Suppose we are given the convolution of two signals, $y = f * g$. When, under which conditions, and how can we reconstruct f and g from the knowledge of y alone if both f and g are unknown? This challenging problem, known as *blind deconvolution*, pervades many areas of science and technology, including astronomy, medical imaging, optics, and wireless communications. A key challenge of this difficult non-convex optimization problem is that it exhibits many local minima. We present an efficient numerical algorithm that is guaranteed to recover the exact solution, when the number of measurements is (up to log-factors) slightly larger than the information-thereotical minimum, and under reasonable conditions on g and f. The proposed regularized gradient descent algorithm converges at a geometric rate and is provably robust in the presence of noise. To the best of our knowledge, our algorithm is the first blind deconvolution algorithm that is numerically efficient, robust against noise, and comes with rigorous recovery guarantees under certain subspace conditions. Moreover, numerical experiments do not only provide empirical verification of our theory, but they also demonstrate that our method yields excellent performance even in situations beyond our theoretical framework.