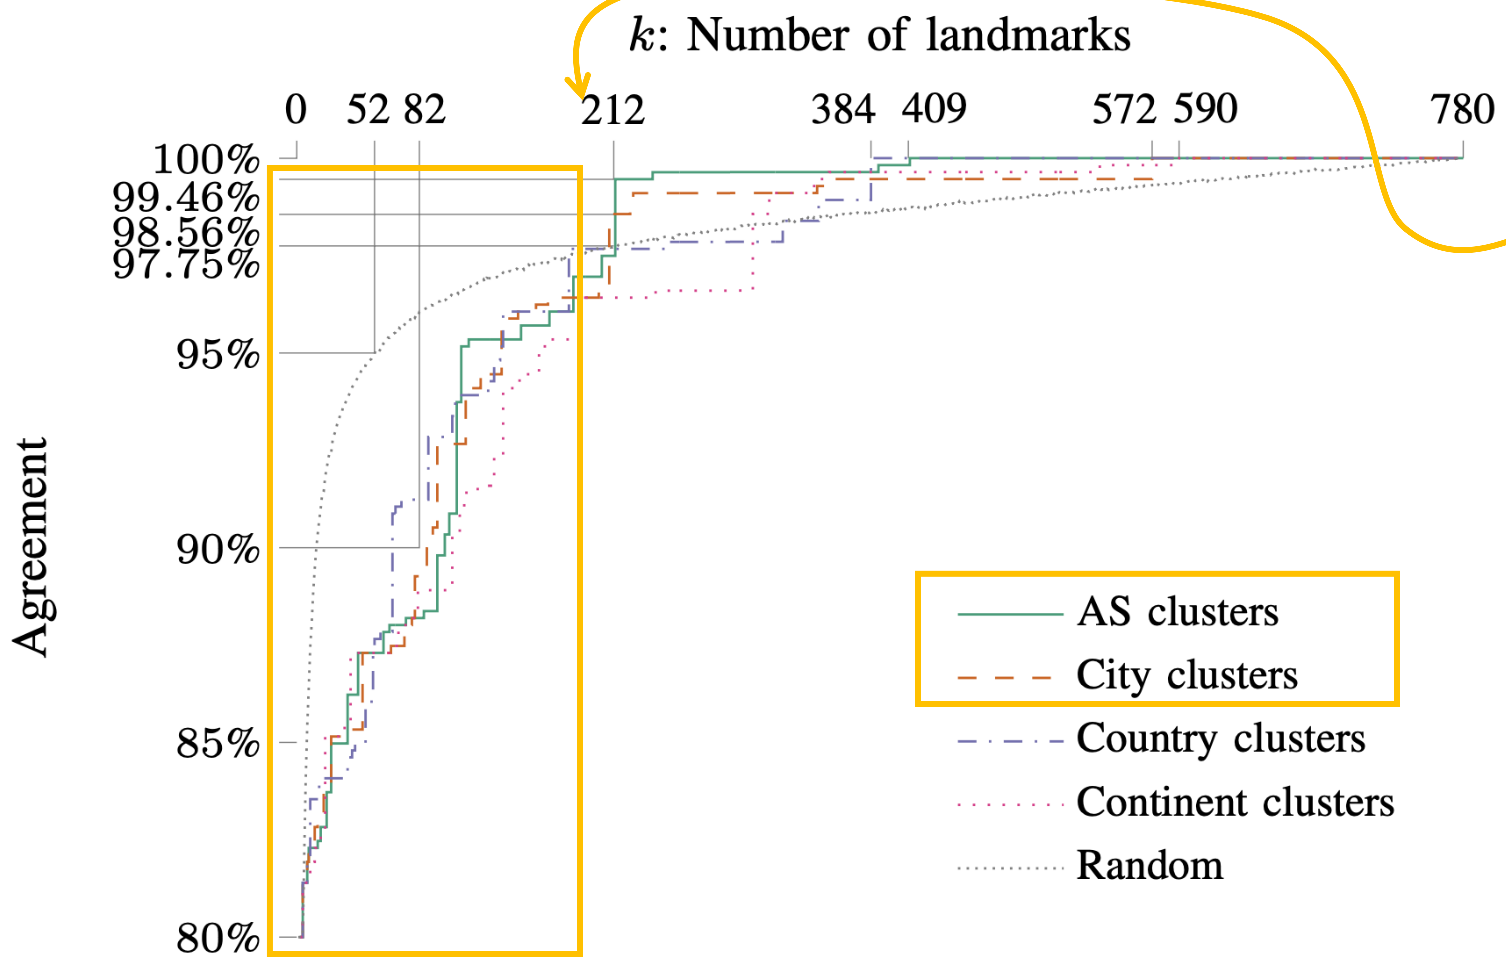

Surpasses random with 210+ landmarks (27% of the pool) using AS or city-based clustering

