Resource-Aware Distributed Loop Closure Detection with Provable Performance Guarantees

Yulun Tian Kasra Khosoussi Jonathan P. How

Massachusetts Institute of Technology

LIDS

Simultaneous Localization and Mapping

Front-end
loop closure
detection

Back-end
optimization

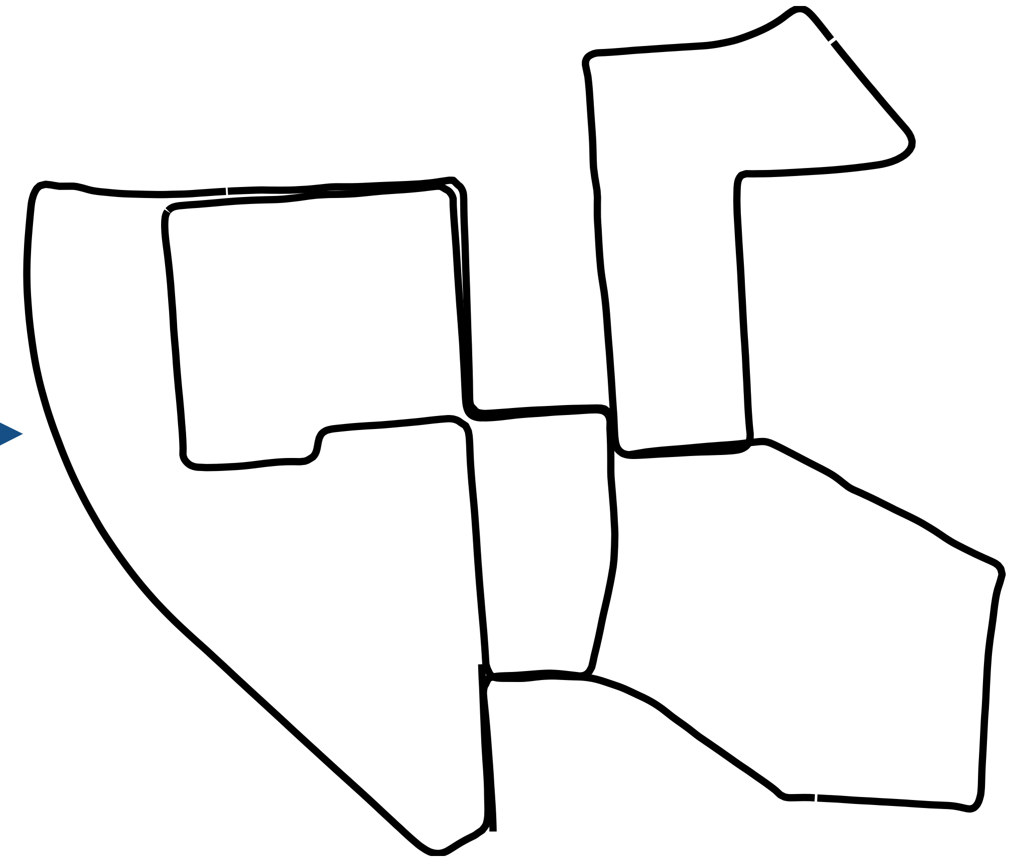

Collaborative Simultaneous Localization and Mapping

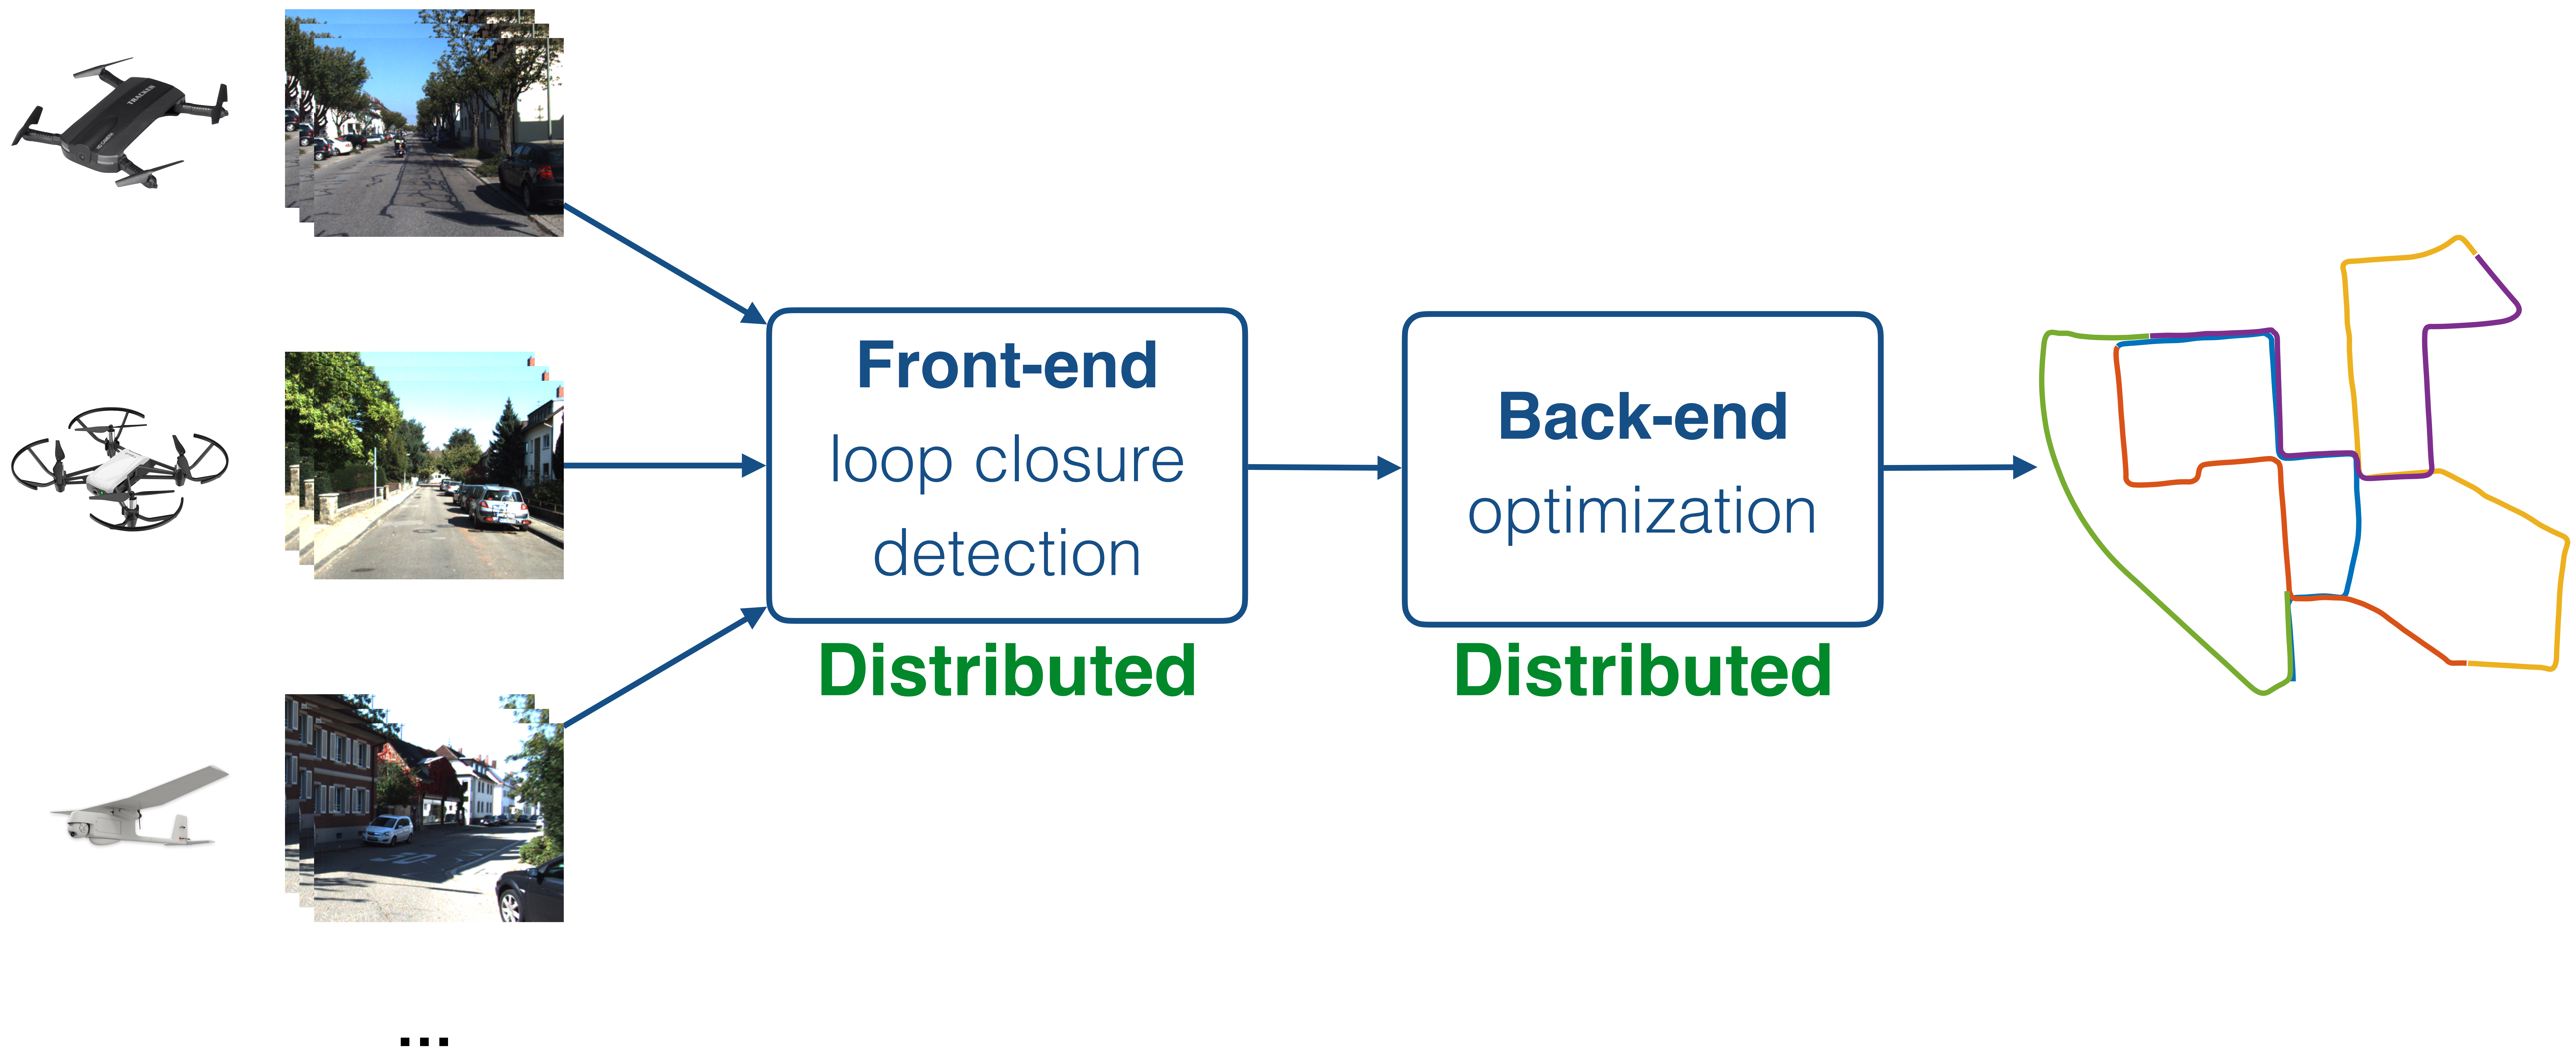

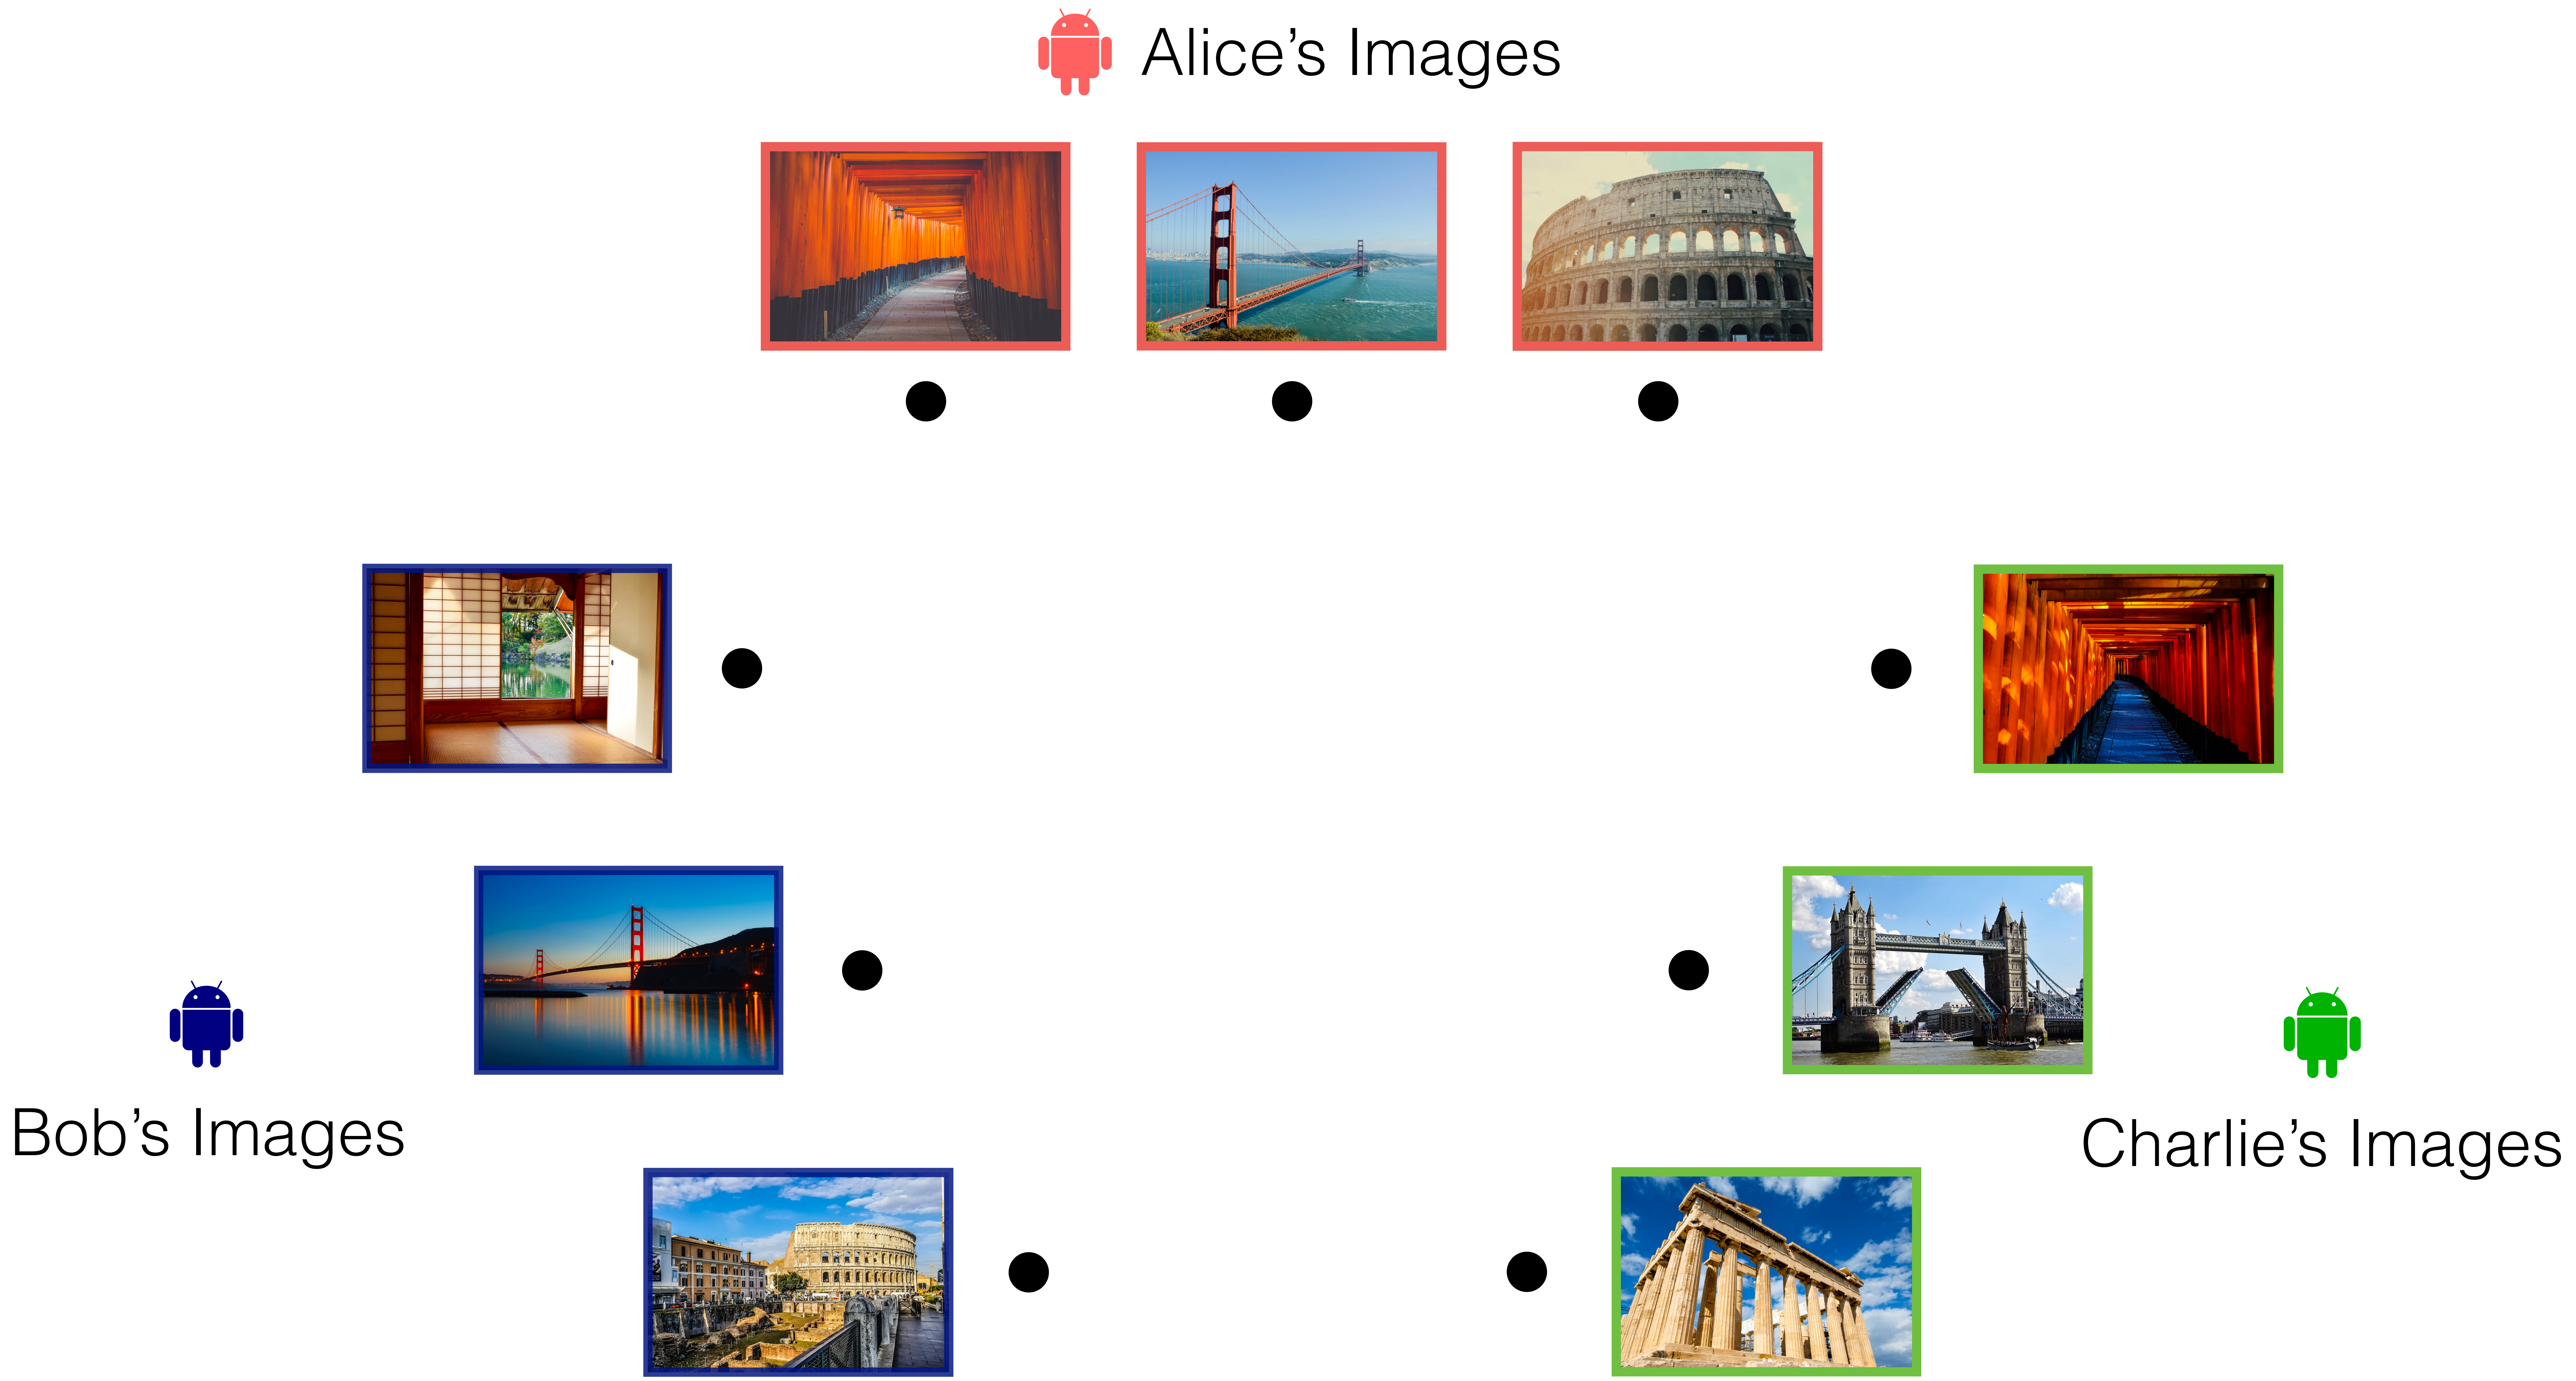
Alice's Images
Bob's Images
Charlie's Images

