

Section Notes 4
Duality, Sensitivity, and the Dual Simplex Algorithm

Applied Math / Engineering Sciences 121

Week of October 10, 2016

TF VERSION - Do Not Distribute

Goals for the week

- understand the relationship between primal and dual programs.
- know how to convert an LP from primal to dual.
- know how to analyze the sensitivity of a BFS to changes in input data.
- understand complementary slackness.

Contents

1	Duality	2
2	The Primal-Dual Relationship	3
3	Sensitivity Analysis	5
4	Understanding Sensitivity Analysis	11
5	The Dual Simplex Algorithm	13
6	Complementary Slackness	16
7	Some More Clarification	18

1 Duality

As we have seen in lecture, for every linear program there is a closely related alternative program called the ‘dual’. The relationship between these two programs is not only fascinating, but also very useful. As we shall see, it can be useful in gaining an intuition for the meaning of a program and to perform ‘sensitivity analysis.’ It will also lead to a new form of the simplex method.

1.1 Two parallel formulations

Given a program in standard equality form:

$$\begin{aligned} \max \quad & c^T x \\ \text{s.t.} \quad & Ax = b \\ & x \geq 0 \end{aligned}$$

Our dual will have the form:

$$\begin{aligned} \min \quad & b^T y \\ \text{s.t.} \quad & A^T y \geq c \\ & y \text{ free} \end{aligned} \tag{1}$$

The constraints in the primal correspond to variables in the dual, and vice versa. In this case, we can draw the following parallels:

Primal	Dual
= Constraints	Free Variables
Non-negative Variables	\geq constraints

The parallel formulations can be extended to LPs that are not in standard equality form. The following table specifies the equivalences between the primal and the dual in a general setting:

Primal (maximize)	Dual (minimize)
i th constraint \leq	i th variable ≥ 0
i th constraint $=$	i th variable free
i th constraint \geq	i th variable ≤ 0
j th variable ≤ 0	j th constraint \leq
j th variable free	j th constraint $=$
j th variable ≥ 0	j th constraint \geq

1.2 Shadow Prices

The dual formulation creates a new set of variables, y . Each variable y is associated with a corresponding constraint from the primal problem. In particular, at the dual optimal solution the value of the y variables tells us how much the objective will increase by relaxing a given constraint by 1 unit. Note that the full value of this increase may not be attainable if another constraint binds first— rather the shadow price pertains to a small ϵ -distance from the current value of the constraint RHS.

2 The Primal-Dual Relationship

2.1 Duality theorems

Theorem 2.1 (Weak Duality Theorem). *The (primal) objective value of any feasible primal solution is no larger than the (dual) objective value of any feasible dual solution:*

$$\text{obj value of feasible primal} \leq \text{obj value of feasible dual}$$

One use of weak duality is to establish the following:

Corollary 2.1 (Certificate of Optimality). *Given primal feasible solution x^* and dual feasible solution y^* where the primal and dual objectives are equal, then x^* is optimal for the primal problem and y^* is optimal for the dual problem*

Do you see why this is a corollary? That means that independent of any algorithm for solving LPs one can verify if a given primal-dual solution pair is in fact optimal. All you need to do is verify that both solutions are feasible, and that the resulting objective values in the two problems are identical.

Strong duality says something special about the case where there is an optimal primal (or optimal dual) solution.

Theorem 2.2 (Strong Duality Theorem). *If an LP has an optimal solution then so does its dual and the two optimal values are equal.*

Refer to the lecture notes for the proof (which appeals to the simplex method). The lemma used to prove this also has a practical implication: it gives us a way to read off the dual solution from the optimal primal tableau.

2.2 Using the duality theorems

Exercise 1

We know that an LP has to be either (a) infeasible, (b) has optimal solution, or (c) unbounded. Based on the duality theorems, answer the following questions.

1. If the primal LP is unbounded, which of (a), (b), and (c) are possible for the dual?
2. If the primal LP has optimal solution, which of (a), (b), and (c) are possible for the dual?
3. If the primal LP is infeasible, which of (a), (b), and (c) are possible for the dual?

