
Tracking Objects with Column Generation

Shaofei Wang
sfwang0928@gmail.com

Steffen Wolf
Heidelberg University, Germany
steffen.wolf@iwr.uni-heidelberg.de

Charless C. Fowlkes
University of California, Irvine
fowlkes@ics.uci.edu

Julian Yarkony
Experian Data Lab, San Diego, CA
julian.e.yarkony@gmail.com

Abstract

We formulate multi-target tracking in video as a maximum weight set packing problem where tracks correspond to sets. We then attack it with column and row generation where pricing is done efficiently using dynamic programming.

1 Introduction

Multi-target tracking in video is often formulated from the perspective of grouping disjoint sets of candidate detections into “tracks” whose underlying trajectories can be estimated using traditional single-target tracking methods such as Kalman filtering. There is a well developed literature on methods for exploring this combinatorial space of possible data associations in order to find collections of low-cost, disjoint tracks.

Our method is most closely related to the Lagrangian relaxation method of [2]. In [2] a large number of short sequences of detections (subtracks) are generated, each of which is associated with a cost. The set of subtracks form the basis from which tracks are constructed. The corresponding objective is attacked via subgradient optimization. In contrast we attack the same problem with column/row generation providing faster inference. We also tighten the bounds by optimizing over additional triplet constraints, which is inspired by [5].

2 Constraint Relaxation for Multi-target Tracking

We now consider our approach. Given a set of candidate detections \mathcal{D} , each with a specified space-time location, our goal is to identify a collection of tracks that describe the trajectories of objects through a scene and the subset of detections associated with each such track.

We denote the set of all possible tracks by \mathcal{P} and use X to denote the detection-track incidence matrix $X \in \{0, 1\}^{|\mathcal{D}| \times |\mathcal{P}|}$ where $X_{dp} = 1$ if and only if track p visits detection d . A solution to the multi-target tracking problem is denoted by the indicator vector $\gamma \in \{0, 1\}^{|\mathcal{P}|}$ where $\gamma_p = 1$ indicates that track p is included in the solution and $\gamma_p = 0$ otherwise. A collection of tracks specified by γ is a valid solution if and only if each detection is associated with at most one active track. Using $\Theta \in \mathbb{R}^{|\mathcal{P}|}$ to denote the costs associated with tracks where Θ_p describes the cost of track p , we express our tracking problem as an integer linear program:

$$\min_{\gamma \in \bar{\Gamma}} \Theta^t \gamma \quad \text{with} \quad \bar{\Gamma} = \{\gamma \in \{0, 1\}^{|\mathcal{P}|} : X\gamma \leq 1\} \quad (1)$$

We note that this is equivalent to finding a maximum-weight set packing which is NP-hard [3].

2.1 Decomposing Track Scores over Subtracks

We consider a general scoring function corresponding to a model in which a track is defined by an ordered sequence of subtracks whose scores in turn depend on detections across several frames. Let \mathcal{S} denote a set of subtracks, each of which contains K detections where K is a user defined modeling parameter that trades off inference complexity and modeling power. For a given subtrack $s \in \mathcal{S}$, let s_k indicate the k 'th detection in the sequence $s = \{s_1, \dots, s_K\}$ ordered by time from earliest to latest. We describe the mapping of subtracks to tracks using $T \in \{0, 1\}^{|\mathcal{S}| \times |\mathcal{P}|}$ where $T_{sp} = 1$ indicates that track p contains subtrack s as a subsequence.

We decompose track costs Θ in terms of the subtrack costs $\theta \in \mathbb{R}^{|\mathcal{S}|}$ where each subtrack s is associated with cost θ_s and use θ_0 to denote a constant cost associated with instancing a track. We define the cost of a track p denoted Θ_p as $\Theta_p = \theta_0 + \sum_{s \in \mathcal{S}} T_{sp} \theta_s$.

2.2 LP Relaxation and Column Generation

We now attack optimization in Eq 1 using the well studied tools of LP relaxations. We use $\Gamma = \{\gamma \in [0, 1]^{|\mathcal{P}|} : X\gamma \leq 1\}$ to denote a convex relaxation of the constraint set $\bar{\Gamma}$.

$$\min_{\gamma \in \Gamma} \Theta^t \gamma \geq \min_{\gamma \in \bar{\Gamma}} \Theta^t \gamma \quad (2)$$

This LP relaxation only contains constraints for collections of tracks that share a common detection. From the view point of maximum-weight set packing, this includes some cliques of conflicting sets but misses many others.