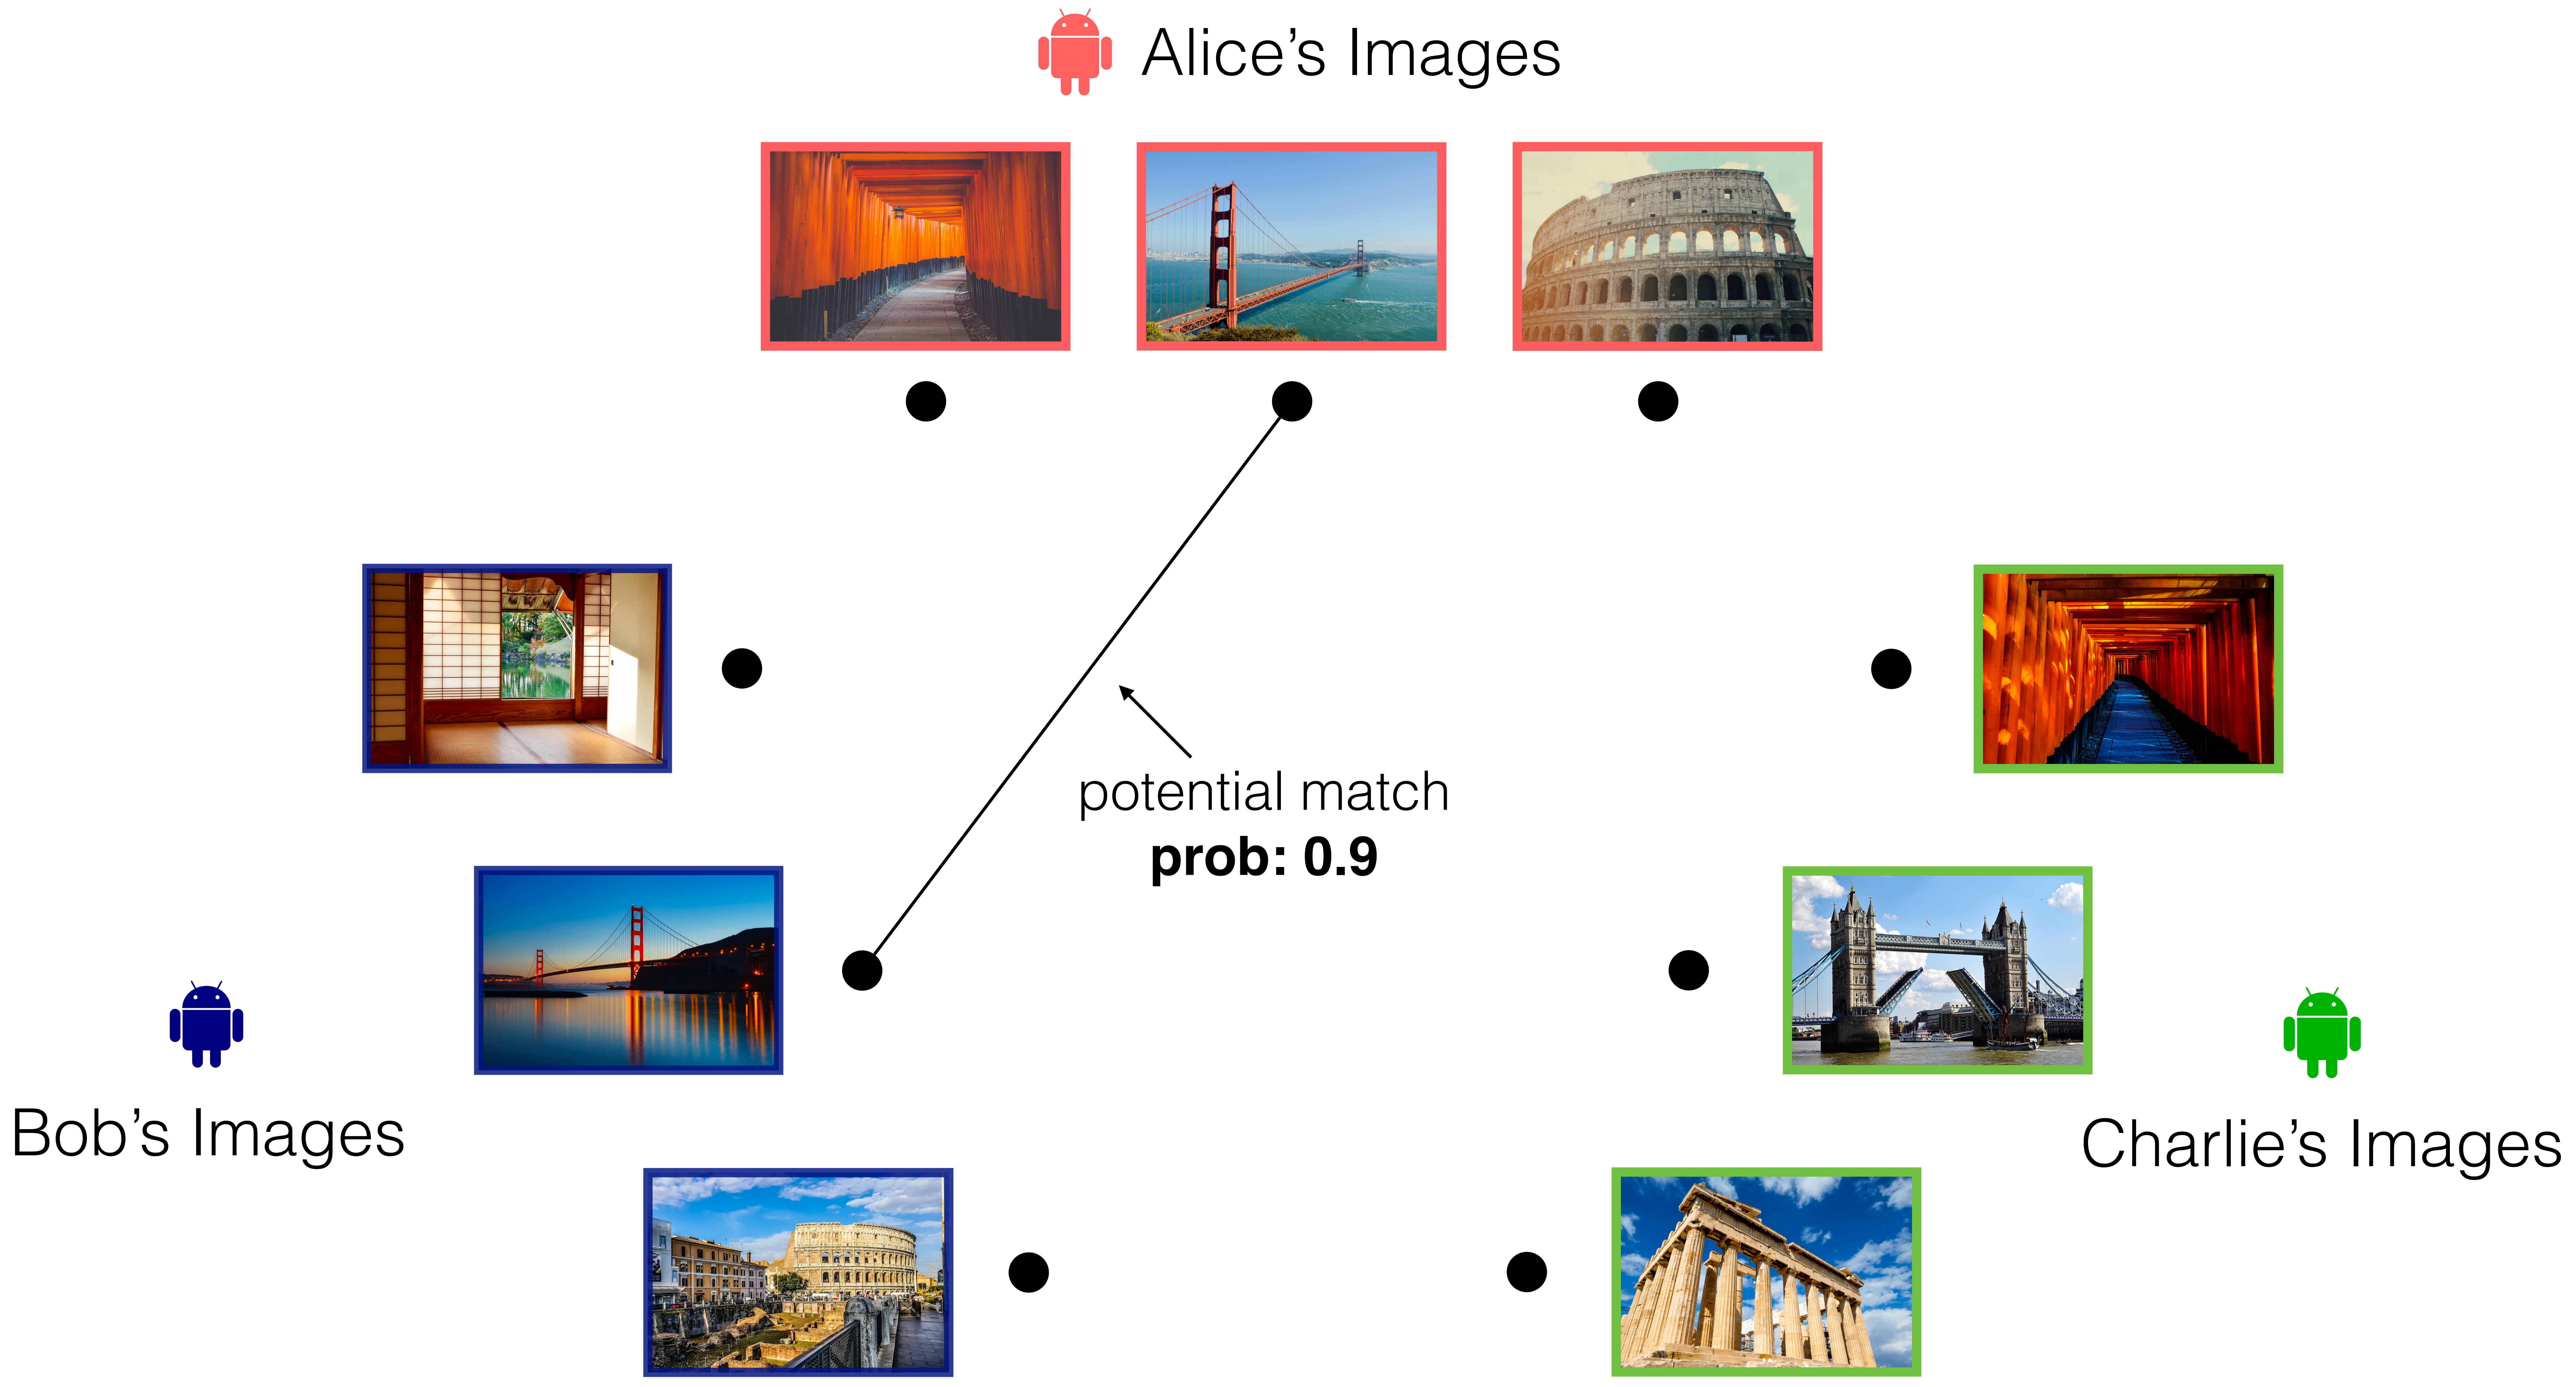
Alice's Images
potential match
prob: 0.9
Bob's Images
Charlie's Images

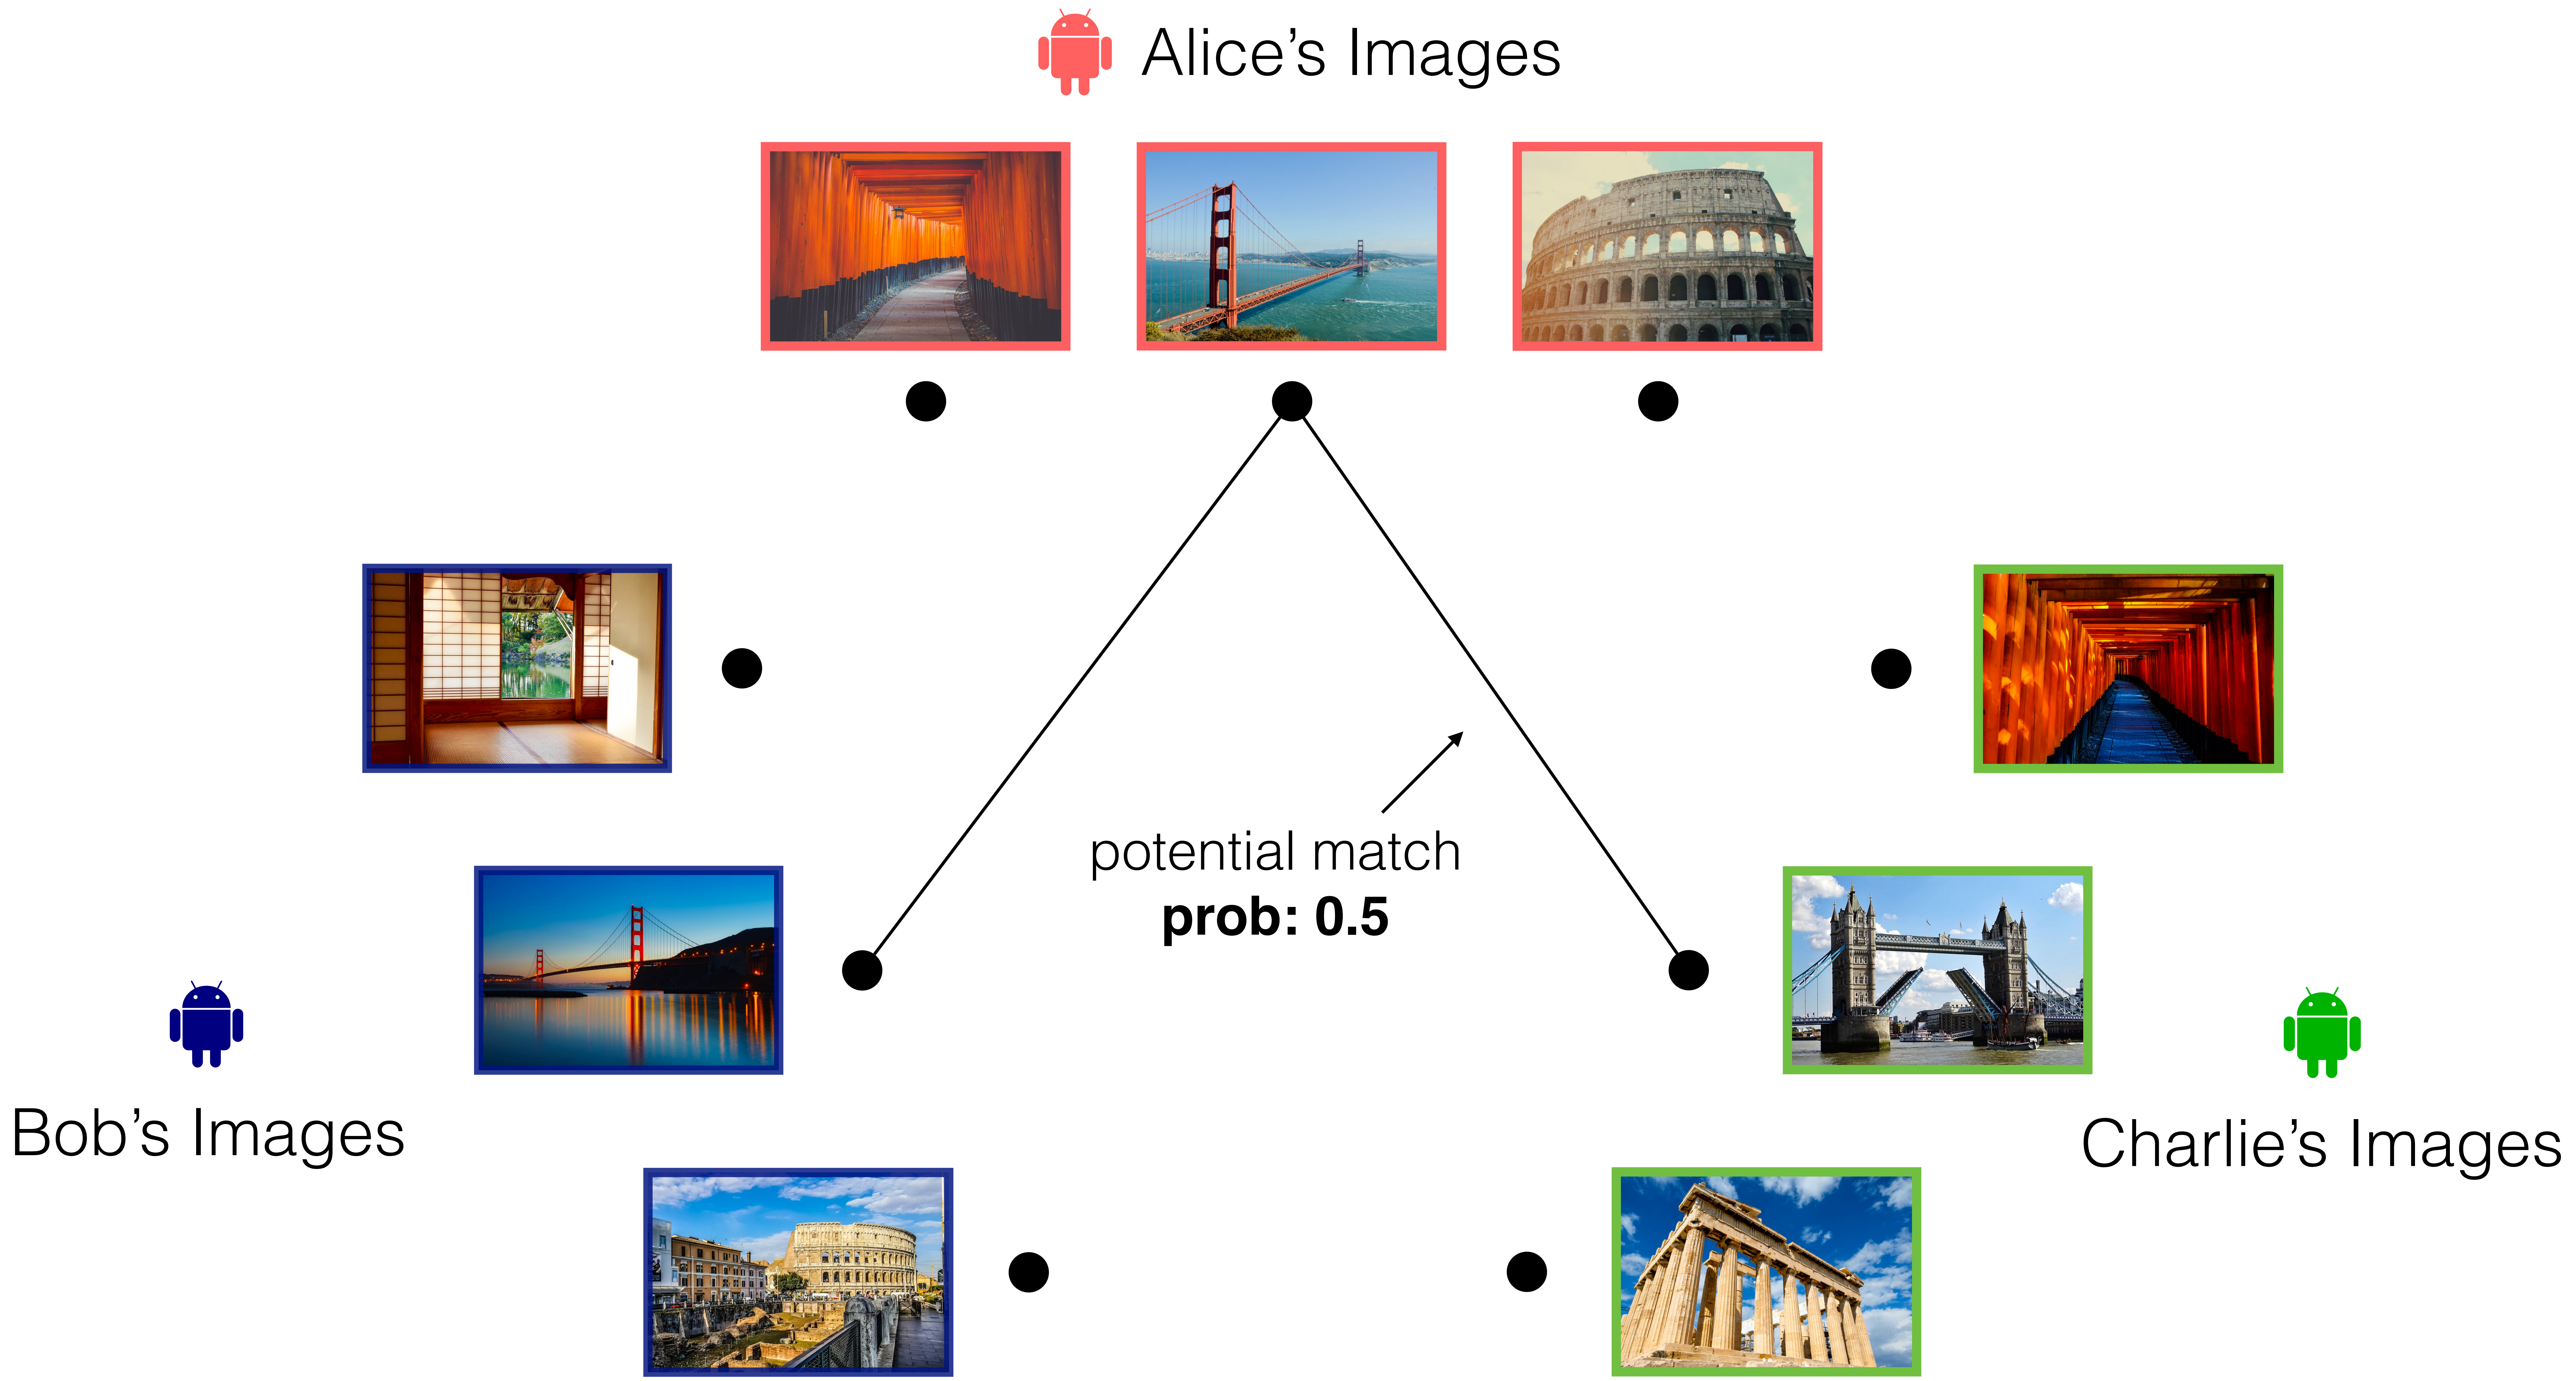
Alice's Images
potential match
prob: 0.5
Bob's Images
Charlie's Images

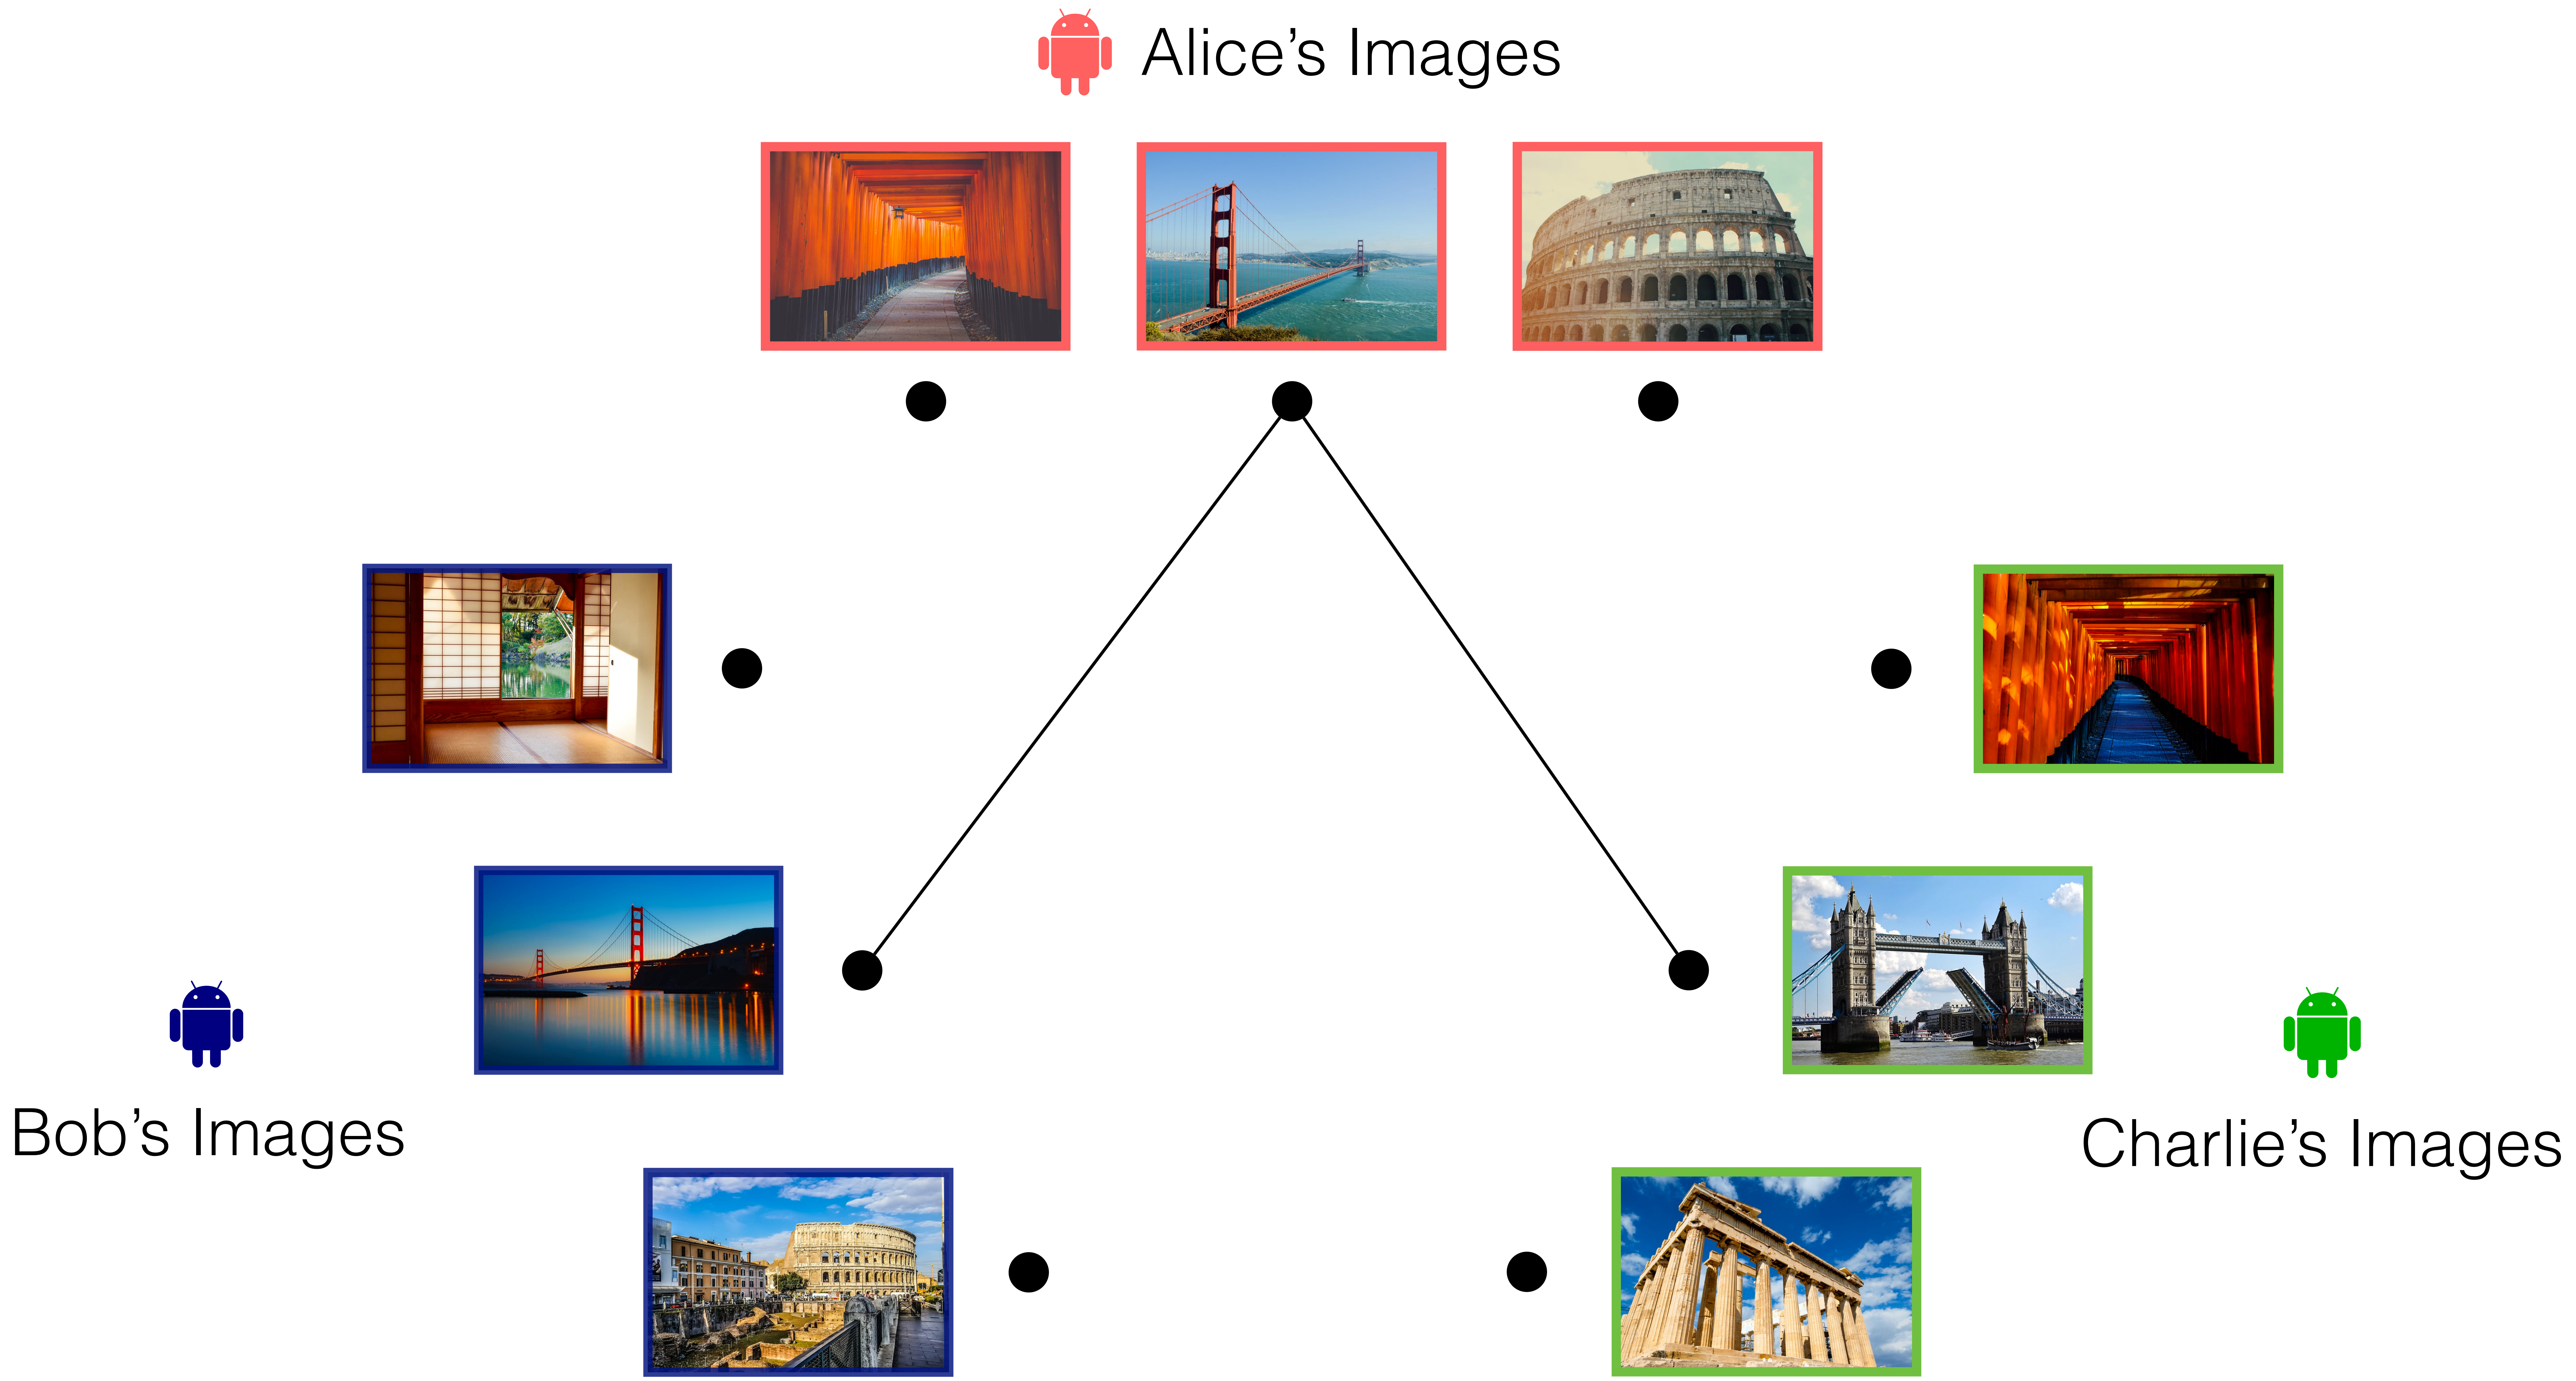
Alice’s Images
Bob’s Images
Charlie’s Images