Worse than random at small subset sizes

LS4: Hybrid 1: Clustering and Great Distance Result

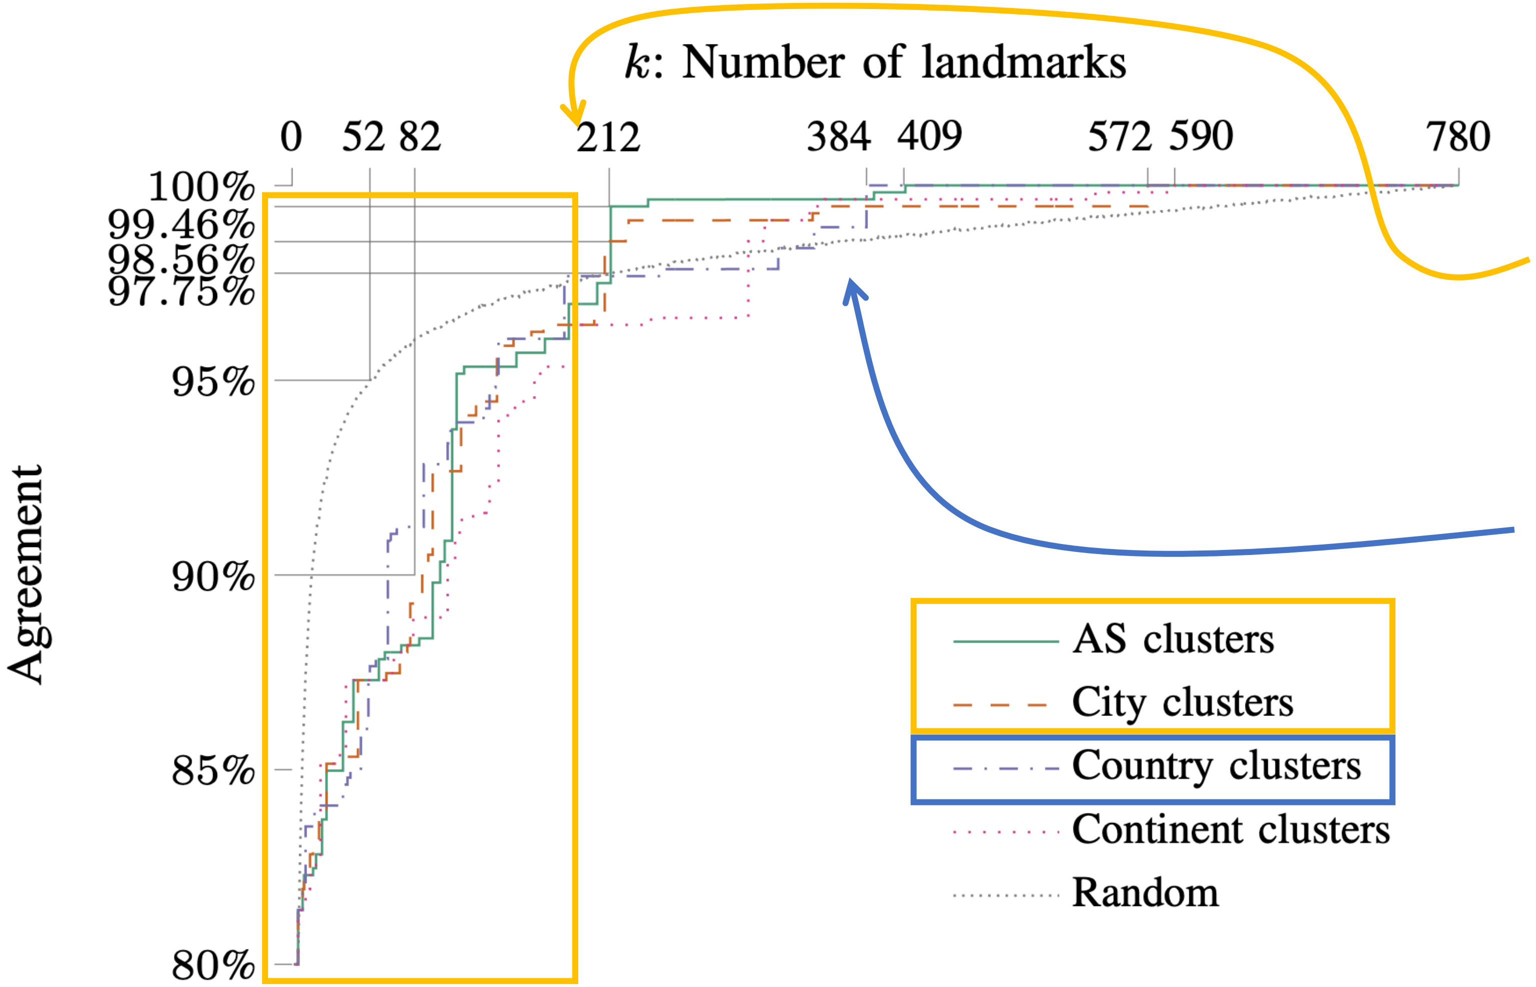

Surpasses random with 210+ landmarks (27% of the pool) using AS or city-based clustering

Achieves perfect agreement first at 384 landmarks (49% of the pool).

