The problem of learning a nonlinear controller directly from experimental data is considered. It is assumed that an existing, unknown controller, able to stabilize the plant, is available, and that input-output measurements can be collected during closed loop operations. A theoretical analysis shows that the error between the input issued by the existing controller and the input given by the learned one shall have low variability in order to achieve closed loop stability. This result is exploited to derive a ℓ_1-norm regularized learning algorithm that achieves the stability condition for a finite number of data points. The approach is completely based on convex optimization. The presented technique is finally tested in real-world experiments to control the flight of a tethered flexible wing, which is characterized by highly nonlinear, unstable and uncertain dynamics and is subject to external disturbances.