
Exploiting skeletal structure in computer vision annotation with Benders decomposition

Shaofei Wang
Beijing A&E Technologies
Beijing China
sfwang0928@gmail.com

Konrad Kording
University of Pennsylvania
Philadelphia PA
koerding@gmail.com

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

Abstract

Many annotation problems in computer vision can be phrased as integer linear programs (ILPs). The use of standard industrial solvers does not exploit the underlying structure of such problems *e.g.*, the skeleton in pose estimation. The leveraging of the underlying structure in conjunction with industrial solvers promises increases in both speed and accuracy. Such structure can be exploited using Bender's decomposition, a technique from operations research, that solves complex ILPs or mixed integer linear programs by decomposing them into sub-problems that communicate via a master problem. The intuition is that conditioned on a small subset of the variables the solution to the remaining variables can be computed easily by taking advantage of properties of the ILP constraint matrix such as block structure. In this paper we apply Benders decomposition to a typical problem in computer vision where we have many sub-ILPs (*e.g.*, partitioning of detections, body-parts) coupled to a master ILP (*e.g.*, constructing skeletons). Dividing inference problems into a master problem and sub-problems motivates the development of a plethora of novel models, and inference approaches for the field of computer vision.

1 Introduction

1.1 Human Pose Estimation

For a wide variety of applications in sports science, training, and rehabilitation, the quantification of human movement is crucial. While this is typically done in specialized labs, it is possible to estimate the three dimensional position of body parts (jointly referred to as a pose for a given person) of subjects from single camera video. Such problems are often modeled using an underlying tree structured model of the human body [12] that models the following two components. The first softly enforces that part locations of a predicted person are supported by evidence in the image as described by local image features [8, 26]. The second component softly enforces that the part locations of a predicted person satisfy the angular and distance relationships consistent with a person [30]. An example of such a relationship is that the head of a person tends to be above the neck. To deal with multiple people in an image, sliding window detectors[25, 8, 11, 29] are employed which check for people across multiple positions, scales, orientations etc. These approaches, however, do not explicitly model the presence and joint likelihood of multiple people outside of tree structured models [29], which is necessary for meaningful multi-person pose estimation(MPPE).

1.2 ILP Formulation for MPPE

MPPE has been modeled as a correspondence problem in [22, 18] where candidate detections (of body parts) are grouped to form each pose. This correspondence problem is formulated as an ILP that labels detections as either inactive or associated with a given part and then clusters the active detections into poses according to the multi-cut clustering criteria. The objective and part detections

for this ILP are produced by a deep neural network[16, 23] and a spring model for the distances between body parts.

1.3 Column Generation

Column generation [4] is an efficient way of solving linear programs (LPs) that have certain structural properties, namely a block structure exhibited by many LPs and ILPs including cutting stock [15], routing vehicles[9], facility location [3], crew scheduling[24], and more recently in computer vision, MPPE[28], hierarchical image segmentation[31], and multi-object tracking [27]. In such cases a new LP relaxation, that is often much tighter than a compact formulation [13, 2], is created that includes an enormous number of variables and a finite number of constraints.

The corresponding optimization problem is solved by iteratively solving the dual LP with a limited set of the dual constraints followed by identifying violated dual constraints using combinatorial optimization (usually dynamic programming) which are then added to the set under consideration. This process terminates when no remaining dual constraints are violated at which point the corresponding primal LP solution is guaranteed to be optimal. The process of identifying violated dual constraints is commonly called “pricing” in operations research.

1.4 Two-Tier Formulation of Multi-Person Pose Estimation

An alternative model of the MPPE problem is the two-tier formulation (TTF)[28]. TTF can be understood as an approximation to [22, 18] such that exact or near exact inference can be achieved efficiently. TTF starts with detections of body parts from a deep convolutional neural network [10, 20] each of which is associated with exactly one body part. Next TTF reasons about the assignment of parts to poses in two tiers: *local assignments*, which correspond to groupings of detections each of which is associated with a single part in a single pose; and *skeletons*, each of which models the outline of a pose and in which each part is associated with no more than one detection.

TTF then solves an ILP formulation over the skeletons and local assignments using costs derived identically to [22, 18]. TTF is designed with inference via column generation in mind and leverages dynamic programming to construct skeletons and explicit enumeration for local assignments.

1.5 Benders Decomposition

In order to exploit the structure of TTF we find further inspiration in operations research in the form of Bender’s decomposition[5]. In Benders decomposition, the variables of the original problem are divided into at least two subsets, one of which is identified with a LP (or ILP) called the master problem with the remainder identified with LPs called sub-problems (also called Benders sub-problems). Conditioned on the solution to the master problem each of the Benders problems is solved independently. Optimization of the master problem relies on transforming the sub-problems to their dual forms and rewriting optimization over only variables associated with the master problem. Solutions to the dual sub-problems form constraints in the master that are added on demand over the course of optimization.

Benders decomposition is especially powerful in domains in which both of the following two conditions hold.

- The LP relaxation of the ILP is inclined to produce fractional values for the variables of the master problem.
- Conditioned on an integer solution to the master problem, the variables of the sub-problems tend not to be fractional or are allowed to be fractional or non-negative.

Such problems are common in the routing literature in which the master variables correspond to pipes, depots, or routes of vehicles while the sub-problem variables correspond to flows of commodities or passengers [7, 6, 14]. An example is the problem of selecting a set of depots to open and jointly allocating flow of commodities from depots to clients. In this case depots are associated with binary variables in the master problem and flows of commodities from depots is considered in the sub-problem. The cost function trades off the fixed cost of opening depots and the variable cost of transporting commodities from depots to clients.