As a concrete example, consider four tracks $\mathcal{P} = \{p_1, p_2, p_3, p_4\}$ over three detections $\mathcal{D} = \{d_1, d_2, d_3\}$ where the first three tracks each contain two of the three detections $\{d_1, d_2\}, \{d_1, d_3\}, \{d_2, d_3\}$, and the fourth track contains all three $\{d_1, d_2, d_3\}$. Suppose the track costs are given by $\Theta_{p_1} = \Theta_{p_2} = \Theta_{p_3} = -4$ and $\Theta_{p_4} = -5$. The optimal integer solution sets $\gamma_{p_4} = 1$, and has a cost of -5 . However the optimal fractional solution sets $\gamma_{p_1} = \gamma_{p_2} = \gamma_{p_3} = 0.5$; $\gamma_{p_4} = 0$ which has cost -6 . Hence the LP relaxation is loose in this case.

A tighter bound can be motivated by the following observation. For any set of three unique detections the number of tracks that pass through two or more members can be no larger than one. We now apply our tighter bound to tracking. We denote the set of groups of three unique detections (which we refer to as triplets) as \mathcal{C} and index it with c . Using $C \in \{0, 1\}^{|\mathcal{C}| \times |\mathcal{P}|}$ we define a tighter space for γ than Γ .

$$\Gamma^C : \{\gamma \in \mathbb{R}^{|\mathcal{P}|} : \gamma \geq 0, \quad X\gamma \leq 1, \quad C\gamma \leq 1\} \quad (3)$$

$$C_{cp} = [\sum_{d \in c} X_{dp} \geq 2] \quad \forall c \in \mathcal{C}, p \in \mathcal{P}$$

3 Optimization over Γ^C

We write tracking as optimization in the primal and dual form below.

$$\min_{\gamma \in \Gamma^C} \Theta^t \gamma = \max_{\substack{\lambda \geq 0 \\ \lambda^c \geq 0 \\ \Theta + X^t \lambda + C^t \lambda^c \geq 0}} -1^t \lambda - 1^t \lambda^c \quad (4)$$

Given that \mathcal{P} and \mathcal{C} are of enormous size we use column and row generation jointly. The nascent subsets of \mathcal{P}, \mathcal{C} are denoted $\hat{\mathcal{P}}, \hat{\mathcal{C}}$ respectively. We write column/row generation optimization given subroutines COLUMN(λ, λ^c), ROW(γ) that identify a group of violated constraints in primal and dual including the most violated in each. We write the column/row generation optimization in Alg 1.

Finding the most violated row consists of the following optimization: $\max_{c \in \mathcal{C}} \sum_{p \in \mathcal{P}} C_{cp} \gamma_p$. We generate its rows as needed by considering only triplets over detections associated with fractional valued tracks.

Algorithm 1 Column/Row Generation

$\hat{\mathcal{P}} \leftarrow \{\}, \hat{\mathcal{C}} \leftarrow \{\}$
repeat
 $\max_{\substack{\lambda \geq 0 \\ \lambda^c \geq 0 \\ \Theta_{\hat{\mathcal{P}}} + X_{(:, \hat{\mathcal{P}})}^t \lambda + C_{(\hat{\mathcal{C}}, \hat{\mathcal{P}})}^t \lambda^c \geq 0}} -1^t \lambda - 1^t \lambda^c$
 Recover γ from λ (provided by LP solver)
 $\hat{\mathcal{P}} \leftarrow \text{COLUMN}(\lambda, \lambda^c) \quad \hat{\mathcal{C}} \leftarrow \text{ROW}(\gamma)$
 $\hat{\mathcal{P}} \leftarrow [\hat{\mathcal{P}}, \hat{\mathcal{P}}] \quad \hat{\mathcal{C}} \leftarrow [\hat{\mathcal{C}}, \hat{\mathcal{C}}]$
until $\hat{\mathcal{P}} = \square$ and $\hat{\mathcal{C}} = \square$

Algorithm 2 Upper Bound Rounding

while $\exists p \in \mathcal{P}$ s.t. $\gamma_p \notin \{0, 1\}$ **do**
 $p^* \leftarrow \arg \min_{\substack{p \in \mathcal{P} \\ \gamma_p > 0}} \Theta_p \gamma_p - \sum_{\hat{p} \in \mathcal{P}_{\perp p}} \gamma_{\hat{p}} \Theta_{\hat{p}}$
 $\gamma_{\hat{p}} \leftarrow 0 \quad \forall \hat{p} \in \mathcal{P}_{\perp p^*}$
 $\gamma_{p^*} \leftarrow 1$
end while
RETURN γ

Figure 1: (Left): Algorithm for dual-optimization of a lower bound on the optimal tracking by column generation where the notation $X_{(:, \hat{\mathcal{P}})}$ denotes selection of a subset of columns of X . (Right) We compute upper-bounds on the optimal tracking using a rounding procedure which greedily selects primal variables γ while removing intersecting tracks.