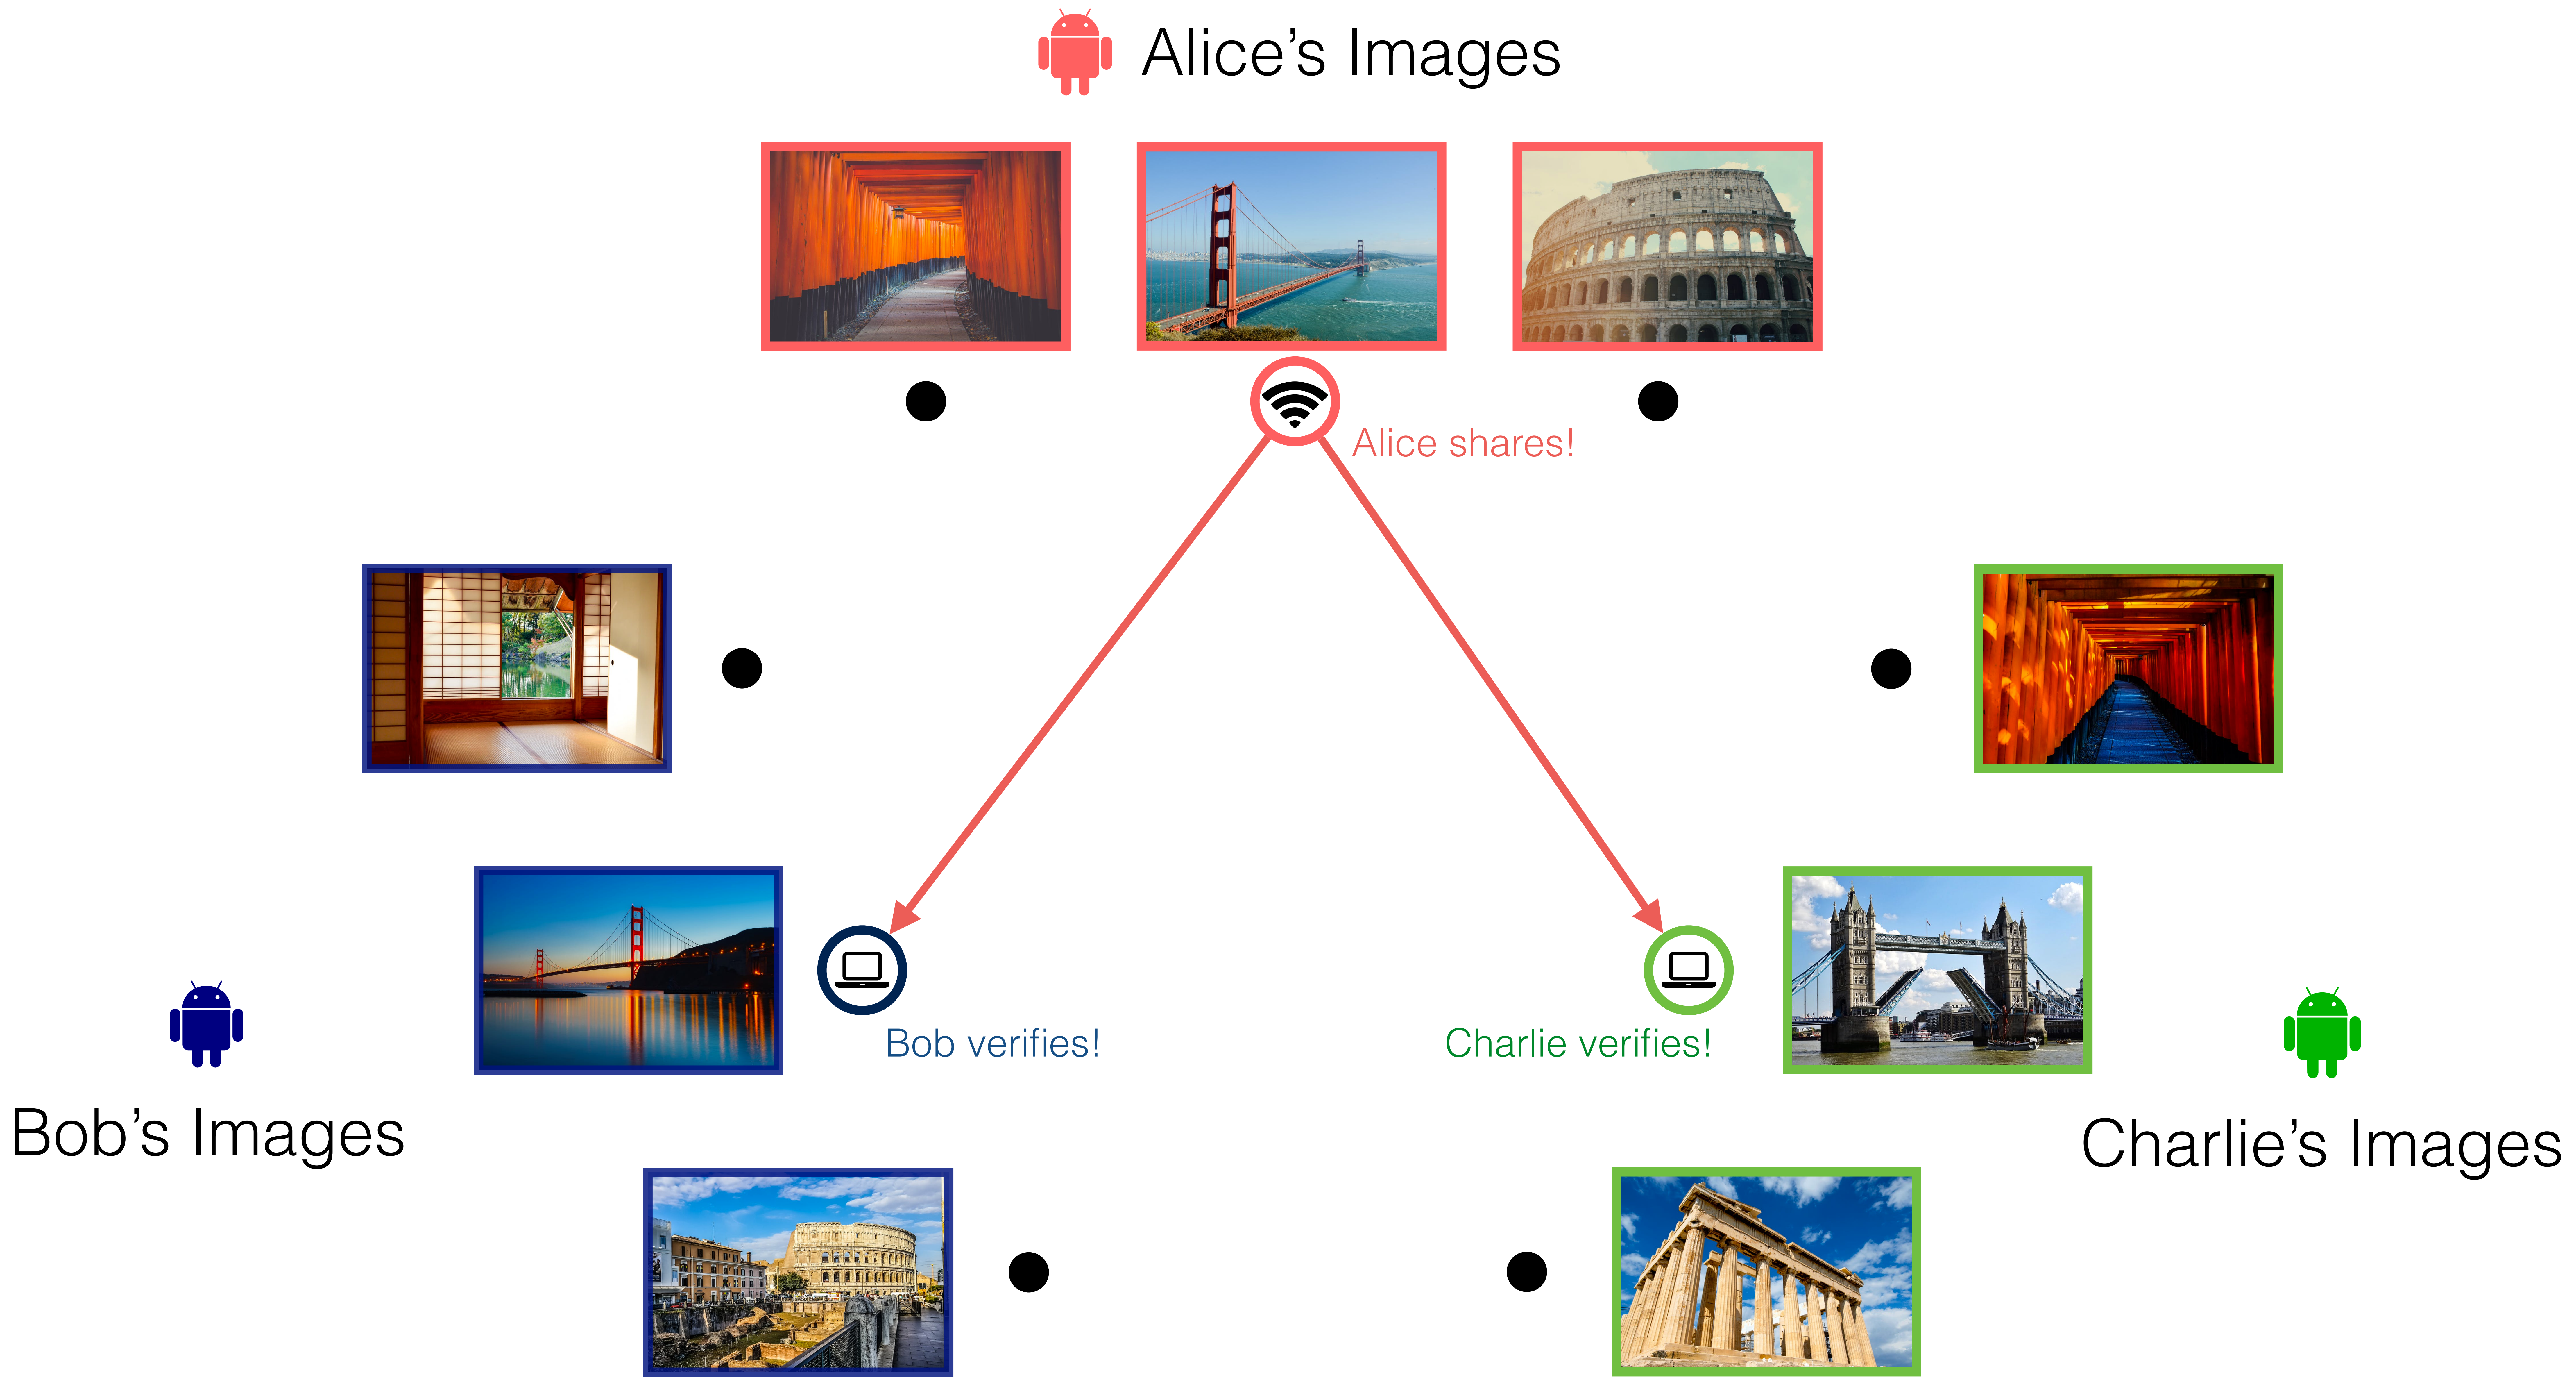

Alice's Images
Alice shares!
Bob verifies!
Charlie verifies!
Bob's Images
Charlie's Images

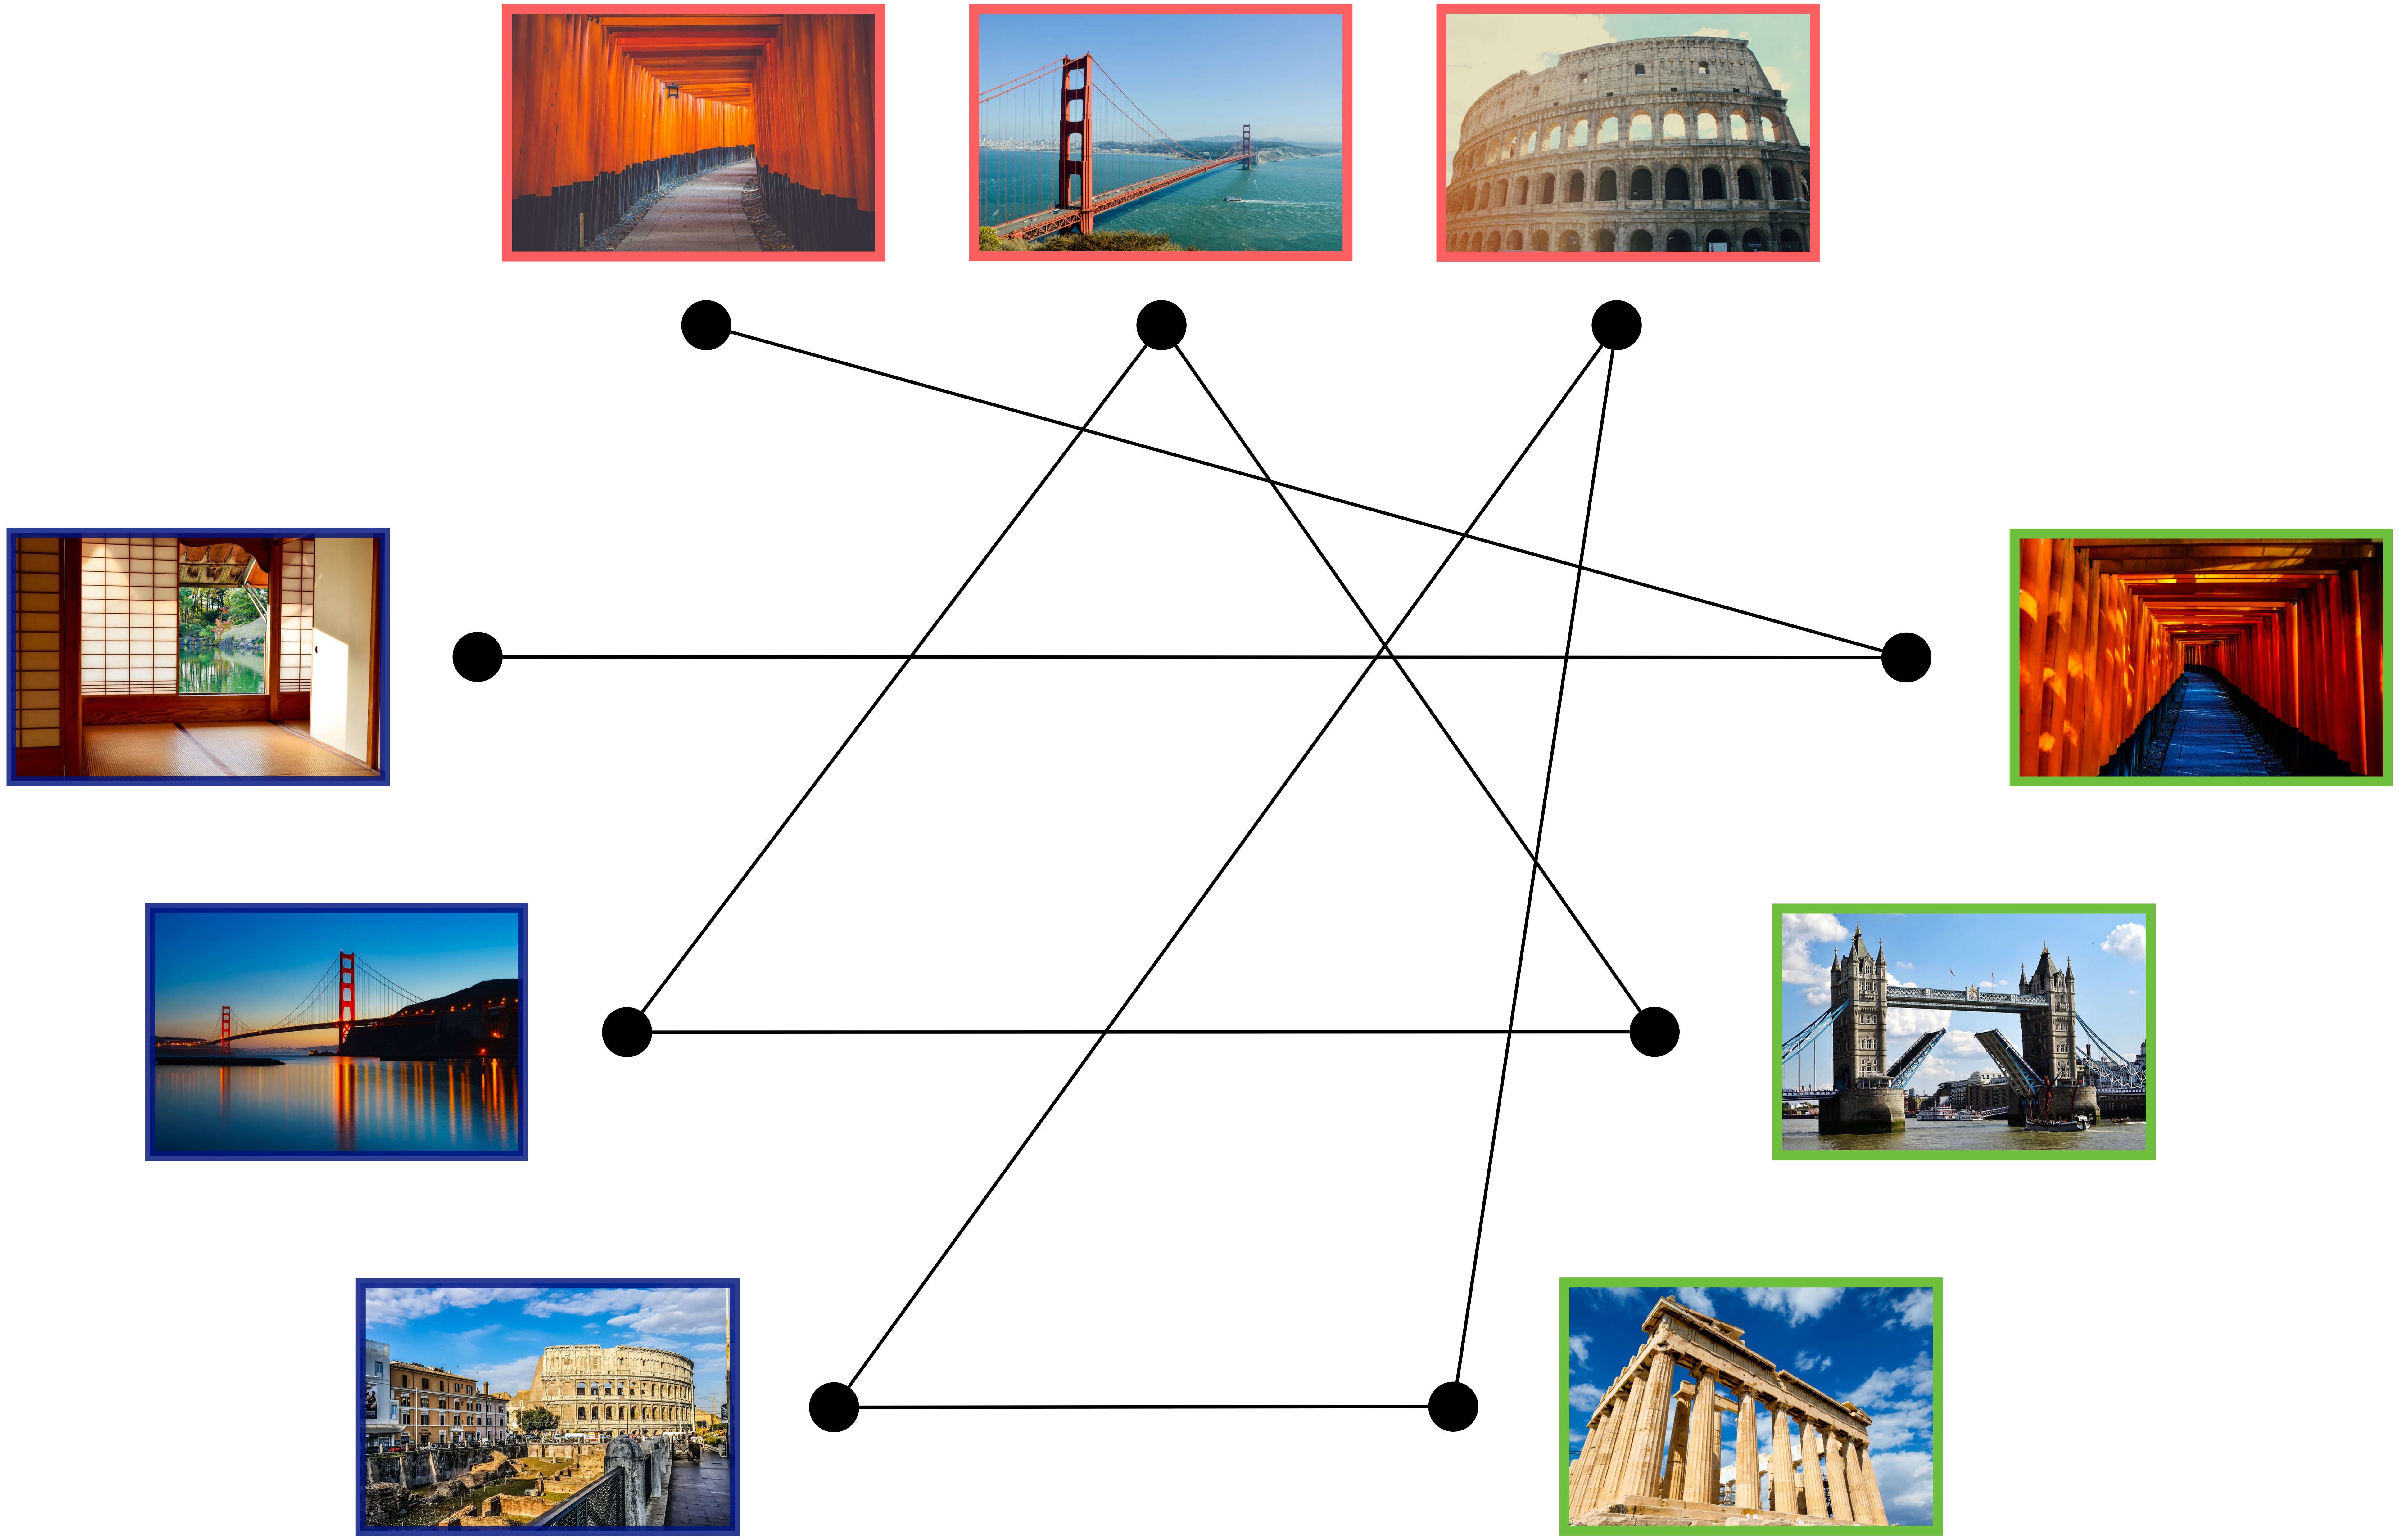

Budgeted Loop Closure Detection Problem

Expected "**information**" gain
(monotone & submodular)

$$
\begin{aligned}
\max_{\mathcal{L} \subseteq \mathcal{L}_{\mathrm{all}}} \quad & f(\mathcal{L}) \\
\mathrm{s.t.} \quad & \mathrm{cost}(\mathcal{L}) \leq b, \\
& \mathrm{cost}(\mathcal{L}) \leq k.
\end{aligned}
$$

