Fusion of Multimodal Partial Estimates

Jiang Xu, Junsong Yuan and Ying Wu

January 2009

Abstract

Fusing partial estimates is a critical and common problem in many computer vision tasks such as part-based detection and tracking. It generally becomes complicated and intractable when there are a large number of multimodal partial estimates, and thus it is desirable to find an effective and scalable fusion method to integrate these partial estimates. This paper presents a novel and effective approach to fusing multimodal partial estimates in a principled way. In this new approach, fusion is related to a computational geometry problem of finding the minimum-volume orthotope, and an effective and scalable branch and bound search algorithm is designed to obtain the global optimal solution. Experiments on tracking articulated objects and occluded objects show the effectiveness of the proposed approach.

Contents

Chapter 1

Introduction

Many computer vision tasks involve the estimation of the unknown $\mathbf{x} \in \mathbb{R}^d$ from many independent estimates $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n\}$, where the individual estimate $\mathbf{y}_i$ may be obtained from various sources (*e.g.* , different views, time and cues), or from partial features. We refer each individual estimate as the *partial estimate* (PE) or *partial belief*, and the final estimation as the *complete estimation* (CE). A PE $\mathbf{y}_i$ gives an individual estimate of the unknown $\mathbf{x}$, and it may only provide the estimate on several specific dimensions of $\mathbf{x}$, so it is called a partial estimate. As the PEs can be quite inaccurate, a critical question is how we can fuse these partial estimates for a better estimation, *i.e.* , how to obtain $\mathbf{x} = \text{fuse}(\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n)$.

One concrete example is the part-based object detection and tracking. The target is represented by its parts and each part is associated with a dedicated detector and tracker, each of which provides a PE of the location and motion of the target. Because a part of the target is generally less discriminative than the entire target, the matching to this part is likely to include many false positives. This is especially true when the target is in a clutter background. Therefore, its PEs tend to have multiple modes, where most of them correspond to false positive matches. We refer one such PE that has multiple modes as a *multimodal partial estimate* (MPE), and our work is focused on the fusion of the MPEs. We want to emphasis that in our work we refer the word *multimodal* to multiple modes in *one* PE, rather than multiple estimates/sensors (Fig. 1.1).

If the PEs $\{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n\}$ are all unimodal, it is possible to obtain a closed-form fusion for the CE $\mathbf{x}$, *e.g.* , through the best linear unbiased estimation (BLUE) [1]. However, when $\mathbf{y}_i$ is multimodal (*e.g.* , modeled as a Gaussian mixture), the fusion for the CE is likely to exhibit an extremely complicated form. If each MPE has m modes, the number of modes in the CE is in the order of $o(m^n)$. In its discrete case, suppose each MPE consists of a set of m discrete estimates, the complexity of searching for the best CE shall be $o(m^n)$. Such an exponential growth of the number of modes (or the combinatorial complexity in the discrete case) makes any form of fuse$(\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n)$ very difficult to be optimized. As the complicated CE has an enormous number of local optima, fusion is likely to end up with a
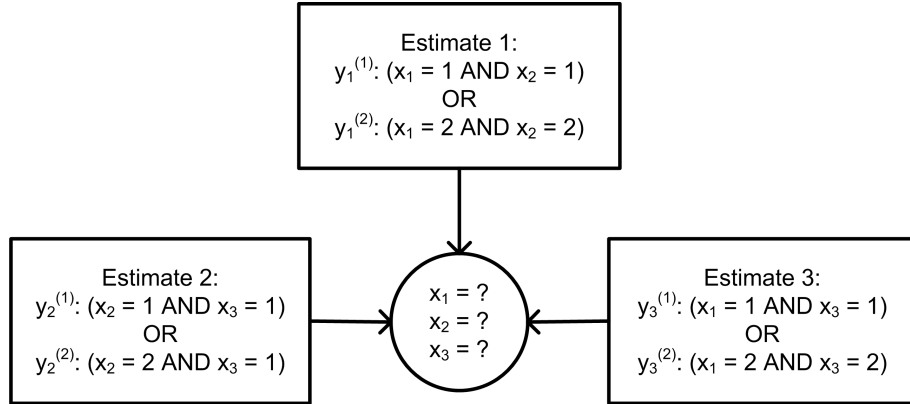
Figure 1.1: Example of fusing three MPEs, and each MPE has two modes. The optimal fusion result CE is the most consistent **x** with all the three MPEs.

low-quality estimation unless an exhaustive search can be performed. Even when an exhaustive search is merely viable when n is small, such a method is not scalable when there are many MPEs to fuse. Thus, new scalable fusion methods are desirable.

As the CE $\mathbf{x} = \text{fuse}(\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n)$ in general may not have good analytical properties, it is difficult to manipulate it directly. In this paper, we convert the error minimization in the fusion problem into a problem that finds a minimum-volume d-orthotope in $\mathbb{R}^d$ subject to some constraints. The *minimum-volume orthotope* problem can also be viewed as a multi-class generalization of the closest-pair problem in computational geometry. We design an effective branch and bound search algorithm to determine the global optimal solution to this problem with a moderate computational complexity.

The novelty of this work includes the following three aspects. (1) The fusion of MPEs is converted to a tractable minimum-volume orthotope problem, in which the intricate CE is exactly optimized in a discrete view, or approximately optimized in a continuous view. This new treatment leads to a tractable solution to fusion. (2) It reveals an interesting connection between probabilistic data fusion and computational geometry. The proposed solution to the minimum-volume orthotope problem provides a non-trivial generalization of the closest-pair problem. (3) The proposed fusion method is very scalable w.r.t. the number of estimates, or information sources, as the complexity is almost constant w.r.t. the number of sources.

The paper is organized as follows: related works are briefly described in Chap. 2. In Chap. 3, we formulate the problem and present the solution. We relate our method to computational geometry in Chap. 4. The experiment results of articulated body tracking and occluded objects tracking by fusing multiple local trackers are shown in Chap. 5, and the conclusion is made in Chap. 6.

Chapter 2

Related Work

There have been extensive studies on distributed estimates fusion. In [1], two criteria for optimal fusion of unimodal Gaussian estimates are summarized. One is the weighted least squares (WLS), and the other is the best linear unbiased estimation (BLUE). These fusion techniques can be applied to some classic computer vision problems such as the optical flow estimation [2].

