Non-monotone Adaptive Submodular Maximization

Alkis Gotovos
ETH Zurich

Amin Karbasi
Yale University

Andreas Krause
ETH Zurich

IJCAI-15
24TH INTERNATIONAL JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE
BUENOS AIRES
JULY 25TH-31TH, 2015

Motivation

- Many AI problems boil down to selecting a number of elements from a large set of options

Motivation

- Many AI problems boil down to selecting a number of elements from a large set of options

- Sequentially make smart choices based on past observations

Motivation

- Many AI problems boil down to selecting a number of elements from a large set of options

- Sequentially make smart choices based on past observations

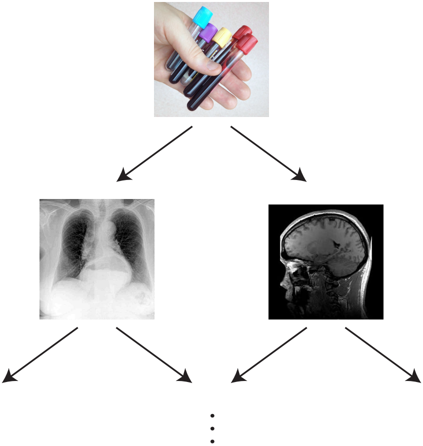

Motivation

Find classes of objective functions that are amenable to efficient sequential optimization with theoretical approximation guarantees