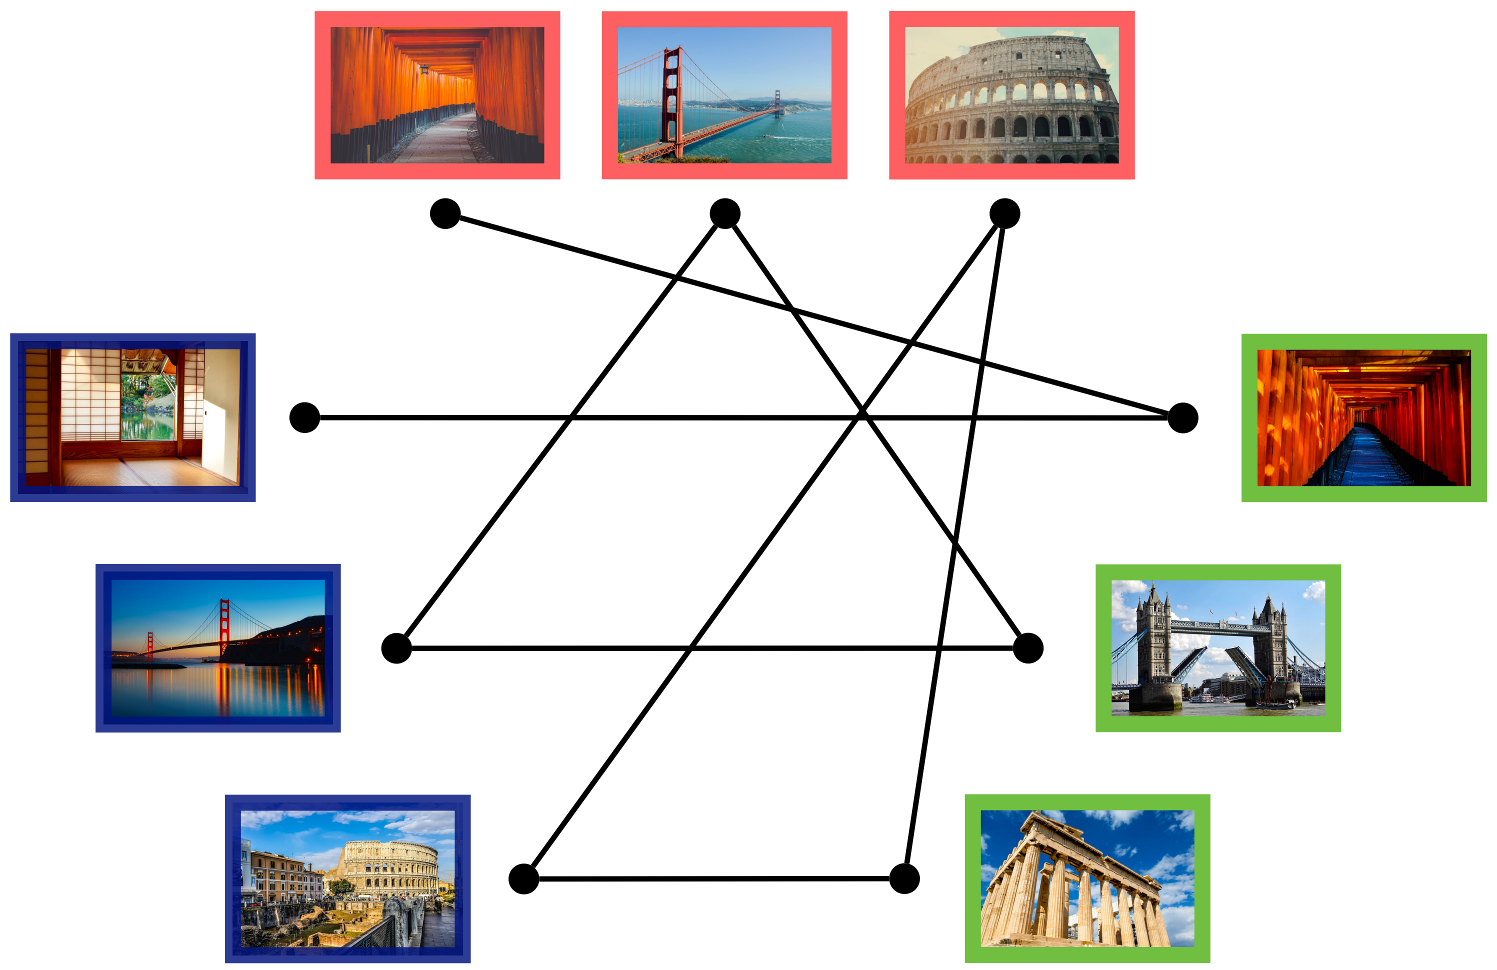

Budgeted Loop Closure Detection Problem

$$
\begin{aligned}
\max_{\mathcal{L} \subseteq \mathcal{L}_{\mathrm{all}},\, \mathcal{V} \subseteq \mathcal{V}_{\mathrm{all}}} \quad & f(\mathcal{L}) \\
\text{s.t.} \quad & |\mathcal{V}| \leq b, \\
& \mathrm{cost}(\mathcal{L}) \leq k.
\end{aligned}
$$

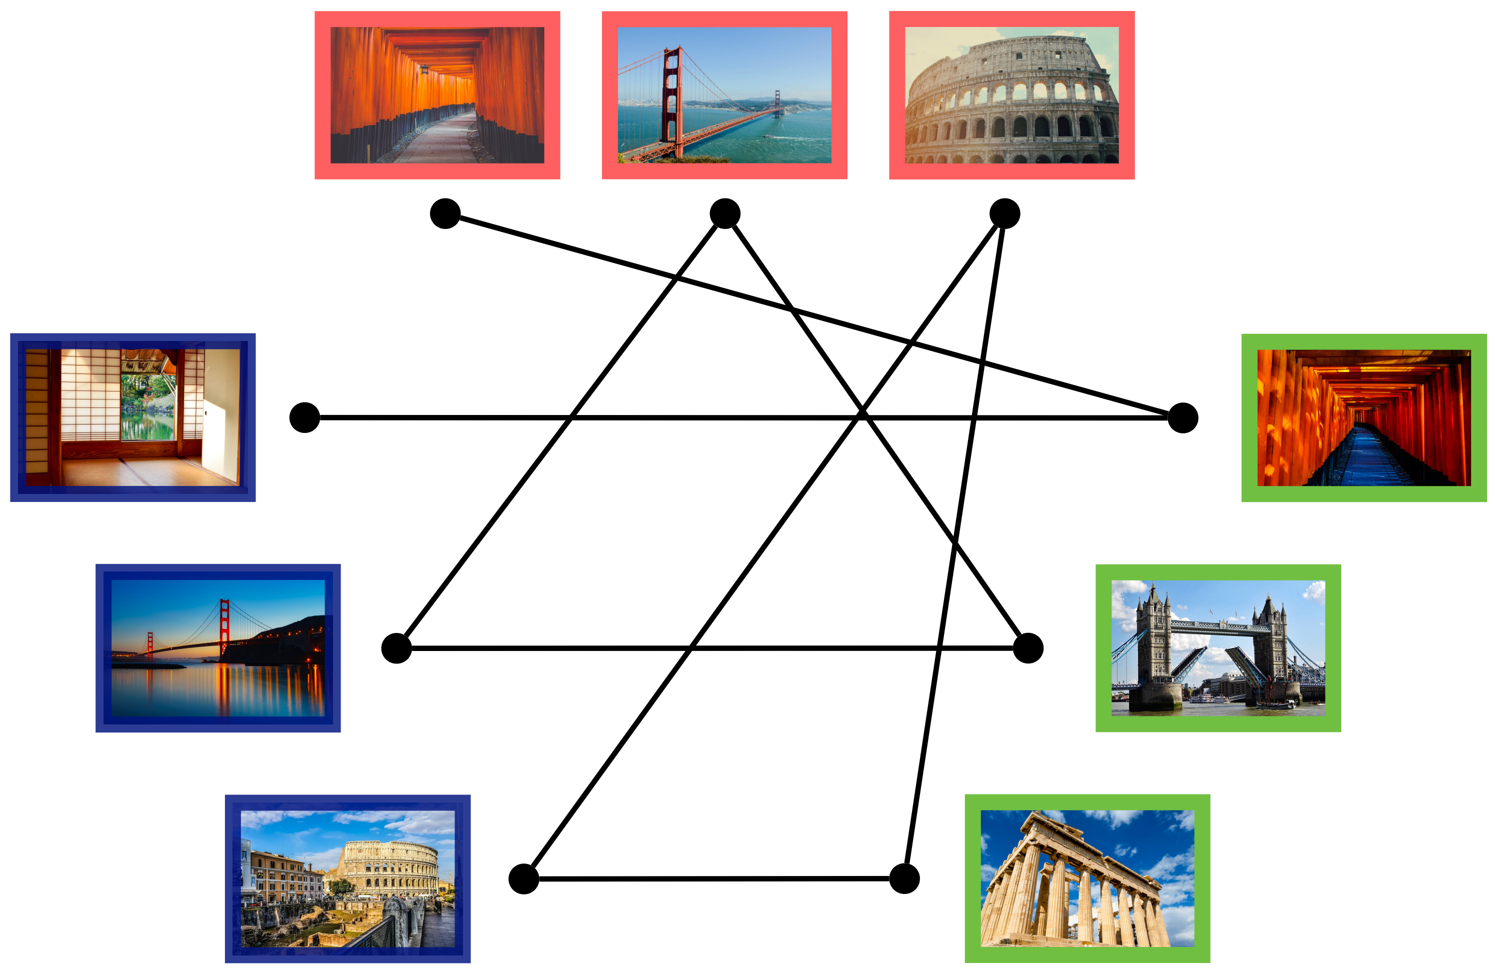

Budgeted Loop Closure Detection Problem

$$\max_{\mathcal{L}\subseteq\mathcal{L}_{\text{all}},\mathcal{V}\subseteq\mathcal{V}_{\text{all}}} \quad f(\mathcal{L})$$

$$\text{s.t.} \quad |\mathcal{V}| \leq b,$$

$$|\mathcal{L}| \leq k.$$

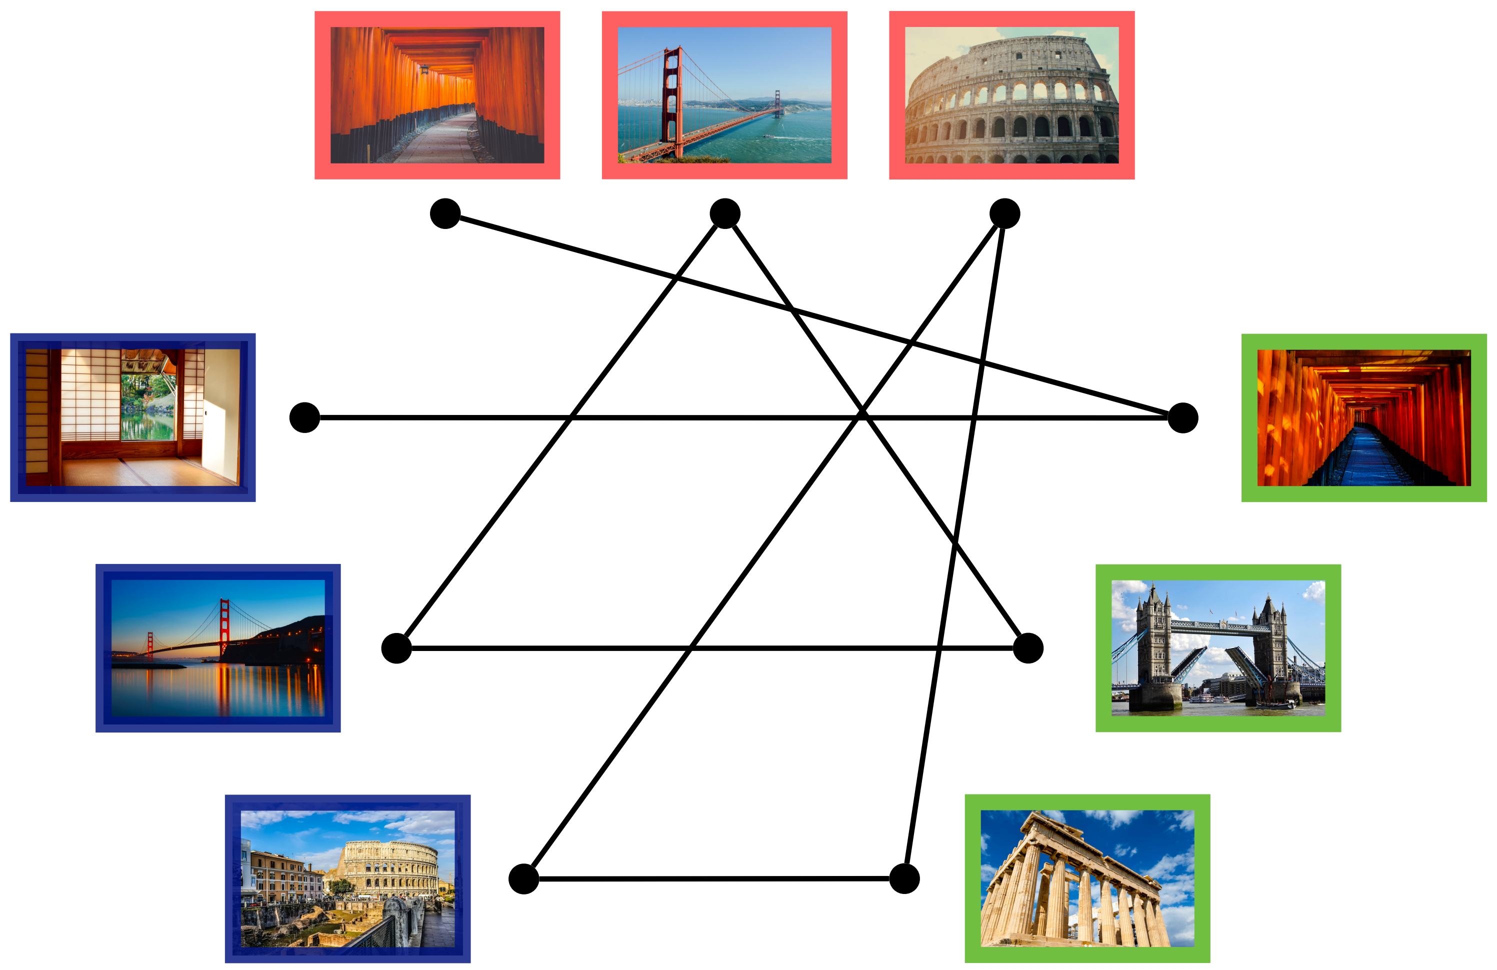

Budgeted Loop Closure Detection Problem

$$\max_{\mathcal{L}\subseteq\mathcal{L}_{\text{all}},\,\mathcal{V}\subseteq\mathcal{V}_{\text{all}}} \quad f(\mathcal{L})$$

NP-hard

need efficient **approximation** algorithm

$$\text{s.t.} \quad |\mathcal{V}| \leq b,$$

$$|\mathcal{L}| \leq k,$$

$$\mathcal{V} \text{ covers } \mathcal{L}.$$

at least one robot must share its measurement

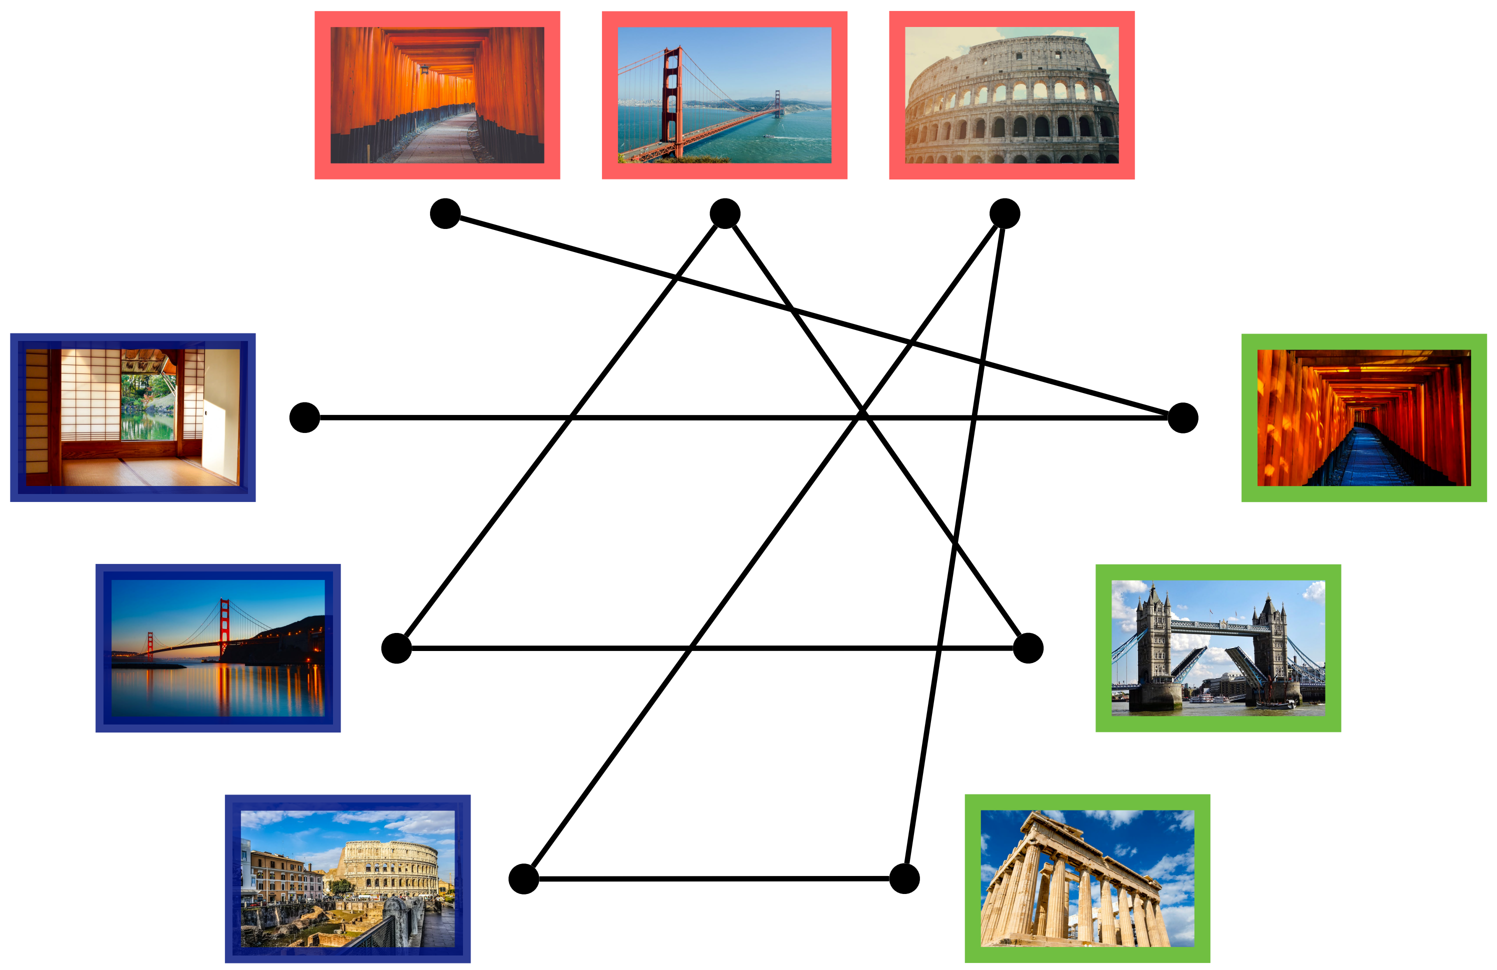

Case I: modular utility

$$
\begin{aligned}
\max_{\mathcal{L} \subseteq \mathcal{L}_{\text{all}}, \mathcal{V} \subseteq \mathcal{V}_{\text{all}}} \quad & \mathbb{E}[\# \text{ of true loop closures in } \mathcal{L}] \\
\text{s.t.} \quad & |\mathcal{V}| \leq b, \\
& |\mathcal{L}| \leq k, \\
& \mathcal{V} \text{ covers } \mathcal{L}.
\end{aligned}
$$

Case I: modular utility

$$
\begin{aligned}
\max_{\mathcal{L} \subseteq \mathcal{L}_{\text{all}}, \mathcal{V} \subseteq \mathcal{V}_{\text{all}}} \quad & \mathbb{E}[\#\text{ of true loop closures in } \mathcal{L}] \\
\text{s.t.} \quad & |\mathcal{V}| \leq b, \\
& |\mathcal{L}| \leq k, \\
& \mathcal{V} \text{ covers } \mathcal{L}.
\end{aligned}
$$

Theorem 1 (constant factor approximation)
A natural greedy algorithm is a $(1 - 1/e)$-approximation.

Case II: submodular utility

$$
\begin{aligned}
\max_{\mathcal{L}\subseteq\mathcal{L}_{\text{all}},\mathcal{V}\subseteq\mathcal{V}_{\text{all}}} \quad & f(\mathcal{L}) \\
\text{s.t.} \quad & |\mathcal{V}| \leq b, \\
& |\mathcal{L}| \leq k, \\
& \mathcal{V} \text{ covers } \mathcal{L}.
\end{aligned}
$$

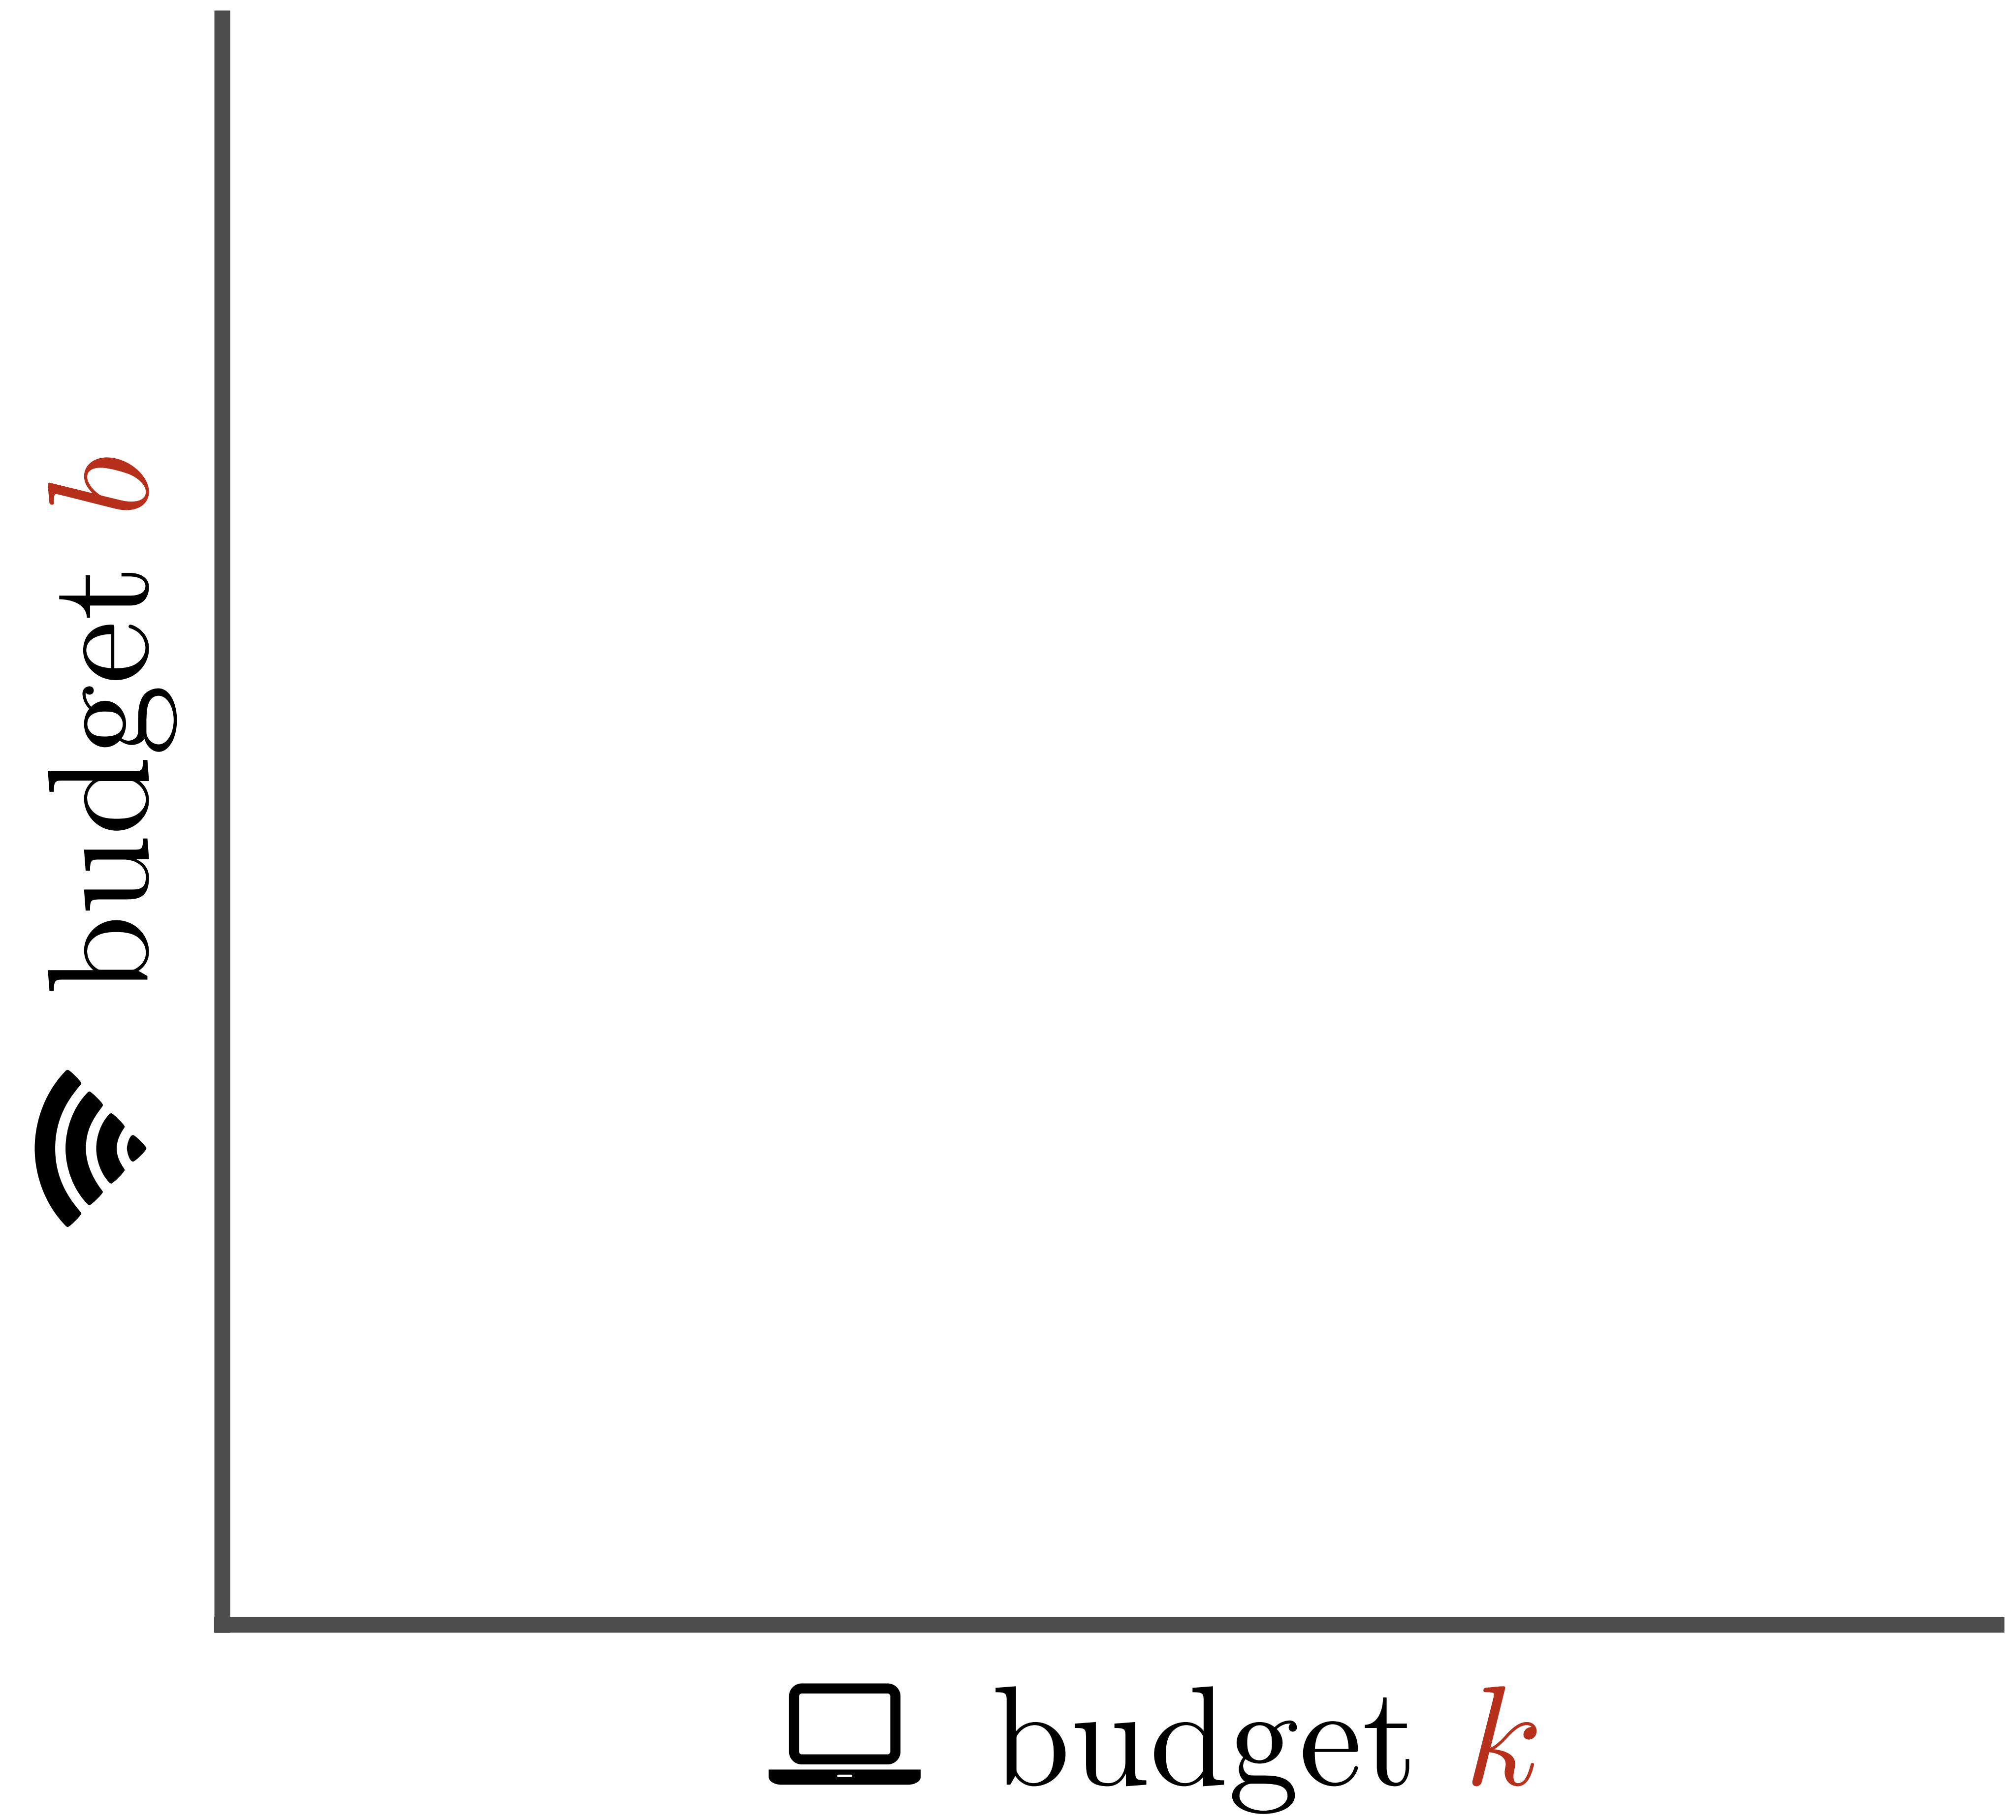

Case II: submodular utility

$$\max_{\mathcal{L}\subseteq\mathcal{L}_{\text{all}},\mathcal{V}\subseteq\mathcal{V}_{\text{all}}} \quad f(\mathcal{L})$$