We now consider a modification of the example above that is particularly amenable to Benders decomposition. Consider that multiple commodities are modeled and that conditioned on the depots opened that routing of the commodities can be done independently. In this case there is one Benders sub-problem for each commodity which is a highly desirable feature that accelerates optimization.

In our Benders decomposition approach we employ a master problem that determines the skeletons and one Benders sub-problem for each body part that associated the local assignments for that body part to the skeletons.

1.6 Our Approach

In this paper we introduce Benders decomposition to computer vision and apply it to optimization over the TTF formulation of MPPE. First we review the TTF formulation of [28]. Next we propose a Benders decomposition formulation for TTF. Then we contrast our approach to the approach of [28]. Next we demonstrate efficiency of inference on problems from the MPPE benchmark [1]. Finally we conclude and discuss extensions.

2 The TTF Model

Body parts and the corresponding detections. Let the sets of detections, parts be denoted \mathcal{D}, \mathcal{R} , and indexed by d, r respectively. We use \mathcal{R}' to denote the subset of \mathcal{R} corresponding to major parts of which at least one must be included in any given pose. We describe the correspondence between detections and parts using a vector $R \in \{\mathcal{R}\}^{|\mathcal{D}|}$ which we index by d where R_d is the part associated with detection d .

Skeletons. We define the set of possible skeletons over \mathcal{D} as \mathcal{G} . Each element of \mathcal{G} contains at least one detection corresponding to a major part and no more than one detection corresponding to any given part. Mappings of detections to skeletons are expressed by a matrix $G \in \{0, 1\}^{|\mathcal{D}| \times |\mathcal{G}|}$, $G_{dg} = 1 \iff$ detection d is included in skeleton g . We maintain an active set of skeletons, $\hat{\mathcal{G}}$, since \mathcal{G} is impractically large.

Local Assignments. The set of possible local assignments over the detections \mathcal{D} is denoted \mathcal{L} , which we index by l . We describe \mathcal{L} using the matrices $L, M \in \{0, 1\}^{|\mathcal{D}| \times |\mathcal{L}|}$, where $L_{dl} = 1 \iff$ detection d is associated with l as a local detection, and $M_{dl} = 1 \iff$ detection d is associated with l as a global detection. We maintain an active set of local assignments, $\hat{\mathcal{L}}$, since \mathcal{L} is impractically large.

Validity Constraints. We use indicator vectors to describe a selection of skeletons and local assignments. We use $\gamma \in \{0, 1\}^{|\mathcal{G}|}$ where $\gamma_g = 1$ indicates that skeleton $g \in \mathcal{G}$ is selected, and $\gamma_g = 0$ otherwise. Similarly, we let $\psi \in \{0, 1\}^{|\mathcal{L}|}$ where $\psi_l = 1$ indicates that local assignment $l \in \mathcal{L}$ is selected, with $\psi_l = 0$ otherwise. We describe the validity of pair (γ, ψ) using the following linear inequalities:

- $-G\gamma + M\psi \leq 0$: The global detection of a selected local assignment is included in some selected skeleton.
- $L\psi + M\psi \leq 1$: Detections can not be shared between selected local assignments.
- $G\gamma + L\psi \leq 1$: A detection can only be global, local, or neither and can not be shared between selected skeletons.

Cost Function. We associate γ, ψ with a cost function which is described using $\theta \in \mathbb{R}^{|\mathcal{D}|}, \phi \in \mathbb{R}^{|\mathcal{D}| \times |\mathcal{D}|}, \omega \in \mathbb{R}$. Here $\theta_d, \phi_{d_1 d_2}$ are respectively the costs of assigning detection d to a pose, and d_1 and d_2 to a common skeleton or local assignment. We use ω to denote the cost of instancing a pose which regularizes the number of people predicted. TTF associates costs to skeletons and local assignments with $\Gamma \in \mathbb{R}^{|\mathcal{G}|}$ and $\Psi \in \mathbb{R}^{|\mathcal{L}|}$ where Γ_g, Ψ_l are the costs of skeleton g and local

assignment l respectively, which are defined below.

$$\begin{aligned}\Gamma_g &= \omega + \sum_{d \in \mathcal{D}} \theta_d G_{dg} + \sum_{\substack{d_1 \in \mathcal{D} \\ d_2 \in \mathcal{D}}} \phi_{d_1 d_2} G_{d_1 g} G_{d_2 g} \\ \Psi_l &= \sum_{d \in \mathcal{D}} \theta_d L_{dl} + \sum_{\substack{d_1 \in \mathcal{D} \\ d_2 \in \mathcal{D}}} \phi_{d_1 d_2} (L_{d_1 l} + M_{d_1 l})(M_{d_2 l} + L_{d_2 l})\end{aligned}\quad (1)$$

Integer Linear Program. TTF formulates MPPE as the following ILP which selects the lowest total cost valid set of skeletons, local assignments.

$$\begin{aligned}\min_{\substack{\gamma \in \{0,1\}^{|\mathcal{G}|} \\ \psi \in \{0,1\}^{|\mathcal{L}|}}} & \Gamma^\top \gamma + \Psi^\top \psi \\ \text{s.t.} & G\gamma + L\psi \leq 1 \\ & L\psi + M\psi \leq 1 \\ & -G\gamma + M\psi \leq 0\end{aligned}\quad (2)$$