End Exercise 1

Solution 1

1. If the primal LP is unbounded, the dual LP must be infeasible. By weak duality thm.
2. If the primal LP has optimal solution, the dual LP must have an optimal solution with the same optimal objective value. BY strong duality thm.

3. If the primal LP is infeasible, the dual LP could either be (a) infeasible or (c) unbounded. By weak duality thm.

End Solution 1

2.3 Converting from Primal to Dual: An Example

Suppose we begin with the following example:

$$\begin{aligned} \max \quad & 2x_1 + 3x_2 - x_3 \\ \text{s.t.} \quad & x_1 + x_2 \leq 8 \\ & x_2 - 3x_3 \leq 6 \\ & x_1, x_2, x_3 \geq 0 \end{aligned}$$

First we construct the following tableau using slack variables x_4 and x_5 :

	x_1	x_2	x_3	x_4	x_5	
z	-2	-3	1	0	0	0
x_4	1	1	0	1	0	8
x_5	0	1	-3	0	1	6

$z=0, x=(0,0,0,8,6), B=(4,5)$

By performing simplex pivot operations, we arrive at the following optimal tableau:

	x_1	x_2	x_3	x_4	x_5	
z	$\frac{2}{3}$	0	0	$2\frac{2}{3}$	$\frac{1}{3}$	$23\frac{1}{3}$
x_3	$\frac{1}{3}$	0	1	$\frac{1}{3}$	$-\frac{1}{3}$	$\frac{2}{3}$
x_2	1	1	0	1	0	8

$z=23\frac{1}{3}, x=(0,8,\frac{2}{3},0,0), B=(3,2)$

Since our reduced costs are all non-negative, we are done.

Exercise 2

Create the dual formulation, either by converting to equality form and using the primal-dual correspondence in (1) or by using the rules for inequality form.

End Exercise 2

Solution 2

Using the rules for inequality form, where non-negative primal variables correspond to \geq dual constraints and \leq primal constraints to non-negative dual variables, we have:

$$\begin{aligned} \min \quad & 8y_1 + 6y_2 \\ \text{s.t.} \quad & y_1 \geq 2 \\ & y_1 + y_2 \geq 3 \\ & -3y_2 \geq -1 \\ & y_1, y_2 \geq 0 \end{aligned}$$

End Solution 2

Exercise 3

Use the optimal dual solution to confirm that the primal objective value is optimal (i.e., the idea of a “certificate of optimality.”)

This can be done in two ways. First, we can read the optimal dual from the reduced costs on slack variables x_4 and x_5 in the optimal primal tableau. Second, we can read A_B^{-1} from the columns corresponding to slack variables x_4 and x_5 in the optimal primal tableau. Given this, we can use the algebraic expression $y^T = c_B^T A_B^{-1}$ to compute the dual. Either way, we verify the dual solution is feasible and that its objective value is equal to that of the primal.

End Exercise 3

Solution 3

Reading off the dual from the shadow costs, we see $y = [2\frac{2}{3}, \frac{1}{3}]$. We can check this is feasible, and also that the dual objective value is objective value of $23\frac{1}{3}$. This confirms the primal solution is optimal. We can also read A_B^{-1} from the optimal primal tableau, using the columns corresponding to the slack variables to get:

$$\begin{bmatrix} \frac{1}{3} & -\frac{1}{3} \\ 1 & 0 \end{bmatrix}$$

Combining with $c_B = [-1, 3]$ (noting the basic variables are in row order 3 then 2), we have $y^T = c_B^T A_B^{-1} = [2\frac{2}{3}, \frac{1}{3}]$.

End Solution 3

By completing the above exercises, you should now:

- understand the relationship between the dual and primal linear program and the utility of the dual
- be comfortable converting the primal problem to the dual problem and vice versa.

3 Sensitivity Analysis

So far we have been happy once we found an optimal solution to an LP. Now we want to look at how robust this solution is to changes in our input data.

Consider an LP in standard equality form:

$$\begin{aligned} \max \quad & c^T x \\ \text{s.t.} \quad & Ax = b \\ & x \geq 0 \end{aligned}$$

We can analyze changes in the cost vector c , in the constraint matrix A and in the RHS of the constraints, b . But before we do so, let's review some of the algebra that's necessary for such an analysis.