$$\text{s.t.} \quad |\mathcal{V}| \leq b,$$

$$|\mathcal{L}| \leq k,$$

$$\mathcal{V} \text{ covers } \mathcal{L}.$$

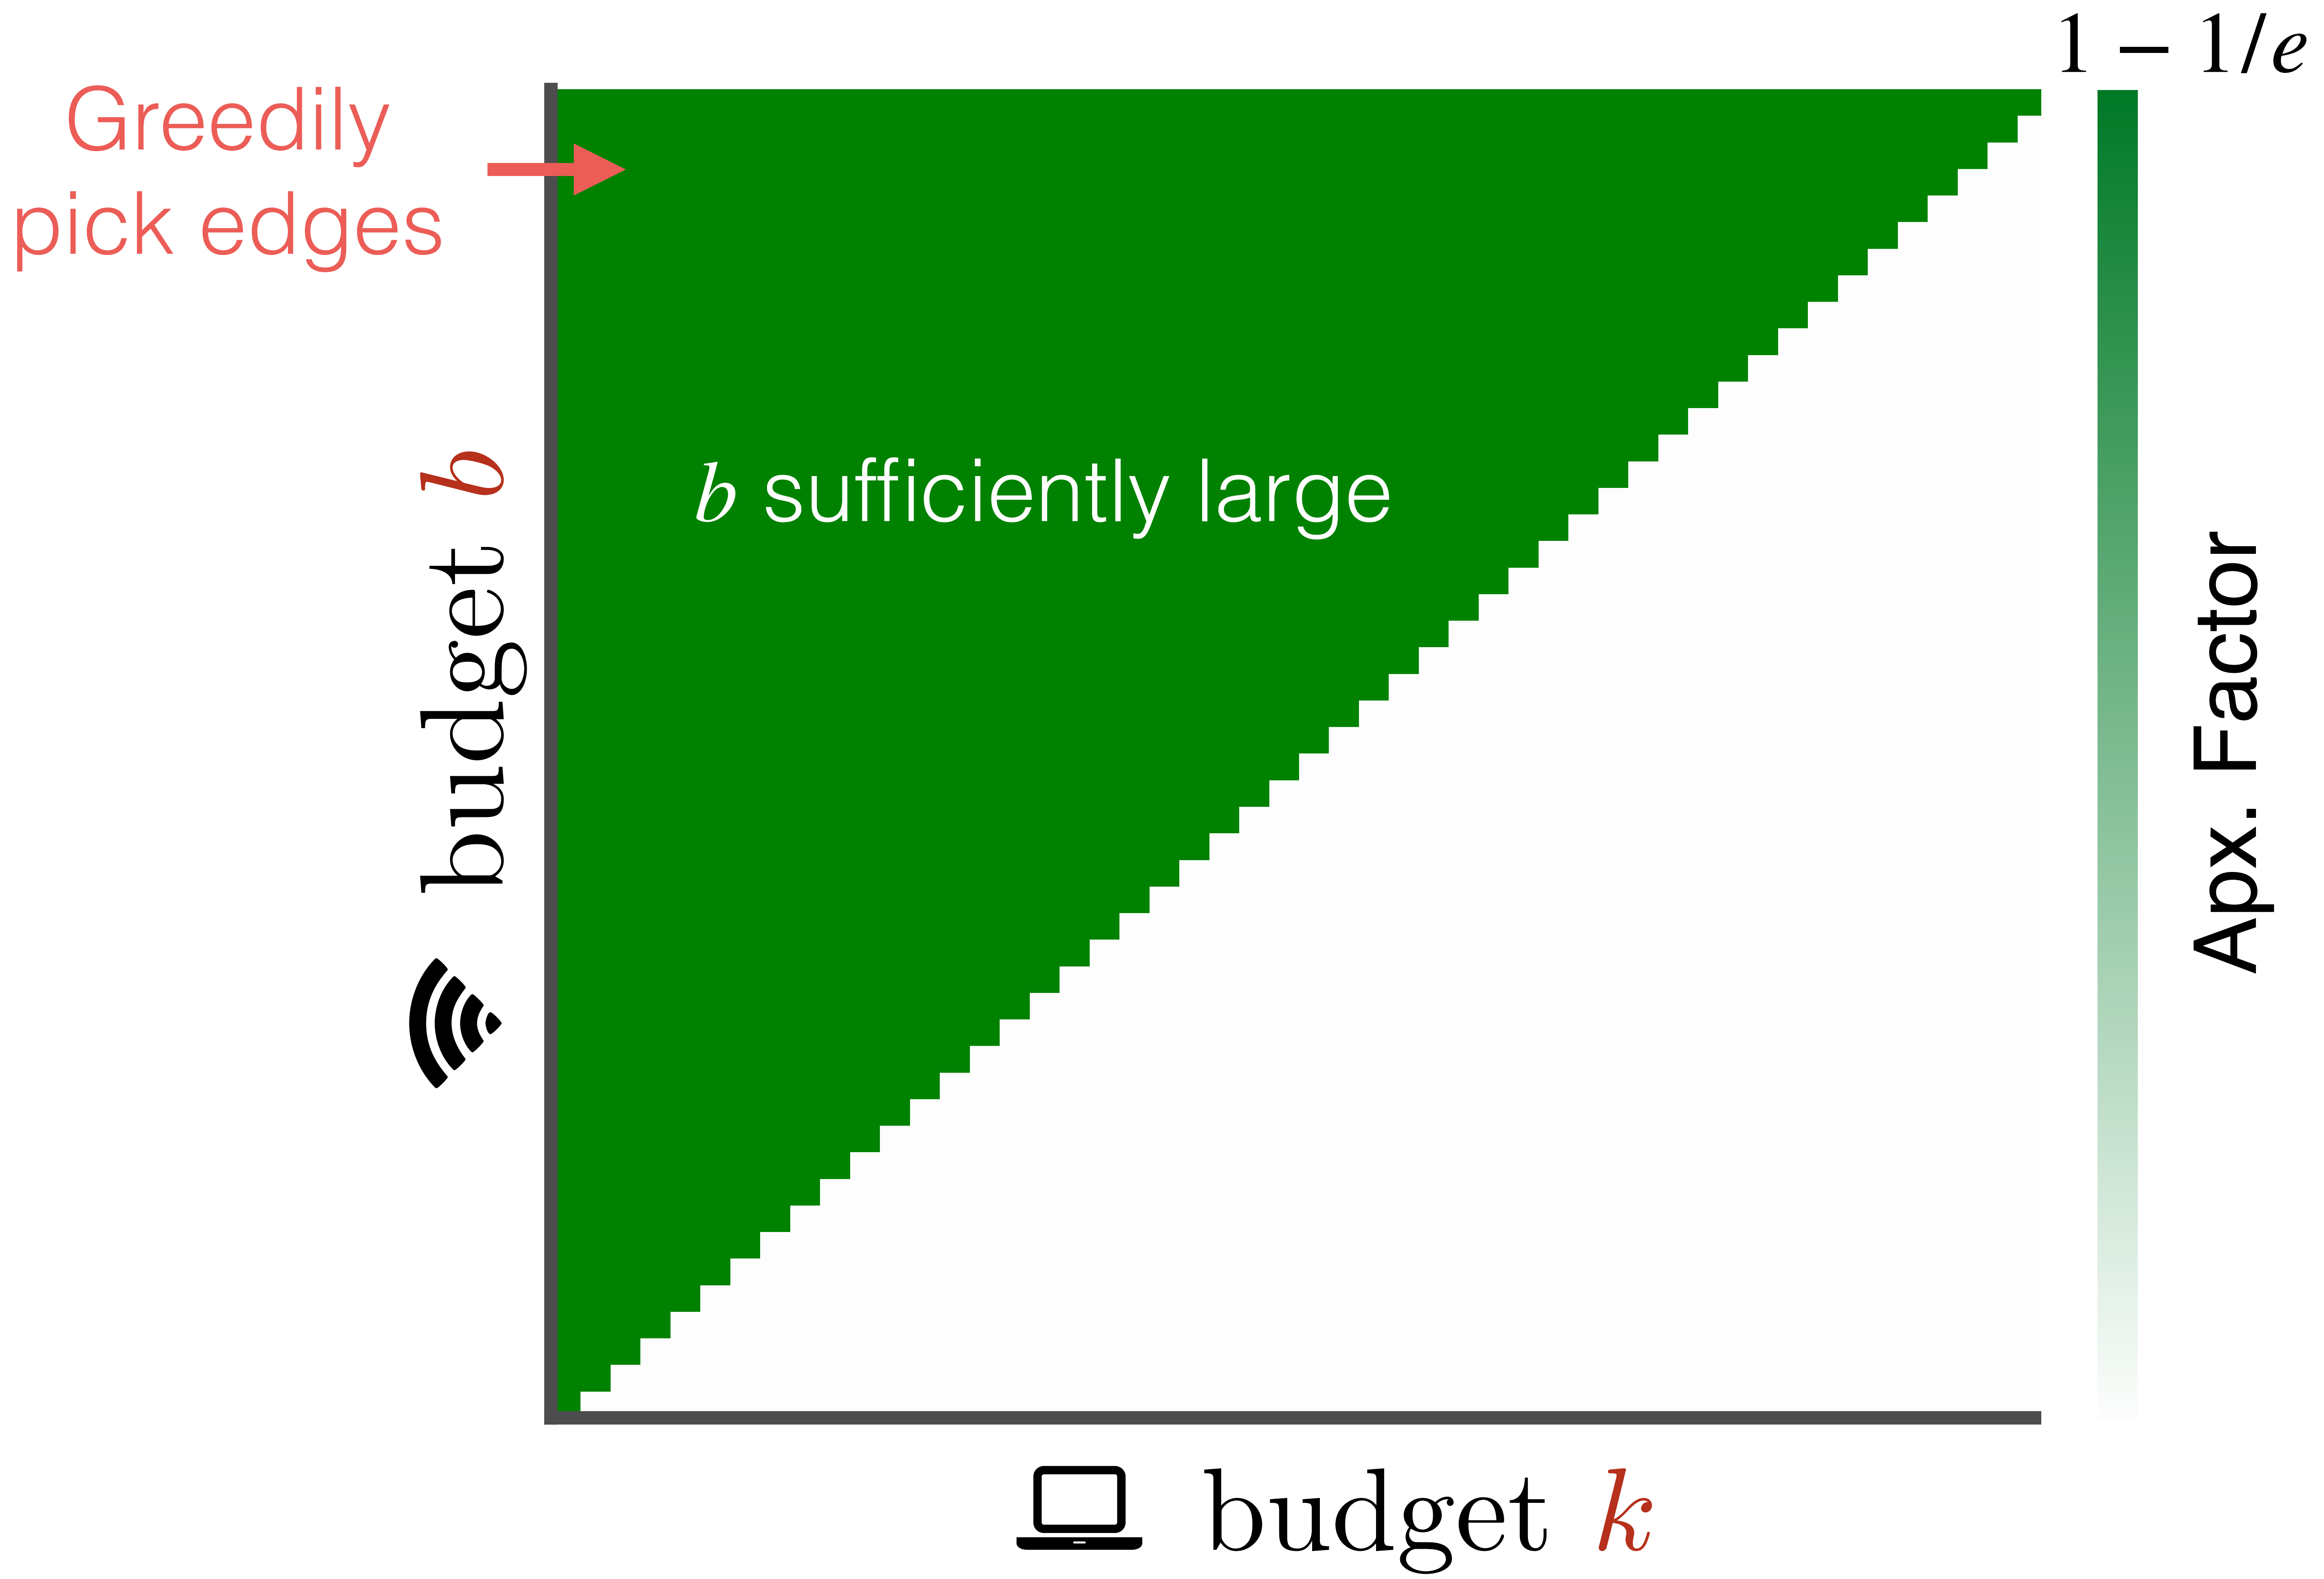

Case II: submodular utility

$$
\begin{aligned}
\max_{\mathcal{L}\subseteq\mathcal{L}_{\text{all}},\mathcal{V}\subseteq\mathcal{V}_{\text{all}}} \quad & f(\mathcal{L}) \\
\text{s.t.} \quad & |\mathcal{V}| \leq b, \\
& |\mathcal{L}| \leq k, \\
& \mathcal{V} \text{ covers } \mathcal{L}.
\end{aligned}
$$

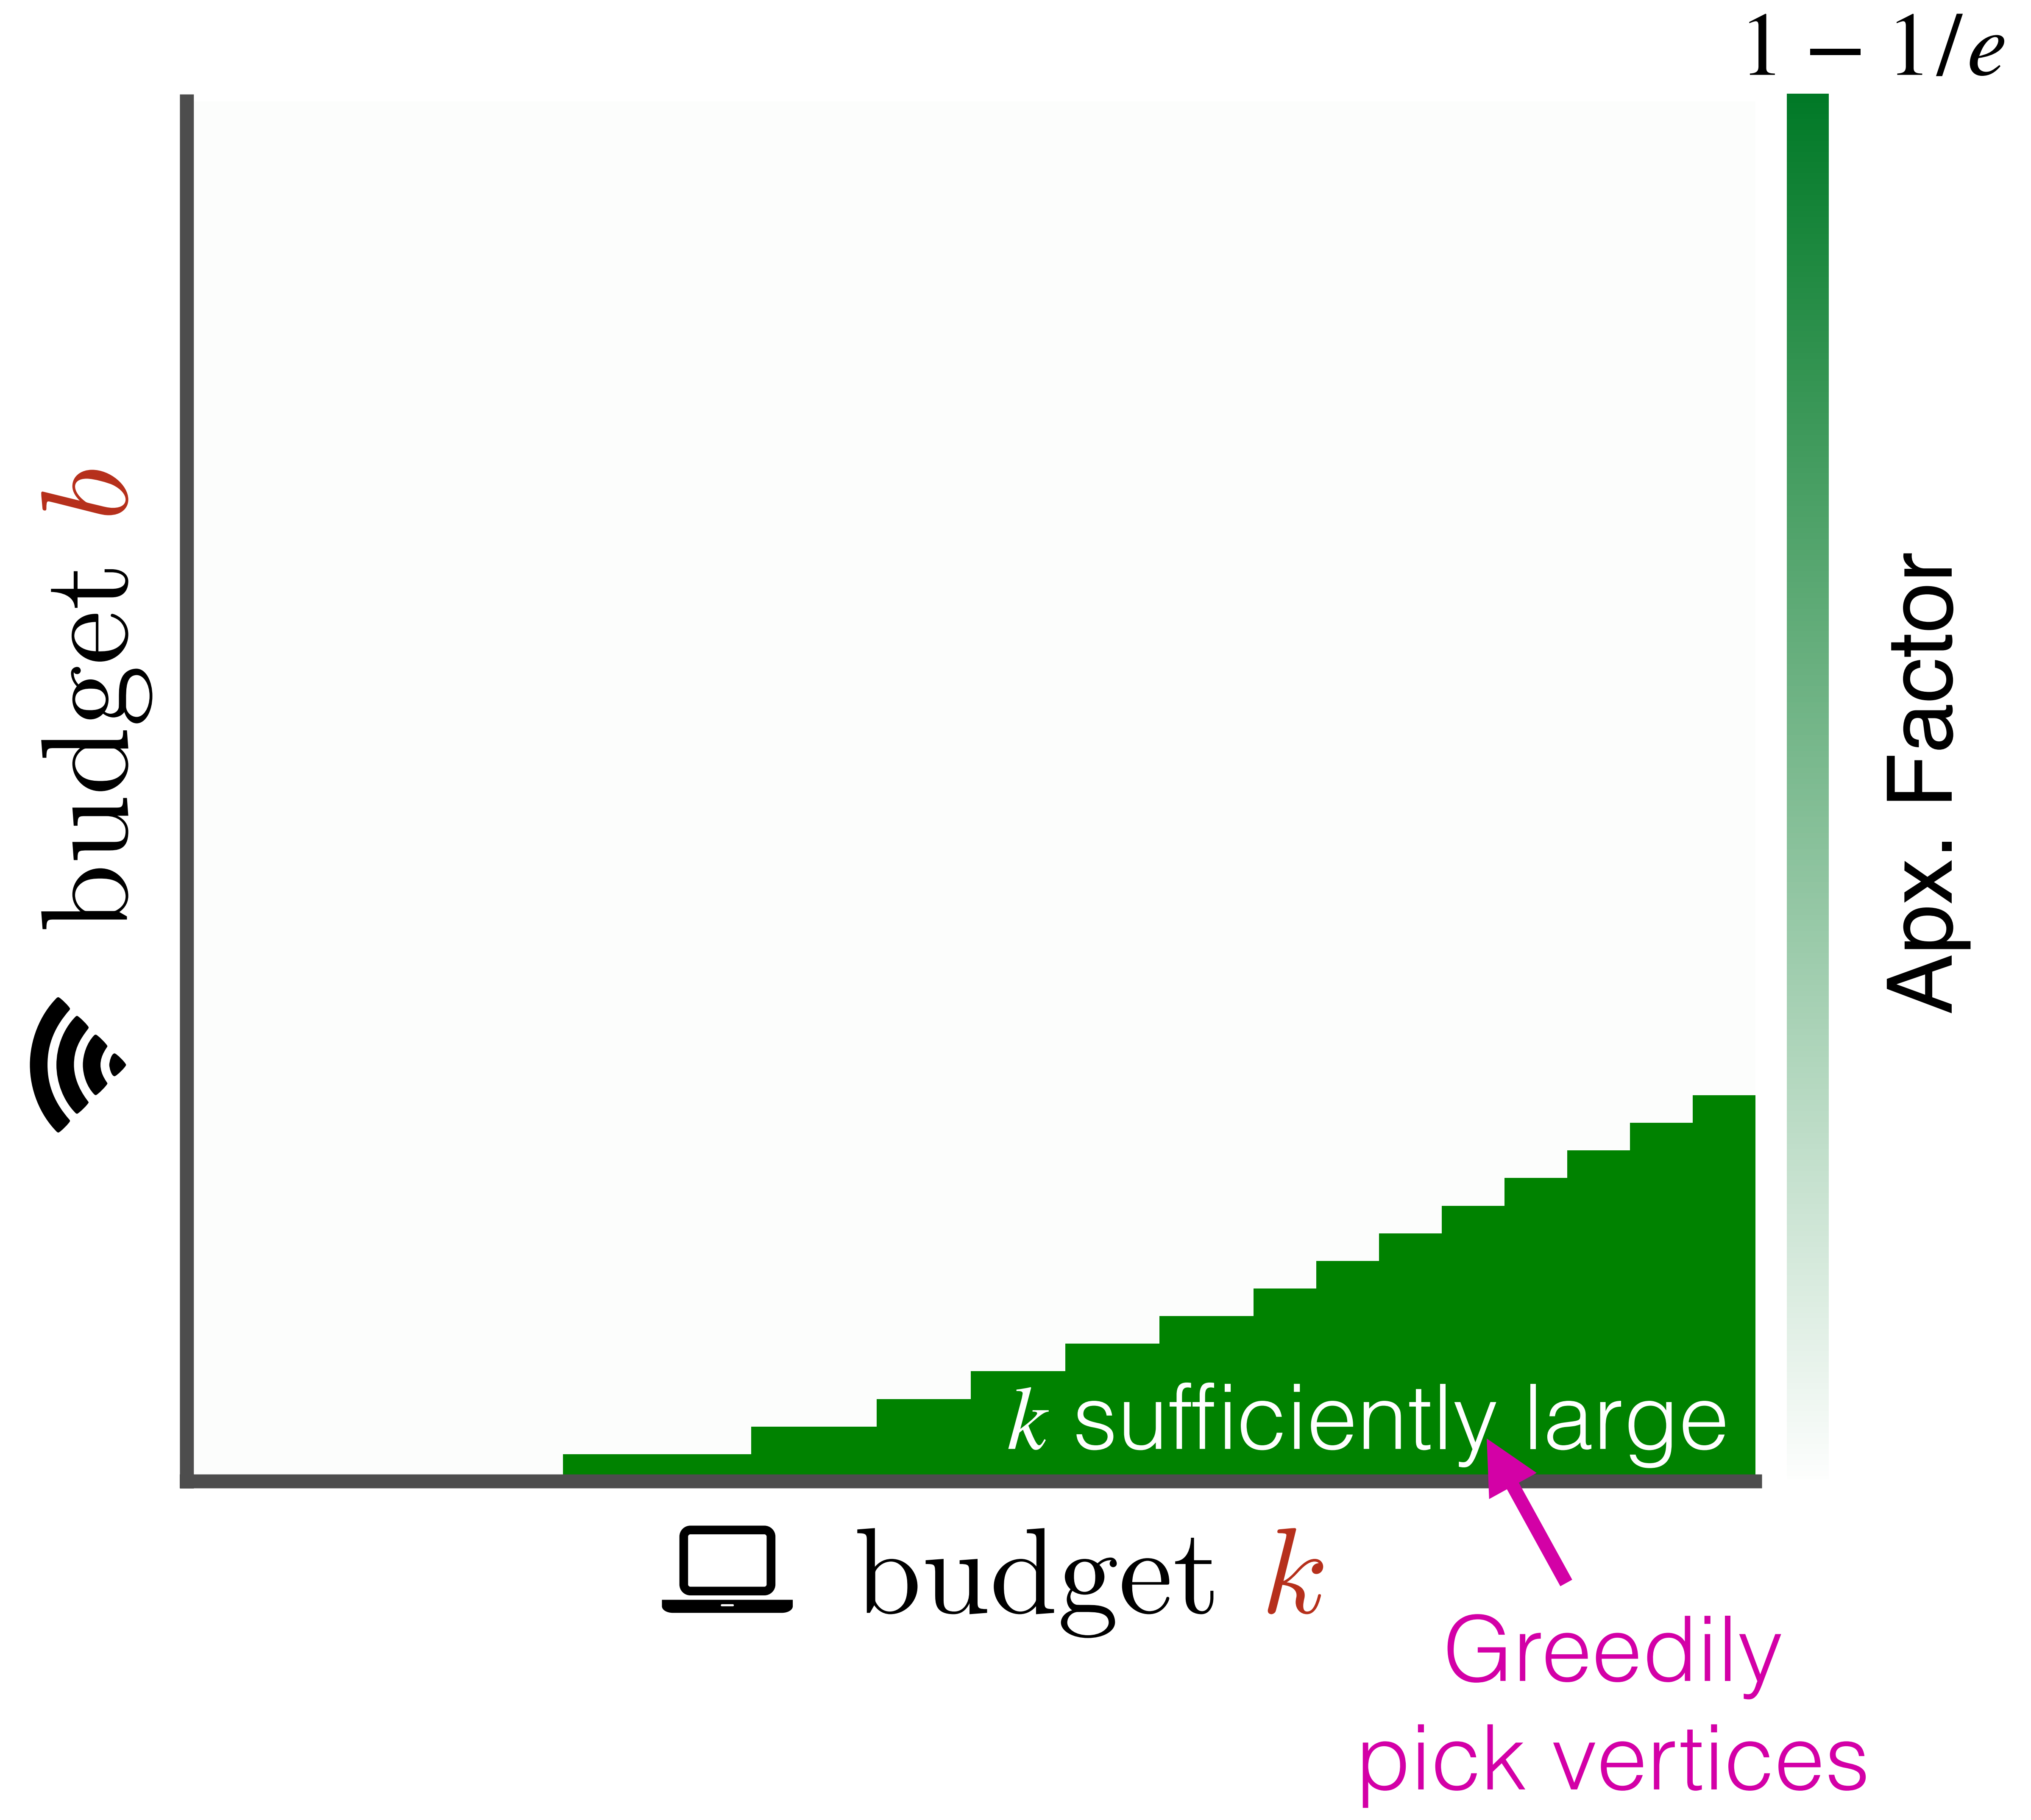

Near-Optimal Budgeted Data Exchange for Distributed Loop Closure Detection, Tian et al., RSS 2018