3 Benders Solution

In this section we develop our Benders decomposition based solution for optimization in the TTF model. We begin by adding the redundant constraint to Eq 2 that no global detection is associated with more than one selected skeleton ($G\gamma \leq 1$) then relaxing integrality.

$$\text{Eq 2} \geq \min_{\substack{\gamma \geq 0 \\ \psi \geq 0 \\ G\gamma \leq 1 \\ G\gamma + L\psi \leq 1 \\ M\psi + L\psi \leq 1 \\ M\psi - G\gamma \leq 0}} \Gamma^\top \gamma + \Psi^\top \psi \quad (3)$$

We use the G^r, L^r, M^r to refer to the matrices corresponding to the subset of the rows of G, L, M over detections of part r . Similarly $\psi^r, \Psi^r, \mathcal{L}^r$ concern all local assignments over part r . We now rewrite optimization using this separation.

$$\text{Eq 3} = \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1}} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} \min_{\substack{\psi^r \geq 0 \\ G^r \gamma + L^r \psi^r \leq 1 \\ M^r \psi^r + L^r \psi^r \leq 1 \\ M^r \psi^r - G^r \gamma \leq 0}} \Psi^{\top r} \psi^r \quad (4)$$

We refer to the optimizations over ψ^r terms as sub-problems and introduce their dual form using dual variables $\lambda^{1r}, \lambda^{2r}, \lambda^{3r}$ each of which lie in $\mathbb{R}_+^{\mathcal{D}}$ for each $r \in \mathcal{R}$.

$$\begin{aligned}\text{Eq 4} &= \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1}} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} \max_{\substack{\lambda^{1r} \geq 0 \\ \lambda^{2r} \geq 0 \\ \lambda^{3r} \geq 0}} -1^\top \lambda^{1r} - 1^\top \lambda^{2r} + \gamma^\top G^{r\top} \lambda^{1r} - \gamma^\top G^{r\top} \lambda^{3r} \\ \text{s.t.} & \Psi^r + L^{r\top} \lambda^{1r} + M^{r\top} \lambda^{2r} + L^{r\top} \lambda^{2r} + M^{r\top} \lambda^{3r} \geq 0\end{aligned}\quad (5)$$

Notice that feasibility of the dual problem is not affected by the value of γ . We now consider the set of all dual feasible solutions for sub-problem r as Λ^r and the union of such sets as Λ and index such sets with i . We use $\lambda^{i1r}, \lambda^{i2r}, \lambda^{i3r}$ to describe a member i of Λ^r . We refer to elements in Λ^r for any r as Benders rows or rows for shorthand. We now rewrite optimization using $-\ell_r$ to denote the objective of the sub-problem r . Notice that the optimal sub-problem objective is always non-positive regardless of the skeletons chosen (thus ℓ_r must be non-negative).

$$\begin{aligned}\text{Eq 5} &= \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1 \\ \ell_r \geq 0}} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} -\ell_r \\ \text{s.t.} & -\ell_r \geq -1^\top \lambda^{i1r} - 1^\top \lambda^{i2r} + \gamma^\top G^{r\top} \lambda^{i1r} - \gamma^\top G^{r\top} \lambda^{i3r} \quad \forall r \in \mathcal{R}, i \in \Lambda^r\end{aligned}\quad (6)$$

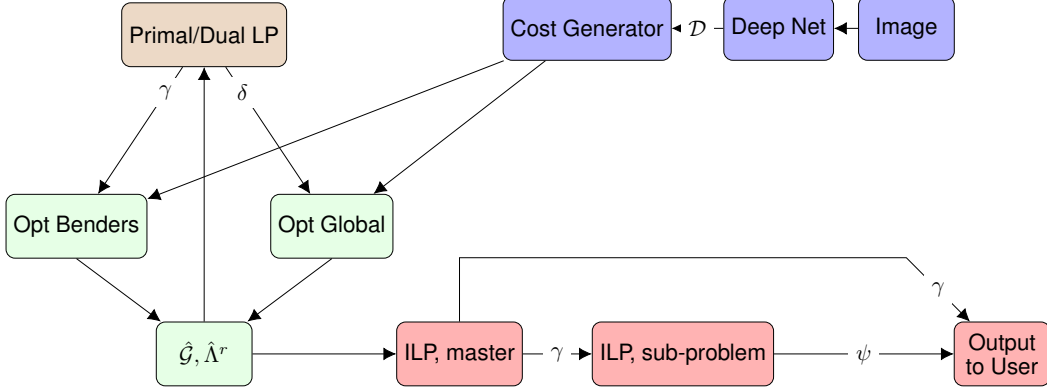


Figure 1: Diagram of our Benders system: blue blocks represent steps for generating unary and pairwise costs. Green blocks represent steps for generating columns/rows. **Opt Benders** and **Opt Global** correspond to the pricing problems. The brown block represents a primal/dual LP solver while red blocks show steps for producing the final integer solutions at termination.