3.1 Algebra Review

Let B denote the optimal basis. We can write the tableau for the initial LP as:

$$\begin{aligned}z - c_B^T x_B - c_{B'} x_{B'} &= 0 \\ A_B x_B + A_{B'} x_{B'} &= b\end{aligned}$$

Now, we can multiply the second row by A_B^{-1} to get

$$I x_B + A_B^{-1} A_{B'} x_{B'} = A_B^{-1} b \quad (2)$$

This tells us that in the final tableau we have:

$$\text{(RHS)} \quad \bar{b} = A_B^{-1} b \quad (3)$$

We can now add c_B^T times (2) to the first row in the initial tableau, to get

$$z + c_B^T A_B^{-1} A_{B'} x_{B'} - c_{B'}^T x_{B'} = c_B^T A_B^{-1} b \quad (4)$$

In particular, since the optimal dual is

$$\text{(DUAL)} \quad y^T = c_B^T A_B^{-1} \quad (5)$$

this means the objective value of the optimal solution is

$$z = y^T b \quad (6)$$

Looking at the structure of (4) we can also see that for non-basic variable $j \in B'$ we have

$$\bar{c}_j = c_B^T A_B^{-1} A_j - c_j \quad (7)$$

More simply, we can write this as

$$\text{(REDUCED COSTS)} \quad \bar{c}_j = y^T A_j - c_j \quad (8)$$

Expressions (3), (5) and (8) are especially important for sensitivity analysis. We will want (8) to remain non-negative for all non-basic variables, which also requires checking what is changing with the dual (5). We will also want the RHS (3) to remain non-negative for primal feasibility.

3.2 Sensitivity: What if something changes?

When analyzing the sensitivity of a basic feasible solution to changes in the input data, we must differentiate the following 6 cases:

1. Changing objective function coefficient of a non-basic variable.
2. Changing objective function coefficient of a basic variable.
3. Changing the RHS.
4. Changing the entries in the column of a non-basic variable.
5. Changing the entries in the column of a basic variable.
6. Adding a variable to the LP.

Exercise 4

Let's write down a brief summary of sensitivity analysis. For each type of change, write down the effect on the optimal solution while the current basis remains feasible and optimal, and what needs to remain true for the current basis to remain feasible and optimal.

1. Changing nonbasic objective function coefficient c_j
2. Changing basic objective function coefficient c_j
3. Changing RHS b_i of a constraint
4. Changing column entries of a nonbasic variable x_j , or adding a new variable

End Exercise 4

Solution 4

1. Changing nonbasic objective function coefficient c_j
 - (a) No change to the optimal solution while in allowable range
 - (b) Need reduced cost $\bar{c}_j \geq 0$ (so just need to check this reduced cost)
2. Changing basic objective function coefficient c_j
 - (a) Objective value may change, but optimal solution remains the same while in allowable range
 - (b) Need reduced cost $\bar{c}_i \geq 0$ for all nonbasic variables x_i (so, need to check effect of change in dual y on these reduced costs)
3. Changing RHS b_i of a constraint
 - (a) Optimal solution as well as the objective value of this solution may change (since \bar{b} is changing)
 - (b) Need RHS $\bar{b} \geq 0$, i.e. RHS for all constraints in final tableau remains non-negative
4. Changing column entries of a nonbasic variable x_j or adding a new variable x_j

- (a) No change to optimal solution while change in allowable range
- (b) Need reduced cost $\bar{c}_j \geq 0$ for this nonbasic variable (and thus just need to check this reduced cost)

End Solution 4

3.3 An example: Changing the objective coefficient of a basic variable

Let's go through one specific example. Consider the following LP:

$$\begin{array}{rllllll}
 \text{maximize} & 5x_1 & + & x_2 & - & 12x_3 & & & & \\
 \text{subject to} & 3x_1 & & & 2x_2 & + & x_3 & & = & 10 \\
 & 5x_1 & + & 3x_2 & & & & + & x_4 & = 16 \\
 & & & & & & & & & x_1, x_2, x_3, x_4 \geq 0
 \end{array}$$

Using the simplex method until all reduced costs are non-positive, we arrive at the following tableau (you can check by yourself!):

	x_1	x_2	x_3	x_4	
z	0	0	2	7	12
x_2	0	1	5	-3	2
x_1	1	0	-3	2	2

$x = \{2, 2, 0, 0\}, z = 12.$

Note that the rows in the final tableau correspond to x_2 and x_1 , in this specific order. This particular order plays a role when we want to read off information from the final tableau.