Case II: submodular utility

$$\max_{\mathcal{L}\subseteq\mathcal{L}_{\text{all}},\mathcal{V}\subseteq\mathcal{V}_{\text{all}}} \quad f(\mathcal{L})$$

$$\text{s.t.} \quad |\mathcal{V}| \le b, \quad |\mathcal{L}| \le k, \quad \mathcal{V} \text{ covers } \mathcal{L}.$$

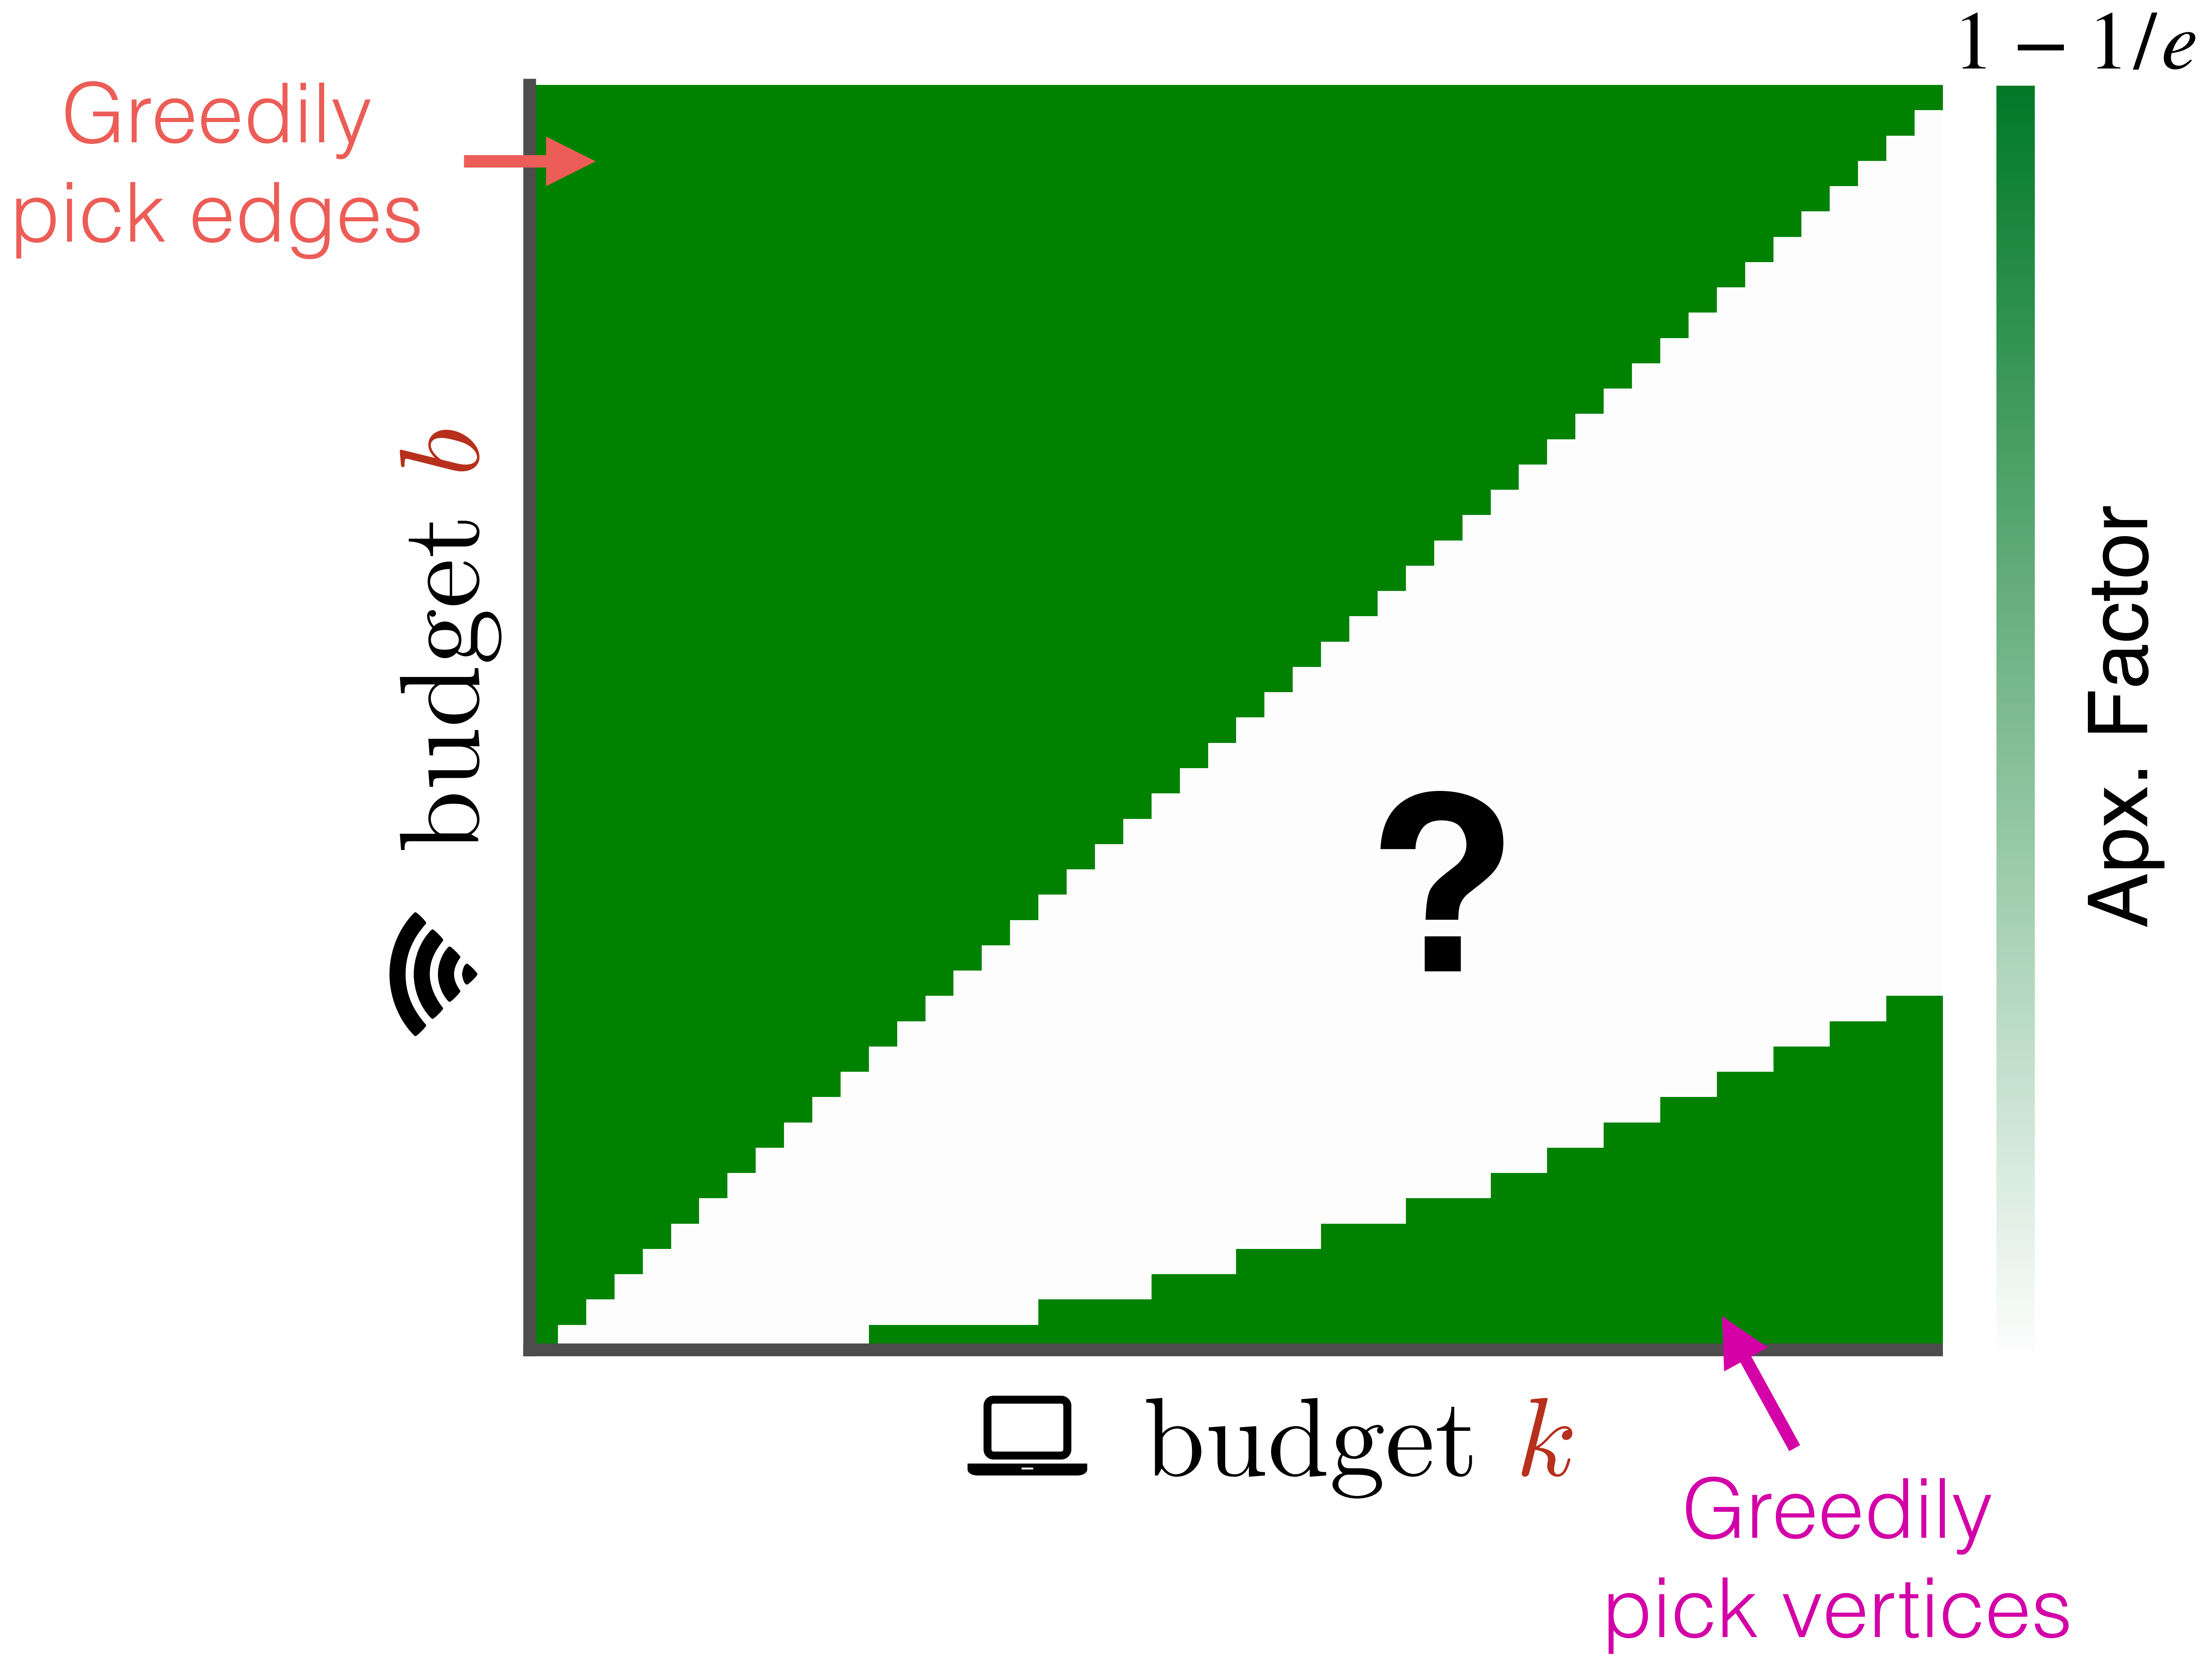

Case II: submodular utility

Theorem 2 (α-approximation)
The better of vertex/edge greedy is an α-approximation where
$\alpha = 1 - \exp\big(-\min\{1, \gamma\}\big)$
in which $\gamma = \max\{b/k, \lfloor k/\Delta \rfloor/b\}$.