Birdwatching

Birdwatching

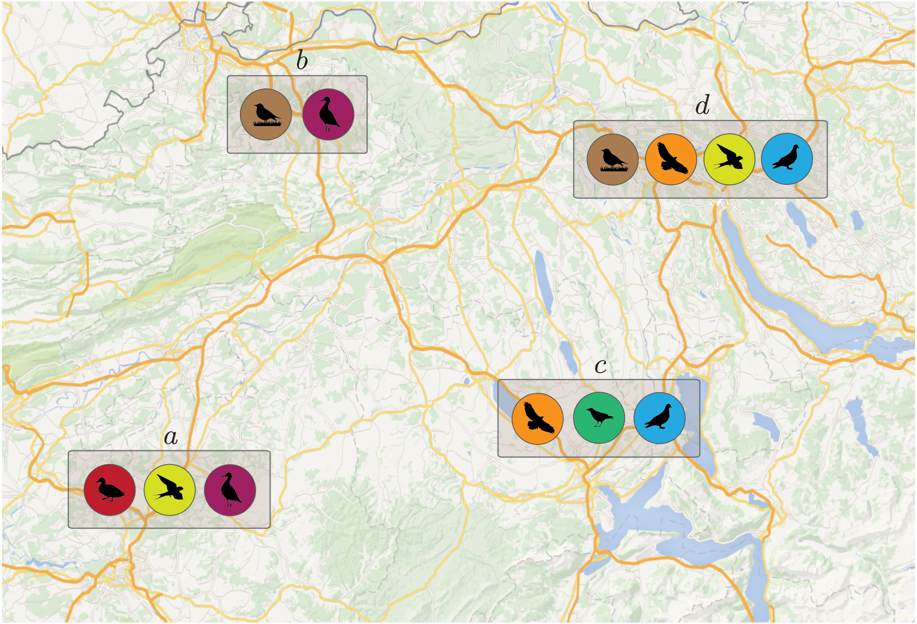

Objective

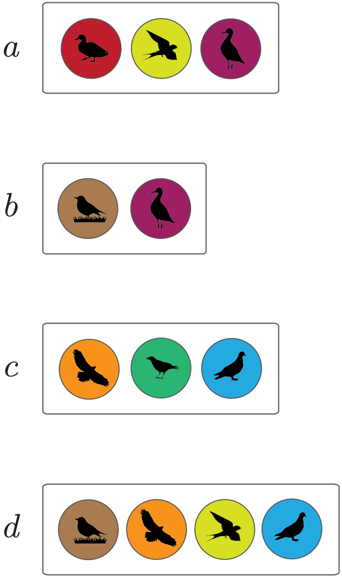

- Ground set $V = \{a, b, c, d\}$

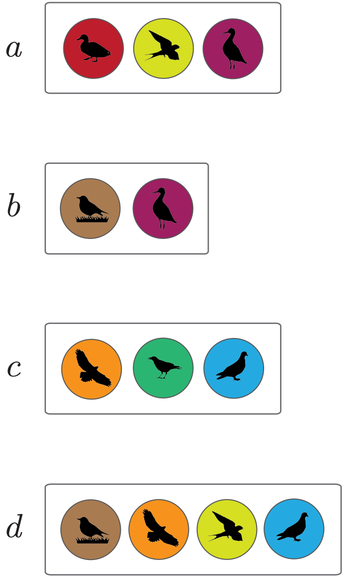

- Ground set $V = \{a, b, c, d\}$
- Objective function $f : 2^V \to \mathbb{R}_{\geq 0}$

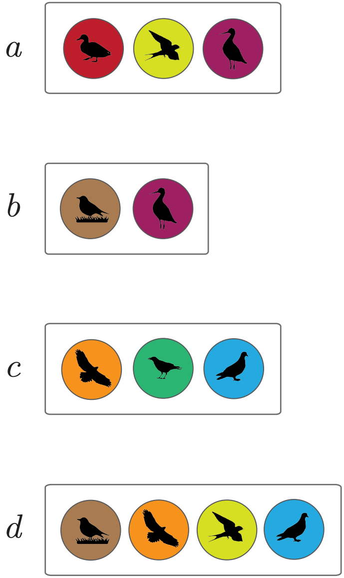

Objective

- Ground set $V = \{a, b, c, d\}$
- Objective function $f : 2^V \to \mathbb{R}_{\geq 0}$
- $f(\{d\}) = 4$

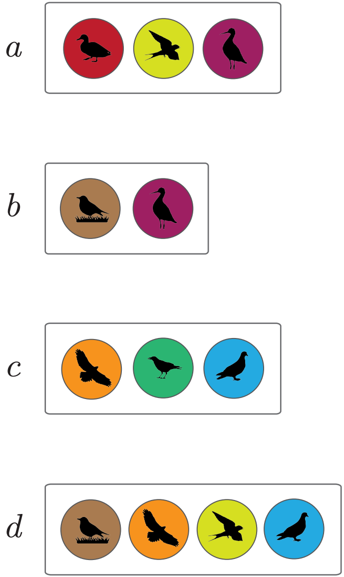

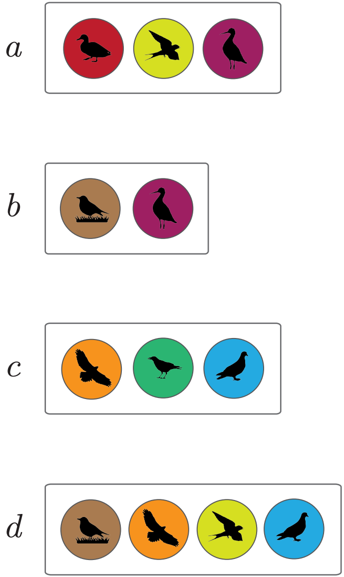

- Ground set $V = \{a, b, c, d\}$
- Objective function $f : 2^V \to \mathbb{R}_{\geq 0}$
- $f(\{d\}) = 4$
- $f(\{c, d\}) = 5$

