Visualizing LP Duality

Manoj Prabhakaran

This is an attempt to geometrically interpret LP duality. For our purposes the LP primal and dual are the following problems:

LP PRIMAL	LP DUAL
$\mathbf{A}\mathbf{x} = \mathbf{b}_m$	
$\mathbf{x} \geq 0$	$\mathbf{A}'\pi \leq \mathbf{c}$
$\min \mathbf{c} \cdot \mathbf{x}$	$\max \mathbf{b}_m \cdot \pi$

Here $\mathbf{A}$ is an $m \times n$ real matrix, $\mathbf{x}, \mathbf{c} \in \mathbb{R}^n$ and $\mathbf{b}_m, \pi \in \mathbb{R}^m$, $m \leq n$. (For illustration we will assume $\mathbf{A} = \left[\mathbf{I}_m | \mathbf{G}_{m \times (n-m)}\right]$ where $\mathbf{I}_m$ is the $m \times m$ identity matrix).

Let $\mathbf{b} \in \mathbb{R}^n$ be defined so that $\mathbf{A}\mathbf{b} = \mathbf{b}_m$. ($\mathbf{A} = \left[\mathbf{I}_m | \mathbf{G}_{m \times (n-m)}\right] \implies \mathbf{b} = \begin{bmatrix} \mathbf{b}_m \\ 0^{n-m} \end{bmatrix}$). Then the LP primal above can be re-written as

$$\mathbf{A}(\mathbf{x} - \mathbf{b}) = 0,\ \mathbf{x} \geq 0,\ \min\ \mathbf{c} \cdot \mathbf{x}$$

Note that the primal and dual problems are not posed in the same space: $\mathbf{x}$ and π lie in different dimensional spaces. We would like to have both in the same space- $\mathbb{R}^n$. For this we transform the dual as follows:

$$\mathbf{y} = \mathbf{c} - \mathbf{A}'\pi,\ \mathbf{y} \geq 0,\ \max\ \mathbf{b}_m \cdot \pi$$

We define $\tilde{\mathbf{A}}$ as the $(n-m) \times n$ matrix with rows being a basis of the complementary orthogonal space of the space spanned by the rows of $\mathbf{A}$. ($\mathbf{A} = \left[\mathbf{I}_m | \mathbf{G}_{m \times (n-m)}\right] \implies \tilde{\mathbf{A}} = [-\mathbf{G}' | \mathbf{I}_{n-m}]$). The constraint that $\mathbf{y} - \mathbf{c}$ is of the form $-\mathbf{A}'\pi$ is the same as $\tilde{\mathbf{A}}(\mathbf{y} - \mathbf{c}) = 0$. Also

$$\mathbf{b}_m \cdot \pi = \mathbf{b} \cdot (\mathbf{c} - \mathbf{y}) \tag{1}$$

Then the above displayed problem can be rephrased as

$$\tilde{\mathbf{A}}(\mathbf{y} - \mathbf{c}) = 0,\ \mathbf{y} \geq 0,\ \min\ \mathbf{b} \cdot \mathbf{y}$$

Now both the problems are posed in $\mathbb{R}^n$. $\mathbf{x} - \mathbf{c}$ lies in the $n - m$ dimensional subspace spanned by the rows of $\tilde{\mathbf{A}}$ and $\mathbf{y} - \mathbf{b}$ in the m dimensional complementary orthogonal subspace, which is spanned by the rows of $\mathbf{A}$. To further simplify the form of the problem we define $\mathbf{a}$ as the unique point in the intersection of the two subspaces. This gives

$$\mathbf{A}(\mathbf{a}-\mathbf{b}) = \tilde{\mathbf{A}}(\mathbf{a}-\mathbf{c}) = 0$$

Then $\mathbf{A}(\mathbf{x}-\mathbf{b}) = \mathbf{A}(\mathbf{x}-\mathbf{a}+\mathbf{a}-\mathbf{b}) = \mathbf{A}(\mathbf{x}-\mathbf{a})$ and $\tilde{\mathbf{A}}(\mathbf{y}-\mathbf{c}) = \tilde{\mathbf{A}}(\mathbf{y}-\mathbf{a}+\mathbf{a}-\mathbf{c}) = \tilde{\mathbf{A}}(\mathbf{y}-\mathbf{a})$. Further minimizing $\mathbf{c}\cdot\mathbf{x}$ under the constraint on $\mathbf{x}$ is the same as minimizing $\mathbf{a}\cdot\mathbf{x}$ under that constraint, and minimizing $\mathbf{b}\cdot\mathbf{y}$ under the constraint on $\mathbf{y}$ is the same as minimizing $\mathbf{a}\cdot\mathbf{y}$. [1]

This gives our final (and very symmetric) primal and dual *Modified Problems*:

MP PRIMAL	MP DUAL
$\mathbf{A}(\mathbf{x}-\mathbf{a}) = 0$	$\tilde{\mathbf{A}}(\mathbf{y}-\mathbf{a}) = 0$
$\mathbf{x} \geq 0$	$\mathbf{y} \geq 0$
$\min \mathbf{a}\cdot\mathbf{x}$	$\min \mathbf{a}\cdot\mathbf{y}$

Note that we did not change the primal feasible set or optimum at all (though the cost was rewritten for aesthetic purposes). The dual problem has undergone a linear-transform: the m-dimensional space in which the LP dual feasible set has been transformed into an m-dimensional (displaced) subspace of $\mathbb{R}^n$ (orthogonal to the (displaced) subspace in which the primal feasible set lives).