Worse than random at small subset sizes

LS4: Hybrid 1: Clustering and Great Distance Result

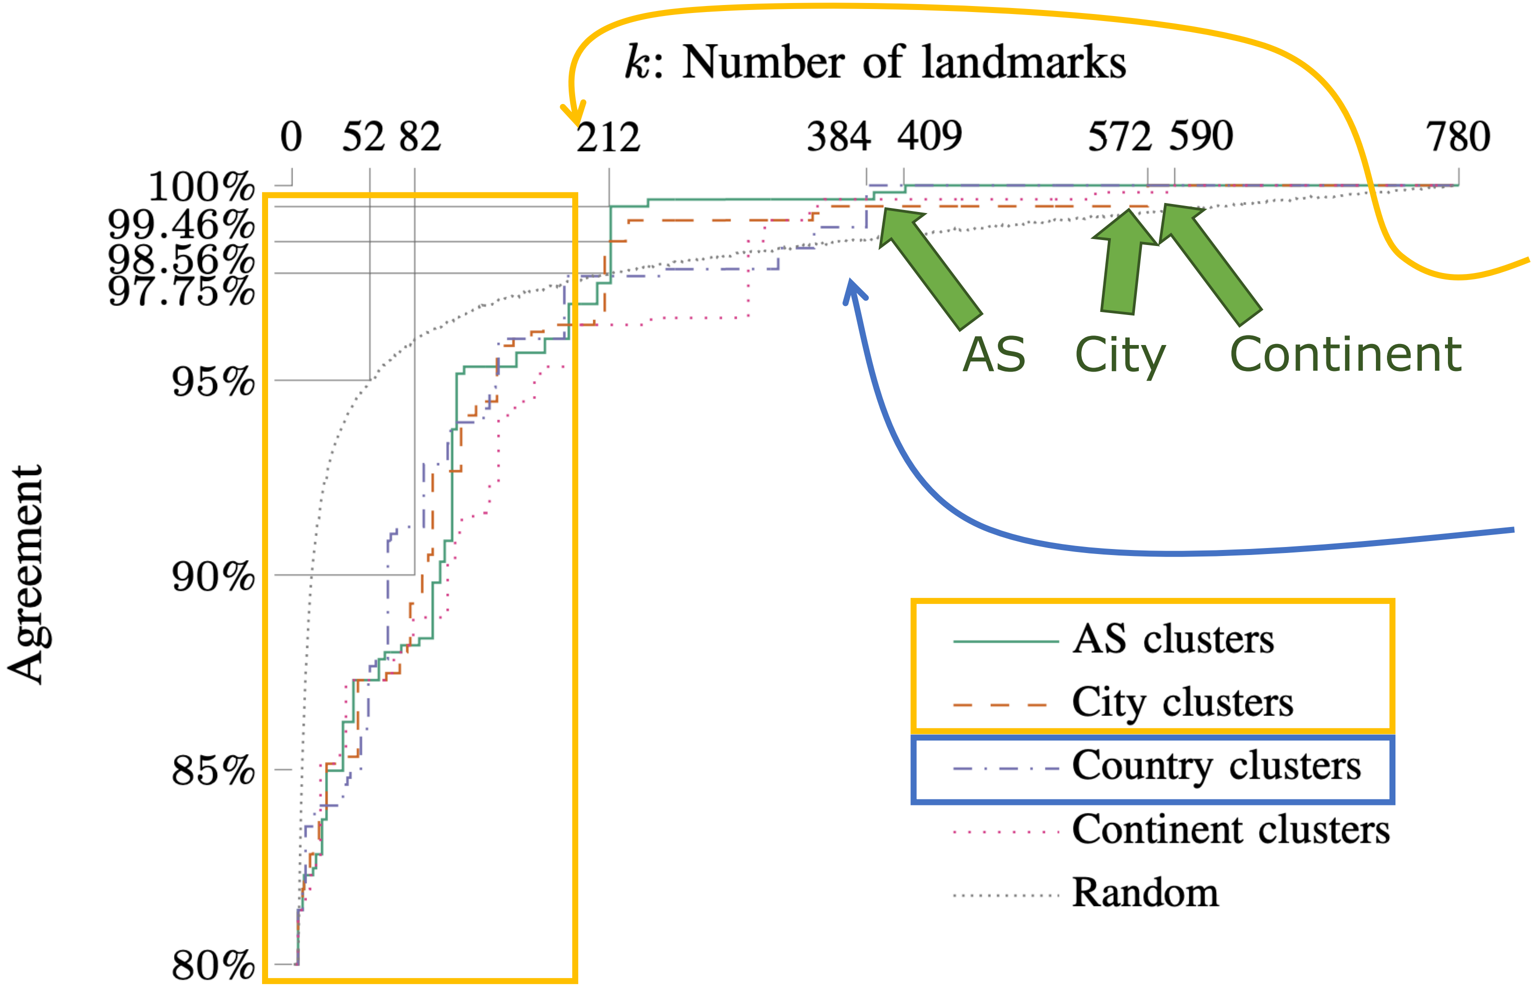

Surpasses random with 210+ landmarks (27% of the pool) using AS or city-based clustering

Achieves perfect agreement first at 384 landmarks (49% of the pool).

AS-based hybrid: at 409
City-based hybrid: at 572
Continent-based hybrid: at 590

Worse than random at small subset sizes

LS4: Hybrid 2: Random, Then Hybrid 1

Observations so far: No algorithm substantially outperforms random selection for small subsets

LS4: Hybrid 2: Random, Then Hybrid 1

Observations so far: No algorithm substantially outperforms random selection for small subsets

Hybrid 2: random, then hybrid 1

- **Initial selection:** Begin by randomly choosing up to 100 landmarks
- **Expansion:** Expand these subsets using the Hybrid 1 approach

LS4: Hybrid 2: Random, Then Hybrid 1 Result

Enhanced Performance: Modification aligns performance closely with random selection across all subset sizes