Now we ask: what happens when the objective function coefficient of x_1 i.e. c_1 , changes by ε , i.e. from 5 to $5 + \varepsilon$?

3.3.1 Sensitivity Analysis by Hand

Note that c_1 is part of c_B because x_1 is a basic variable. Thus, let's look back at the algebraic expressions that we have derived to see in which expressions c_B shows up.

Exercise 5

In which of the algebraic expressions does c_B show up?

End Exercise 5

Solution 5

- Only in the expression for the reduced cost of non-basic variables:
- Non-basic reduced cost: $\bar{c}_j = y^T A_j - c_j$, since $y = c_B^T A_B^{-1}$ is changing.

End Solution 5

We have to check the reduced costs of all non basic variables to see they remain non-negative. We will use $\bar{c}_j = y^T A_j - c_j$, since $y = c_B^T A_B^{-1}$ and check for the reduced costs on non-basic variables c_3 and c_4 .

First, $c_B^T = (1, 5 + \varepsilon)$; second, $A_B^{-1} = \begin{pmatrix} 5 & -3 \\ -3 & 2 \end{pmatrix}$, which we can read off from the final tableau. (The order of the basic variables in the two matrices must match the order of the rows in the final tableau.)

Exercise 6

Compute the new reduced costs \bar{c}_3 and \bar{c}_4 using the algebraic expression $\bar{c}_j = y^T A_j - c_j$ and $y = c_B^T A_B^{-1}$

End Exercise 6

Solution 6

First, we compute the dual solution $y^T = c_B^T A_B^{-1}$ as follows:

$$\begin{aligned} y^T &= c_B^T A_B^{-1} \\ &= (1, 5 + \varepsilon) \begin{pmatrix} 5 & -3 \\ -3 & 2 \end{pmatrix} \\ &= (-15 - 3\varepsilon + 5, 10 + 2\varepsilon - 3) \\ &= (-10 - 3\varepsilon, 7 + 2\varepsilon) \end{aligned}$$

Now, we can use our algebraic expression for the reduced costs \bar{c}_3 and \bar{c}_4 and compute those. Thus, we get:

$$\begin{aligned} \bar{c}_3 &= y^T A_3 - c_3 \\ &= (-10 - 3\varepsilon, 7 + 2\varepsilon) \cdot (1, 0)^T + 12 \\ &= -10 - 3\varepsilon + 12 \\ &= 2 - 3\varepsilon \end{aligned}$$

and

$$\begin{aligned} \bar{c}_4 &= y^T A_4 - c_4 \\ &= (-10 - 3\varepsilon, 7 + 2\varepsilon) \cdot (0, 1)^T - 0 \\ &= +7 + 2\varepsilon - 0 \\ &= 7 + 2\varepsilon \end{aligned}$$

End Solution 6

To maintain optimality, even after c_1 has changed by ε , the reduced costs must be non-negative.

Exercise 7

Compute the conditions on ε such that the optimal solution does not change even if c_1 has been increased by ε .

Hint: Basically, we want to ensure that the new reduced costs computed in the previous exercise are always non-negative.