Following is the relation between the LP and MP solutions: Consider any feasible $\mathbf{x}$ in primal LP and MP, and any feasible $\mathbf{y}$ in dual MP with π as the corresponding feasible solution in dual LP. Note that $\mathbf{x}-\mathbf{b}$ and $\mathbf{y}-\mathbf{c}$ are in complementary orthogonal spaces. Thus $(\mathbf{x}-\mathbf{b})\cdot(\mathbf{y}-\mathbf{c}) = 0$. The dual LP value is (see (1))

$$\mathbf{b}_m\cdot\pi = \mathbf{b}\cdot(\mathbf{c}-\mathbf{y}) = \mathbf{c}\cdot\mathbf{x} - \mathbf{x}\cdot\mathbf{y} \leq \mathbf{c}\cdot\mathbf{x} \tag{2}$$

where as the primal LP cost is $\mathbf{c}\cdot\mathbf{x}$. The inequality above follows from the fact that $\mathbf{x},\mathbf{y} \geq 0$. This is the weak-duality theorem. The strong duality theorem says that there exist $\mathbf{x}^*, \pi^*$ feasible for LP primal and dual respectively such that $\forall \mathbf{x}, \pi$ feasible for LP primal and dual resp.,

$$\mathbf{b}_m\cdot\pi \leq \mathbf{b}_m\cdot\pi^* = \mathbf{c}\cdot\mathbf{x}^* \leq \mathbf{c}\cdot\mathbf{x}$$

This can be rephrased as (see 2) $\exists\ \mathbf{x}^*, \mathbf{y}^*$ feasible for primal, dual resp. such that

$$\mathbf{x}^*\cdot\mathbf{y}^* = 0$$

This solving the MP amounts to finding two orthogonal points in the feasible regions of the primal and dual.

[1] $\mathbf{c}\cdot\mathbf{x} = (\mathbf{c}-\mathbf{a}+\mathbf{a})\cdot(\mathbf{x}-\mathbf{b}+\mathbf{b}) = \mathbf{a}\cdot\mathbf{x} + (\mathbf{c}-\mathbf{a})\cdot(\mathbf{x}-\mathbf{b}) + (\mathbf{c}-\mathbf{a})\cdot\mathbf{b}$
$\mathbf{A}(\mathbf{x}-\mathbf{b}) = 0, \tilde{\mathbf{A}}(\mathbf{c}-\mathbf{a}) = 0 \implies (\mathbf{c}-\mathbf{a})\cdot(\mathbf{x}-\mathbf{b}) = 0$. Thus $\mathbf{c}\cdot\mathbf{x} = \mathbf{a}\cdot\mathbf{y} + \text{const.}$
Similarly, $\mathbf{b}\cdot\mathbf{y} = (\mathbf{b}-\mathbf{a}+\mathbf{a})\cdot(\mathbf{y}-\mathbf{c}+\mathbf{c}) = \mathbf{a}\cdot\mathbf{y} + (\mathbf{b}-\mathbf{a})\cdot(\mathbf{y}-\mathbf{c}) + (\mathbf{b}-\mathbf{a})\cdot\mathbf{c}$,
and $\tilde{\mathbf{A}}(\mathbf{y}-\mathbf{c}) = 0, \mathbf{A}(\mathbf{b}-\mathbf{a}) = 0 \implies (\mathbf{b}-\mathbf{a})\cdot(\mathbf{y}-\mathbf{c}) = 0$.

Remarks

Remark 1 We visualize the strong duality theorem: consider two orthogonal flats in $\mathbb{R}^n$, defined by $\mathbf{A}(\mathbf{x}-\mathbf{a})=0$ and $\tilde{\mathbf{A}}(\mathbf{x}-\mathbf{a})=0$, of dimensions $n-m$ and m respectively. Suppose $\mathbf{a}$ is both primal and dual feasible (i.e, all co-ordinates of $\mathbf{a}$ are ≥ 0). Then, consider starting at $\mathbf{a}$ and following the projection of the vector to the origin, but confined to the (primal or dual) plane. The question is which is the farthest point to which one can get this way staying in non-negative co-ordinates. Consider the primal plane. Suppose $\mathbf{x}^*$ is the final point in this journey. $\mathbf{x}^*$ lies on certain $n-m$ axes-planes and will have the corresponding co-ordinates zero. The farthest point for the dual- call it $\mathbf{y}^*$ - is zero in the other co-ordinates. In other words every axes plane has at least one of $\mathbf{x}^*, \mathbf{y}^*$ lying on it (and so removing any one axes plane will let either the primal or dual plane escape into the *negative space*).

Remark 2 Solving the LP amounts to separating the two feasible points $\mathbf{x}$ and $\mathbf{y}$ by as large an angle as possible. It may be interesting to visualize the primal-dual algorithms in this set-up, as now we can see both the primal and the dual variables moving in the same space.

Acknowledgement

I would like to thank Bernard Chazelle and Andy Yao for reading through a preliminary version of this write-up.