Conventional solution and proof

For $x \neq 0$, the chain rule applied to $(x^{\mathsf{T}}x)^{1/2}$ gives

$$\boxed{\nabla f(x) = x/\|x\|_2.}$$

The denominator is positive, so the expression is well defined. At zero, along any unit vector v, $f(tv) = |t|$ has unequal one-sided derivatives. Therefore no derivative exists at zero. Equivalently, a putative linear differential L would satisfy $L(v) = L(-v) = 1$, contradicting linearity.

Structural or analytic derivation

The subdifferential at zero is exactly the closed unit ball. Indeed, g is a subgradient if $\|y\|_2 \geq g^{\mathsf{T}}y$ for every y. Cauchy–Schwarz proves this for $\|g\|_2 \leq 1$; choosing $y = g$ proves the converse. At nonzero x, equality in Cauchy–Schwarz makes the subgradient unique and equal to $x/\|x\|_2$. A norm contains conjugation in complex coordinates and is not holomorphic.

Comparative cost and preferred method

Computing a dense n-coordinate norm gradient costs $\Theta(n)$: the formula reads every coordinate and writes n outputs, giving a matching input/output lower bound. Coordinatewise numerical differentiation requires $\Theta(n^2)$ work if each function call scans the full vector. Against a competent analytic or automatic-differentiation baseline there is no asymptotic improvement. Complex-step differentiation of the norm is invalid.