End Exercise 7

Solution 7

$$\begin{aligned}\bar{c}_3 &\geq 0 \\ \Leftrightarrow 2 - 3\varepsilon &\geq 0 \\ \Leftrightarrow 2 &\geq 3\varepsilon \\ \Leftrightarrow \varepsilon &\leq \frac{2}{3}\end{aligned}$$

and

$$\begin{aligned}\bar{c}_4 &\geq 0 \\ \Leftrightarrow +7 + 2\varepsilon &\geq 0 \\ \Leftrightarrow 2\varepsilon &\geq -7 \\ \Leftrightarrow \varepsilon &\geq -\frac{7}{2}\end{aligned}$$

Putting the two together, we get:

$$-\frac{7}{2} \leq \varepsilon \leq \frac{2}{3}$$

which completes the sensitivity analysis for the objective coefficient of x_1 .

End Solution 7

3.3.2 Sensitivity Analysis in AMPL/CPLEX

Now that we have done an example for a sensitivity analysis by hand, let's see how this works in AMPL. See the file `sensitivity.mod`. We first write down the formulation of the LP in AMPL:

```

var x1 >= 0;
var x2 >= 0;
var x3 >= 0;
var x4 >= 0;

```

```

maximize objective: 5*x1 + x2 -12*x3;

```

```

subject to constraint1: 3*x1 + 2*x2 + x3 = 10;

```

```

subject to constraint2: 5*x1 + 3*x2 + x4 = 16;

```

We then solve the model, setting the solves as CPLEX and telling CPLEX to provide sensitivity information and use the Simplex algorithm:

```

model sensitivity.mod;
option solver cplex;
option cplex_options 'sensitivity primalopt';
solve;

```

We obtain the sensitivity report by querying the bounds on the cost coefficient for x_1 :

```

ampl: display x1.current; x1.current = 5

```

```

ampl: display x1.up; x1.up = 5.66667

```

```

ampl: display x1.down; x1.down = 1.5

```

which agrees with our findings that c_1 can increase by at most $\frac{2}{3}$ and decrease by at most 3.5.

Note, the bounds on the RHS constraints (which we have not analyzed by hand!) could be displayed with the following commands:

```

ampl: display constraint1.current; constraint1.current = 10

```

```

ampl: display constraint1.up; constraint1.up = 10.6667

```

```

ampl: display constraint1.down; constraint1.down = 9.6

```

4 Understanding Sensitivity Analysis

Consider the following LP:

$$\begin{array}{rcllcl}
 \text{maximize} & 5x_1 & + & x_2 & - & 12x_3 & & & & \\
 \text{subject to} & 3x_1 & & 2x_2 & + & x_3 & & & = & 10 \\
 & 5x_1 & + & 3x_2 & & & + & x_4 & = & 16 \\
 & x_1, x_2, x_3, x_4 & \geq & 0 & & & & & &
 \end{array}$$

Solving with AMPL/CPLEX, we obtain the following solution for this LP:

$$x = (2, 2, 0, 0) \quad \text{and} \quad z = 12 \tag{9}$$

Furthermore, let's assume we get the following sensitivity analysis output from AMPL (see `sensitivity.mod` and `sensitivity.run`):

```

%% : _varname      _var.rc      _var.down _var.current  _var.up      :=
1  x1             -7.10543e-15  1.5       5             5.66667
2  x2             -3.55271e-15  0.6       1             3.33333
3  x3             -2            -1e+20    -12          -10
4  x4             -7            -1e+20    0            7 ;

:   _conname      _con.dual  _con.down  _con.current  _con.up      :=
1  constraint1    -10        9.6        10           10.6667
2  constraint2     7          15         16           16.6667 ;

```

Exercise 8

For the following questions, either answer them using the information from the sensitivity analysis or state if you need more information to answer it.

1. Assume the the RHS of constraint 1 increases by 0.5. Is the current basis still optimal? What happens to the optimal value?
2. What happens to the optimal value if the RHS of constraint 2 decreases by 1?
3. What happens to the optimal value if the RHS of constraint 2 increases by 1?

End Exercise 8

Solution 8

1. Looking at the sensitivity analysis for constraint 1, we see that an increase of 0.5 is within the allowable range of change (+2/3 is OK). Thus, the basis will not change. Now, recall the equation for the objective value: $z = y^T b$. Thus, we only have to look at the corresponding dual variables in the sensitivity analysis to compute the change of the objective value. The total cost changes by $0.5 \cdot (-10) = -5$. Thus, the new cost is 7.
2. This change is also in the allowable range of change. Thus, looking at the dual variables we know that the total cost changes by $-1 \cdot 7 = -7$. Thus, the new cost is 5.
3. Looking at the sensitivity analysis, we see that an increase of 1 is outside of the allowable range and the optimal basis will change. Thus, we cannot answer this question directly with the information present.