Objective

- f is monotone

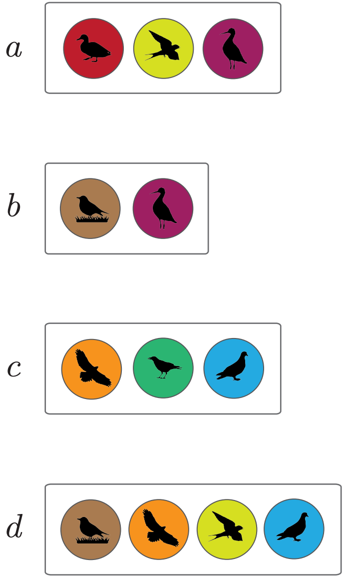

Objective

- f is monotone
- f is submodular

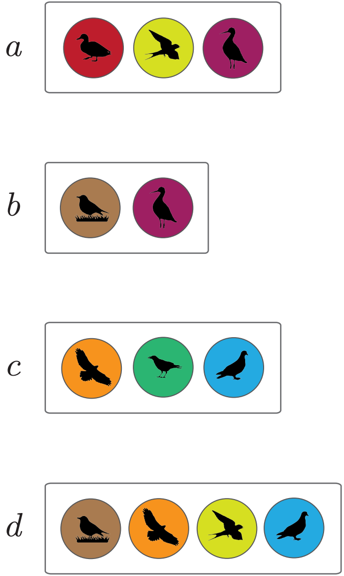

Objective

- f is monotone
- f is submodular
- Benefit of visiting c, given that…

a

b

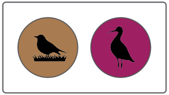

c

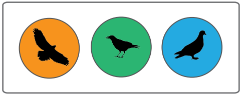

d

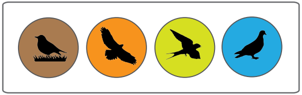

Objective

- f is monotone
- f is submodular
- Benefit of visiting c, given that…
  - …it is the first place we visit:

$$f(\{c\}) = 3$$

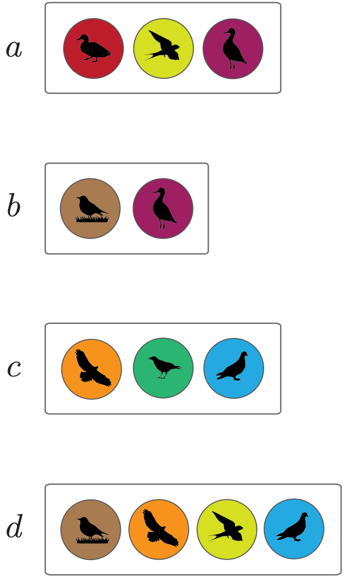

Objective

- f is monotone
- f is submodular
- Benefit of visiting c, given that…
  - …it is the first place we visit:

$$f(\{c\}) = 3$$

  - …we have already visited d:

$$f(\{c, d\}) - f(\{d\}) = 5 - 4 = 1$$

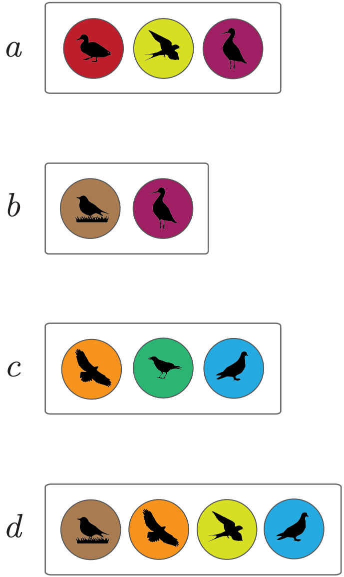

- Unconstrained problem:

$$\text{maximize} \quad f(S)$$

Monotone submodular maximization

- Unconstrained problem:

maximize $f(S)$ $\longrightarrow$ Trivial OPT $= f(V)$

Monotone submodular maximization

- Unconstrained problem:

$$\text{maximize} \quad f(S) \qquad \longrightarrow \qquad \text{Trivial} \quad \text{OPT} = f(V)$$

- Cardinality-constrained problem:

$$\begin{aligned} \text{maximize} \quad & f(S) \\ \text{subject to} \quad & |S| \leq k \end{aligned}$$

Monotone submodular maximization

- Unconstrained problem:

$$\text{maximize} \quad f(S) \qquad \longrightarrow \qquad \text{Trivial} \quad \text{OPT} = f(V)$$

- Cardinality-constrained problem:

$$\begin{aligned} \text{maximize} \quad & f(S) \\ \text{subject to} \quad & |S| \leq k \end{aligned} \qquad \longrightarrow \qquad \text{NP-hard}$$

Monotone submodular maximization

- Unconstrained problem:

$$\text{maximize} \quad f(S) \qquad \longrightarrow \qquad \text{Trivial} \quad \text{OPT} = f(V)$$

- Cardinality-constrained problem:

$$\begin{aligned} \text{maximize} \quad & f(S) \\ \text{subject to} \quad & |S| \leq k \end{aligned} \qquad \longrightarrow \qquad \text{NP-hard}$$

- More general constraints: matroid, knapsack, etc.

