Through the Lens of Sequence Submodularity

Sara Bernardini, Fabio Fagnani, Chiara Piacentini

Department of Computer Science
Royal Holloway University of London
sara.bernardini@rhul.ac.uk

30th International Conference on Automated Planning and Scheduling
October 26 - 30, 2020

Outstanding Paper Award Honorable Mention

Classical Submodularity

Ω finite set, $F : 2^{\Omega} \to \mathbb{R}$ submodular

$$F(A \cup \{\sigma\}) - F(A) \geq F(B \cup \{\sigma\}) - F(B), \quad A \subseteq B \subseteq \Omega,\ \sigma \in \Omega$$

Common property that expresses a diminishing return condition
(e.g. coverage, entropy, mutual information, detection probability)

Maximization via greedy algorithm has provable good performance:
$A^n = \{\sigma_1, \ldots, \sigma_n\}$, $\sigma_n = \operatorname{argmax}_{\sigma \in \Omega} F(A^{n-1} \cup \{\sigma\})$

Theorem

$F : 2^{\Omega} \to \mathbb{R}$ is monotonic and submodular. Let O^T be a maximizing set for F among the sets of cardinality T. Then,

$$F(A^T) \geq \left(1 - \frac{1}{e}\right) F(O^T)$$

Extension to Sequences

Ω finite set, $\mathbb{H}(\Omega)$ set of sequences over Ω, $F : \mathbb{H}(\Omega) \to \mathbb{R}$

- *backward monotonic* if $F(\sigma \perp S) \geq F(S), \quad S \in \mathbb{H}(\Omega),\ \sigma \in \Omega$
- *forward submodular* if

$$F(S \perp \sigma) - F(S) \geq F(S \perp R \perp \sigma) - F(S \perp R), \quad R, S \in \mathbb{H}(\Omega),\ \sigma \in \Omega$$

Greedy algorithm: $S^n = (\sigma_1, \ldots, \sigma_n)$, $\sigma_n = \text{argmax}_{\sigma \in \Omega} F(S^{n-1} \perp \sigma)$

Theorem

$F : \mathbb{H}(\Omega) \to \mathbb{R}$ is backward monotonic and forward submodular. Let O^T be a maximizing sequence for F among the sequences of length T. Then,

$$F(S^T) \geq \left(1 - \frac{1}{e}\right) F(O^T)$$

Search &Tracking

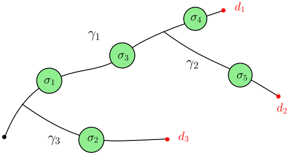

UAV is tasked with searching for a moving target and tracking it to destination. To find the target, the UAV selects a set of flight search patterns Ω based on the possible paths Γ that the target might follow.

$$\sigma \in \Omega \rightarrow \begin{cases} t(\sigma) & \text{reference time} \\ \phi_\sigma & \text{detection probability} \end{cases}$$

$S = (S_1, \ldots, S_n)$, $S_i \in \Omega$, $F(S)$ detection probability

Problem: Minimize the average detection time

$$\min_{|S|\leq T} \sum_{k=1}^{T} t(S_k) \left[F(S|_1^k) - F(S|_1^{k-1})\right] + K(1 - F(S))$$

$$\max_{|S|\leq T} \sum_{k=1}^{T} (K - t(S_k)) \left[F(S|_1^k) - F(S|_1^{k-1})\right]$$

Recommender Systems

Ω set of movies, $g : \Omega \to [0, 1]$ user's satisfaction probability

$\mathcal{T}$ set of topics, $t(\sigma) \subseteq \mathcal{T}$ subset of topics covered by $\sigma \in \Omega$

S sequence of movies generated by the recommender system
$F(S) = | \cup_k \, t(S_k)|$ number of topics covered by S

User chooses a topic t in $\mathcal{T}$ uniformly at random, picks the first item S_k in the sequence for which $t \in t(S_k)$ and is satisfied with probability $g(S_k)$

Problem: Maximization of the user's satisfaction probability

$$\max_{|S| \leq T} \frac{1}{|\mathcal{T}|} \sum_{k=1}^{T} g(S_k) \left[F(S|_1^k) - F(S|_1^{k-1}) \right]$$