When bad or fault estimates exist, WLS or BLUE cannot work well. To better handle noisy estimates, which refer to bad or outlier estimates, one possible solution is to allocate large variances to the bad estimates, such as in the methods of Covariance Intersection/Union (CI/CU) [3]. In [4], the fusion problem when the measurement errors are heteroscedastic is addressed, and the problem is solved in a WLS way. Another solution to handle noisy estimates is to keep good estimates while discarding bad ones. Variable-Bandwidth Density-based Fusion (VBDF) [5] falls into such a category, which performs globally, and attempts to alleviate the influence of the outliers by gradually reducing the bandwidth of the modes. By applying VBDF, a tracking method is presented in [6]. However, it cannot guarantee the global optimality in fusion. Another method to alleviate the effect of bad estimates is to measure the goodness of the estimates locally. For example in [7], a principle to estimate the fidelity of each measurement in a localized calculation is presented.

Unfortunately, all of the above fusion methods are not designed for multimodal cases, namely multiple modes in one estimate. To handle the multimodal estimates fusion problem, there have been two types of solutions: using distributed algorithms [8, 9], or using randomized algorithms [10]. If the estimates can be represented in a loosely-connected graph, several techniques can be applied, such as the variational methods [11], Belief Propagation (BP) [8], or Nonparametric Belief Propagation (NBP) [9]. However, if the graph is densely connected, these methods are easily trapped by local minima, or cannot even converge, due to the loops in the densely connected graph. To avoid local minima and to guarantee the convergence, randomized algorithms can be applied. For example, RANSAC [10] has the ability to obtain a robust estimation from noisy MPEs. Even if only one mode

is correct, and all of the others are outliers in each MPE, RANSAC may still obtain the global optimum with some probability, but the performance of RANSAC deteriorates when the number of modes in each MPE or the number of the MPEs increases.

Chapter 3

Problem Formulation and Solution

Given a collection of MPEs $\{\mathbf{y}_1, \mathbf{y}_2, ..., \mathbf{y}_n\}$, we want to obtain the CE of the high dimensional unknown $\mathbf{x} \in \mathbb{R}^d$. To better explain our idea, we first examine the problem of fusing discrete MPEs, then present the solution of fusing continuous MPEs in Chap. 4.2.

For discrete MPEs, each MPE $\mathbf{y}_i$ contains multiple point estimations (*i.e.* modes):

$$\mathbf{y}_i = \{\mathbf{y}_i^1, \cdots, \mathbf{y}_i^{\nu_i}\},$$

where $\mathbf{y}_i^j \in \mathbb{R}^{h_i}$ is the j-th mode of $\mathbf{y}_i$ in the h_i-dimensional subspace ($h_i \leq d$), and ν_i is the number of modes in $\mathbf{y}_i$. Given two modes belonging to two different subspaces, $\alpha \in \mathbb{R}^A$ and $\beta \in \mathbb{R}^B$, we define their arithmetic operations (summation, subtraction and maximization) as below.

Computation Rule 1. Addition and Subtraction

In addition or subtraction, we only perform the calculation in space $\mathbb{R}^A \bigcap \mathbb{R}^B$, i.e., if $\gamma = \alpha \pm \beta$, then

$$\gamma_i = \begin{cases} \alpha_i \pm \beta_i & \text{if } i \in \mathbb{R}^A \bigcap \mathbb{R}^B \\ \text{undefined} & \text{otherwise} \end{cases}$$

Computation Rule 2. Max and Min

In maximization or minimization, we perform the calculation in space $\mathbb{R}^A \bigcup \mathbb{R}^B$, i.e., if $\gamma = \max(\alpha, \beta)$, then

$$\gamma_i = \begin{cases} \max(\alpha_i, \beta_i) & \text{if } i \in \mathbb{R}^A \bigcap \mathbb{R}^B \\ \alpha_i & \text{if } i \in \mathbb{R}^A \bigcap \mathbb{R}^{\overline{B}} \\ \beta_i & \text{if } i \in \mathbb{R}^{\overline{A}} \bigcap \mathbb{R}^B \\ \text{undefined} & \text{otherwise} \end{cases}$$

An illustrative example is shown in Fig. 3.1.

| W	α	β	$\alpha + \beta$	$\max(\alpha, \beta)$
1		β_1		β_1
2	α_2	β_2	$\alpha_2 + \beta_2$	$\max(\alpha_2, \beta_2)$
3	α_3			α_3

Figure 3.1: An illustration of arithmetic operations over two vectors belonging to two different subspaces. Suppose there are two vectors $\alpha \in \mathbb{R}^A$ and $\beta \in \mathbb{R}^B$. We use α_i to denote the value at α's i-th dimension. For example, here $A = \{2, 3\}$, $B = \{1, 2\}$, which means α is in the 2nd and 3rd dimensions of the whole space, and β is in the 1st and 2nd dimensions of the whole space. As the figure shows, $\alpha = [11, 12]$, $\beta = [13, 14]$, then $\alpha_2 = 11$, $\alpha_3 = 12$, $\beta_1 = 13$, $\beta_2 = 14$, while α_1 and β_3 are undefined. The addition is performed at the intersection of the subspaces, while the maximization is performed at the union of the subspaces.

3.1 The Objective Function

A good fusion result should be consistent with the MPEs. One natural objective is to minimize the average estimation error, *i.e.* ,

$$\min_{\mathbf{x}} \frac{1}{n} \sum_i \Psi(\mathbf{x}, \mathbf{y}_i), \tag{3.1}$$

where $\Psi(\mathbf{x}, \mathbf{y}_i)$ is the measurement of the inconsistency. For example, we can choose $\Psi(\mathbf{x}, \mathbf{y}_i)$ as:

$$\Psi(\mathbf{x}, \mathbf{y}_i) = \min_j \|\mathbf{x} - \mathbf{y}_i^j\|_\infty,$$

where the final estimation $\mathbf{x}$ is expected to be consistent with at least one of the modes of $\mathbf{y}_i$. Here the L_∞ norm of a vector α is $\|\alpha\|_\infty = \max_i |\alpha_i|$. Although other types of measurements are possible, we will show later that the L_∞ norm leads to an elegant global optimal solution.