Monotone submodular maximization

- Unconstrained problem:

$$\text{maximize} \quad f(S) \qquad \longrightarrow \qquad \text{Trivial} \quad \text{OPT} = f(V)$$

- Cardinality-constrained problem:

$$\begin{aligned} \text{maximize} \quad & f(S) \\ \text{subject to} \quad & |S| \leq k \end{aligned} \qquad \longrightarrow \qquad \text{NP-hard}$$

- More general constraints: matroid, knapsack, etc.

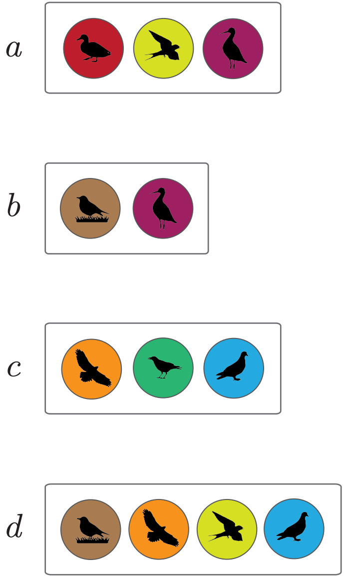
a
b
c
d

Greedy

- $k = 2$

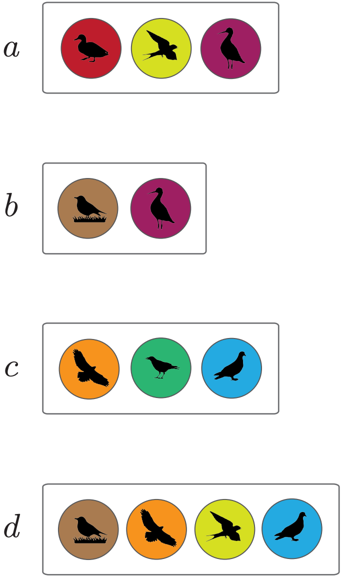

- $k = 2$
- $S_0 = \varnothing \longrightarrow f(S_0) = 0$

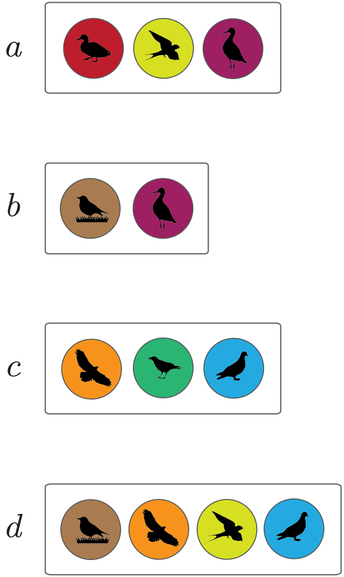

- $k = 2$
- $S_0 = \varnothing \longrightarrow f(S_0) = 0$
- $S_1 = \{d\} \longrightarrow f(S_1) = 4$

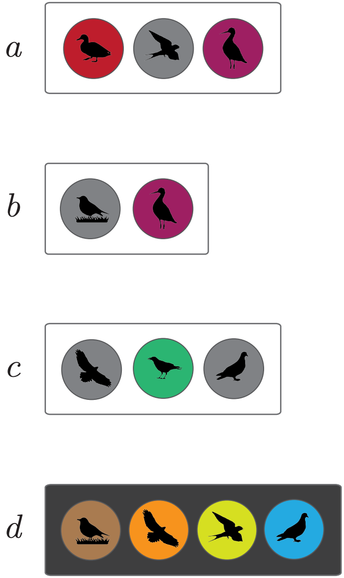

- $k = 2$
- $S_0 = \emptyset \longrightarrow f(S_0) = 0$
- $S_1 = \{d\} \longrightarrow f(S_1) = 4$
- $S_2 = \{d, a\} \longrightarrow f(S_2) = 6$

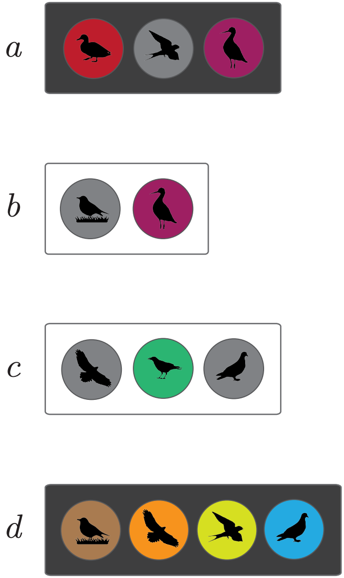

Greedy

Theorem [Nemhauser *et al.*, 1978]

If f is monotone submodular, then greedy gives a $(1 - 1/e)$-approximation.