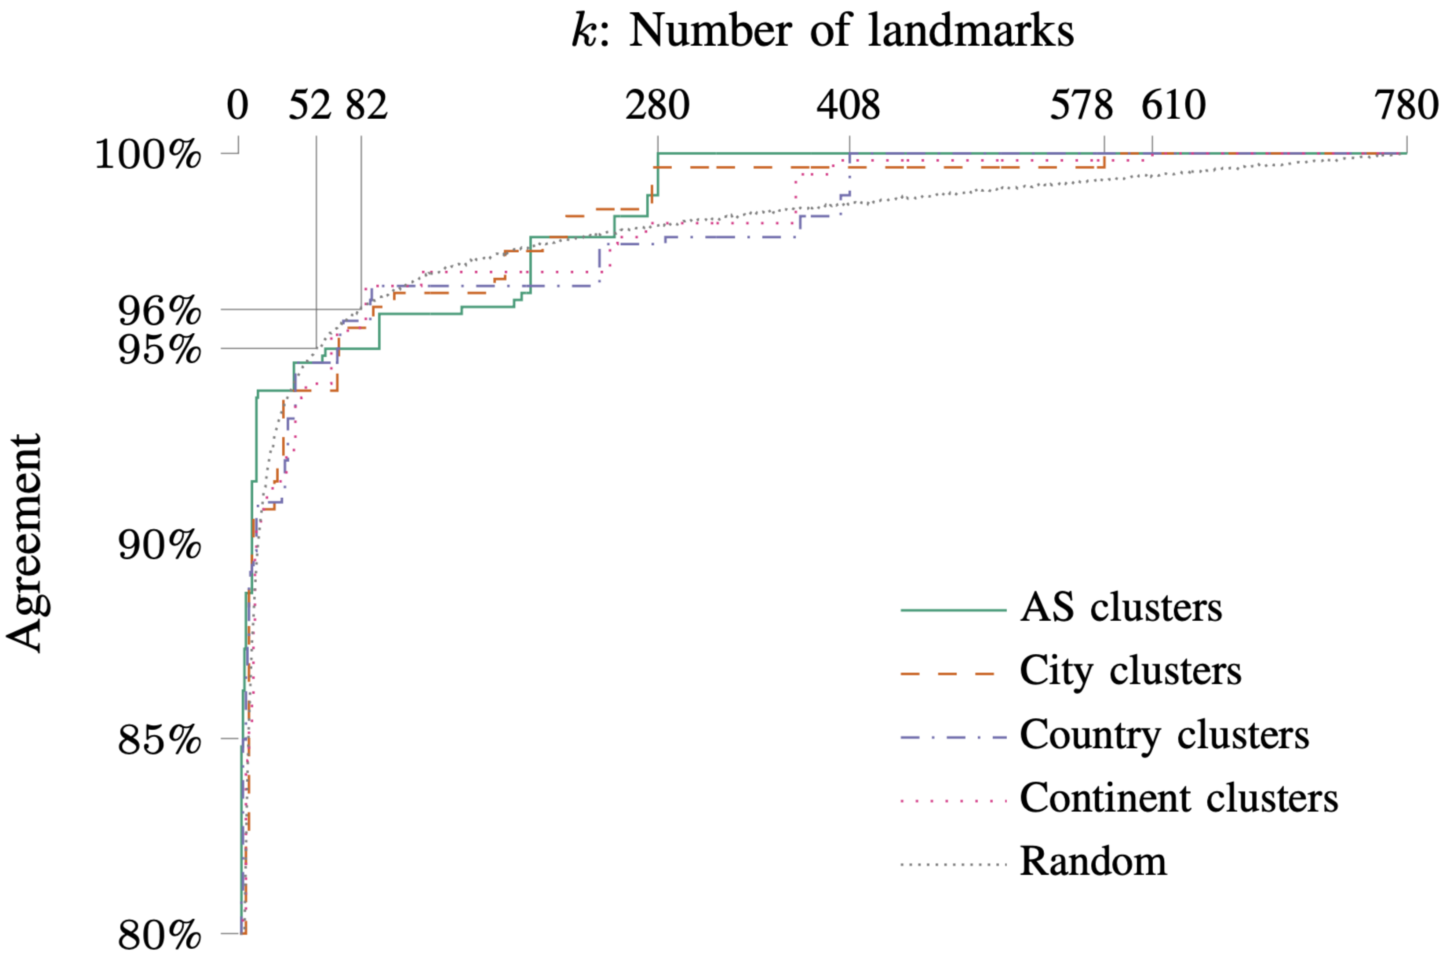

LS4: Hybrid 2: Random, Then Hybrid 1 Result

Enhanced Performance:
Modification aligns performance closely with random selection across all subset sizes

AS Clustering:
Reaches full agreement with 280 landmarks, 130 fewer than Hybrid 1

