Slater Strong Duality: Finiteness Conditions for Attainment

2026-08-24

Verdict: false-source (resolved).

1 Source passage

Let V and W be finite-dimensional real inner-product spaces, let $K \subseteq V$ be a closed pointed convex cone with nonempty interior, and let $T : V \to W$ be linear. For $c \in V$ and $b \in W$, Chapter 4 of Ref. [Wol12] defines the primal and dual values by

$$C_p = \inf\{\langle c|x\rangle : x \in K,\ T(x) = b\}, \tag{1}$$
$$C_d = \sup\{\langle b|y\rangle : c - T^*(y) \in K^*\}. \tag{2}$$

At lines 72–78 of the local source, primal strict feasibility is said to imply $C_p = C_d$ and attainment of the dual supremum, while dual strict feasibility is said to imply attainment of the primal infimum [Wol12, Chapter 4].

The semidefinite specialization at lines 100–105 of the same source contains the same missing finiteness condition and a separate directional error [Wol12, Chapter 4]. It assumes either a positive-definite primal matrix X satisfying $\mathrm{tr}(F_i X) = b_i$, or a dual vector y whose slack $F_0 - \sum_i y_i F_i$ is positive definite, and then states in both cases that the dual extremum is attained. Primal strict feasibility yields dual attainment when the primal value is finite. Dual strict feasibility instead yields primal attainment when the dual value is finite; it does not generally yield dual attainment.

2 Missing boundary conditions

Strict feasibility on one side does not ensure that the opposite feasible set is nonempty or that the relevant optimum is finite. Both failures occur for $V = \mathbb{R}$, $K = \mathbb{R}_{\geq 0}$, and $T = 0$.

First, set $b = 0$ and $c = -1$. The point $x = 1$ is strictly primal feasible, but

$$\inf_{x \geq 0}(-x) = -\infty. \tag{3}$$

Dual feasibility would require

$$-1 \in \mathbb{R}_{\geq 0}, \tag{4}$$

which is impossible. Thus the dual feasible set is empty and has no optimizer.

Second, set $b = 1$ and $c = 1$. The dual slack is 1 for every y, so the dual problem is strictly feasible. The primal equality constraint becomes

$$0 = 1, \tag{5}$$

so the primal feasible set is empty and has no optimizer.

These examples are also semidefinite programs on 1×1 Hermitian matrices with one equality constraint. For the first example, take $F_0 = -1$, $F_1 = 0$, and $b_1 = 0$. The matrix $X = 1$ is positive definite and satisfies $\operatorname{tr}(F_1 X) = b_1$, but the primal objective is unbounded below and the dual slack satisfies

$$F_0 - y_1 F_1 = -1 \tag{6}$$

for every y_1, so it is never positive semidefinite. For the second example, take $F_0 = 1$, $F_1 = 0$, and $b_1 = 1$. Every dual slack is the positive-definite matrix 1, while the primal equality

$$\operatorname{tr}(F_1 X) = 1 \tag{7}$$

is impossible.

Therefore, the attainment statements at lines 72–78 and their semidefinite specialization at lines 100–105 require a finiteness condition, or an equivalent condition combining opposite-side feasibility and boundedness [Wol12, Chapter 4].

A finite-valued semidefinite program isolates the directional error at line 105. In the convention of Ref. [Wol12], set

$$F_0 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tag{8}$$

$$F_1 = \begin{pmatrix} -1 & 0 \\ 0 & 0 \end{pmatrix}, \tag{9}$$

$$F_2 = \begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix}, \tag{10}$$

$$b = (-1, 0). \tag{11}$$

Writing the dual variable as $y = (x, z)$ gives

$$F_0 - xF_1 - zF_2 = \begin{pmatrix} x & 1 \\ 1 & z \end{pmatrix}, \tag{12}$$

$$\langle b|y\rangle = -x. \tag{13}$$

The dual problem therefore maximizes $-x$ subject to positivity of the matrix in (12). It is strictly feasible at $(x,z)=(2,2)$. Feasibility implies

$$x > 0, \tag{14}$$

$$z > 0, \tag{15}$$

$$xz \geq 1. \tag{16}$$

For every $t > 0$, the point $(x,z)=(t,t^{-1})$ is feasible. Hence

$$\sup\{-x : (x,z) \text{ is dual feasible}\} = 0, \tag{17}$$

but the supremum is approached only as $t \downarrow 0$ and is not attained.

The associated primal constraints are

$$-X_{11} = -1, \tag{18}$$

$$-X_{22} = 0. \tag{19}$$

Positive semidefiniteness then forces

$$X = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \tag{20}$$

and its objective value is

$$\operatorname{tr}(F_0 X) = 0. \tag{21}$$

Thus the primal infimum is attained and equals the dual supremum, exactly as the correctly oriented Slater theorem asserts.