Considering that it is difficult to minimize the average estimation error in Eq. 3.1, we slightly change the original formulation by replacing the average error with the maximum error among the MPEs:

$$\mathbf{x}^* = \arg_{\mathbf{x}} \min_{\mathbf{x}} \max_i \Psi(\mathbf{x}, \mathbf{y}_i), \tag{3.2}$$

or the median error among the MPEs:

$$\mathbf{x}^* = \arg_{\mathbf{x}} \min_{\mathbf{x}} median_i \Psi(\mathbf{x}, \mathbf{y}_i). \tag{3.3}$$

We call Eq.3.2 as the *maximum fusion* while Eq. 3.3 as the *median fusion*. The median fusion is less sensitive to noise while the maximum fusion may be influenced by an outlier MPE, in which all its modes are outliers. For clarity, we mainly discuss how to solve the maximum fusion in this section. The solution to the median fusion follows the same strategy and will be briefly discussed in Chap. 3.4 as well.

3.2 Equivalence to Orthotope Search

Figure 3.2: An illustration of orthotope search for partial estimation fusion. There are three MPEs, $\mathbf{y}_1 = \{y_1^1, y_1^2\}$, $\mathbf{y}_2 = \{y_2^1, y_2^2\}$, $\mathbf{y}_3 = \{y_3^1, y_3^2\}$. The orthotope V contains a mode if and only if this mode is contained in V's projection to this mode's subspace, e.g, $y_1^1 \in V$ and $y_3^1 \notin V$. Minimizing estimation error (Eq. 3.2) is equivalent to minimizing the volume of the orthotope V, which must contain at least one mode from each MPE (Eq. 3.4). In this figure, the orthotope V contains y_1^1, y_2^1, y_3^2, so $\mathcal{W}(V) = 1$.

The minimization of Eq. 3.2 can be converted into a min-volume orthotope search problem, as explained in Figure 3.2. Our task is to find an orthotope (a

high-dimensional bounding box) that can cover at least one mode from every MPE $\mathbf{y}_i$. To minimize the maximum error in Eq. 3.2, we require the longest edge of the orthotope has the minimum length. Based on the above definitions, we propose the following optimization problem

$$\begin{aligned} \min \quad & \|V\|_\infty \\ \text{s.t} \quad & \mathcal{W}(V) = 1. \end{aligned} \tag{3.4}$$

Here we denote by V a d-dimensional axis-aligned orthotope. An orthotope V's **volume** is related to the length of its longest edge, denoted by $\|V\|_\infty$. $\mathcal{W}(V)$ is the predicate function of the orthotope V:

$$\mathcal{W}(V) = \begin{cases} 1 & \forall\, i, \exists\, j, \text{such that } \mathbf{y}_i^j \in V \\ 0 & \text{otherwise} \end{cases} \tag{3.5}$$

When a lower-dimensional mode $\mathbf{y}_i^j$ is inside the subspace projection of the d-dimensional axis-aligned orthotope V, the orthotope V **contains** the mode $\mathbf{y}_i^j$, and we denote by $\mathbf{y}_i^j \in V$ (Figure 3.2).

To justify our formulation as an orthotope search problem, we prove the equivalence between Eq. 3.2 and Eq. 3.4 in Theorem 1. We further derive the property under the condition of unique optimal solution in Theorem 2. The proof of both theorems are in the appendix.

Max Fusion Theorem 1. *The equivalence of the optimization in Eq. 3.2 and Eq. 3.4*

Let

$$v_1 = \min_{\mathbf{x}} \max_i \Psi(\mathbf{x}, \mathbf{y}_i)$$

and

$$v_2 = \min_V \|V\|_\infty, \qquad s.t. \quad \mathcal{W}(V) = 1.$$

Then

$$v_1 = v_2/2.$$

Max Fusion Theorem 2. *If Eq. 3.2 has a unique optimal $\mathbf{x}^*$ and V^* is the optimal solution to Eq. 3.4, then $\mathbf{x}^*$ is the center of V^*.*

3.3 A Branch and Bound Solution

According to Theorem 2, we solve Eq. 3.2 by optimizing Eq. 3.4, which is to find a minimum-volume orthotope satisfying the predicate. In order to obtain the global optimal solution in the high-dimensional space, we propose a branch and bound search algorithm to find the best orthotope efficiently. As an efficient search method, branch and bound has been applied to object detection [12] and action detection [13]. Our solution is related to [12, 13], but works in a high-dimensional discretized space.

Algorithm 1: Maximum Fusion of MPEs

---

input : Multimodal partial estimates (MPEs) $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n\}$
output: Complete estimation (CE) $\mathbf{x}$

Initialize $\mathbb{V}$ as the collection of all orthotope candidates in the d-dimensional space.
Initialize an empty priority queue $\mathcal{Q}$, in which the element with the smallest key value pops first.

repeat
    split $\mathbb{V}$ into $\mathbb{V}'$ and $\mathbb{V}''$
    if $\mathcal{W}(\overline{\mathbb{V}'}) = 1$ **then**
        $\lfloor$ $\mathbb{V}' \to \mathcal{Q}$ by the key value $\|\underline{\mathbb{V}'}\|_\infty$
    if $\mathcal{W}(\overline{\mathbb{V}''}) = 1$ **then**
        $\lfloor$ $\mathbb{V}'' \to \mathcal{Q}$ by the key value $\|\underline{\mathbb{V}''}\|_\infty$
    retrieve the top element $\mathbb{V}$ from $\mathcal{Q}$
until $\mathbb{V}$ *contains only one element*
retrieve the only element V^* of $\mathbb{V}$
return $\mathbf{x}$ as the center point of V^*

Our branch and bound search algorithm is presented in Algorithm 1. Let $\mathbb{V} = \{V_i\}$ be an *orthotope-set*, where each V_i is an orthotope in the d-dimensional space. The union of $\mathbb{V}$, denoted by $\overline{\mathbb{V}}$, is the minimum orthotope which satisfies $\forall V \in \mathbb{V}, \overline{\mathbb{V}} \supseteq V$. The intersection of $\mathbb{V}$, denoted by $\underline{\mathbb{V}}$, is the maximum orthotope which satisfies $\forall V \in \mathbb{V}, \underline{\mathbb{V}} \subseteq V$. We provide an illustrative example in Figure 3.3.