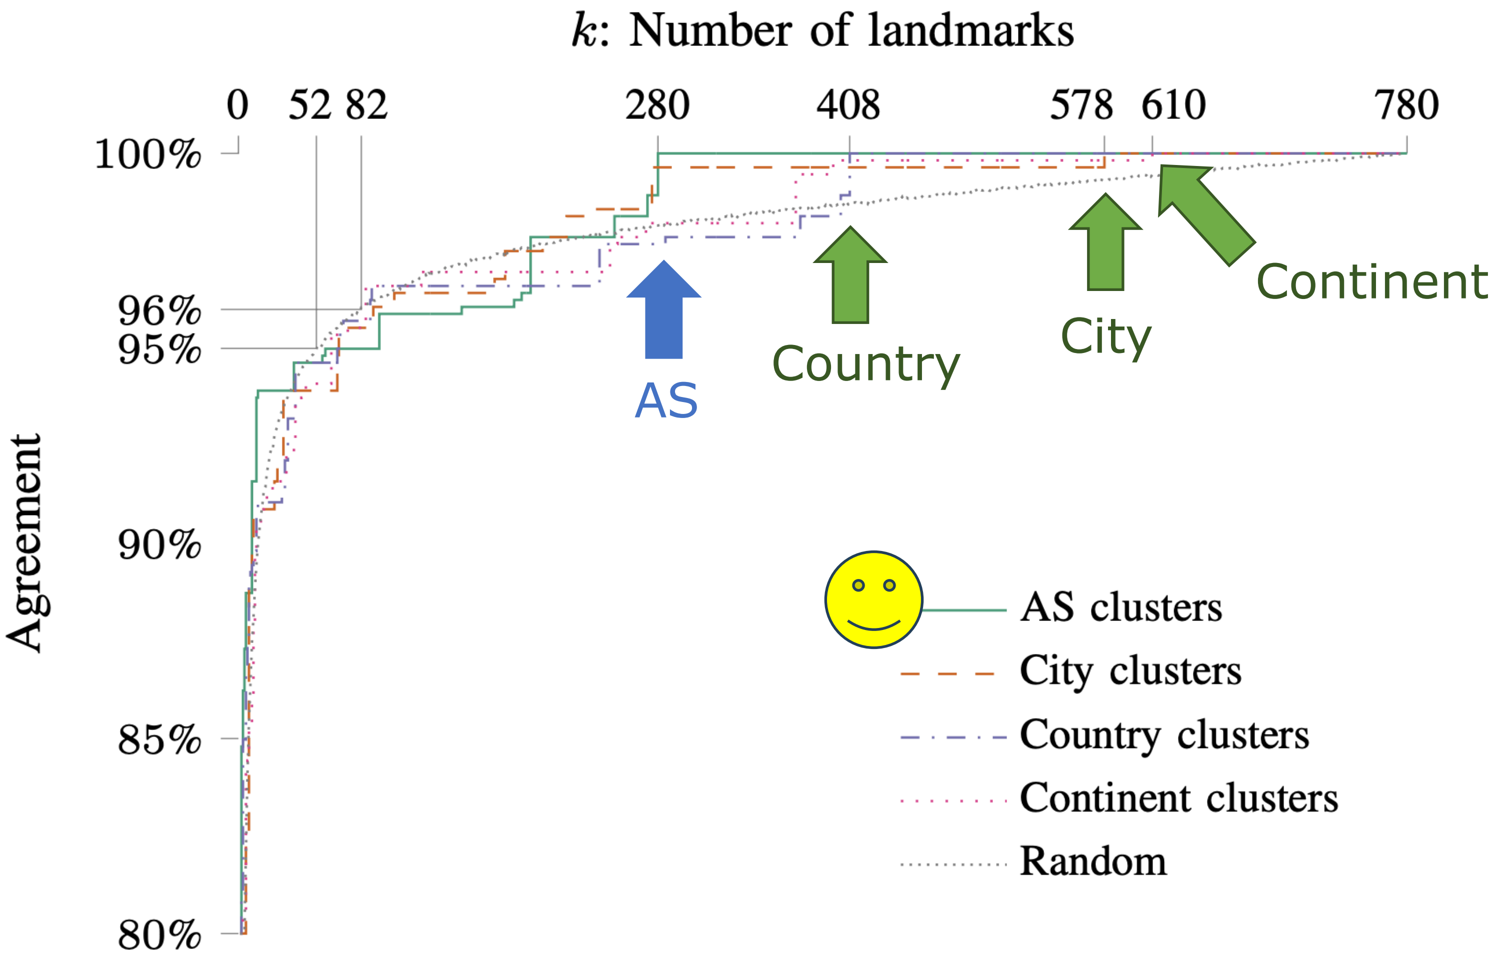

Analysis

Shorthand	Metric	Cluster by	First 100 random?	# landmarks to… beat random	perfect agreement
CLUSTER-CITY		Cities		179	683
H2-CONTINENT	Geodesic	Continents	Yes	85	610
CLUSTER-AS		ASes		254	605
DIST-GEO	Geodesic			305	590
H1-CONTINENT	Geodesic	Continents		305	590
H2-CITY	Geodesic	Cities	Yes	179	578
DIST-RTT	Travel time			547	547
H1-AS	Geodesic	ASes		213	410
H2-COUNTRY	Geodesic	Countries	Yes	88	408
H1-COUNTRY	Geodesic	Countries		182	384
H2-AS	Geodesic	ASes	Yes	195	280