3 Corrected statement

In the signs and adjoint convention of Ref. [Wol12], the finite-dimensional Slater theorem has the following two directions:

$$\text{primal strict feasibility and } C_p \in \mathbb{R} \Longrightarrow C_p = C_d \text{ and dual attainment,} \tag{22}$$

$$\text{dual strict feasibility and } C_d \in \mathbb{R} \Longrightarrow C_p = C_d \text{ and primal attainment.} \tag{23}$$

The formal development proves both implications in `QICLean/Analysis/ConicProgram`.[1] The first pair assumes primal strict feasibility and $C_p = p \in \mathbb{R}$; the second assumes dual strict feasibility and $C_d = p \in \mathbb{R}$. The declarations `ConicProgram.primalValue_eq_coe_iff_isGLB` and `ConicProgram.dualV`

[1] The primal-strict direction is given by `ConicProgram.exists_isDualOptimizer_of_isPrimalStrictlyFeasible_of_primalValue_eq_coe` and `ConicProgram.values_eq_of_isPrimalStrictlyFeasible_of_primalValue_eq_coe`. The dual-strict direction is given by `ConicProgram.exists_isPrimalOptimizer_of_isDualStrictlyFeasible_of_dualValue_eq_coe` and `ConicProgram.values_eq_of_isDualStrictlyFeasible_of_dualValue_eq_coe`.

`alue_eq_coe_iff_isLUB` identify these hypotheses with the corresponding real greatest-lower-bound and least-upper-bound conditions whenever the relevant feasible set is nonempty.

The direct semidefinite specializations are proved in `QICLean/Analysis/SemidefiniteDuality`.[2] A positive-definite primal point together with $C_p = p \in \mathbb{R}$ gives dual attainment and value equality. A positive-definite dual slack together with $C_d = p \in \mathbb{R}$ gives primal attainment and value equality.

4 Separation argument

Let C be a cone, let A be an equality map, let d be an objective vector, and let r be the right-hand side. The proof separates the lifted epigraph

$$D = \{(A(x), \langle d|x\rangle + s) : x \in C,\ s \geq 0\} \subseteq \operatorname{ran}(A) \times \mathbb{R}. \tag{24}$$

Restricting the codomain to $\operatorname{ran}(A)$ makes the lifted linear map surjective and therefore open in finite dimension. Strict feasibility gives D nonempty interior. If p is the finite greatest lower bound, then

$$(r, p) \in \overline{D} \setminus \operatorname{int} D. \tag{25}$$

A supporting functional f therefore satisfies

$$f(r, p) = 0, \tag{26}$$
$$f(z) \leq 0 \qquad \text{for every } z \in D. \tag{27}$$

Its vertical coefficient

$$\alpha = f(0, 1) \tag{28}$$

is strictly negative. If $\alpha = 0$, strict feasibility and openness of the nonzero functional would make 0 an interior point of $(-\infty, 0]$, which is impossible. Normalization by $-\alpha$ gives the required dual multiplier and its attained objective value.

For the converse direction, write the dual problem as a primal problem on $W \times K^*$ with equality map and objective

$$(y, z) \longmapsto T^*(y) + z, \tag{29}$$
$$(y, z) \longmapsto \langle -b|y\rangle. \tag{30}$$

The multiplier returned by the same separation argument satisfies

$$T(-x) = b, \tag{31}$$
$$-x \in K^{**} = K. \tag{32}$$

[2] The primal-strict declarations are `SemidefiniteProgram.exists_dualOptimizer_of_primalStrict_of_primalValue_eq_coe` and `SemidefiniteProgram.values_eq_of_primalStrict_of_primalValue_eq_coe`. The dual-strict declarations are `SemidefiniteProgram.exists_primalOptimizer_of_dualStrict_of_dualValue_eq_coe` and `SemidefiniteProgram.values_eq_of_dualStrict_of_dualValue_eq_coe`.

These are precisely the primal feasibility conditions in Wolf's signs. The argument never assumes that a linear image of a closed cone is closed. It uses only the closure membership in (25), which follows from the greatest-lower-bound property.

5 Verdict

Verdict: false as printed; corrected theorem resolved. The two one-dimensional examples show that neither the general assertions at lines 72–78 nor their semidefinite specialization at lines 100–105 follows from strict feasibility alone. The 2×2 example also shows that dual strict feasibility and a finite dual supremum do not imply dual attainment, so line 105 names the wrong optimizer in the dual-strict case. With the missing finite-real optimum hypothesis, value equality and attainment on the opposite side are now formalized by separate declarations. This resolves QICLean issue #110 while retaining the source correction. The same correction governs the semidefinite specialization tracked by QICLean issue #111.

References

[Wol12] Michael M. Wolf. Quantum channels & operations: Guided tour. Lecture notes, 2012.