Figure 3.3: An example showing the upper and lower bounds of $\mathbb{V}$. The two black rectangles are V_1 and V_2, and $\mathbb{V} = \{V_1, V_2\}$. $\overline{\mathbb{V}}$ is the red rectangle which contains V_1 and V_2, and $\underline{\mathbb{V}}$ is the blue rectangle which is contained by V_1 and V_2.

Given the original orthotope-set $\mathbb{V}$, our task is to find a minimum-volume $V^* \in \mathbb{V}$ satisfying the predicate, and the optimal CE $\mathbf{x} \in \mathbb{R}^d$ can be uniquely determined by V^*. In each iteration in Algorithm 1, we split $\mathbb{V}$ into two parts $\mathbb{V}'$ and $\mathbb{V}''$, and the splitting point is the middle point of the longest dimension of $\mathbb{V}$ in the orthotope-set space.

10

We use the bound properties of $\overline{\mathbb{V}}$ and $\underline{\mathbb{V}}$: If the union of $\mathbb{V}$ cannot satisfy the predicate, it impossible for any $V \in \mathbb{V}$ to satisfy the predicate. As a result, only if $\overline{\mathbb{V}}$ satisfies the predicate, it is worth to perform a further check on $\mathbb{V}$. Otherwise this $\mathbb{V}$ can be safely pruned.

We index $\mathbb{V}$ with a key value $\|\underline{\mathbb{V}}\|_\infty$. This key value provides a lower bound, *i.e.* $\forall V \in \mathbb{V}, \|\underline{\mathbb{V}}\|_\infty \leq \|V\|_\infty$. We use a priority queue $\mathcal{Q}$ to store the orthotope-sets by their key values. Each time we retrieve from $\mathcal{Q}$ a candidate orthotope-set $\mathbb{V}$ with the smallest key value. The retrieving process keeps going until the retrieved $\mathbb{V}$ contains only one orthotope V^*, then V^* has to be the optimal solution because V^* satisfies the predicate and has a minimum volume compared with all other possible orthotopes in $\mathcal{Q}$.

3.4   Median Fusion

As mentioned earlier, the limitation of the maximum fusion is that it is sensitive to noisy MPEs. To address this issue, we can modify our objective function to a robust form by using the median fusion in Eq.3.3. The corresponding orthotope search problem can still be formulated in Eq. 3.4, but with a different predicate function:

$$\mathcal{W}(V) = \begin{cases} 1 & \text{for at least half of } i, \exists\, j, \text{such that } \mathbf{y}_i^j \in V \\ 0 & \text{otherwise} \end{cases} \tag{3.6}$$

Using $\mathcal{W}(V)$ in Eq. 3.6 and the same branch and bound procedure as in Algorithm 1, we can obtain the optimal solution to the median fusion. The theorems in max fusion can be extended in median fusion.

Chapter 4

Beyond Basic Formulation

Although we obtain the global optimal solution under the discrete MPE case, the MPE fusion is more difficult when each MPE provides a continuous estimation. In this section, we firstly show the connection between our algorithm and the bichromatic pair problem in computational geometry, then we extend our solution to the continuous MPE fusion and provide a probabilistic interpretation of our approach.

4.1 Link to Computational Geometry

The bichromatic pair problem [14] is formulated as

$$\min_{j(1),j(2)} |\mathbf{y}_1^{j(1)} - \mathbf{y}_2^{j(2)}|, \tag{4.1}$$

where the objective is to find the closest pair $\mathbf{y}_1^{j(1)} \in \mathbf{y}_1$ and $\mathbf{y}_2^{j(2)} \in \mathbf{y}_2$ from different classes $\mathbf{y}_1$ and $\mathbf{y}_2$.

We extend this problem to multiple classes, as well as multiple subspaces, where $\mathbf{y}_i$ and $\mathbf{y}_k$ are two MPEs and can belong to different subspaces. The *multichromatic pair problem* is similar to Eq. 4.1:

$$\min_{j(\cdot) \text{ for all } i,\, k} \max \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty, \tag{4.2}$$

where the goal is to find a mode from each MPE, such that the maximum distance among all mode-pairs is minimized. Accroding to the following Theorem, the multichromatic pair problem is equivalent to Eq. 3.4, therefore it can be solved by our proposed branch and bound method as well.

Max Fusion Theorem 3. *The equivalence of Eq. 3.4 and Eq. 4.2.*

Let

$$v_2 = \min \|V\|_\infty, \qquad s.t. \quad \mathcal{W}(V) = 1.$$

and

$$v_3 = \min_{j(\cdot) \text{ for all } i,\, k} \max \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty.$$

Then

$$v_2 = v_3.$$

The proof of Theorem 3 is in the appendix.

Corollary 1. *Optimizing Eq. 3.2, Eq. 3.4 and Eq. 4.2 are equivalent.*

In summary, fusing discrete MPEs can be converted to finding a minimum orthotope containing at least one mode from each MPE, and is also equivalent to the multichromatic pair problem.

4.2   MPE Fusion in a Probabilistic View

Now we consider the fusion of continuous MPEs. Suppose that each MPE $\mathbf{y}_i$ generates a multimodal distribution $p_i(\mathbf{x}|\mathbf{y}_i)$:

$$p_i(\mathbf{x}|\mathbf{y}_i) = \sum_j p(\mathbf{y}_i^j) k(\mathbf{x} - \mathbf{y}_i^j), \tag{4.3}$$

where $p(\mathbf{y}_i^j)$ is the prior of mode $\mathbf{y}_i^j$. If $k_i(\cdot)$ is the Gaussian kernel, then $p_i(\mathbf{x}|\mathbf{y}_i)$ is a Gaussian Mixture (GM). In our definition of $k_i(\cdot)$, we call $p_i(\mathbf{x}|\mathbf{y}_i)$ an Infinity Mixture (IM), as $k_i(\cdot)$ uses the L_∞ norm:

$$k_i(\alpha) = C_i \exp(-\frac{\|\alpha\|_\infty}{\sigma}), \tag{4.4}$$

where σ is the kernel bandwidth, and C_i is the normalization term. This IM justifies our previous optimization method in a probabilistic view.