Birdwatching with costs

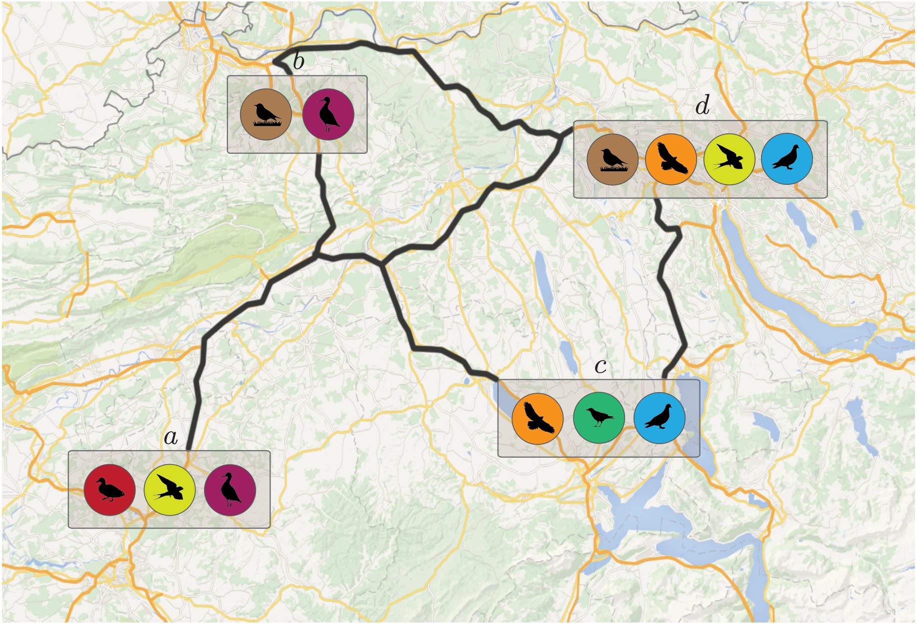

Non-monotonicity

- $g(A) = \underbrace{f(A)}_{\text{monotone submodular}} - \underbrace{c(A)}_{\text{cost term}}$

Non-monotonicity

- $g(A) = \underbrace{f(A)}_{\text{monotone submodular}} - \underbrace{c(A)}_{\text{cost term}}$

- Greedy has no guarantees for non-monotone functions

Non-monotonicity

- $g(A) = \underbrace{f(A)}_{\text{monotone submodular}} - \underbrace{c(A)}_{\text{cost term}}$
- Greedy has no guarantees for non-monotone functions
- Introduce randomization $\longrightarrow$ *random greedy* algorithm

Theorem [Buchbinder *et al.*, 2014]

If f is submodular, then random greedy gives a $(1/e)$-approximation*.

* In expectation over the randomness of the algorithm.

Random greedy

Theorem [Buchbinder *et al.*, 2014]

If f is submodular, then random greedy gives a $(1/e)$-approximation*.

If f is also monotone, then random greedy gives a $(1-1/e)$-approximation*.

* In expectation over the randomness of the algorithm.