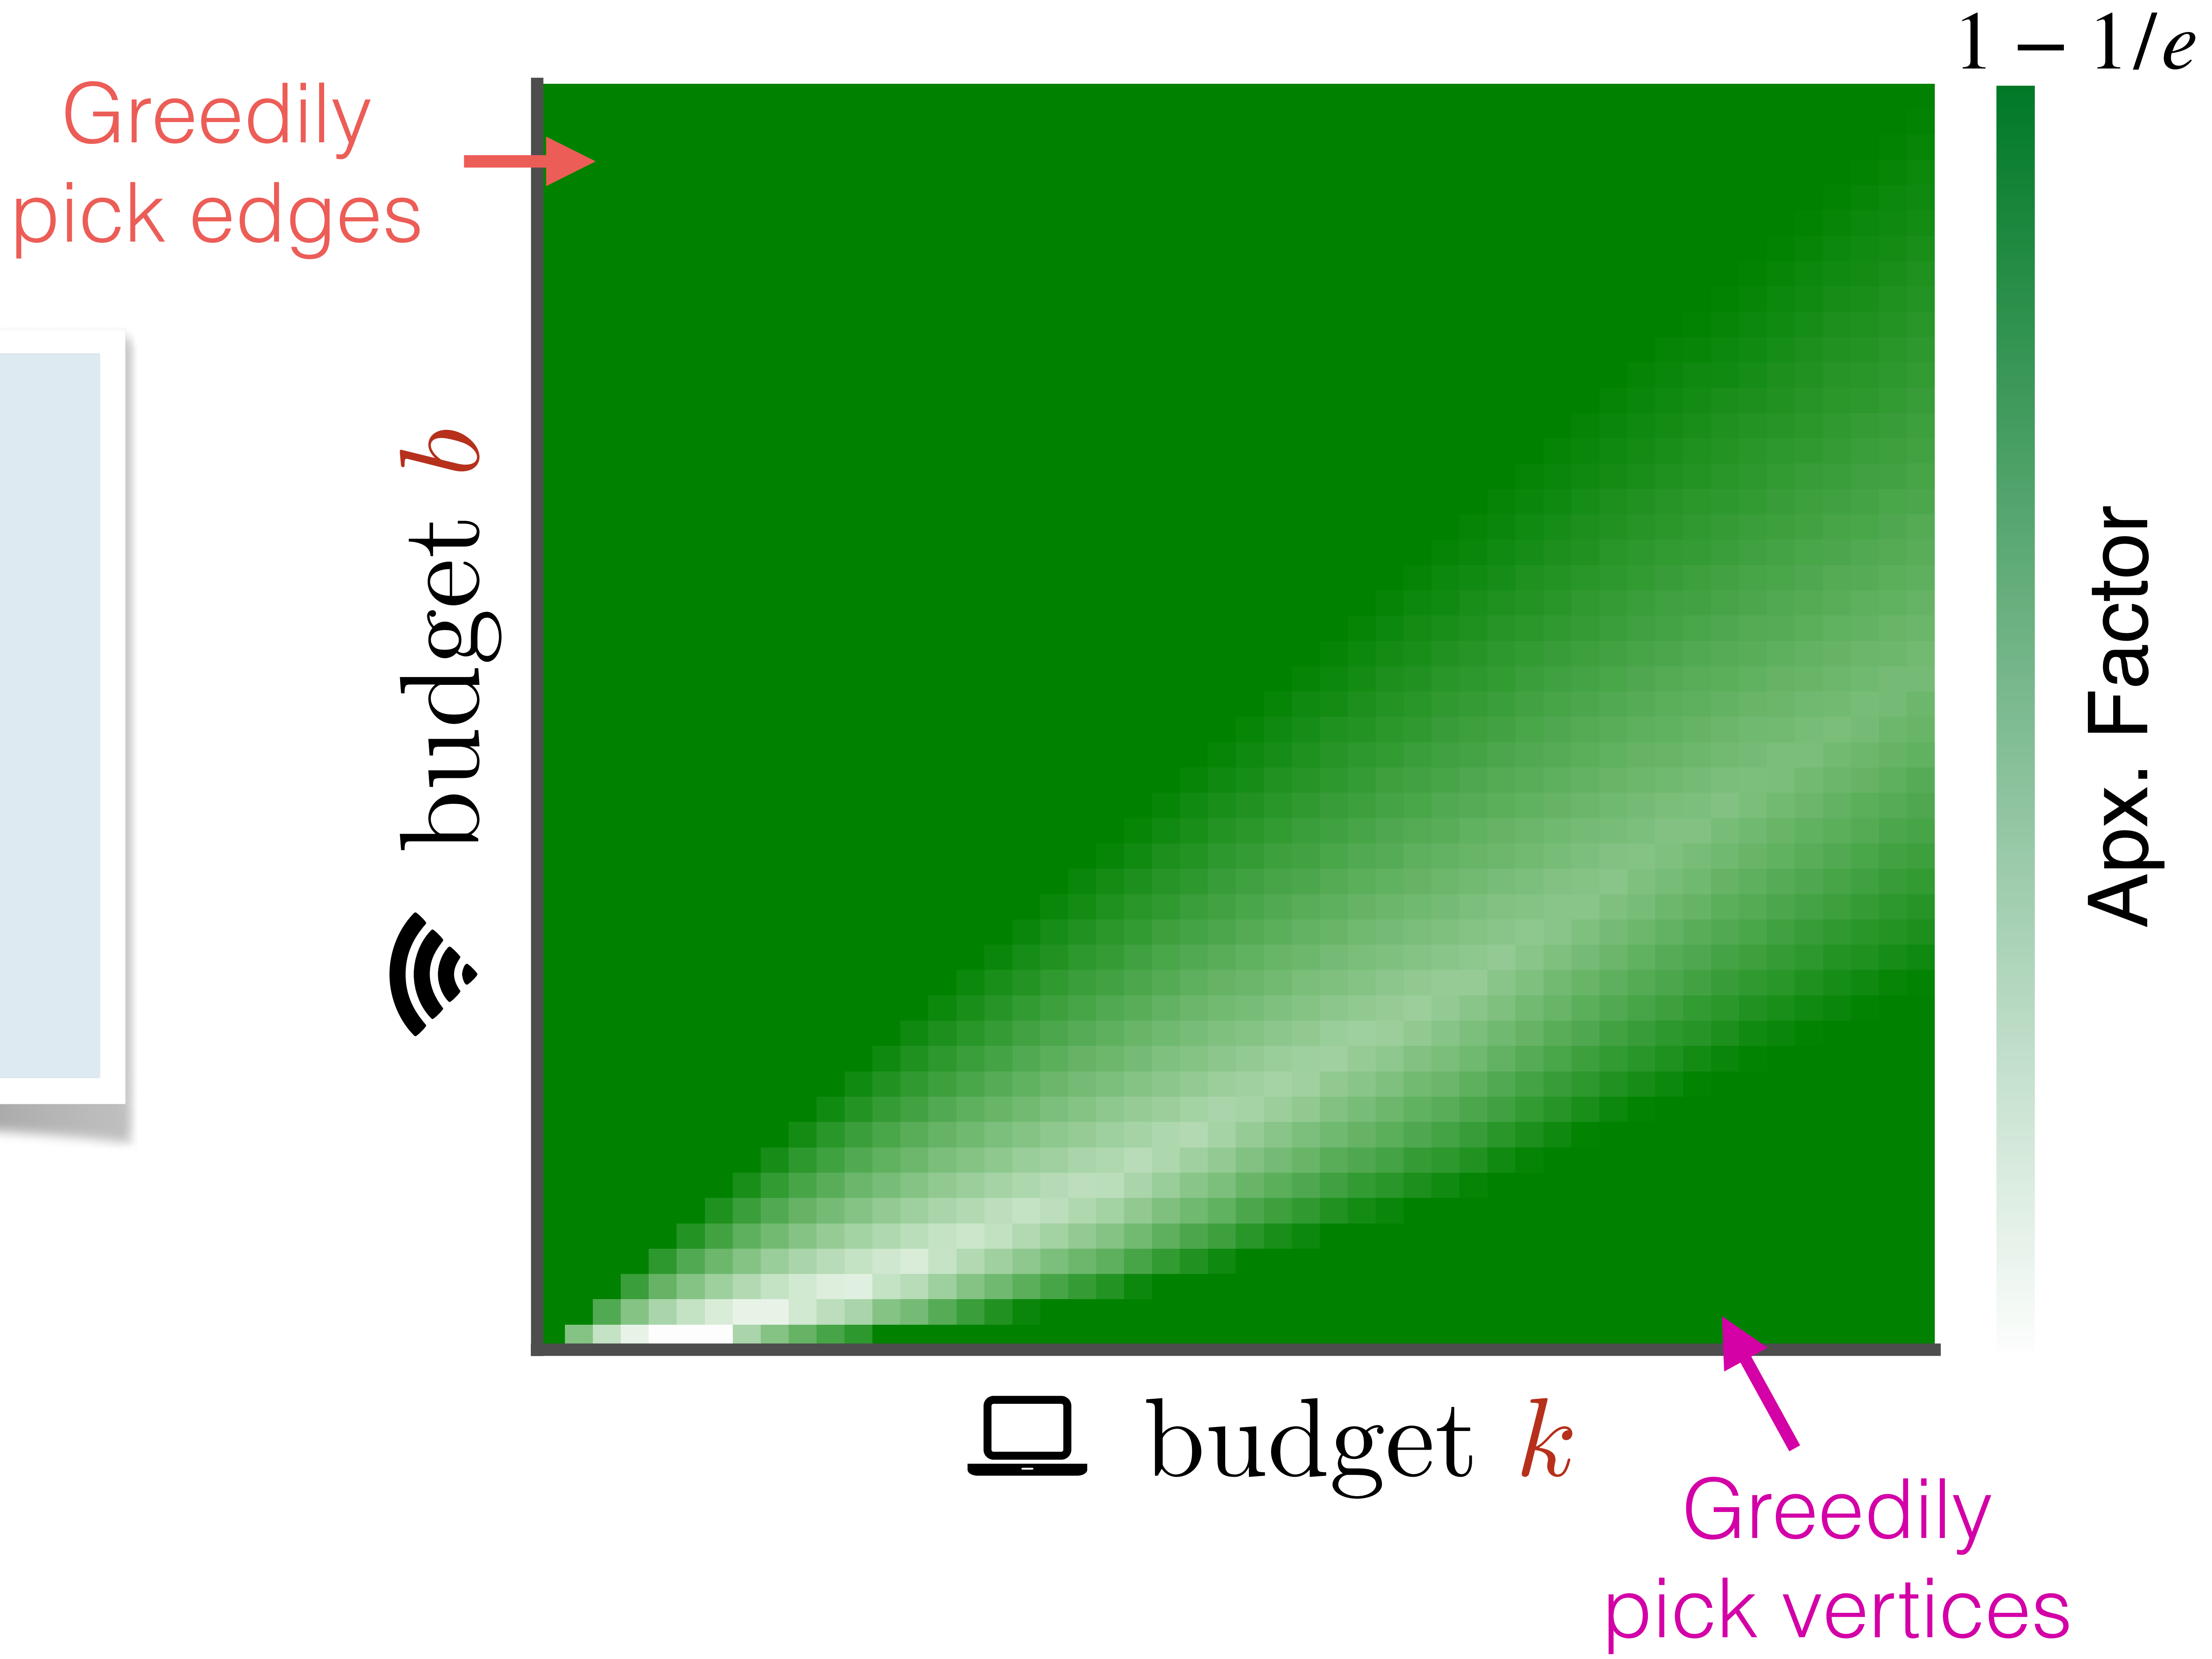

Experiments

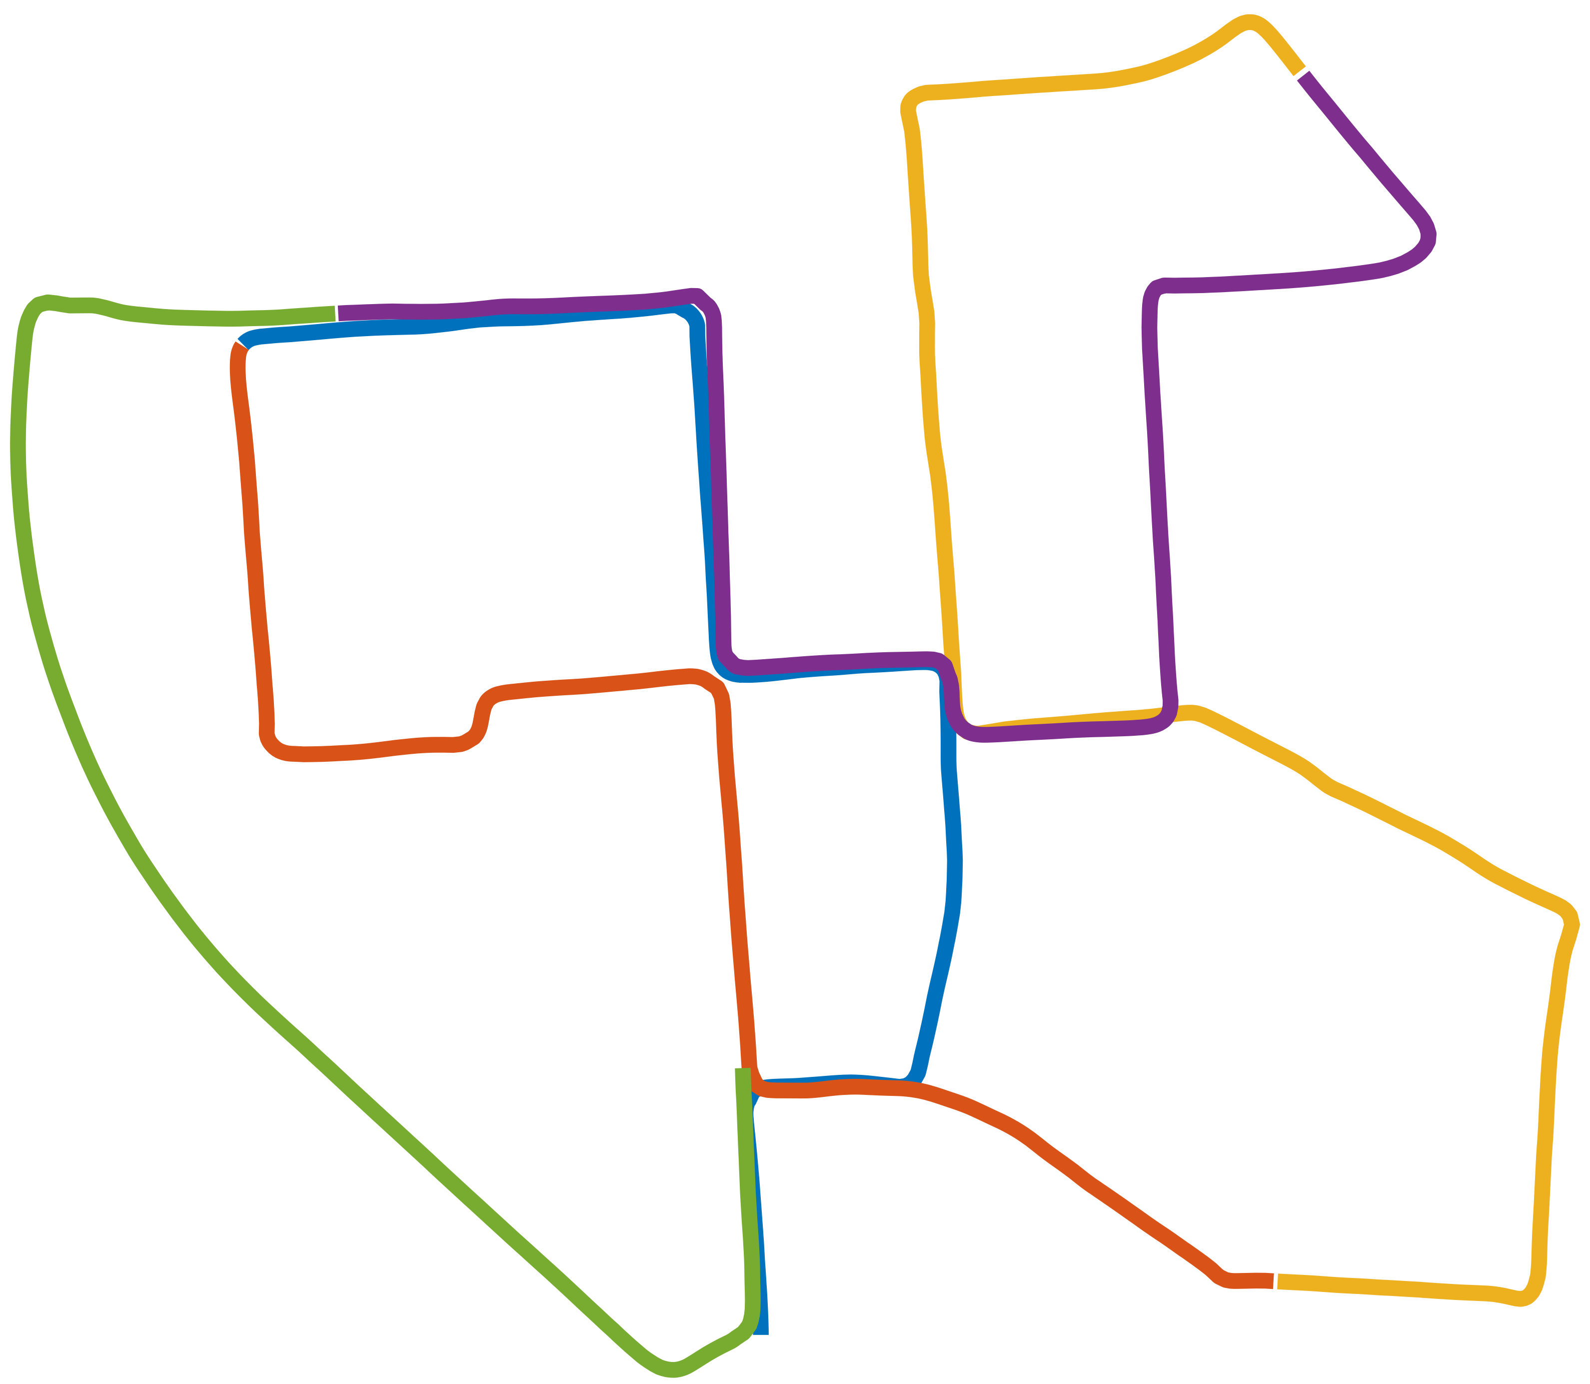

KITTI

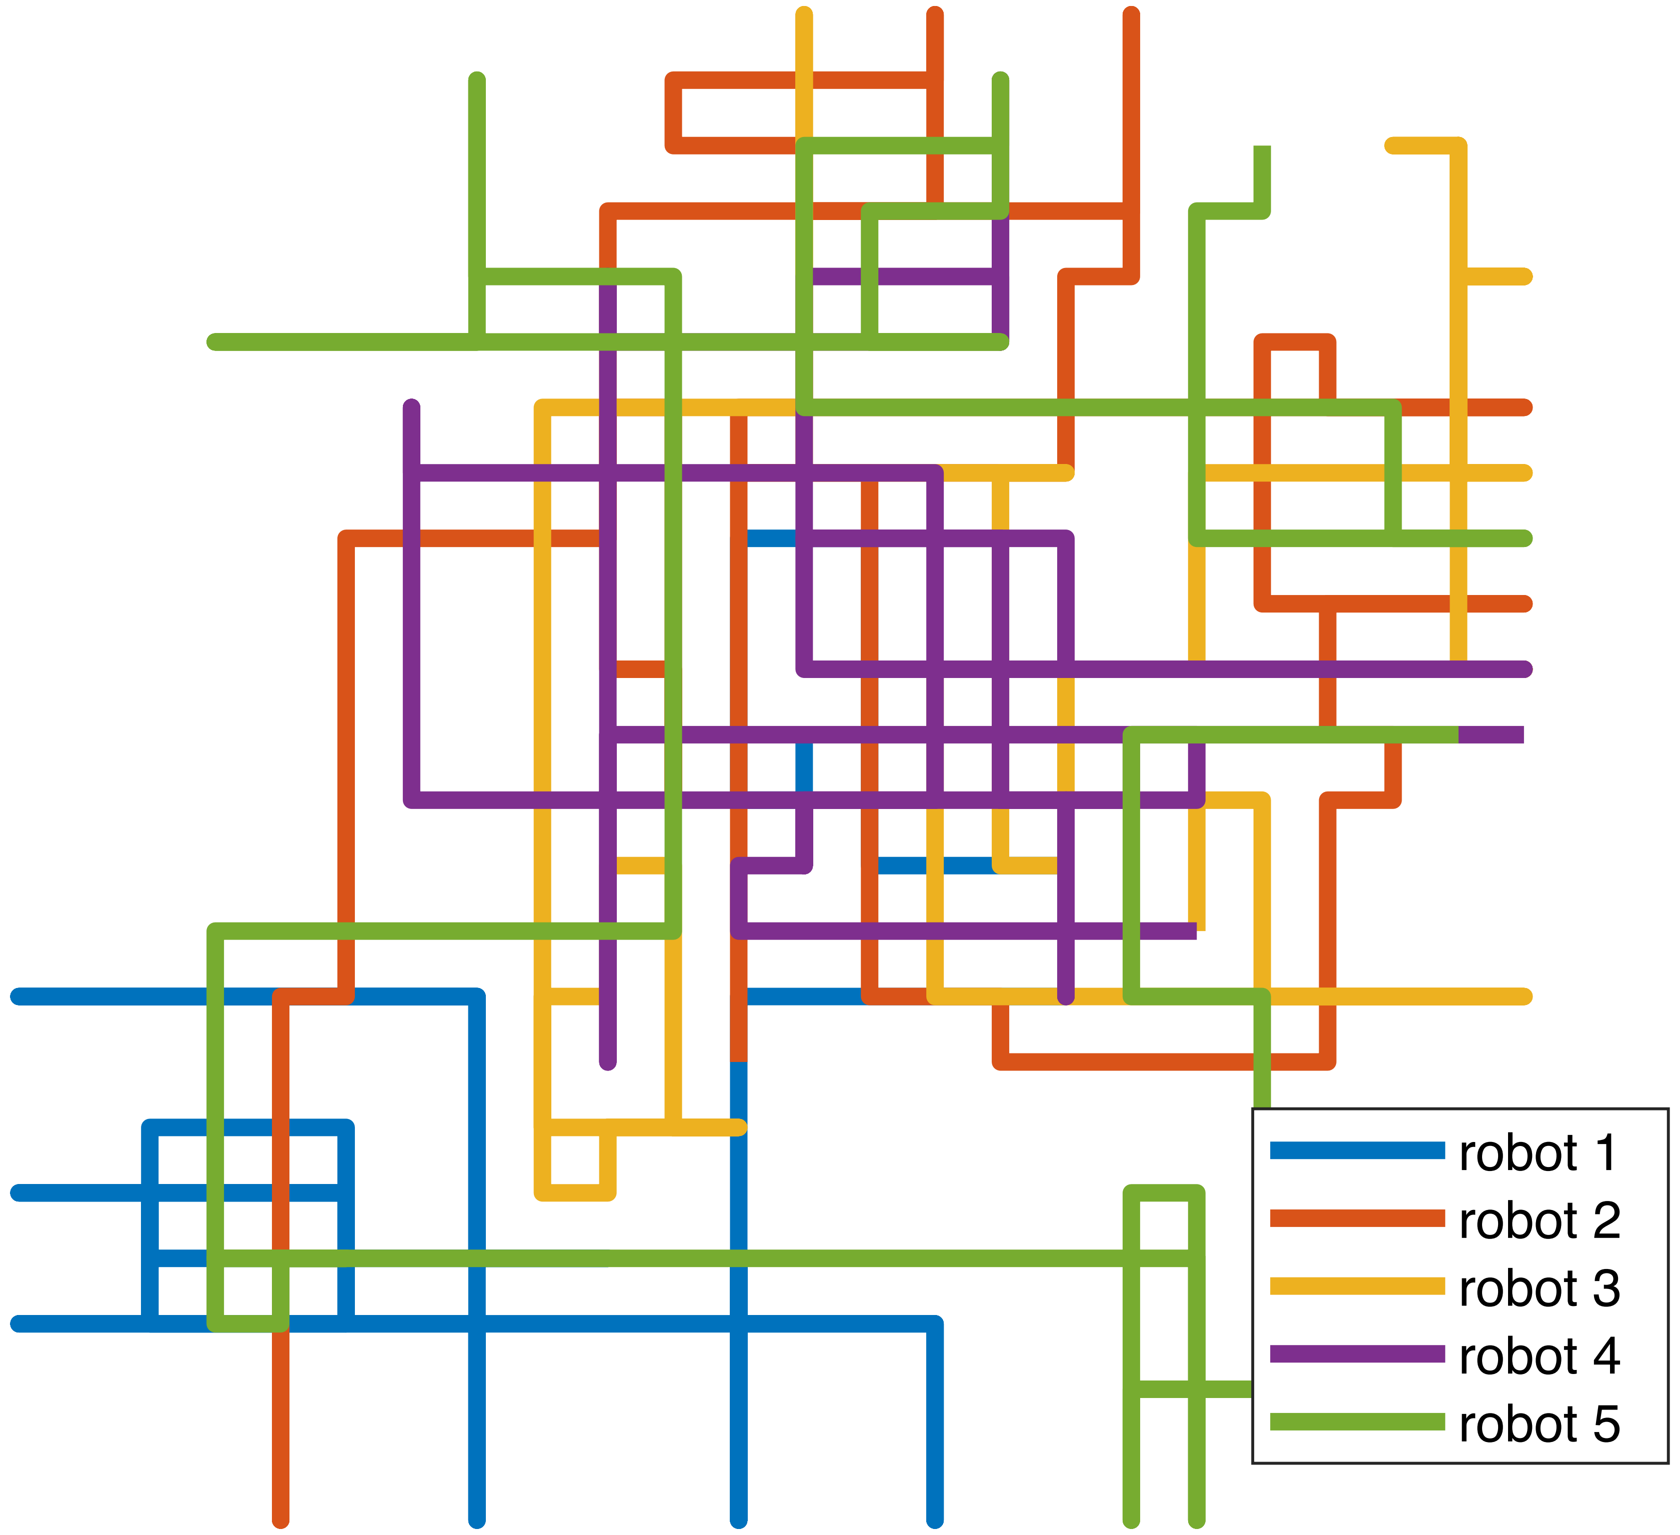

Atlas

Modular experiments

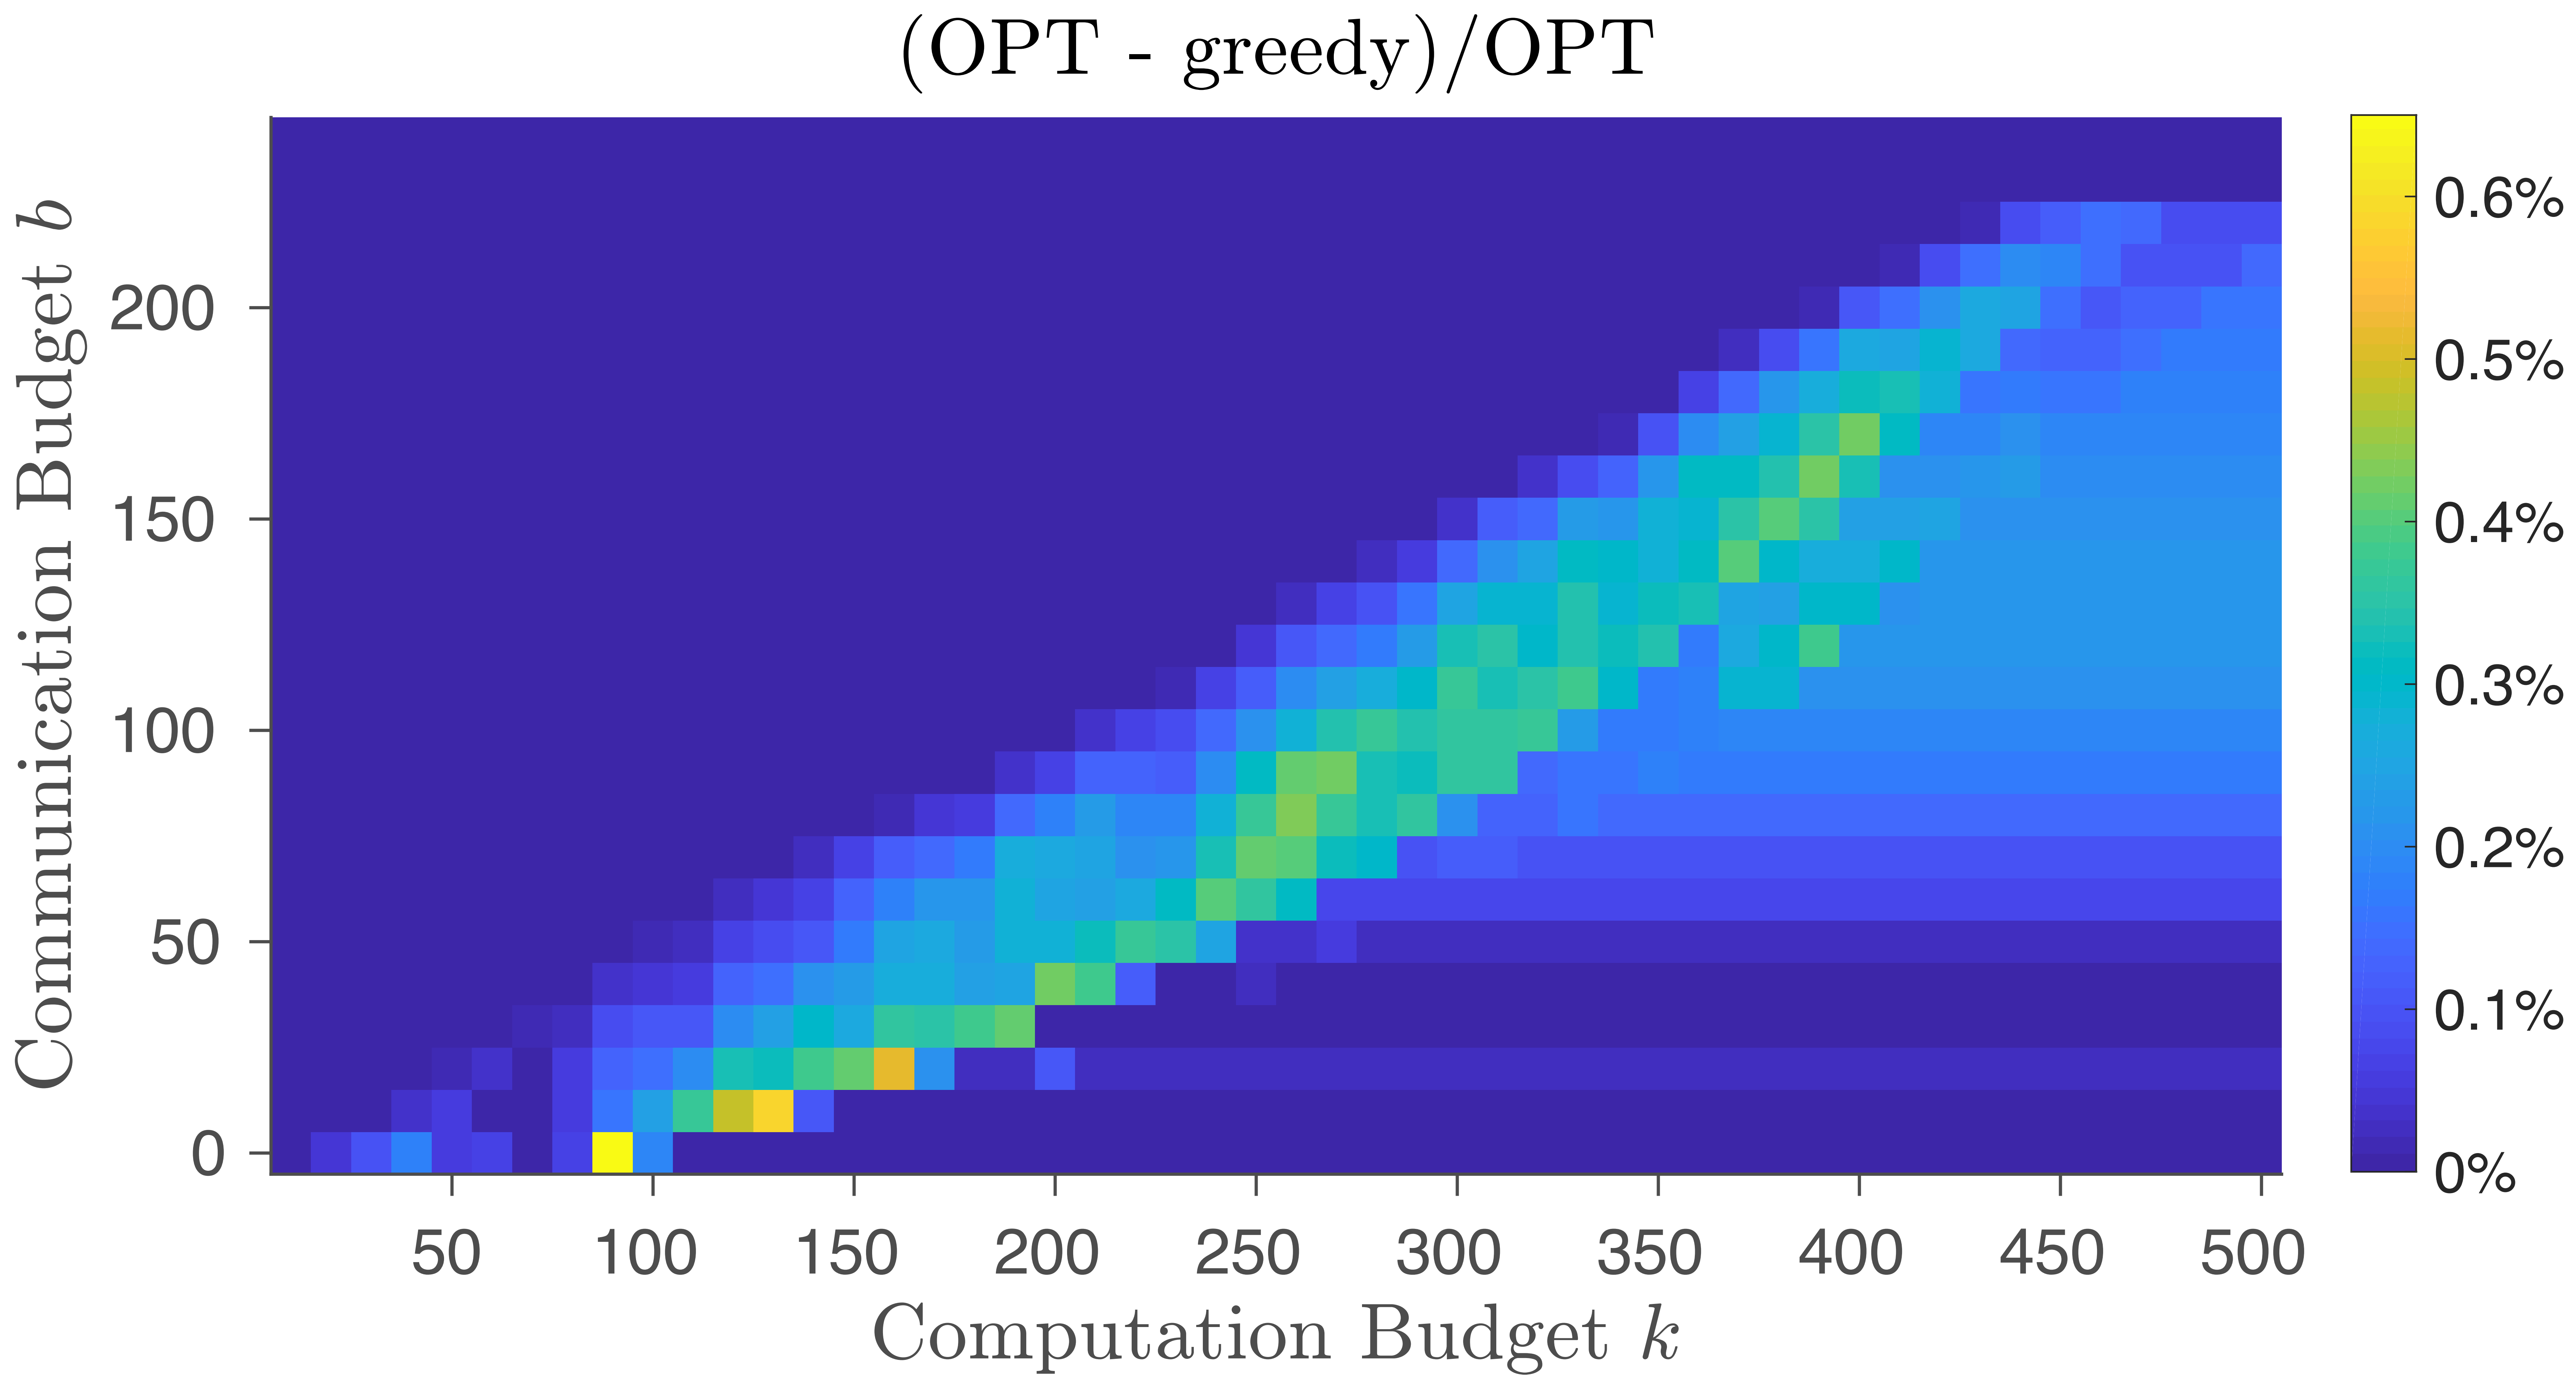

- Worst case: miss 6/1000 loop closures in expectation
- Near-optimal performance >> theoretical guarantee