Denote by $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_n\}$. Suppose $p(\mathbf{x}|\mathbf{Y})$ follows the Products of Experts (PoE) model [15], the distribution becomes:

$$p(\mathbf{x}|\mathbf{Y}) \propto \prod_i p_i(\mathbf{x}|\mathbf{y}_i), \tag{4.5}$$

where $p_i(\mathbf{x}|\mathbf{y}_i)$ is the partial estimation, or partial belief of $\mathbf{x}$ from $\mathbf{y}_i$, and we assume that $p_i(\mathbf{x}|\mathbf{y}_i)$ are independent. Our objective is to find an estimate $\mathbf{x} \in \mathbb{R}^d$ with the highest probability:

$$\mathbf{x}^* = \arg\max_{\mathbf{x}} p(\mathbf{x}|\mathbf{Y}), \qquad \mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_n\}. \tag{4.6}$$

Searching $p(\mathbf{x}|\mathbf{Y})$ in the high dimensional space is an extremely difficult problem. For an arbitrary $\mathbf{x} \in \mathbb{R}^d$, we consider the orthotope V **centered at x**. If we only count the modes located inside the orthotope, while ignoring the modes outside the orthotope, we obtain the following lower bound of $p(\mathbf{x}|\mathbf{Y})$ by combining

13

Eq. 4.3, 4.4 and 4.5:

$$p(\mathbf{x}|\mathbf{Y}) = \mathcal{C}_1 \prod_i \sum_j p(\mathbf{y}_i^j) k(\mathbf{x} - \mathbf{y}_i^j)$$

$$\geq \mathcal{C}_1 \prod_i \sum_{\hat{j}} p(\mathbf{y}_i^j) k(\mathbf{x} - \mathbf{y}_i^j)$$

$$= \mathcal{C}_1 \prod_i \sum_{\hat{j}} p(\mathbf{y}_i^j) C_i \exp(-\frac{\|\mathbf{x} - \mathbf{y}_i^j\|_\infty}{\sigma})$$

$$\geq \mathcal{C}_1 \prod_i \sum_{\hat{j}} p(\mathbf{y}_i^j) C_i \exp(-\frac{\|V\|_\infty}{2\sigma})$$

where $\hat{j}$ only counts the modes inside V, and $\mathcal{C}_1$ is a constant. The first inequality is obtained by ignoring the contribution from the modes outside the orthotope, and the second inequality is obtained from $\|\mathbf{x} - \mathbf{y}_i^j\|_\infty \leq \|V\|_\infty/2$ when $\mathbf{y}_i^j \in V$.

By taking the logarithm of the above equation, we obtain

$$\log p(\mathbf{x}|\mathbf{Y}) \geq \mathcal{C}_2 - n\frac{\|V\|_\infty}{2\sigma} + \sum_i \log(\sum_{\hat{j}} p(\mathbf{y}_i^j)) \tag{4.7}$$

$$= \mathcal{L}(V),$$

where $\mathcal{C}_2$ is a constant, n is the number of MPEs, then $\mathcal{L}(V)$ is the lower bound of $\log p(\mathbf{x}|\mathbf{Y})$. Searching for the optimal $\mathbf{x}^*$ now amounts to finding V^* with the largest $\mathcal{L}(V^*)$, *i.e.* maximizing the lower bound of $\log p(\mathbf{x}|\mathbf{Y})$. When σ is very small (the extreme case is $\sigma \to 0$, and it is degenerated to the discrete case), the $\frac{\|V\|_\infty}{2\sigma}$ term is dominant. Under this condition, maximizing $\mathcal{L}(V)$ is equivalent to minimizing $\|V\|_\infty$, which is equivalent to our discrete solution.

We maximize Eq. 4.7 by using a similar branch and bound method as in Algorithm 1. To further speed up the branch and bound process, we derive the lower and upper bounds of $\mathcal{L}(V) - \mathcal{C}_2$, respectively:

$$f^+(V) = \sum_i \log(\sum_{\hat{j}} p(\mathbf{y}_i^j))$$

$$f^-(V) = -n\frac{\|V\|_\infty}{2\sigma} + \sum_i \log(\sum_{\hat{j}} p(\mathbf{y}_i^j)).$$

Here $f^+(V)$ is obtained by putting all of the modes inside V at the center of the orthotope V, and $f^-(V)$ is obtained by putting all the modes inside V at the boundary of the orthotope V. Neglecting the constant terms, $f^+(V)$ and $f^-(V)$ provide upper and lower bounds of $\mathcal{L}(V)$, respectively. The branch and bound technique can be further accelerated by using these upper and lower bounds for more efficient pruning.

Chapter 5

Experiments

We evaluate our new MPE fusion methods in two tracking scenarios: one is to track articulated objects (testing the max fusion), and the other is to track occluded objects (testing the median fusion).

5.1 Tracking Articulated Objects

To track an articulated object, the object is decomposed into several parts, and each part is tracked by an individual part-tracker, as explained in Figure 5.1. As the part-trackers are connected and influence each other, the final tracking result is obtained by fusing the results from the set of part-trackers.


Figure 5.1: Example of a two-part articulated body. The CE is $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5, \mathbf{x}_6]$. MPE 1 provides the estimation of $[\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4]$, and MPE 2 provides the estimation of $[\mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5, \mathbf{x}_6]$.