Analysis

				# landmarks to...	
Shorthand	Metric	Cluster by	First 100 random?	beat random	perfect agreement
CLUSTER-CITY		Cities		179	683
H2-CONTINENT	Geodesic	Continents	Yes	85	610
CLUSTER-AS		ASes		254	605
DIST-GEO	Geodesic			305	590
H1-CONTINENT	Geodesic	Continents		305	590
H2-CITY	Geodesic	Cities	Yes	179	578
DIST-RTT	Travel time			547	547
H1-AS	Geodesic	ASes		213	410
H2-COUNTRY	Geodesic	Countries	Yes	88	408
H1-COUNTRY	Geodesic	Countries		182	384
H2-AS	Geodesic	ASes	Yes	195	280

Analysis

Shorthand	Metric	Cluster by	First 100 random?	# landmarks to… beat random	# landmarks to… perfect agreement
CLUSTER-CITY		Cities		179	683
H2-CONTINENT	Geodesic	Continents	Yes	85	610
CLUSTER-AS		ASes		254	605
DIST-GEO	Geodesic			305	590
H1-CONTINENT	Geodesic	Continents		305	590
H2-CITY	Geodesic	Cities	Yes	179	578
DIST-RTT	Travel time			547	547
H1-AS	Geodesic	ASes		213	410
H2-COUNTRY	Geodesic	Countries	Yes	88	408
H1-COUNTRY	Geodesic	Countries		182	384
Top Performer H2-AS	Geodesic	ASes	Yes	195	280

Analysis

Shorthand	Metric	Cluster by	First 100 random?	# landmarks to… beat random	perfect agreement
CLUSTER-CITY		Cities		179	683
H2-CONTINENT	Geodesic	Continents	Yes	85	610
CLUSTER-AS		ASes		254	605
DIST-GEO	Geodesic			305	590
H1-CONTINENT	Geodesic	Continents		305	590
H2-CITY	Geodesic	Cities	Yes	179	578
DIST-RTT	Travel time			547	547
H1-AS	Geodesic	ASes		213	410
H2-COUNTRY	Geodesic	Countries	Yes	88	408
H1-COUNTRY	Geodesic	Countries		182	384
Top Performer H2-AS	Geodesic	ASes	Yes	195 (36%)	280

ICMP Echo Request packets **1,308,060**

469,560 (36%)

Analysis

	Shorthand	Metric	Cluster by	First 100 random?	# landmarks to... beat random	perfect agreement	ICMP Echo Request packets **1,308,060**
	CLUSTER-CITY		Cities		179	683	
	H2-CONTINENT	Geodesic	Continents	Yes	85	610	
	CLUSTER-AS		ASes		254	605	
	DIST-GEO	Geodesic			305	590	
	H1-CONTINENT	Geodesic	Continents		305	590	
	H2-CITY	Geodesic	Cities	Yes	179	578	
	DIST-RTT	Travel time			547	547	
	H1-AS	Geodesic	ASes		213	410	
	H2-COUNTRY	Geodesic	Countries	Yes	88	408	
Consistent Runner-Up	H1-COUNTRY	Geodesic	Countries		182 (49%)	384	643,968 (49%)
Top Performer	H2-AS	Geodesic	ASes	Yes	195 (36%)	280	469,560 (36%)

Summary

- Demonstrated that it is possible to **reduce landmarks by 2/3** with no change in the overall results
- **City/AS-based clusters outperform** country/continent-based clusters
    - Highlighting the need for fine-grained diversity
- **Geographic distance is a better metric** than RTT for selecting landmarks close to distant targets
- **Future directions:** Combine selection rules with incremental geolocation algorithms to further reduce landmarks and leverage RIPE Atlas probes for greater diversity

Thank You

scho@smith.edu