Optimization Problem

- Ω finite set, $\mathbb{H}(\Omega)$ set of sequences on Ω
- $g : \Omega \to \mathbb{R}^+$ any function, $F : \mathbb{H}(\Omega) \to \mathbb{R}$ permutation invariant, monotonic, and submodular
- $F_g(S) = \sum_{k=1}^{n} g(S_k) \left[F(S|_1^k) - F(S|_1^{k-1}) \right]$ for $S = (S_1, \ldots, S_n)$

$$\max_{S \in \mathbb{H}(\Omega), |S| \leq T} F_g(S)$$

Optimization Problem

- Ω finite set, $\mathbb{H}(\Omega)$ set of sequences on Ω
- $g : \Omega \to \mathbb{R}^+$ any function, $F : \mathbb{H}(\Omega) \to \mathbb{R}$ permutation invariant, monotonic, and submodular
- $F_g(S) = \sum_{k=1}^{n} g(S_k) \left[F(S|_1^k) - F(S|_1^{k-1}) \right]$ for $S = (S_1, \dots, S_n)$

$$\max_{S \in \mathbb{H}(\Omega), |S| \leq T} F_g(S) = \max_{S \in \mathcal{I}(g), |S| \leq T} F_g(S)$$

where $\mathcal{I}(g) = \{ S \in \mathbb{H}(\Omega) \,|\, g(S_i) \geq g(S_{i+1}) \, \forall i \}$

Optimization Problem

- Ω finite set, $\mathbb{H}(\Omega)$ set of sequences on Ω
- $g : \Omega \to \mathbb{R}^+$ any function, $F : \mathbb{H}(\Omega) \to \mathbb{R}$ permutation invariant, monotonic, and submodular
- $F_g(S) = \sum_{k=1}^{n} g(S_k) \left[F(S|_1^k) - F(S|_1^{k-1}) \right]$ for $S = (S_1, \ldots, S_n)$

$$\max_{S \in \mathbb{H}(\Omega), |S| \leq T} F_g(S) = \max_{S \in \mathcal{I}(g), |S| \leq T} F_g(S)$$

where $\mathcal{I}(g) = \{S \in \mathbb{H}(\Omega) \,|\, g(S_i) \geq g(S_{i+1}) \, \forall i\}$

Remark: other choices for the domain are possible (e.g. $S_i \neq S_j$ for $i \neq j$)

Optimization Problem

- Ω finite set, $\mathbb{H}(\Omega)$ set of sequences on Ω
- $g : \Omega \to \mathbb{R}^+$ any function, $F : \mathbb{H}(\Omega) \to \mathbb{R}$ permutation invariant, monotonic, and submodular
- $F_g(S) = \sum_{k=1}^{n} g(S_k) \left[F(S|_1^k) - F(S|_1^{k-1}) \right]$ for $S = (S_1, \ldots, S_n)$

$$\max_{S \in \mathbb{H}(\Omega), |S| \leq T} F_g(S) = \max_{S \in \mathcal{I}(g), |S| \leq T} F_g(S)$$

where $\mathcal{I}(g) = \{ S \in \mathbb{H}(\Omega) \mid g(S_i) \geq g(S_{i+1}) \, \forall i \}$

Remark: other choices for the domain are possible (e.g. $S_i \neq S_j$ for $i \neq j$)

F_g is forward submodular but not backward monotonic!

Generalized Greedy Algorithm

$J : \mathbb{H}(\Omega) \to \mathbb{R}$, $\mathcal{I} \subseteq \mathbb{H}(\Omega)$, $\max_{S \in \mathcal{I}, |S| \leq T} J(S)$

The algorithm produces recursively an ordered sequence $S^T = (S^T_1, \ldots, S^T_T) \in \mathcal{I}$ as follows:

- $S^1 = (S^1_1)$ where $S^1_1 \in \operatorname{argmax}_{\sigma \in \Omega} J(\sigma)$;
- Given $S^k = (S^k_1, \ldots, S^k_k) \in \mathcal{I}$, we define

$$S^{k+1} = \operatorname*{argmax}_{\substack{U = (S^k_1, \ldots, S^k_h, \sigma, S^k_{h+1}, \ldots, S^k_k) \\ U \in \mathcal{I}}} J(U)$$

Instead of simply augmenting the sequence on the right hand side, we allow each new element to be placed in any position among the elements that are already present in the sequence

Example: Search & Tracking

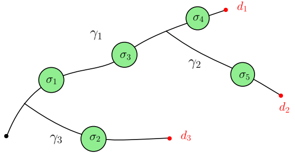

$\Omega = \{\sigma_1, \sigma_2, \sigma_3, \sigma_4, \sigma_5\}$, $t(\sigma_i) = i$, $K = 10$

$\phi_{\sigma_1} = 0.2, \phi_{\sigma_2} = \phi_{\sigma_4} = 0.5, \phi_{\sigma_3} = 0.3, \phi_{\sigma_5} = 0.95$

- $T = 1$: $\operatorname*{argmax}_{\sigma \in \Omega}(K - t(\sigma))F(\sigma) = \{\sigma_5\}$

  Classical greedy stops here! $\Rightarrow$ $S = (\sigma_5)$, Detection time$= 8,40s$

- $T = 2$:
  $\operatorname*{argmax}_{\sigma \in \Omega}[(K - t(\sigma))F(\sigma) + (K - t(\sigma_5))(F(\sigma \perp \sigma_5) - F(\sigma)] = \{\sigma_1\}$

  Generalized greedy: $\Rightarrow$ $S = (\sigma_1, \sigma_5)$, Detection time$= 7,73s$

Technical Results: $\mathcal{I}$-Monotonicity and $\mathcal{I}$-Submodularity

Definition

Let $\mathcal{I} \subseteq \mathbb{H}(\Omega)$ and $J : \mathcal{I} \to \mathbb{R}$.