The flowchart of our tracking is shown in Figure 5.2. In our experiments, each part-tracker is manually initialized by a fixed size rectangle which covers one part of the object. During the tracking process, each part-tracker randomly samples image patches in its neighborhood regions. These image patches are of the same size as the initialized rectangle. To track these patches, we check if their appearances are similar to their initialized appearance, *e.g.* using the sum-of-squared-differences (SSD) measurement. When the similarity score is higher than a predefined threshold, we treat the corresponding coordinates of the matched location as one mode in the partial estimate of the location of the part-tracker. The collection
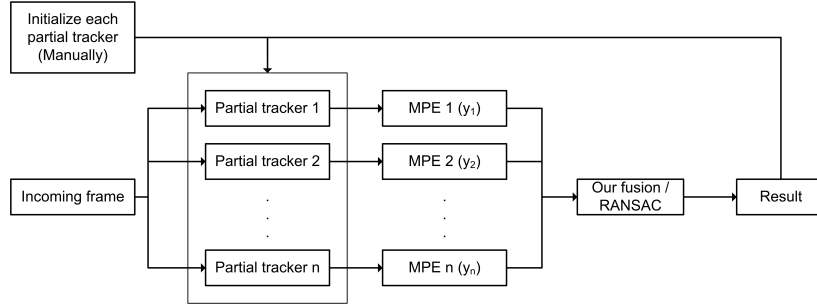
15

Figure 5.2: The flow chart of our tracking experiments.

of all such modes gives the MPE of the corresponding tracker. To tolerate appearance variations, we use a less rigid matching criterion that leads to many false positive modes (Figure 5.3). In our approach, we resample about 1,000 patches and obtain about 20~200 modes for each part-tracker.


Figure 5.3: The left figure is the input frame, and the right figure shows the modes of two part-trackers. Two part-trackers handle the upper and lower arms, respectively. Each part-tracker generates 20~200 modes (shown in green rectangles) in one frame.

After the MPEs are obtained, we apply our new fusion method and compare its performance to RANSAC. In the RANSAC approach, we iterate $1,000,000 \sim 5,000,000$ times. In each iteration, we randomly select one mode from each MPE to obtain the CE by averaging the selections. Then we choose the closest mode to CE from each MPE and calculate the SSD. The experiment settings are the same in our approach as in RANSAC, except for the fusion step. By increasing the number of iterations in RANSAC, it can give good results on tracking a two-part arm. However, for an articulated object that has more than two parts, the RANSAC method performs poorly even if we increase the number of iterations in RANSAC. The execution time of our algorithm is almost fixed when we increase the number of modes in each part-tracker. This shows the good scalability of our algorithm to the number of modes from each MPE.

16

Figure 5.4: Tracking articulated objects. From top to bottom, the articulated objects are split into 3, 4, 5 parts, respectively.

We test our new fusion method on tracking different articulated objects, and some sample results are shown in Figure 5.4 and 5.5. Figure 5.4 shows the experiment results of tracking articulated objects. By optimizing globally, our algorithm can keep tracking the structure of the articulated object. From top to bottom, the articulated objects have 3, 4, 5 parts, respectively. Figure 5.5 shows the comparison between our algorithm and RANSAC. The 1st and 2nd row of the figure shows the tracking results of a two-part articulated arm, and both our algorithm and RANSAC can provide good results. The 3rd and 4th row of the figure shows the tracking results of a three-part articulated finger, our algorithm is able to track the finger successfully, but RANSAC fails to give correct results: it begins to drift after several frames. From further experiments, we observe that our fusion method is able to successfully find a global optimum, and outperforms RANSAC. In general, we observe that the more parts we have in the fusion, the better our new fusion method achieves comparing with RANSAC.

5.2   Tracking Occluded Objects

We evaluate the median fusion in tracking an occluded object. The experiment setting keeps the same as that in the articulated objects tracking. The only difference is that each part-tracker follows a certain part of the object, rather than an articulated part. For example, in the face sequence shown in Figure 5.6, the face is modeled by eight overlapping parts. Instead of tracking the whole face, we track the eight overlapping face patches. Although every individual part-tracker tracks one of the eight parts and induces many false estimates, the fusion of all the part-trackers leads to a strong tracker which is very robust to partial occlusion. As long as half of the eight patches are visible, the median fusion is able to successfully

Figure 5.5: Comparison between our algorithm and RANSAC. Odd rows: our results; even rows: RANSAC results.

Figure 5.6: Tracking occluded face (sequence from [16]).

handle the severe occlusion.

Chapter 6

Conclusions

Fusing partial estimates from different sources is challenging because of the multimodal nature of the partial estimates: a multimodal objective function can make the optimization process easily trapped in local minima. Generally, it is difficult to obtain the global optimal estimation, especially in a high-dimensional parameter space. By revealing the connection between the probabilistic data fusion and computational geometry, we present a novel approach to the above challenges. We relate the error minimization problem of MPE fusion to a computational geometry problem of finding the minimum-volume orthotope in the parameter space. A branch and bound search algorithm is designed to obtain the global optimal solution. Our proposed new fusion method is scalable w.r.t. the number of estimates and its complexity is almost constant w.r.t. the number of partial estimates. Our proposed algorithm can be applied to a wide variety of applications (*e.g.* articulated objects tracking, occluded objects tracking), where effective information fusion from separate sources is needed.

Bibliography

[1] X. Li, Y. Zhu, J. Wang, and C. Han, "Optimal linear estimation fusion-part I: unified fusion rules," *IEEE Transactions on Information Theory*, vol. 49, no. 9, pp. 2192–2208, 2003.

[2] E. Simoncelli, E. Adelson, and D. Heeger, "Probability distributions of optical flow," in *CVPR*, 1991, pp. 310–315.

[3] J. Uhlmann, "Covariance consistency methods for fault-tolerant distributed data fusion," *Information Fusion*, vol. 4, no. 3, pp. 201–215, 2003.

[4] X. Zhou, A. Gupta, and D. Comaniciu, "An information fusion framework for robust shape tracking," *PAMI*, vol. 27, no. 1, pp. 115–129, 2005.

[5] D. Comaniciu, "Nonparametric information fusion for motion estimation," in *CVPR*, vol. 1, June 2003.

[6] B. Han and L. Davis, "Probabilistic fusion-based parameter estimation for visual tracking," *CVIU*, 2008.

[7] G. Hua and Y. Wu, "Measurement integration under inconsistency for robust tracking," in *CVPR*, 2006.

[8] J. Yedidia, W. Freeman, and Y. Weiss, "Understanding belief propagation and its generalizations," *Exploring artificial intelligence in the new millennium*, pp. 239–236, 2003.

[9] E. Sudderth, E. Ihler, W. Freeman, and A. Willsky, "Nonparametric belief propagation," in *CVPR*, 2002, pp. 605–612.

[10] M. Fischler and R. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.

[11] M. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul, "An introduction to variational methods for graphical models," *Machine Learning*, vol. 37, no. 2, pp. 183–233, 1999.

