Direct evaluation uses n^2 weighted terms. Two forward FFTs, one inverse and linear preparation yield $O(n \log n)$ arithmetic. For $|z| \neq 1$, chirp magnitudes can grow or decay exponentially in n^2, so precision, overflow and bit complexity require separate control. On the DFT specialization, optimized SciPy FFT is a stronger conventional baseline because it already uses Bluestein for poorly factorizable sizes (reference S1). The zero-parameter branch is essential to the word "any" in the question.