3.1 Computing $\text{COLUMN}(\lambda, \lambda^c)$ using Dynamic Programming Without Triplets

We now discuss how $\text{COLUMN}(\lambda, \lambda^c)$ is computed efficiently for our track cost model using dynamic programming when λ^c is zero valued. We later show how to use this when λ^c is not zero valued. We specify that a subtrack s may be preceded by a subtrack \hat{s} if and only if the least recent $K - 1$ detections in s correspond to the most recent $K - 1$ detections in \hat{s} . We denote the set of valid subtracks that may precede a subtrack s as $\{\Rightarrow s\}$. We use ℓ_s to denote the cost of the cheapest track that terminates at subtrack s .

$$\ell_s \leftarrow \theta_s + \lambda_{s_K} + \min_{\hat{s} \in \{\Rightarrow s\}} \left\{ \min_{k=0}^{K-1} \ell_{\hat{s}_k}, \theta_0 + \sum_{k=0}^{K-1} \lambda_{s_k} \right\} \quad (5)$$

Empirically we observe large speed ups by adding the most violated track terminating at each subtrack (if one exists) which is easy to extract from the dynamic program. Thus we add many columns per solution to dynamic programming.

3.2 Generating Columns under Triplet Constraints

We denote the value of the slack corresponding to an arbitrary column p as $V(\Theta, \lambda, \lambda^c, p)$ and the most violated as $V^*(\Theta, \lambda, \lambda^c)$ which we define below.

$$V(\Theta, \lambda, \lambda^c, p) = \Theta_p + \sum_{d \in \mathcal{D}} \lambda_d X_{dp} + \sum_{c \in \hat{\mathcal{C}}} \lambda_c^c C_{cp} \quad V^*(\Theta, \lambda, \lambda^c) = \min_{p \in \mathcal{P}} V(\Theta, \lambda, \lambda^c, p) \quad (6)$$

Solving for $V^*(\Theta, \lambda, \lambda^c)$ can not be directly attacked using dynamic programming as in Section 3.1. However dynamic programming can be applied if we ignore the triplet term $\sum_{c \in \hat{\mathcal{C}}} \lambda_c^c C_{cp}$, providing a lower bound.

This invites a branch and bound (B&B) approach. The set of branches in our B&B tree is denoted \mathcal{B} . Each branch $b \in \mathcal{B}$ is defined by two sets \mathcal{D}_{b+} and \mathcal{D}_{b-} . These correspond to detections that must be included on the track and those that must not be included on the track respectively. We write the set of all tracks that are consistent with a given \mathcal{D}_{b-} , \mathcal{D}_{b+} or consistent with both \mathcal{D}_{b-} and \mathcal{D}_{b+} as \mathcal{P}_{b-} , \mathcal{P}_{b+} and $\mathcal{P}_{b\pm}$ respectively. The initial branch b is defined by $\mathcal{D}_{b+} = \mathcal{D}_{b-} = \{\}$.

Bounding Operation: Let $V^b(\Theta, \lambda, \lambda^c)$ denote the value of the most violating slack over columns in $\mathcal{P}_{b\pm}$. We can compute a lower-bound for this value, denoted V_{lb}^b by independently optimizing the




Figure 2: When training on a subset of motion features on MOT dataset we get instances with loose bound. For the two examples we plot the gap (absolute value of the difference) between the bounds and the final lower bound as a function of time. We indicate each time that a triplet is added with a black dot on the lower bound plot. In all examples the bound of [2] is loose and at least one triplet is needed to produce a tight bound which results in visually compelling trackings.

dynamic program and the triplet penalty.

$$\begin{aligned}
 V^b(\Theta, \lambda, \lambda^c) &= \min_{p \in \mathcal{P}_{b\pm}} V(\Theta, \lambda, \lambda^c, p) \geq \min_{p \in \mathcal{P}_{b-}} \Theta_p + \sum_{d \in \mathcal{D}} \lambda_d X_{dp} + \min_{p \in \mathcal{P}_{b+}} \sum_{c \in \hat{\mathcal{C}}} \lambda_c^c C_{cp} \\
 &\geq \min_{p \in \mathcal{P}_{b-}} \Theta_p + \sum_{d \in \mathcal{D}} \lambda_d X_{dp} + \sum_{c \in \hat{\mathcal{C}}} \lambda_c^c [\sum_{d \in \mathcal{D}} [d \in \mathcal{D}_{b+}] \geq 2] = V_{lb}^b(\Theta, \lambda, \lambda^c)
 \end{aligned}$$

Observe that dynamic programming can be used to efficiently search over \mathcal{P}_{b-} to minimize the first term.