[12] C. Lampert, M. Blaschko, T. Hofmann, and S. Zurich, "Beyond Sliding Windows: Object Localization by Efficient Subwindow Search," in *CVPR*, 2008.

[13] J. Yuan, Z. Liu, and Y. Wu, "Discriminative Subvolume Search for Efficient Action Detection," in *CVPR*, 2009.

[14] S. Khuller and Y. Matias, "A Simple Randomized Sieve Algorithm for the Closest-Pair Problem," *Information and Computation*, vol. 118, no. 1, pp. 34–37, 1995.

[15] G. Hinton, "Products of experts," in *ICANN*, vol. 1, 1999.

[16] A. Adam, E. Rivlin, and I. Shimshoni, "Robust Fragments-based Tracking using the Integral Histogram," in *CVPR*, 2006.

Appendix A

Max Fusion Theorems

A.1 Max Fusion Theorem 1

We prove $v_1 \leq v_2/2$ first. Assume v_2's corresponding optimum orthotope is V^*. Then for all i, there exists $j^*(i)$, such that $\mathbf{y}_i^{j^*(i)} \in V^*$. We construct $\mathbf{x}^*$ located at the center of V^*, then

$$\|\mathbf{x}^* - \mathbf{y}_i^{j^*(i)}\|_\infty \leq \|V^*\|_\infty/2 = v_2/2,$$

and we obtain

$$
\begin{aligned}
v_1 &= \min_{\mathbf{x}} \max_i \Psi(\mathbf{x}, \mathbf{y}_i) \\
&\leq \max_i \Psi(\mathbf{x}^*, \mathbf{y}_i) \\
&= \max_i \min_{j(\cdot)} \|\mathbf{x}^* - \mathbf{y}_i^{j(i)}\|_\infty \\
&\leq \max_i \|\mathbf{x}^* - \mathbf{y}_i^{j^*(i)}\|_\infty \\
&\leq v_2/2.
\end{aligned}
$$

We then prove $v_1 \geq v_2/2$. Assume

$$
\begin{aligned}
v_1 &= \min_{\mathbf{x}} \max_i \Psi(\mathbf{x}, \mathbf{y}_i) \\
&= \max_i \|\mathbf{x}' - \mathbf{y}_i^{j'(i)}\|_\infty,
\end{aligned}
$$

we then consider the orthotope V' centered at $\mathbf{x}'$ and with volume $\|V'\|_\infty = 2v_1$. For any i,

$$\|\mathbf{x}' - \mathbf{y}_i^{j'(i)}\|_\infty \leq v_1 \quad \Rightarrow \quad \mathbf{y}_i^{j'(i)} \in V'.$$

Therefore, V' satisfies Eq. 4, and

$$v_2 \leq 2v_1 \quad \Rightarrow \quad v_1 \geq v_2/2.$$

Therefore, $v_1 = v_2/2$.

A.2 Max Fusion Theorem 2

We proof this theorem by contradiction. Assume $\mathbf{x}^*$ is not the center of V^*, then we denote the center of V^* $\mathbf{x}'$ and $\mathbf{x}' \neq \mathbf{x}^*$. From the proof of Theorem 1, $\mathbf{x}'$ is the optimal solution of Eq. 2, this contradicts the condition that $\mathbf{x}^*$ is the unique optimal solution. Therefore, the assumption does not hold, and $\mathbf{x}' = \mathbf{x}^*$.

A.3 Max Fusion Theorem 3

We prove $v_2 \geq v_3$ first. Assume v_2's corresponding optimal orthotope is V^*. Then for all i, there exists $j^*(i)$, such that $\mathbf{y}_i^{j^*(i)} \in V^*$. Then for all i,k, it satisfies

$$\|\mathbf{y}_i^{j^*(i)} - \mathbf{y}_k^{j^*(k)}\|_\infty \leq \|V^*\|_\infty = v_2,$$

and we obtain

$$v_3 = \min_{j(\cdot)} \max_{\text{for all } i,\, k} \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty$$

$$\leq \max_{\text{for all } i,\, k} \|\mathbf{y}_i^{j^*(i)} - \mathbf{y}_k^{j^*(k)}\|_\infty$$

$$\leq v_2.$$

We then prove $v_2 = v_3$ by contradiction. Assume $v_2 \neq v_3$, then $v_2 > v_3$. Let

$$v_3 = \min_{j(\cdot)} \max_{\text{for all } i,\, k} \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty$$

$$= \max_{\text{for all } i,\, k} \|\mathbf{y}_i^{j'(i)} - \mathbf{y}_k^{j'(k)}\|_\infty.$$

We construct the orthotope V', and its two opposite vertex are $\min_{\text{all } i} \mathbf{y}_i^{j'(i)}$ and $\max_{\text{all } i} \mathbf{y}_i^{j'(i)}$. Then

$$\|V'\|_\infty = \|\max_{\text{all } i} \mathbf{y}_i^{j'(i)} - \min_{\text{all } k} \mathbf{y}_k^{j'(k)}\|_\infty$$

$$= \max_{\text{for all } i,\, k} \|\mathbf{y}_i^{j'(i)} - \mathbf{y}_k^{j'(k)}\|_\infty$$

$$= v_3.$$

Therefore, $\|V'\|_\infty = v_3 < v_2$, this contradicts v_2 is the optimal solution of Eq. 3.4. Thus, the assumption $v_2 \neq v_3$ is not true, and $v_2 = v_3$.

Appendix B

Median Fusion Theorems

B.1 Notations

All the theorems in maximum fusion can be extended to median fusion. To make the presentation clear, we define the validity function for median fusion as

$$\mathcal{W}(V) = \begin{cases} 1 & \text{for at least half of } i, \exists\, j, \text{such that } \mathbf{y}_i^j \in V \\ 0 & \text{otherwise,} \end{cases} \tag{B.1}$$

and the corresponding optimizations become

$$\min_{\mathbf{x}} median_i \Psi(\mathbf{x}, \mathbf{y}_i), \tag{B.2}$$

$$\begin{aligned} \min \quad & \|V\|_\infty \\ \text{s.t} \quad & \mathcal{W}(V) = 1. \end{aligned} \tag{B.3}$$

$$\min_{j(\cdot)\ \text{for half } i, k \in \{1,2,\cdots,n\}} \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty, \tag{B.4}$$

Similar to max fusion, optimizing Eq. B.2, Eq. B.3 and Eq. B.4 are equivalent, and related theorems and proofs are shown below.