Stochastic birdwatching

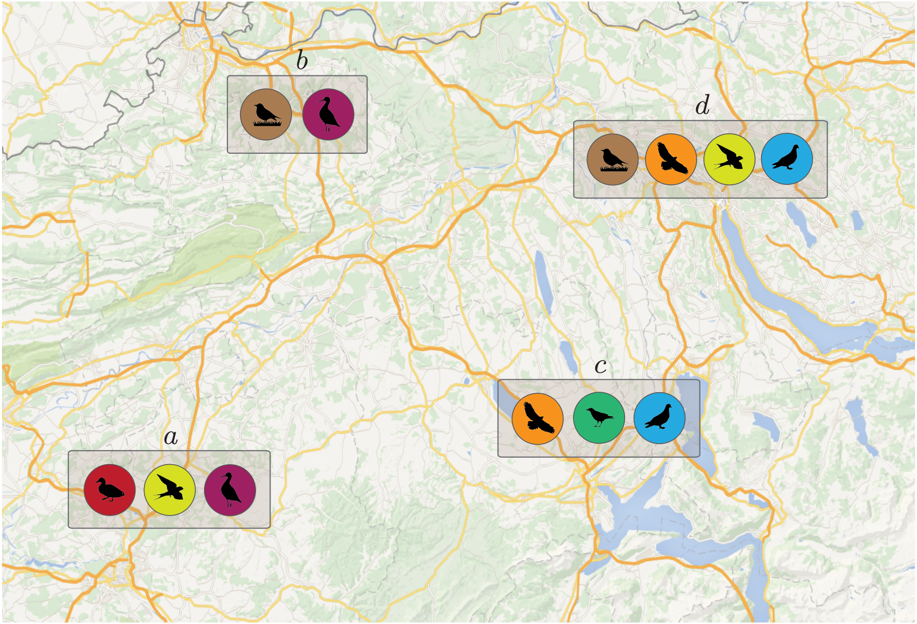

Stochastic birdwatching

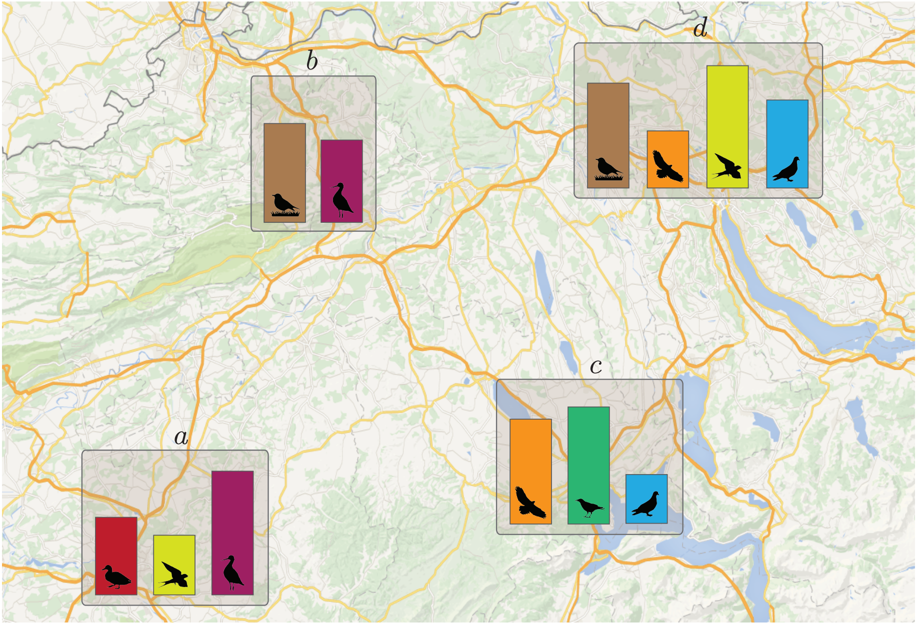

Adaptivity

- Non-adaptive: choose set of locations in advance without looking at outcomes

Adaptivity

- Non-adaptive: choose set of locations in advance without looking at outcomes
- Adaptive: sequentially make choices based on past outcomes

Adaptivity

- Non-adaptive: choose set of locations in advance without looking at outcomes
- Adaptive: sequentially make choices based on past outcomes
- Monotonicity and submodularity $\longrightarrow$ adaptive monotonicity and adaptive submodularity

Adaptivity

- Non-adaptive: choose set of locations in advance without looking at outcomes
- Adaptive: sequentially make choices based on past outcomes
- Monotonicity and submodularity $\longrightarrow$ adaptive monotonicity and adaptive submodularity
- Greedily select the most promising location in conditional expectation $\longrightarrow$ *adaptive greedy* algorithm

Adaptive greedy

Theorem [Golovin and Krause, 2011]

If f is adaptive monotone submodular, then adaptive greedy gives a $(1 - 1/e)$-approximation*.