End Solution 8

Exercise 9

For the following questions, either answer them using the information from the sensitivity analysis or state if you need more information to answer it.

1. Assume we increase the objective coefficient of variable x_1 by 0.5. Is the current basis still optimal? What happens to the optimal value?
2. Assume we increase the objective coefficient of variable x_4 by 3. Is the current basis still optimal? What happens to the optimal value?
3. Assume we increase the objective coefficient of variable x_3 by 3. Is the current basis still optimal? What happens to the optimal value?

End Exercise 9

Solution 9

1. By looking at the sensitivity analysis for variable x_1 we see that an increase of 0.5 is within the allowable range of change (x_1 can be between 1.5 and 5.6666667). Thus, the current basis remains optimal. The new optimal value changes by $0.5 \cdot 2 = 1$. Thus the new total value is 13.
2. The cost increase is within the allowable range for variable x_4 . Thus, the current basis remains optimal. Because x_4 is a non-basic variable, i.e. it is set equal to 0, this change in the cost coefficient does not change the objective value.
3. This cost increase is outside of the allowable range for variable x_3 (var.up was -10). Thus, the optimal basis will change and we cannot answer this question fully without doing more computation.

End Solution 9

By completing the above exercises, you should now:

- understand how robustness of a solution can be analyzed using sensitivity analysis and be able to conduct sensitivity analysis by hand.
- understand sensitivity information in AMPL/CPLEX and be comfortable determining how changes in the input affect the optimal solution using sensitivity analysis output from AMPL.

5 The Dual Simplex Algorithm

5.1 Review

We have already seen how we can use duality to check optimality of a basic solution at hand. Now, we will see a very practical algorithm, called the Dual Simplex Algorithm, that uses duality and can be applied in situations where the usual simplex algorithm, primal simplex, cannot.

We will see in lecture that there is a direct correspondence between variables in the primal simplex algorithm and the variables in the dual simplex algorithm. This correspondence holds because we insist that complementary slackness holds between the primal tableau and dual tableau.

In particular:

1. For every primal tableau, there is a corresponding dual tableau

Exercise 10

For the following exercises, remember the algebraic relations listed at the beginning of the section notes and refer to the sensitivity analysis from exercise 1 above. Remember from the sensitivity analysis that when the RHS of constraint 2 is increased by 1, we are outside the allowable range and the optimal basis changes. We can use the dual simplex algorithm to proceed:

1. Assume, the RHS of the constraint 2 is increased by 1. Compute the new RHS, i.e. \bar{b} .
2. Compute the new objective value z for the (dual feasible) tableau.
3. Construct the new tableau corresponding to the current basis but considering the change in b .
4. By performing one iteration of the dual simplex algorithm, we arrive at the following tableau:

	x_1	x_2	x_3	x_4	
z	0	$\frac{7}{3}$	$13\frac{2}{3}$	0	$16\frac{2}{3}$
x_4	0	$-\frac{1}{3}$	$-\frac{5}{3}$	1	$\frac{1}{3}$
x_1	1	$\frac{2}{3}$	$\frac{1}{3}$	0	$3\frac{1}{3}$

Explain what are the entering and leaving variables, why are they chosen, and what is the pivot operation? Explain why the dual simplex algorithm is finished. Given the optimal solution and its objective value.