B.2 Median Fusion Theorem 1

The equivalence of Eq. B.2 and Eq. B.3
Let

$$v_1 = \min_{\mathbf{x}} median_i \Psi(\mathbf{x}, \mathbf{y}_i)$$

and

$$v_2 = \min \|V\|_\infty, \qquad \text{s.t.} \quad \mathcal{W}(V) = 1.$$

Then

$$v_1 = v_2/2.$$

We prove $v_1 \leq v_2/2$ first. Assume v_2's corresponding optimum orthotope is V^*. For at least half of i, there exists $j^*(i)$, such that $\mathbf{y}_i^{j^*(i)} \in V^*$. We then construct $\mathbf{x}^*$ located at the center of V^*, then for at least half of i,

$$\|\mathbf{x}^* - \mathbf{y}_i^{j^*(i)}\|_\infty \leq \|V^*\|_\infty/2 = v_2/2 \rightarrow median_i \|\mathbf{x}^* - \mathbf{y}_i^{j^*(i)}\|_\infty \leq v_2/2,$$

so

$$\begin{aligned}
v_1 &= \min_{\mathbf{x}} median_i \Psi(\mathbf{x}, \mathbf{y}_i) \\
&\leq median_i \Psi(\mathbf{x}^*, \mathbf{y}_i) \\
&= median_i \min_{j(\cdot)} \|\mathbf{x}^* - \mathbf{y}_i^{j(i)}\|_\infty \\
&\leq median_i \|\mathbf{x}^* - \mathbf{y}_i^{j^*(i)}\|_\infty \\
&\leq v_2/2.
\end{aligned}$$

We then prove $v_1 \geq v_2/2$. Assume

$$\begin{aligned}
v_1 &= \min_{\mathbf{x}} median_i \Psi(\mathbf{x}, \mathbf{y}_i) \\
&= median_i \|\mathbf{x}' - \mathbf{y}_i^{j'(i)}\|_\infty,
\end{aligned}$$

we then consider the orthotope V' centered at $\mathbf{x}'$ and with volume $\|V'\|_\infty = 2v_1$. For at least half of i,

$$\|\mathbf{x}' - \mathbf{y}_i^{j'(i)}\|_\infty \leq v_1 \quad \Rightarrow \quad \mathbf{y}_i^{j'(i)} \in V'.$$

Therefore, V' satisfies Eq. B.3, and

$$v_2 \leq 2v_1 \quad \Rightarrow \quad v_1 \geq v_2/2.$$

Therefore, $v_1 = v_2/2$.

B.3   Median Fusion Theorem 2

If Eq. B.2 has a unique optimal $\mathbf{x}^*$ and V^* is the optimization result of Eq. B.3, then $\mathbf{x}^*$ is the center of V^*.

The proof is identical with proof of Theorem 2 of Max Fusion.

B.4   Median Fusion Theorem 3

The equivalence of Eq. B.3 and Eq. B.4.

Let

$$v_2 = \min \|V\|_\infty, \qquad \text{s.t.} \quad \mathcal{W}(V) = 1.$$

and

$$v_3 = \min_{\substack{j(\cdot) \text{ for half}}} \max_{i,k \in \{1,2,\cdots,n\}} \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty.$$

Then

$$v_2 = v_3.$$

We prove $v_2 \geq v_3$ first. Assume v_2's corresponding optimum orthotope is V^*. Then there exists $\mathbf{U} \subseteq \{1, 2, \cdots, n\}$ such that $|\mathbf{U}| \geq n/2$ and for all $i \in \mathbf{U}$, there exists $j^*(i)$, such that $\mathbf{y}_i^{j^*(i)} \in V^*$. Then for all $i, k \in \mathbf{U}$, it satisfies

$$\|\mathbf{y}_i^{j^*(i)} - \mathbf{y}_k^{j^*(k)}\|_\infty \leq \|V^*\|_\infty = v_2,$$

and we obtain

$$v_3 = \min_{\substack{j(\cdot) \text{ for half}}} \max_{i,k \in \{1,2,\cdots,n\}} \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty$$

$$\leq \min \max_{\text{for all } i, k \in \mathbf{U}} \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty$$

$$\leq \max_{\text{for all } i, k \in \mathbf{U}} \|\mathbf{y}_i^{j^*(i)} - \mathbf{y}_k^{j^*(k)}\|_\infty$$

$$\leq v_2.$$

We then prove $v_2 = v_3$ by contradiction. Assume $v_2 \neq v_3$, then $v_2 > v_3$. Let

$$v_3 = \min_{\substack{j(\cdot) \text{ for half}}} \max_{i,k \in \{1,2,\cdots,n\}} \|\mathbf{y}_i^{j(i)} - \mathbf{y}_k^{j(k)}\|_\infty$$

$$= \max_{\text{for all } i, k} \|\mathbf{y}_i^{j'(i)} - \mathbf{y}_k^{j'(k)}\|_\infty.$$

We construct the orthotope V', and its two opposite vertex are $\min_{\text{all } i \in \mathbf{U}} \mathbf{y}_i^{j'(i)}$ and $\max_{\text{all } i \in \mathbf{U}} \mathbf{y}_i^{j'(i)}$. Then

$$\|V'\|_\infty = \|\max_{\text{all } i \in \mathbf{U}} \mathbf{y}_i^{j'(i)} - \min_{\text{all } k \in \mathbf{U}} \mathbf{y}_k^{j'(k)}\|_\infty$$

$$= \max_{\text{for all } i, k \in \mathbf{U}} \|\mathbf{y}_i^{j'(i)} - \mathbf{y}_k^{j'(k)}\|_\infty$$

$$= v_3.$$

Therefore, $\|V'\|_\infty = v_3 < v_2$, this contradicts v_2 is the optimal solution of Eq. 3.4. Thus, the assumption $v_2 \neq v_3$ is not true, and $v_2 = v_3$.