Submodular experiments

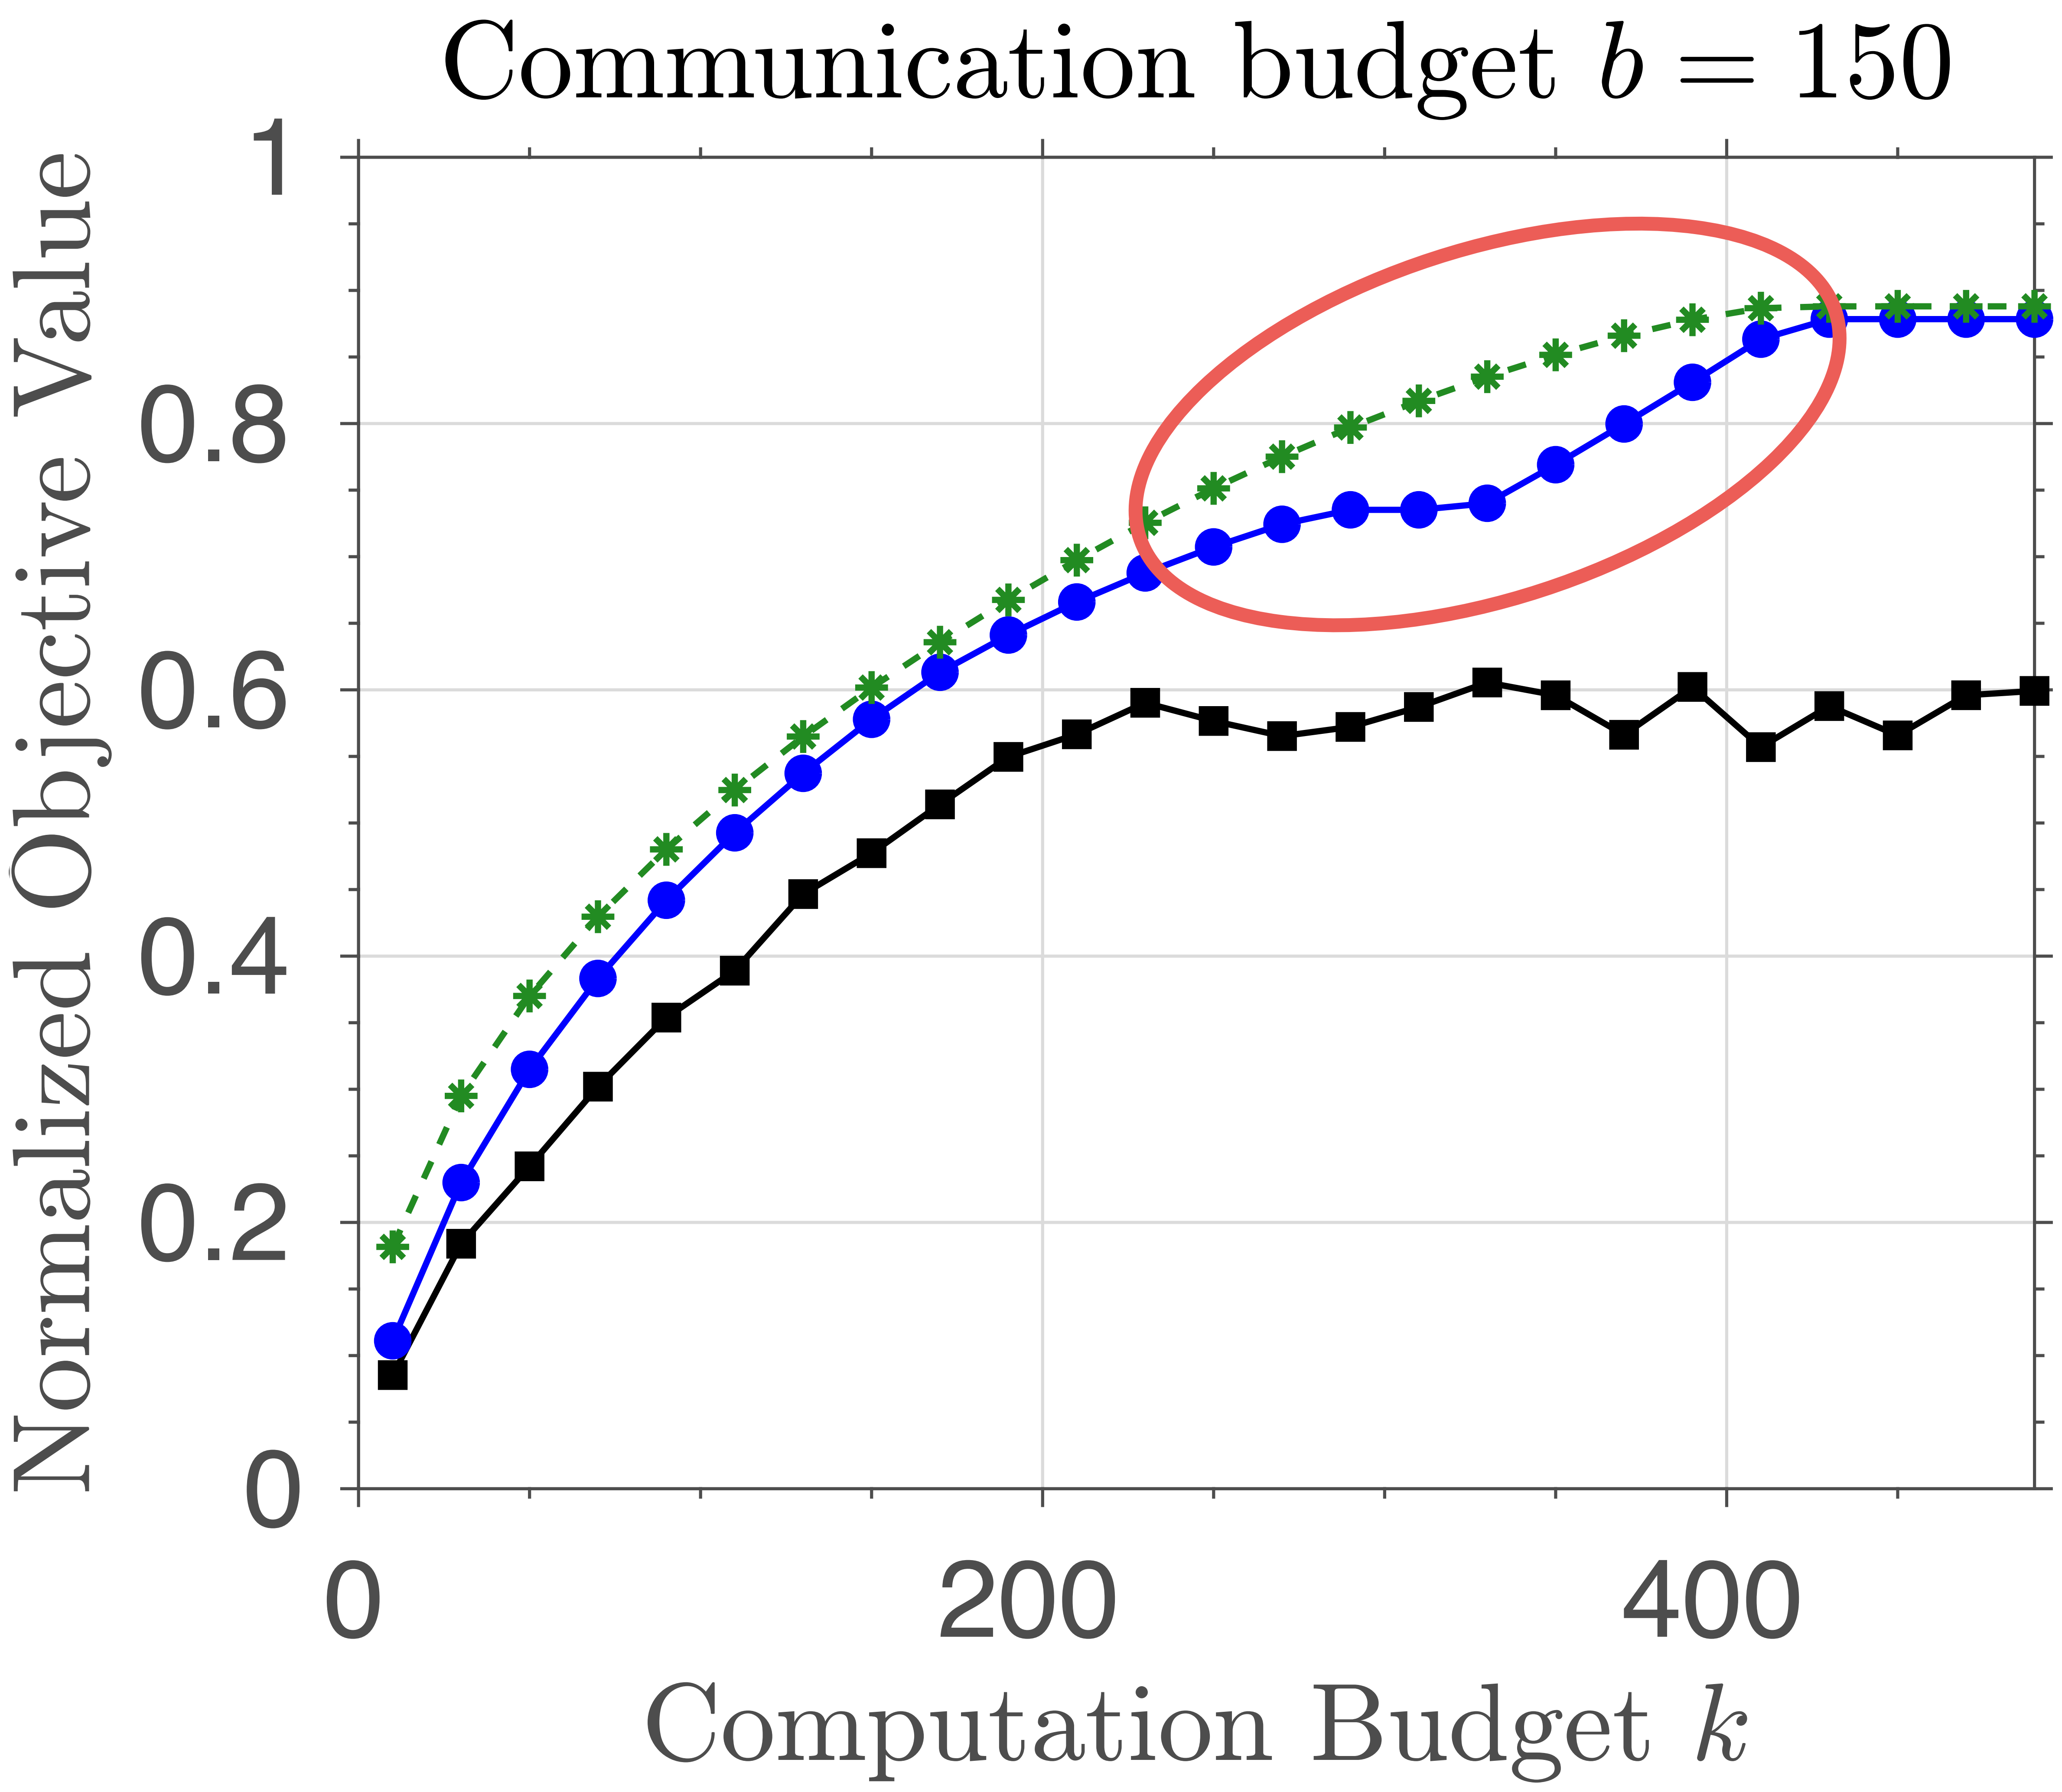

- Near-optimal performance
- Phase transition matches theoretical guarantee

Conclusion

- Ubiquitous resource constraints
- This paper: resource-aware loop closure for collaborative SLAM
- Select “best” feasible subset of potential loop closures
- NP-hard ⟶ efficient approximation algorithm with near-optimality guarantee

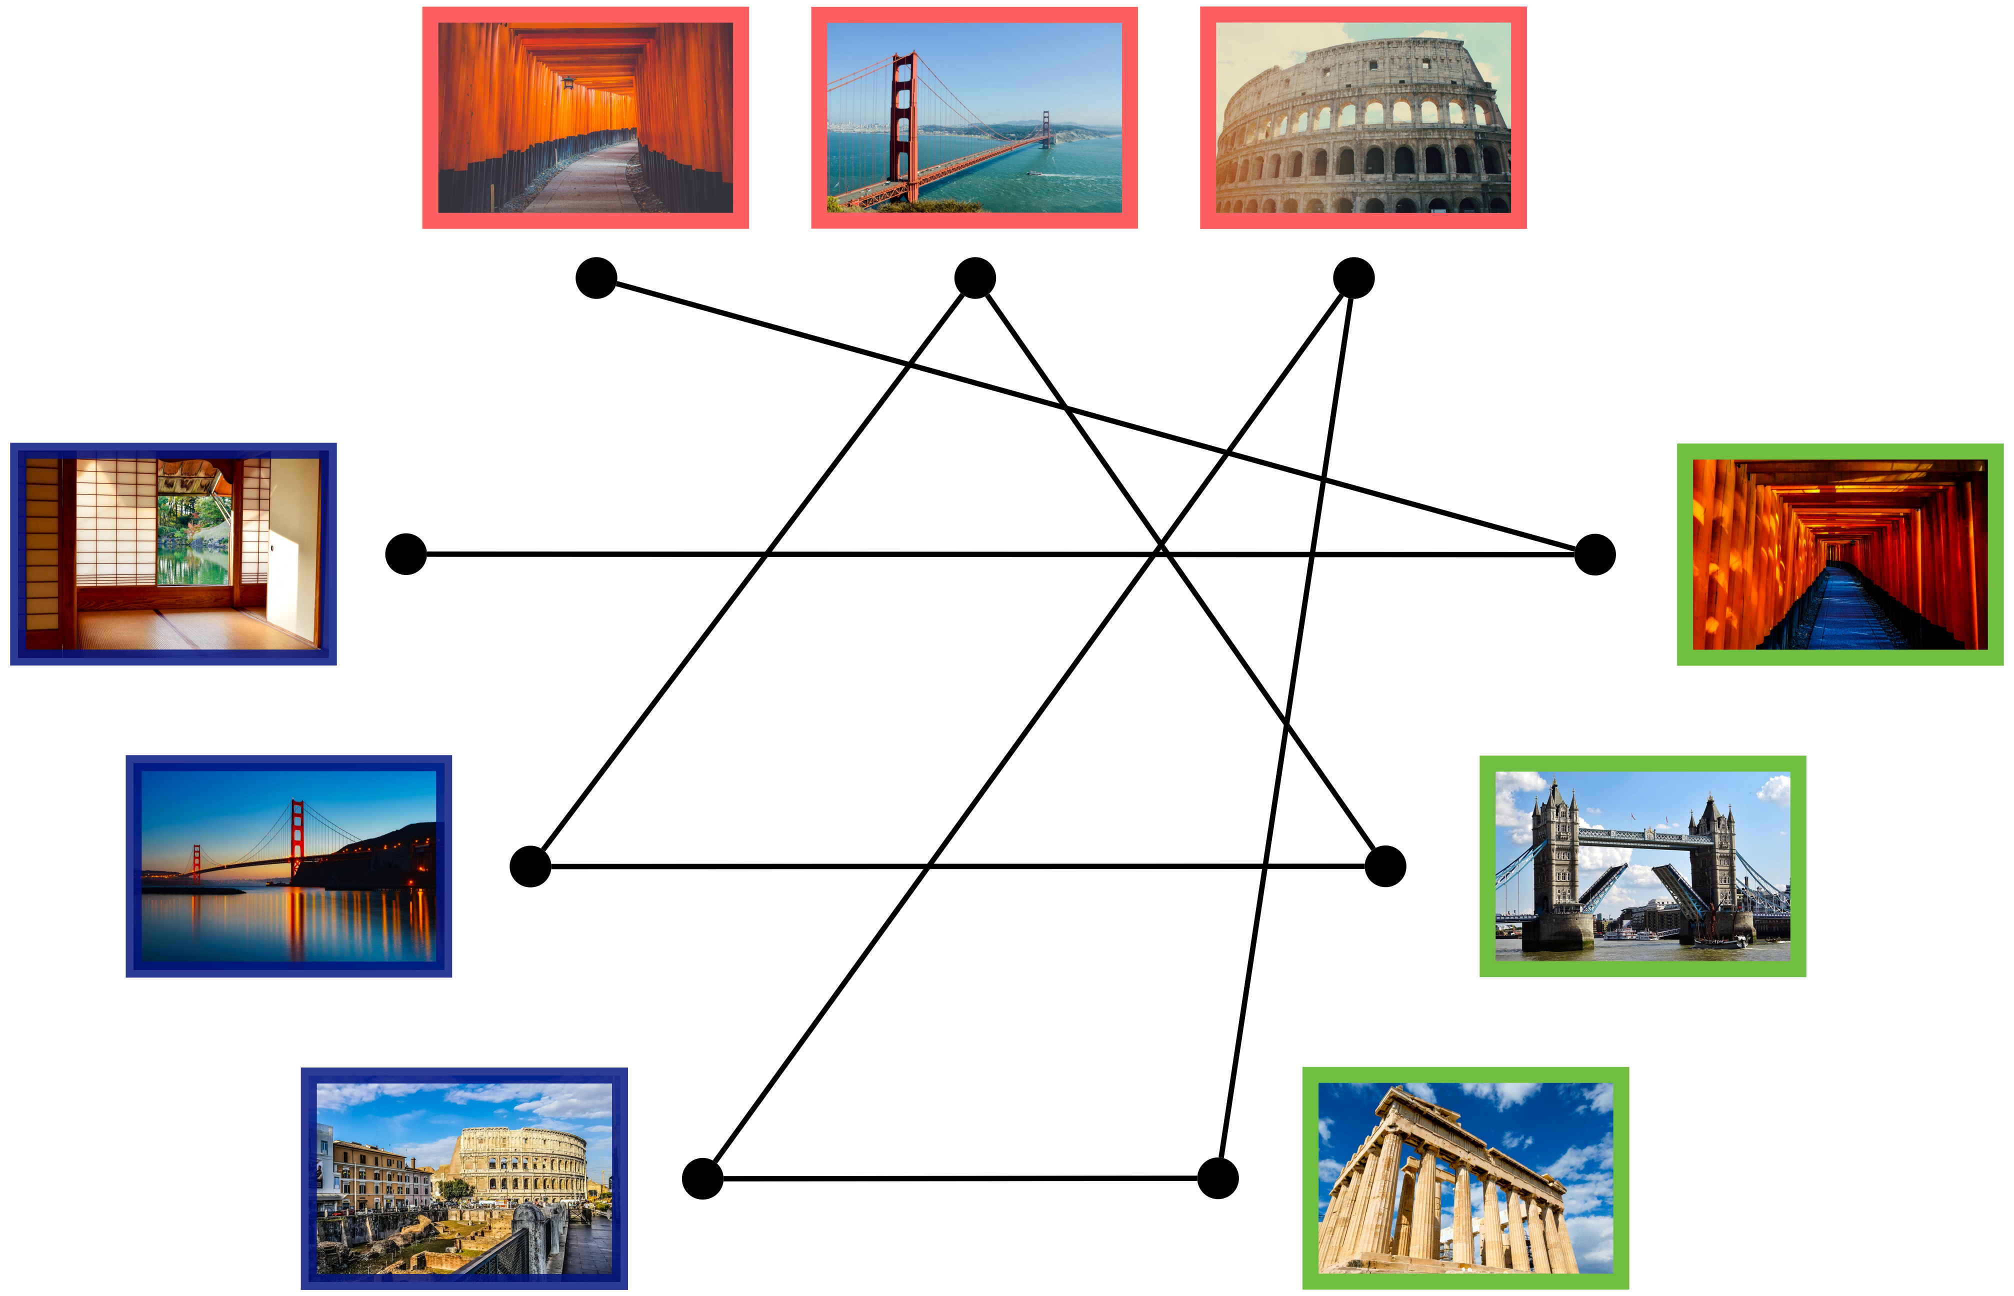

Dirty Laundry

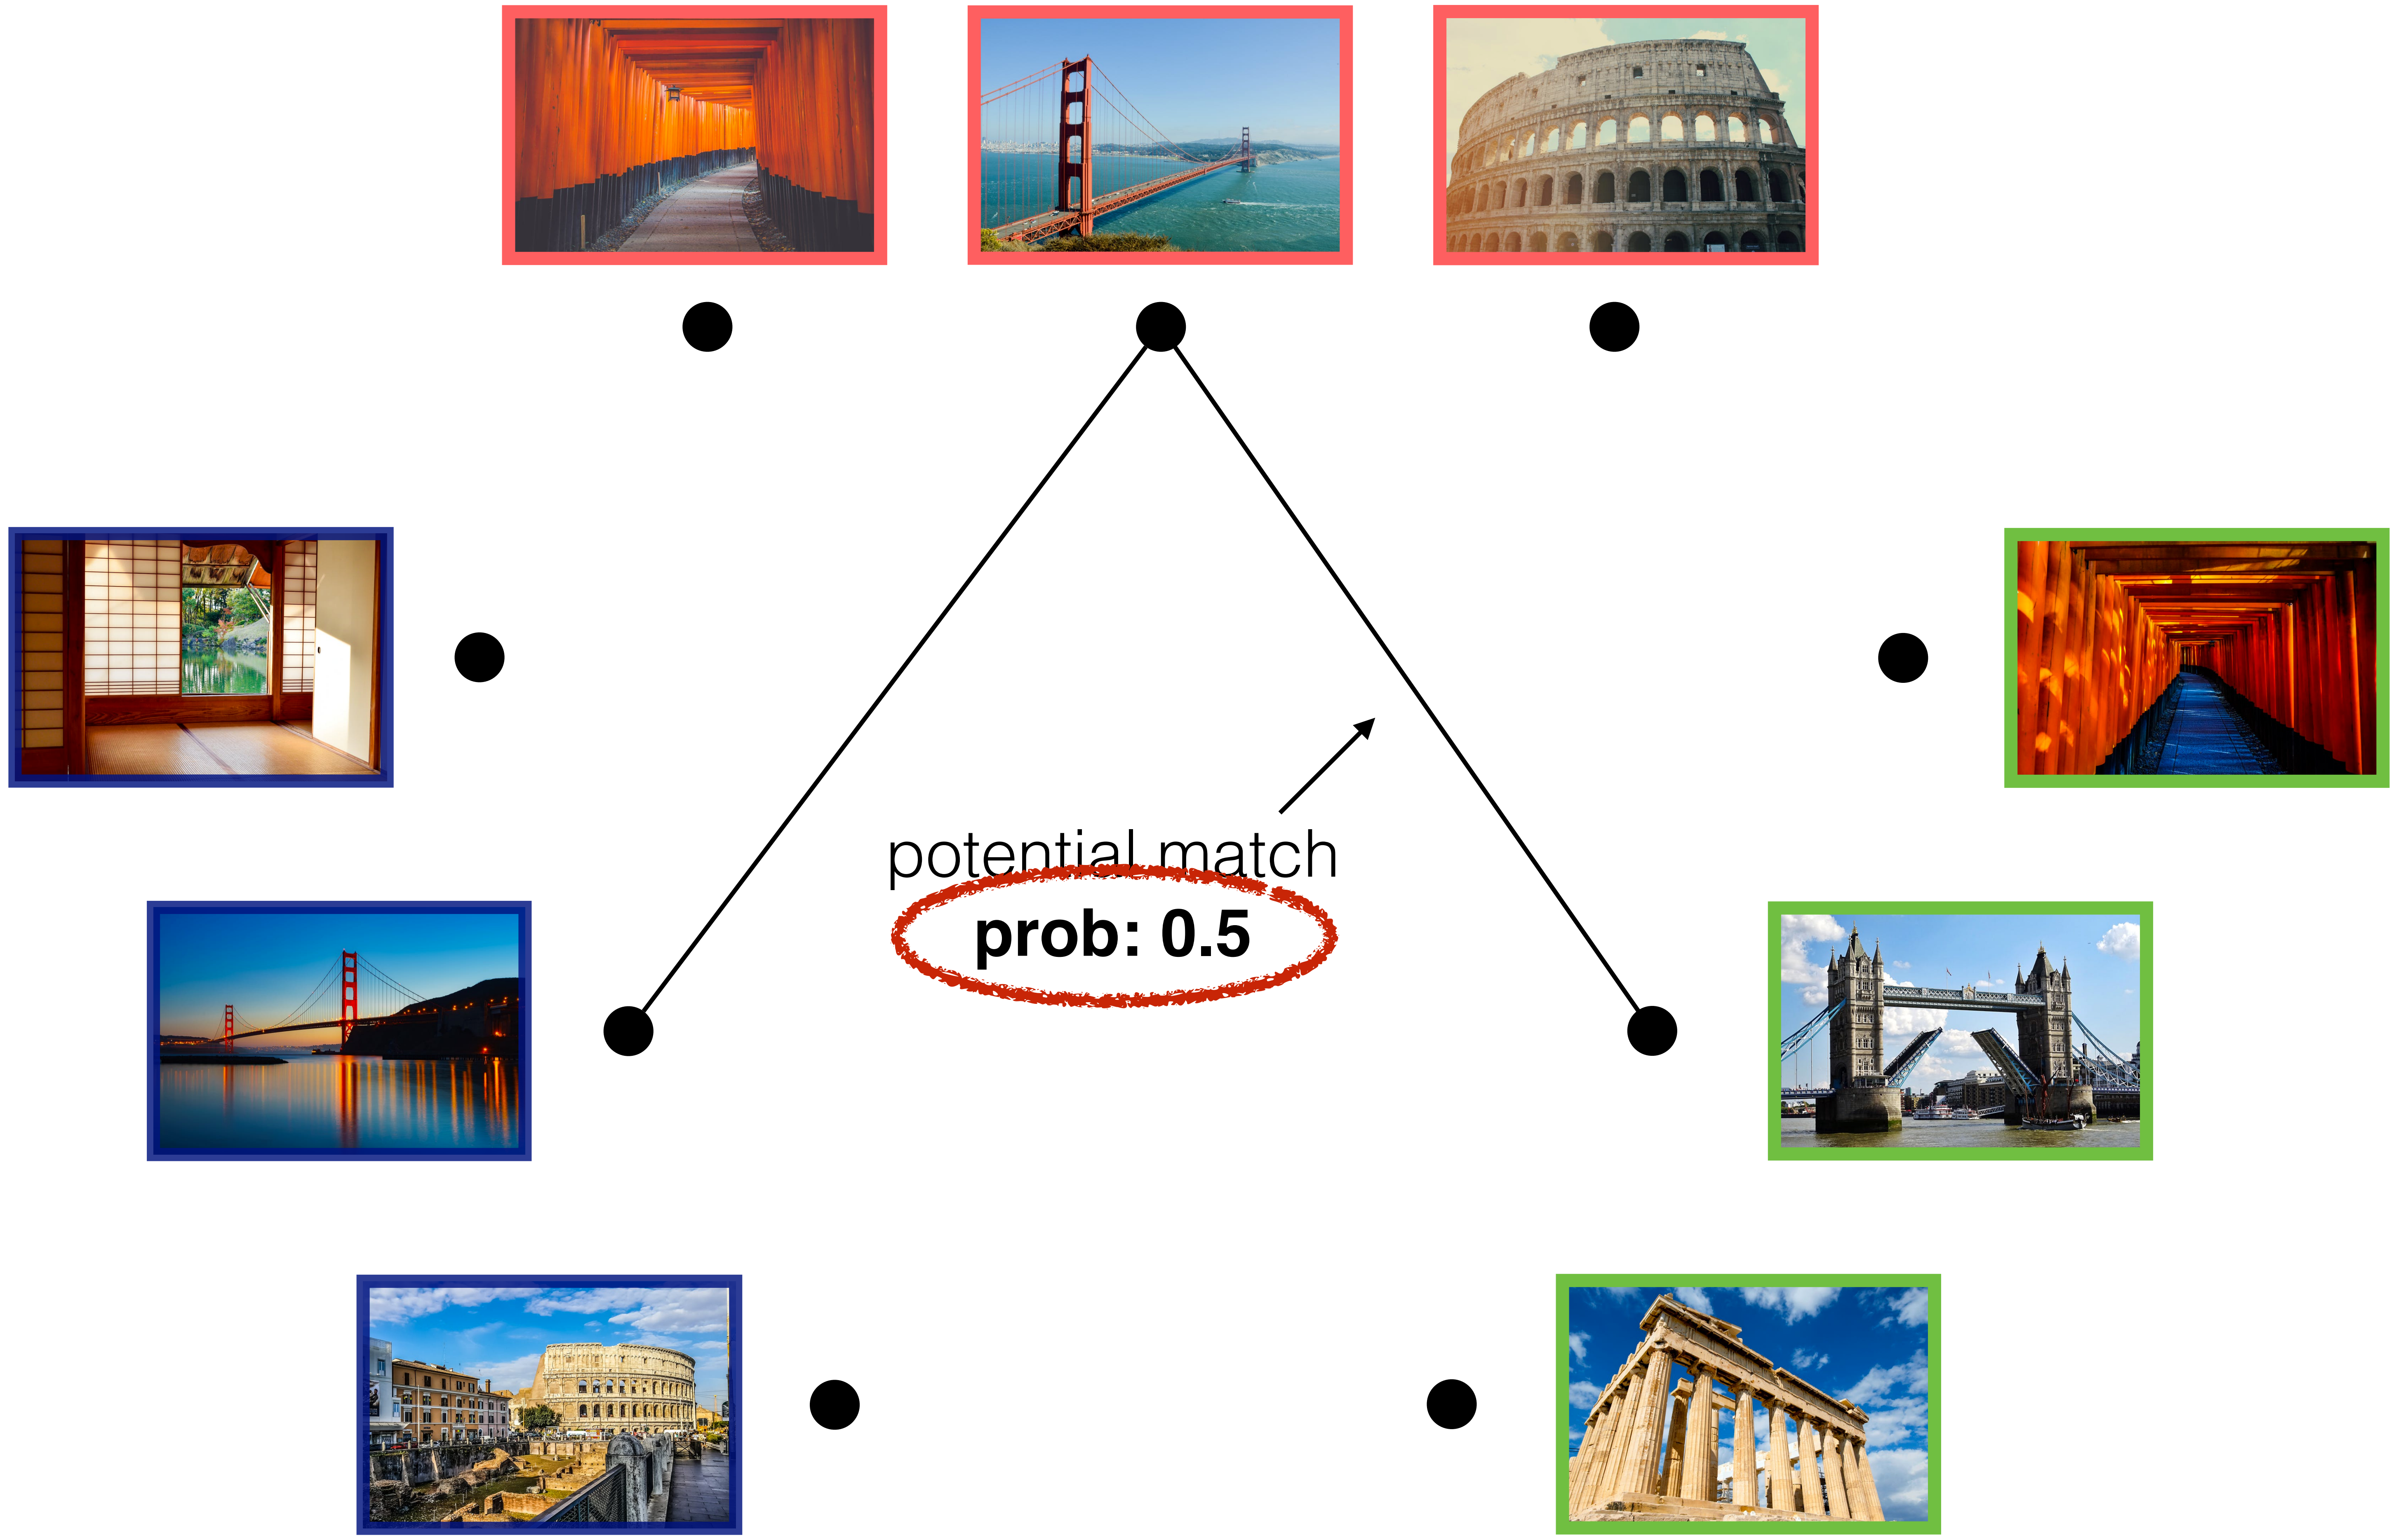

- Approximate probability as normalized visual similarity score
- Assume edges occur *independently*

Thank you!