We write the dual form of Eq 6 below using Lagrange multipliers μ . Here μ has one index for each d (denoted μ_d^0) and one index for each (i,r) pair (denoted μ_{ir}).

$$\text{Eq 6} = \max_{\substack{\mu_d^0 \geq 0 \\ \mu_{ir} \geq 0 \\ \sum_{i \in \Lambda^r} \mu_{ir} \geq 1}} \sum_{d \in \mathcal{D}} -\mu_d^0 + \sum_{\substack{r \in \mathcal{R} \\ i \in \Lambda^r}} \mu_{ir} (-1^\top \lambda^{i1r} - 1^\top \lambda^{i2r}) \quad (7)$$

$$\text{s.t.} \quad \Gamma_g + \sum_{d \in \mathcal{D}} \mu_d^0 G_{dg} + \sum_{\substack{r \in \mathcal{R} \\ i \in \Lambda^r}} \mu_{ir} \sum_{d \in \mathcal{D}} (\lambda_d^{i1r} - \lambda_d^{i3r}) G_{dg}^r \geq 0 \quad \forall g \in \mathcal{G}$$

3.1 Solving Benders Formulation

Since we can not enumerate all possible Benders rows or skeletons we iteratively construct small subsets sufficient for optimality. We use $\hat{\Lambda}^r, \hat{\mathcal{G}}$ to denote the nascent subsets of Λ^r, \mathcal{G} and write the corresponding optimization below.

$$\begin{aligned} & \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1 \\ \gamma_g = 0 \forall g \notin \hat{\mathcal{G}} \\ \ell_r \geq 0}} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} -\ell_r \quad (8) \\ & \text{s.t.} \quad -\ell_r \geq -1^\top (\lambda^{i1r} + \lambda^{i2r}) + \gamma^\top G^{r^\top} \lambda^{i1r} - \gamma^\top G^{r^\top} \lambda^{i3r} \quad \forall r \in \mathcal{R}, i \in \hat{\Lambda}^r \\ = & \max_{\substack{\mu_d^0 \geq 0 \\ \mu_{ir} \geq 0 \\ \sum_{i \in \Lambda^r} \mu_{ir} \geq 1 \\ \mu_{ir} = 0 \forall r \in \mathcal{R}, i \notin \hat{\Lambda}^r}} \sum_{d \in \mathcal{D}} \mu_d^0 + \sum_{\substack{r \in \mathcal{R} \\ i \in \Lambda^r}} \mu_{ir} (-1^\top \lambda^{i1r} - 1^\top \lambda^{i2r}) \\ & \text{s.t.} \quad \Gamma_g + \sum_{d \in \mathcal{D}} \mu_d^0 G_{dg} + \sum_{\substack{r \in \mathcal{R} \\ i \in \Lambda^r}} \mu_{ir} \sum_{d \in \mathcal{D}} (\lambda_d^{i1r} - \lambda_d^{i3r}) G_{dg}^r \geq 0 \quad \forall g \in \hat{\mathcal{G}} \end{aligned}$$

We now formally study the solution procedure for the primal and dual form in Eq 8 which are solved jointly. We diagram our procedure in Fig 1. For each $r \in \mathcal{R}$ we initialize $\hat{\Lambda}^r$ with the solution corresponding to $\gamma_g = 0$ for all g and initialize $\hat{\mathcal{G}}$ with the empty set. We then proceed by solving the LP in Eq 8 producing both the primal solution and the dual solution. We use the dual solution to find violated dual constraints corresponding to \mathcal{G} and the primal solution to find violated constraints in Λ^r for all $r \in \mathcal{R}$. We repeat these steps until convergence. We study the generation of members of $\hat{\mathcal{G}}$ and $\hat{\Lambda}^r$ in the succeeding subsections.

Algorithm 1 Benders Decomposition Dual Optimization

```

1:  $\hat{\mathcal{G}} \leftarrow \emptyset$ 
2:  $\hat{\mathcal{L}}^r \leftarrow \emptyset, \forall r \in \mathcal{R}$ 
3:  $\hat{\Lambda}^r \leftarrow \{ i \leftarrow \max_{\substack{\lambda^{1r} \geq 0 \\ \lambda^{2r} \geq 0 \\ \lambda^{3r} \geq 0}} -1^\top \lambda^{1r} - 1^\top \lambda^{2r} \text{ over dual feasible space } \}$ 
4: repeat
5:   did_change  $\leftarrow$  0
6:    $[\lambda, \gamma] \leftarrow$  Maximize primal/dual in Eq. (14) over column sets  $\hat{\mathcal{G}}, \hat{\Lambda}^r$ 
7:   for  $d_* \in \mathcal{D}$  s.t.  $R_{d_*} \in \mathcal{R}'$  do
8:      $g_* \leftarrow \arg \min_{\substack{g \in \mathcal{G} \\ G_{d_* g} = 1}} \Gamma_g + \sum_{d \in \mathcal{D}} G_{dg} \delta_d$ 
9:     if  $\Gamma_{g_*} + \sum_{d \in \mathcal{D}} G_{dg_*} \delta_d < 0$  then
10:       $\hat{\mathcal{G}} \leftarrow [\hat{\mathcal{G}} \cup g_*]$ 
11:      did_change  $\leftarrow$  1
12:    end if
13:  end for
14:  for  $r \in \mathcal{R}$  do
15:    repeat
16:       $\lambda_*^r \leftarrow \max_{\substack{\lambda^{1r} \geq 0 \\ \lambda^{2r} \geq 0 \\ \lambda^{3r} \geq 0}} -1^\top \lambda^{1r} - 1^\top \lambda^{2r} + \gamma^\top G^{r\top} \lambda^{1r} - \gamma^\top G^{r\top} \lambda^{3r}$  over dual feasible space
17:       $\dot{\mathcal{L}}^r \leftarrow \emptyset$ 
18:      for  $d_* \in \mathcal{D}$  s.t.  $R_{d_*} = r$  do
19:         $l_* \leftarrow \arg \min_{\substack{l \in \mathcal{L}^r \\ M_{d_* l}^r = 1}} (\lambda_{d_*}^{3r} + \lambda_{d_*}^{2r}) M_{d_* l}^r + \sum_{d \in \mathcal{D}} (\lambda_d^{1r} + \lambda_d^{2r}) L_{dl}^r + \Psi_l^r$ 
20:        if  $(\lambda_{d_*}^{3r} + \lambda_{d_*}^{2r}) M_{d_* l_*}^r + \sum_{d \in \mathcal{D}} (\lambda_d^{1r} + \lambda_d^{2r}) L_{dl_*}^r + \Psi_{l_*}^r < 0$  then
21:           $\dot{\mathcal{L}} \leftarrow [\dot{\mathcal{L}} \cup l_*]$ 
22:          did_change  $\leftarrow$  1
23:        end if
24:      end for
25:       $\dot{\mathcal{L}}^r \leftarrow [\dot{\mathcal{L}}^r, \dot{\mathcal{L}}^r]$ 
26:    until  $|\dot{\mathcal{L}}^r| = 0$ 
27:     $i \leftarrow$  index of  $\Lambda^r$  corresponding to final  $\lambda_*^r$ 
28:     $\hat{\Lambda}^r \leftarrow [\hat{\Lambda}^r \cup i]$ 
29:  end for
30: until did_change=0