- J is called *$\mathcal{I}$-monotonic* if for every $Q, R \in \mathcal{I}$ and $\sigma \in \Omega$ such that $Q \perp \sigma \perp R \in \mathcal{I}$, it holds:

$$J(Q \perp \sigma \perp R) \geq J(Q \perp R) \tag{1}$$

- J is called *$\mathcal{I}$-submodular* if for every $Q, R, S \in \mathcal{I}$ and $\sigma_1, \sigma_2 \in \Omega$ such that $Q \perp \sigma_1 R \perp \sigma_2 \perp S \in \mathcal{I}$, it holds:

$$\begin{aligned} &J(Q \perp \sigma_1 \perp R \perp \sigma_2 \perp S) - J(Q \perp \sigma_1 \perp R \perp S) \\ &\leq J(Q \perp R \perp \sigma_2 \perp S) - J(Q \perp R \perp S) \end{aligned} \tag{2}$$

Technical results: Fully Extendable Sets

$Q, R \in \mathbb{H}(\Omega)$, $Q \leq R$ means that Q can be obtained from R eliminating some elements

Definition

A subset $\mathcal{I} \subseteq \mathbb{H}(\Omega)$ is called *fully extendable* if the following conditions are satisfied:

1. For every $\sigma \in \Omega$, $(\sigma) \in \mathcal{I}$;
2. If $R \in \mathcal{I}$ and $Q \leq R$, then $Q \in \mathcal{I}$;
3. If $Q, R \in \mathcal{I}$, there exists $U \in \mathcal{I}$ such that $Q, R \leq U$ and $|U| \leq |Q| + |R|$

Examples:
$$\begin{aligned} \mathcal{I}(g) &= \{S \in \mathbb{H}(\Omega) \,|\, g(S_i) \geq g(S_{i+1}) \,\forall i\} \\ \mathcal{I}^d(g) &= \{S \in \mathbb{H}(\Omega) \,|\, g(S_i) > g(S_{i+1}) \,\forall i\} \end{aligned}$$

Technical Results: Main Theorems

Theorem

Function $J : \mathbb{H}(\Omega) \to \mathbb{R}$ *and fully extendable set* $\mathcal{I} \subseteq \mathbb{H}(\Omega)$. J *is* $\mathcal{I}$*-monotonic and* $\mathcal{I}$*-submodular. Let* O^T *be a maximizing sequence for* J *among the sequences in* $\mathcal{I}$ *of length* T *and let* S^T *be the result of the generalized greedy algorithm. Then,*

$$J(S^T) \geq \left(1 - \frac{1}{e}\right) J(O^T)$$

Theorem

Functional $F : \mathbb{H}(\Omega) \to \mathbb{R}$ *permutation invariant, monotonic, and submodular. Function* $g : \Omega \to \mathbb{R}^+$ *and fully extendable set* $\mathcal{I} \subseteq \mathcal{I}(g)$. *Then,* F_g *is* $\mathcal{I}$*-monotonic and* $\mathcal{I}$*-submodular.*

Experimental Results: Search & Tracking

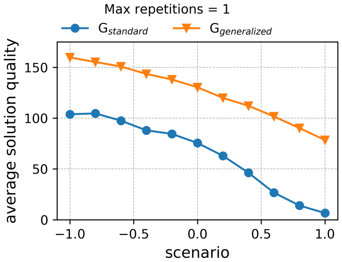

Figure: Average objective values obtained by the standard and the generalized greedy algorithms

Conclusions

- Several applicative problems involve the maximization of a recursive functional with the following structure:
$$F_g(S) = \sum_{k=1}^{n} g(S_k) \left[F(S|_1^k) - F(S|_1^{k-1}) \right] \text{ for } S = (S_1, \ldots, S_n)$$
- Existing greedy algorithms do not yield strong theoretical guarantees for this functional
- We introduce an efficient generalized greedy algorithm that ensures finding solutions that are $(1 - \frac{1}{e})$ from the optimal
- Experiments directly show the power of the new algorithm

Acknowledgements

We thank the anonymous reviewers for their insightful comments.

This work has been supported by EPSRC Grant EP/S016473/1, Leverhulme Trust Grant VP1-2019-037 and MIUR Grant "Dipartimenti di Eccellenza 2018- 2022".