End Exercise 10

Solution 10

1. We can read off A_B^{-1} from the final tableau, i.e. $A_B^{-1} = \begin{pmatrix} 5 & -3 \\ -3 & 2 \end{pmatrix}$. We know that $\bar{b} = A_B^{-1}b$, thus $\bar{b} = \begin{pmatrix} 5 & -3 \\ -3 & 2 \end{pmatrix} \cdot \begin{pmatrix} 10 \\ 17 \end{pmatrix} = \begin{pmatrix} -1 \\ 4 \end{pmatrix}$.
2. From the AMPL output for the sensitivity analysis we know that the dual variables are: $y = (-10, 7)$. Thus, $z = y^T b = (-10, 7) \cdot \begin{pmatrix} 10 \\ 17 \end{pmatrix} = 19$.
3. Thus, the new tableau looks as follows:

	x_1	x_2	x_3	x_4	
z	0	0	2	7	19
x_2	0	1	5	-3	-1
x_1	1	0	-3	2	4

4. Because x_2 is now negative, it has to leave the basis. Thus we pivot on row 1. The only negative entry in this row correspond to x_4 , thus x_4 will enter the basis.

The new tableau is dual optimal. Thus, we are done. We found a new basic feasible solution $x = (3\frac{1}{3}, 0, 0, \frac{1}{3})$ with optimal objective value $z = 16\frac{2}{3}$.

End Solution 10

By completing the above exercises, you should now:

- understand how the dual simplex algorithm works.
- be comfortable applying the dual simplex algorithm to solve an LP.

6 Complementary Slackness

There is a strong correspondence between primal and dual problems. Among these, we have seen that the constraints in one problem are complementary to the variables in the other. Stemming from this, *complementary slackness* refers to an association between the slack in primal inequality constraints and positivity of their corresponding dual variable.

It is stated for the standard inequality form.

Definition. (Complementary slackness) *Given a feasible primal solution x and a feasible dual solution y , complementary slackness insists that there is not slack in BOTH a variable and a corresponding constraint.*

Slack in a variable means it is strictly positive. Slack in a constraint means it is not binding.

We see that comp slackness between a primal and dual pair can equivalently be stated as

- If $x_j > 0$, then the j th constraint in (D) is binding
- If the j th constraint in (D) is not binding, then $x_j = 0$

For dual variable primal constraint pairs, we have

- If $y_i > 0$, then the i th constraint in (P) is binding
- If the i th constraint in (P) is not binding, then $y_i = 0$

Theorem 6.1. *A feasible primal x and a feasible dual y are both optimal if and only if (x, y) satisfy complementary slackness.*

See the lecture notes for the proof.

Complementary Slackness is useful not only because it helps us to interpret the meaning of dual variables (e.g., dual values are positive only if primal constraints are binding!), but also because it provides a convenient way to find the dual optimal solution knowing the primal optimal solution.

6.1 Practice

Suppose we have the following problem:

$$\begin{aligned} \max \quad & 2x_1 + 4x_2 + 3x_3 + x_4 \\ \text{s.t.} \quad & 3x_1 + x_2 + x_3 + 4x_4 \leq 12 \\ & x_1 - 3x_2 + 2x_3 + 3x_4 \leq 7 \end{aligned}$$

$$2x_1 + x_2 + 3x_3 - x_4 \leq 10$$

$$x_1, x_2, x_3, x_4 \geq 0$$

If we solve this problem using Simplex, we will find that Optimal Value of 42 can be obtained at the solution vector $(0, 10.4, 0, .4)$.

Now, we will use the complementary slackness conditions to solve for the dual optimal solution without using the simplex algorithm.

Exercise 11

From the optimal basic solution $x = (0, 10.4, 0, .4)$, we know that x_2 and x_4 are positive in the optimal solution. What does this imply about the 2nd and 4th dual constraints based on the complementary slackness conditions? The dual is as follows:

$$\begin{aligned} \min \quad & 12y_1 + 7y_2 + 10y_3 \\ \text{s.t.} \quad & 3y_1 + y_2 + 2y_3 \geq 2 \\ & y_1 - 3y_2 + y_3 \geq 4 \\ & y_1 + 2y_2 + 3y_3 \geq 3 \\ & 4y_1 + 3y_2 - y_3 \geq 1 \\ & y_1, y_2, y_3 \geq 0 \end{aligned}$$