```

3.2 Constructing Nascent Set $\hat{\mathcal{G}}$

We consolidate the dual terms as follows using $\delta \in \mathbb{R}^{|\mathcal{D}|}$ defined as follows for index d .

$$\delta_d = \mu_d^0 + \sum_{i \in \Lambda^{R_d}} \mu_{ir} (\lambda_d^{i1r} - \lambda_d^{i3r}) \quad (9)$$

For each detection d_* such that $R_{d_*} \in \mathcal{R}'$ (i.e., d_* corresponds to a major part), we compute the most violated constraint corresponding to a skeleton that includes detection d_* as follows. We use indicator vector $x \in \{0, 1\}^{|\mathcal{D}|}$, and define a new column g to be included in $\hat{\mathcal{G}}$, defined by $G_{dg} = x_d^*$ for all $d \in \mathcal{D}$, where x^* is the solution to:

$$\begin{aligned} & \min_{\substack{g \in \mathcal{G} \\ G_{d_* g} = 1}} \Gamma_g + \sum_{d \in \mathcal{D}} G_{dg} \delta_d \quad (10) \\ & = \omega + \min_{\substack{x \in \{0, 1\}^{|\mathcal{D}|} \\ x_{d_*} = 1 \\ \sum_{d \in \mathcal{D}} R_{dr} x_d \leq 1 \quad \forall r \in \mathcal{R}}} \sum_{d \in \mathcal{D}} (\theta_d + \delta_d) x_d + \sum_{d_1 d_2 \in \mathcal{D}} \phi_{d_1 d_2} x_{d_1} x_{d_2} \end{aligned}$$

We employ a tree structured model augmented with connections from the neck (which is the only major part) to all other non-adjacent body parts. Conditioned on the global detection associated with the neck, the conditional model is tree-structured and can be optimized using dynamic programming.

3.3 Constructing Nascent Set $\hat{\Lambda}^r$

Generating the most violated Benders row conditioned on γ for a given r is a simple linear programming problem.

$$\begin{aligned} \max \quad & -1^\top \lambda^{1r} - 1^\top \lambda^{2r} + \gamma^\top G^{r\top} \lambda^{1r} - \gamma^\top G^{r\top} \lambda^{3r} \\ & \lambda^{1r} \geq 0 \\ & \lambda^{2r} \geq 0 \\ & \lambda^{3r} \geq 0 \\ \text{s.t.} \quad & \Psi^r + L^{r\top} \lambda^{1r} + M^{r\top} \lambda^{2r} + L^{r\top} \lambda^{2r} + M^{r\top} \lambda^{3r} \geq 0 \end{aligned} \quad (11)$$

We apply this optimization for each r and apply column generation to produce local assignments to solve the LP. Notice that this decouples from the master problem any need to consider the operations in the sub-problems.

The set of local assignments for any part r is too large to consider in optimization but not too large to enumerate in advance. Thus to solve Eq 11 we employ an explicit column generation approach where we iteratively solve the dual LP with a limited set of constraints and then identify violated constraints, which are added to the set under consideration.

3.3.1 Crucial Point in Optimization

When solving optimization in Eq 11 we add a tiny negative bias to objective corresponding to terms $\lambda^{1r}, \lambda^{3r}$. This ensures that the corresponding terms do not increase beyond what is needed to produce an optimal dual solution.

The additional small biases can be understood as ensuring that $\lambda_d^{1r} - \lambda_d^{3r}$ becomes the marginal cost for adding a local assignment to the solution in which d is the global detection. This addition does not affect the optimality of the solution if the bias is small enough in magnitude.

3.4 Anytime Lower Bounds

In the Appendix we consider the production of anytime lower bounds in the master problem and Benders sub-problems. We write the corresponding bounds below and leave their derivations to the Appendix.