Branch Operation: We now consider the branch operation. We describe an upper bound on $V^b(\Theta, \lambda, \lambda^c)$ as $V_{ub}^b(\Theta, \lambda, \lambda^c)$. This is constructed by adding in the active λ^c terms ignored when constructing $V_{lb}^b(\Theta, \lambda, \lambda^c)$. Let $p_b = \arg \min_{p \in \mathcal{P}_{b-}} \Theta_p + \sum_{d \in \mathcal{D}} \lambda_d X_{dp}$.

$$\begin{aligned}
 V_{ub}^b(\Theta, \lambda, \lambda^c) &= V_{lb}^b(\Theta, \lambda, \lambda^c) + \sum_{c \in \hat{\mathcal{C}}} \lambda_c^c C_{cp_b} [\sum_{d \in \mathcal{D}} [d \in \mathcal{D}_{b+}] < 2] \quad (7) \\
 &= \Theta_{p_b} + \sum_{d \in \mathcal{D}} \lambda_d X_{dp_b} + \sum_{c \in \hat{\mathcal{C}}} \lambda_c^c [\sum_{d \in \mathcal{D}} [d \in \mathcal{D}_{b+}] \geq 2] + \sum_{c \in \hat{\mathcal{C}}} \lambda_c^c C_{cp_b} [\sum_{d \in \mathcal{D}} [d \in \mathcal{D}_{b+}] < 2] \\
 &\geq V(\Theta, \lambda, \lambda^c, p_b) \geq V^b(\Theta, \lambda, \lambda^c) \quad (8)
 \end{aligned}$$

Now consider the largest triplet constraint term λ_c^c that is included in $V_{ub}^b(\Theta, \lambda, \lambda^c, p_b)$ but not $V_{lb}^b(\Theta, \lambda, \lambda^c)$. $c^* \leftarrow \arg \max_{c \in \hat{\mathcal{C}}} \lambda_c^c C_{cp_b} [\sum_{d \in \mathcal{D}} [d \in \mathcal{D}_{b+}] < 2]$

We create eight child branches of b where there is one for each of the eight different ways of splitting the detections in the triplet corresponding to c^* between the include (+) and exclude (-) sets.

We compute upper bounds using a fast principled method that avoids resolving the LP. We round fractional γ via a greedy iterative approach described in Alg 2. We rely on the technique of [6] (supplement Section 4) to produce the following lower bound: $\min_{\gamma \in \Gamma} \Theta^t \gamma \geq -1^t \lambda + \sum_{d \in \mathcal{D}} \min\{0, \min_{s \in \mathcal{S}} \ell_s\}$.

4 Experiments

We use a part of MOT 2015 training set [4] to train and evaluate real-world tracking models. For $K = 2$ we observe 48.5% Multiple Object Tracking Accuracy [1], 11 identity switches and 9 track fragments or for short hand (48.5,11,9). However when setting $K = 3, 4$ the performance is (49,10,7), and (49.9,9,7) which constitutes noticeable improvements over all three metrics. In Fig 2 we compared the timing/cost performance of our algorithm with the baseline algorithm of [2].

References

- [1] K. Bernardin and R. Stiefelhagen. Evaluating multiple object tracking performance: The clear mot metrics. *Journal on Image Video Processing*, 2008.
- [2] A. Butt and R. Collins. Multi-target tracking by lagrangian relaxation to min-cost network flow. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 1846–1853, June 2013.
- [3] R. M. Karp. Reducibility among combinatorial problems. The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972.
- [4] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler. MOTChallenge 2015: Towards a benchmark for multi-target tracking. *arXiv:1504.01942 [cs]*, Apr. 2015. arXiv: 1504.01942.
- [5] D. Sontag, T. Meltzer, A. Globerson, T. Jaakkola, and Y. Weiss. Tightening lp relaxations for map using message passing. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 503–510, july 2008.
- [6] J. Yarkony and C. Fowlkes. Planar ultrametrics for image segmentation. In *Neural Information Processing Systems*, 2015.