End Exercise 11

Solution 11

What do the complementary slackness conditions tell us? They tell us the dual constraints for positive primal variables will bind, and thus in fact be satisfied as equalities. Because we know x_2 and x_4 are positive, we 2nd and 4th dual constraints should be binding (equalities):

$$\begin{aligned} y_1 - 3y_2 + y_3 &= 4 \\ 4y_1 + 3y_2 - y_3 &= 1 \end{aligned}$$

End Solution 11

Exercise 12

Note that at our primal solution, the first and third primal constraints bind. But the second one does not – there is slack in this constraint:

$$\begin{aligned} x_1 - 3x_2 + 2x_3 + 3x_4 &\leq 7 \\ 1(0) - 3(10.4) + 2(0) + 3(.4) &\leq 7 \\ -30 &\leq 7 \end{aligned}$$

What does this imply about the 2nd dual variable y_2 ?

End Exercise 12

Solution 12

The second primal constraint has a slack of $7 - (-30) = 37 > 0$. Complementary slackness then implies that the corresponding dual variable, y_2 should be 0. Giving us the following set of equations:

$$\begin{aligned}y_1 - 3y_2 + y_3 &= 4 \\4y_1 + 3y_2 - y_3 &= 1 \\y_2 &= 0\end{aligned}$$

Or:

$$\begin{aligned}y_1 + y_3 &= 4 \\4y_1 - y_3 &= 1 \\y_2 &= 0\end{aligned}$$

End Solution 12

Exercise 13

Solve the system of equations to find the dual optimal solution. What is the optimal objective value?

End Exercise 13

Solution 13

The system of equations

$$\begin{aligned}y_1 + y_3 &= 4 \\4y_1 - y_3 &= 1 \\y_2 &= 0\end{aligned}$$

tells us that $y_1 = 1$ and $y_3 = 3$. This gives us a dual solution of $y = (1, 0, 3)$. We can check that this solution is feasible for the other dual constraints (it is). And we can check that the value of this solution in the dual problem is 42, which is equal to the value of our primal solution and thus optimal.

End Solution 13

By completing the above exercises, you should now:

- understand the concept of complementary slackness.
- be comfortable using complementary slackness to solve the dual LP.

7 Some More Clarification

In earlier sections, we saw that we can read some useful values from our final tableaux. But why does this work?

For example, recall Exercise 3, where we were given the following initial tableau:

	x_1	x_2	x_3	x_4	x_5	
z	-2	-3	1	0	0	0
x_4	1	1	0	1	0	8
x_5	0	1	-3	0	1	6

$z=0, x=(0,0,0,8,6), B=(4,5)$

and the following optimal tableau:

	x_1	x_2	x_3	x_4	x_5	
z	$\frac{2}{3}$	0	0	$2\frac{2}{3}$	$\frac{1}{3}$	$23\frac{1}{3}$
x_3	$\frac{1}{3}$	0	1	$\frac{1}{3}$	$-\frac{1}{3}$	$\frac{2}{3}$
x_2	1	1	0	1	0	8

$z=23\frac{1}{3}, x=(0,8,\frac{2}{3},0,0), B=(3,2)$

We wanted to isolate the matrix A_B^{-1} , which by definition we can find by taking the inverse of the columns corresponding to the optimal basis in the initial tableau. In other words, since the optimal basis is (3,2) we can read the following matrix from the initial tableau:

$$A_B = \begin{bmatrix} 0 & 1 \\ -3 & 1 \end{bmatrix}$$

Taking the inverse of this, we then have

$$A_B^{-1} = \begin{bmatrix} \frac{1}{3} & -\frac{1}{3} \\ 1 & 0 \end{bmatrix}$$

Now we clearly see that these same values can be found in the final tableau in the columns corresponding to our initial basis (4,5). Note that in the initial tableau, these variables' columns form an identity matrix (true for any current basic variables). Also note that the pivot operations needed to isolate the final basic variables equate to multiplying matrix A_B by its inverse (to obtain a new identity matrix). Since the operations applied hold over the entirety of the tableau rows, these operations would transform the matrix coefficients of the initial basic variables into those of A_B^{-1} .

Thus we can read off the desired coefficients without actually having to compute the inverse of any matrices!

See the lecture notes for a proof of this observation.