* In expectation over the randomness of the environment.

What's missing?

	Non-adaptive	Adaptive
Monotone		
Non-monotone		

What's missing?

	Non-adaptive	Adaptive
Monotone	Greedy $(1 - 1/e)$	
Non-monotone		

What's missing?

	Non-adaptive	Adaptive
Monotone	Greedy $(1 - 1/e)$	
Non-monotone	Random greedy $(1/e)$	

What's missing?

	Non-adaptive	**Adaptive**
Monotone	Greedy $(1 - 1/e)$	Adaptive greedy $(1 - 1/e)$
Non-monotone	Random greedy $(1/e)$	

What's missing?

	Non-adaptive	Adaptive
Monotone	Greedy $(1 - 1/e)$	Adaptive greedy $(1 - 1/e)$
Non-monotone	Random greedy $(1/e)$	?

What's missing?

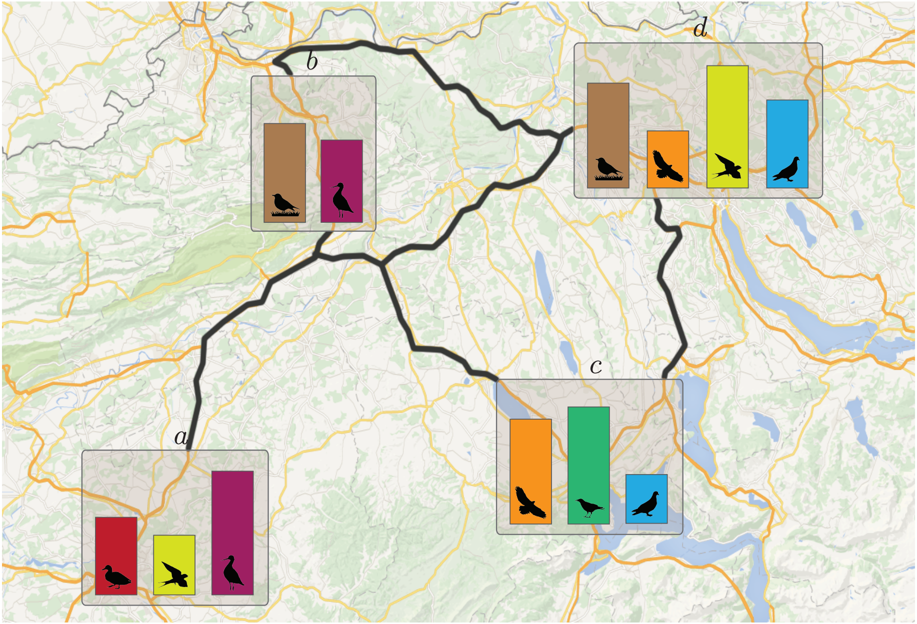

What's missing?

How do we maximize a non-monotone adaptive submodular function subject to a cardinality constraint?

What's missing?

How do we maximize a non-monotone adaptive submodular function subject to a cardinality constraint?

Adaptive random greedy

Theorem [Our contribution]

If f is adaptive submodular, then adaptive random greedy gives a $(1/e)$-approximation*.

* In expectation over the randomness of the algorithm and the environment.

Adaptive random greedy

Theorem [Our contribution]

If f is adaptive submodular, then adaptive random greedy gives a $(1/e)$-approximation*.

If f is also adaptive monotone, then adaptive random greedy gives a $(1-1/e)$-approximation*.

* In expectation over the randomness of the algorithm and the environment.

Conclusion

	Non-adaptive	**Adaptive**
Monotone	Greedy $(1 - 1/e)$	Adaptive greedy $(1 - 1/e)$
Non-monotone	Random greedy $(1/e)$	?

Conclusion

	Non-adaptive	Adaptive
Monotone	Greedy $(1 - 1/e)$	Adaptive greedy $(1 - 1/e)$
Non-monotone	Random greedy $(1/e)$	**Adaptive random greedy** $(1/e)$

Conclusion

More in our poster! (Panel 40)

- Details on algorithm
- Classes of non-monotone objectives
- Experimental evaluation on social networks