We write below a lower bound corresponding to Benders sub-problem r given any non-negative λ .

$$\lambda^{1\top} (-1 + G^r \gamma) - \lambda^{2\top} 1 - \lambda^{3\top} G^r \gamma + \sum_{\substack{d \in \mathcal{D} \\ R_d = r}} \min[0, \min_{\substack{l \in \mathcal{L} \\ M_{dl} = 1}} \Psi_l] \quad (12)$$

We write below a lower bound corresponding to the master problem given any set of Benders rows and Lagrange multipliers λ .

$$-\mu^{0\top} 1 - \sum_{\substack{r \in \mathcal{R} \\ i \in \Lambda^r}} \mu_{ir} (1^\top (\lambda^{i1r} + \lambda^{i2r})) + \sum_{\substack{\hat{d} \in \mathcal{R}' \\ G_{\hat{d}g} = 1}} \min_{g \in \mathcal{G}} \gamma_g (\Gamma_g + \sum_{d \in \mathcal{D}} G_{dg} (\mu_d^0 + \sum_{\substack{r \in \mathcal{R} \\ i \in \Lambda^r}} \mu_{ir} (\lambda_d^{i1r} - \lambda_d^{i3r}))) \quad (13)$$

4 Comparison to Column Generation Formulation for TTF

Our approach can be contrasted with the approach of [28] which employs an optimization approach that we refer to as pure column generation (PCG) in contrast to our approach which we refer to as Benders Column generation (BCG). PCG relaxing the integrality constraints on γ, ψ producing the




Figure 2: Scatter plot of the normalized gap vs time

following primal/dual optimization over $\lambda^1, \lambda^2, \lambda^3 \in \mathbb{R}_{0+}^{|\mathcal{D}|}$:

$$\begin{array}{ll}
 \min & \Gamma^\top \gamma + \Psi^\top \psi = \\
 \begin{array}{l}
 \gamma \geq 0 \\
 \psi \geq 0 \\
 G\gamma + L\psi \leq 1 \\
 L\psi + M\psi \leq 1 \\
 -G\gamma + M\psi \leq 0
 \end{array}
 & \begin{array}{l}
 \max \\
 \lambda^1 \geq 0 \\
 \lambda^2 \geq 0 \\
 \lambda^3 \geq 0 \\
 \Gamma + G^\top (\lambda^1 - \lambda^3) \geq 0 \\
 \Psi + L^\top \lambda^1 + (M^\top + L^\top) \lambda^2 + M^\top \lambda^3 \geq 0
 \end{array}
 \end{array} \quad -1^\top \lambda^1 - 1^\top \lambda^2 \quad (14)$$

Optimization in Eq 14 consists of iteratively solving the dual form of Eq 14 using a limited set of the primal variables followed by identifying violated dual constraints. This identification is done using dynamic programming and explicit column generation as in BCG.

The BCG and PCG strategies attack the same LP relaxation and thus achieve the same objective. BCG however is distinct in that it solves the optimizations associated with local assignments entirely separately from the master problem. In contrast PCG reasons about both local assignments and skeletons along with all constraints in a common LP. Both however use identical mechanisms to identify skeletons and local assignments that are added to consideration during optimization.

5 Experiments

We evaluate our approach against the baseline of [28] using identical problem instances. These problem instances are derived from that of [18] and provided generously by its authors of [28].

We have capped computation time at 300 seconds of dual optimization. Afterwards we solve the ILP over the primal variables generated.

We show a scatter plot of the gap between the upper and lower bound vs termination time in Fig 2. To normalize these gaps we divide by the minus one times maximum lower bound produced at termination for each problem instance. We also plot a histogram of the computation time and gap in Fig 3 and Fig 4 respectively.

We have results on 688 problem instances currently. We observed some numerical issues in the LP solver that we do not see in PCG. All experiments are conducted with MATLAB 2016 with its built in linprog LP solver using a 2014 MacBook Pro with 2.6 GHz Intel Core i5 chip.




Figure 3: Cumulative histogram displaying proportion of problems with consuming less than a given amount of time.




Figure 4: Cumulative histogram of proportion of problems with gap less than a given value.

6 Conclusions

To meaningfully use the structural information contained in the skeleton of humans, we have attacked optimization in the TTF using Benders decomposition [5]. We reformulate the TTF LP relaxation to separate optimization over of skeletons and local assignments. Using our Benders decomposition approach we produce one separate optimization for each part as opposed to one single sub-problem as is common in the Benders decomposition literature [21, 7, 6]. We produce tight bounds and fast inference.

The introduction of Bender’s decomposition opens a number of new opportunities for computer vision. The decomposition allows the separation of the integer and non-integer components in mixed integer programs, which promises to help circumvent difficulties induced by relaxing the integer variables. This could help for solving tracking problems in which splits occur (as in cell division in cell tracking). It also allows problems with heavy nesting as in planar ultrametrics [31] in which the master problem of one layer is the subproblem of the next.

Similar approaches based on the nested Benders decomposition could allow for dynamic programs to be solved in which the state space for individual variables is tractable but not the vector product. This could have use when inference in high tree width problems are considered . Such a case would be in pose estimation where hyper-nodes (nodes in the high tree width representation) correspond to parts and their components are detections.

Lastly, the use of Benders decomposition promises to enable parallelization of tracking and annotation by permitting the sub-problems to be solved independently while at the same time not relying on a Lagrangian relaxation with sub-gradient optimization [19, 17].

References

- [1] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3686–3693, 2014.
- [2] A. P. Armacost, C. Barnhart, and K. A. Ware. Composite variable formulations for express shipment service network design. *Transportation science*, 36(1):1–20, 2002.
- [3] F. Barahona and D. Jensen. Plant location with minimum inventory. *Mathematical Programming*, 83(1):101–111, 1998.
- [4] C. Barnhart, E. L. Johnson, G. L. Nemhauser, M. W. P. Savelsbergh, and P. H. Vance. Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, 46:316–329, 1996.
- [5] J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische mathematik*, 4(1):238–252, 1962.
- [6] J.-F. Cordeau, G. Stojković, F. Soumis, and J. Desrosiers. Benders decomposition for simultaneous aircraft routing and crew scheduling. *Transportation science*, 35(4):375–388, 2001.
- [7] A. M. Costa. A survey on benders decomposition applied to fixed-charge network design problems. *Computers & operations research*, 32(6):1429–1450, 2005.
- [8] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 886–893. IEEE, 2005.
- [9] G. Desaulniers, J. Desrosiers, and M. M. Solomon. *Column generation*, volume 5. Springer Science & Business Media, 2006.
- [10] C. Farabet, C. Couprie, L. Najman, and Y. LeCun. Learning hierarchical features for scene labeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1915–1929, 2013.
- [11] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9):1627–1645, 2010.
- [12] P. F. Felzenszwalb and D. P. Huttenlocher. Pictorial structures for object recognition. *International journal of computer vision*, 61(1):55–79, 2005.
- [13] A. M. Geoffrion. Lagrangian relaxation for integer programming. *50 Years of Integer Programming 1958-2008*, pages 243–281, 2010.

- [14] A. M. Geoffrion and G. W. Graves. Multicommodity distribution system design by benders decomposition. *Management science*, 20(5):822–844, 1974.
- [15] P. C. Gilmore and R. E. Gomory. A linear programming approach to the cutting-stock problem. *Operations research*, 9(6):849–859, 1961.
- [16] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- [17] D. Huisman, R. Jans, M. Peeters, and A. P. Wagelmans. Combining column generation and lagrangian relaxation. *Column generation*, pages 247–270, 2005.
- [18] E. Insafutdinov, L. Pishchulin, B. Andres, M. Andriluka, and B. Schiele. Deepercut: A deeper, stronger, and faster multi-person pose estimation model. In *ECCV*, 2016. (Accepted).
- [19] N. Komodakis, N. Paragios, and G. Tziritas. Mrf energy minimization and beyond via dual decomposition. *IEEE transactions on pattern analysis and machine intelligence*, 33(3):531–552, 2011.
- [20] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [21] T. L. Magnanti and R. T. Wong. Accelerating benders decomposition: Algorithmic enhancement and model selection criteria. *Operations research*, 29(3):464–484, 1981.
- [22] L. Pishchulin, E. Insafutdinov, S. Tang, B. Andres, M. Andriluka, P. Gehler, and B. Schiele. DeepCut: Joint subset partition and labeling for multi person pose estimation. In *CVPR*, 2016.
- [23] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Parallel distributed processing: Explorations in the microstructure of cognition, vol. 1. chapter Learning Internal Representations by Error Propagation, pages 318–362. MIT Press, Cambridge, MA, USA, 1986.
- [24] P. H. Vance, C. Barnhart, E. L. Johnson, and G. L. Nemhauser. Airline crew scheduling: A new formulation and decomposition algorithm. *Operations Research*, 45(2):188–200, 1997.
- [25] P. Viola and M. J. Jones. Robust real-time face detection. *International journal of computer vision*, 57(2):137–154, 2004.
- [26] C. Vondrick, A. Khosla, T. Malisiewicz, and A. Torralba. Hoggles: Visualizing object detection features. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1–8, 2013.
- [27] S. Wang, S. Wolf, C. Fowlkes, and J. Yarkony. Tracking objects with higher order interactions using delayed column generation. *Artificial Intelligence and Statistics*, 2017.
- [28] S. Wang, C. Zhang, M. A. Gonzalez-Ballester, and J. Yarkony. Efficient pose and cell segmentation using column generation. *arXiv preprint arXiv:1612.00437*, 2016.
- [29] Y. Yang, S. Baker, A. Kannan, and D. Ramanan. Recognizing proxemics in personal photos. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3522–3529. IEEE, 2012.
- [30] Y. Yang and D. Ramanan. Articulated pose estimation with flexible mixtures-of-parts. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1385–1392. IEEE, 2011.
- [31] J. Yarkony and C. Fowlkes. Planar ultrametrics for image segmentation. In *Neural Information Processing Systems*, 2015.

A Anytime Lower Bounds on the Master Problem

Given any set of the Benders rows we produce any anytime lower bound the optimum of the LP relaxation of the master problem as follows.

$$\begin{aligned} \text{Eq 5} &\geq \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1 \\ \ell_r \geq 0}} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} -\ell_r \quad (15) \\ -\ell_r &\geq -1^\top \lambda^{i1r} - 1^\top \lambda^{i2r} + \gamma^\top G^{r\top} \lambda^{i1r} - \gamma^\top G^{r\top} \lambda^{i3r} \quad \forall r \in \mathcal{R}, i \in \hat{\Lambda}^r \end{aligned}$$

We now convert the constraints into dual variables and but retain $G\gamma \leq 1$ in the master problem.

$$\begin{aligned} \text{Eq 15} &= \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1 \\ \ell_r \geq 0}} \max_{\mu \geq 0} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} -\ell_r \quad (16) \\ &+ \mu^{0\top} (G\gamma - 1) + \sum_{r \in \mathcal{R}, i \in \hat{\Lambda}^r} \mu_{ir} (\ell_r - 1^\top \lambda^{i1r} - 1^\top \lambda^{i2r} + \gamma^\top G^{r\top} \lambda^{i1r} - \gamma^\top G^{r\top} \lambda^{i3r}) \end{aligned}$$

We now relax the optimality of μ and consider any dual feasible μ which implies μ is non negative and that for each $r \in \mathcal{R}$ that $\sum_{i \in \hat{\Lambda}^r} \mu_{ir} \geq 1$.

$$\begin{aligned} \text{Eq 15} &\geq \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1 \\ \ell_r \geq 0}} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} -\ell_r \quad (17) \\ &+ \mu^{0\top} (G\gamma - 1) + \sum_{r \in \mathcal{R}, i \in \hat{\Lambda}^r} \mu_{ir} (\ell_r - 1^\top \lambda^{i1r} - 1^\top \lambda^{i2r} + \gamma^\top G^{r\top} \lambda^{i1r} - \gamma^\top G^{r\top} \lambda^{i3r}) \\ &= \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1 \\ \ell_r \geq 0}} \Gamma^\top \gamma + \sum_{r \in \mathcal{R}} -\ell_r - \mu^{0\top} 1 + \sum_{r, i \in \hat{\Lambda}^r} \mu_{ir} (\ell_r - 1^\top (\lambda^{i1r} + \lambda^{i2r})) \\ &\quad + \sum_{g \in \mathcal{G}} \gamma_g (\Gamma_g + \sum_{d \in \mathcal{D}} G_{dg} (\mu_d^0 + \sum_{r \in \mathcal{R}, i \in \hat{\Lambda}^r} \mu_{ir} (\lambda_d^{i1r} - \lambda_d^{i3r}))) \end{aligned}$$

Since $\sum_{i \in \hat{\Lambda}^r} \mu_{ir} \geq 1$ we know that ℓ_r is zero valued.

$$\begin{aligned} \text{Eq 17} &= \min_{\substack{\gamma \geq 0 \\ G\gamma \leq 1}} \Gamma^\top \gamma - \mu^{0\top} 1 - \sum_{r \in \mathcal{R}, i \in \hat{\Lambda}^r} \mu_{ir} 1^\top (\lambda^{i1r} + \lambda^{i2r}) \quad (18) \\ &+ \sum_{g \in \mathcal{G}} \gamma_g (\Gamma_g + \sum_{d \in \mathcal{D}} G_{dg} (\mu_d^0 + \sum_{r \in \mathcal{R}, i \in \hat{\Lambda}^r} \mu_{ir} (\lambda_d^{i1r} - \lambda_d^{i3r}))) \end{aligned}$$

We now relax optimization to permit the selection of one skeleton associated with each detection associated with a global part. Note that in the case of multiple global parts a given detection may be include more than once.

$$\begin{aligned} \text{Eq 17} &\geq -\mu^{0\top} 1 - \sum_{\substack{r \in \mathcal{R} \\ i \in \hat{\Lambda}^r}} \mu_{ir} (1^\top (\lambda^{i1r} + \lambda^{i2r})) \quad (19) \\ &+ \sum_{d \in \mathcal{R}'} \min_{\substack{g \in \mathcal{G} \\ G_{dg}=1}} \gamma_g (\Gamma_g + \sum_{d \in \mathcal{D}} G_{dg} (\mu_d^0 + \sum_{\substack{r \in \mathcal{R} \\ i \in \hat{\Lambda}^r}} \mu_{ir} (\lambda_d^{i1r} - \lambda_d^{i3r}))) \end{aligned}$$

Note that the minimization over g is simply the minimization done for generating columns and is known to be a dynamic program.

B Anytime Lower Bounds on a Benders Sub-Problem

We now consider the production of anytime lower bounds on the solution the the LP corresponding to a given Benders sub-problem. The production of these allows for an anytime dual feasible solution to

be generated prior to convergence of the LP of the dual sub-problem and hence anytime generation of Benders rows. Recall the dual optimization corresponding to part r given master problem solution γ .

$$\min_{\substack{\psi^r > 0 \\ G^r \gamma + L^r \psi^r \leq 1 \\ M^r \psi^r + L^r \psi^r \leq 1 \\ M^r \psi^r - G^r \gamma \leq 0}} \Psi^{\top r} \psi^r \quad (20)$$

We begin by augment our primal ILP with the constraint that $M^r \psi^r \leq 1$ and then add the Lagrange multipliers $\lambda^1, \lambda^2, \lambda^3$ to optimization.

$$\min_{\substack{\psi^r > 0 \\ M^r \psi^r \leq 1}} \max_{\lambda \geq 0} \Psi^{\top r} \psi^r + \lambda^{1\top} (-1 + G^r \gamma + L^r \psi^r) + \lambda^{2\top} (-1 + G^r \gamma + L^r \psi^r) + \lambda^{3\top} (M^r \psi^r - G^r \gamma) \quad (21)$$

We now relax optimality in λ producing the following bound.

$$\begin{aligned} \min_{\substack{\psi^r > 0 \\ M^r \psi^r \leq 1}} \Psi^{\top r} \psi^r + \lambda^{1\top} (-1 + G^r \gamma) - \lambda^{2\top} 1 - \lambda^{3\top} G^r \gamma + (\lambda^1 + \lambda^2)^{\top} L^r \psi^r + (\lambda^2 + \lambda^3)^{\top} M^r \psi^r \\ = \lambda^{1\top} (-1 + G^r \gamma) - \lambda^{2\top} 1 - \lambda^{3\top} G^r \gamma + \sum_{\substack{d \in \mathcal{D} \\ R_d = r}} \min[0, \min_{\substack{l \in \mathcal{L} \\ M_{dl} = 1}} \Psi_l] \end{aligned} \quad (22)$$

Observe that the minimization computed is the lowest reduced cost local assignment given a fixed global detection which is computed during each iteration of optimization.