The estimation of local surface orientation (slant and tilt) is fundamental to the recovery of the three-dimensional structure of the environment. Although a great deal of research has been devoted to understanding how visual systems solve this problem, most previous work has focused on performance with artificial stimuli. Here, we study human surface tilt estimation with both natural and artificial stimuli. Natural stimuli are sampled from a stereo-image database of natural scenes with precisely co-registered range data at each pixel; groundtruth tilt, slant, and distance information is obtained directly from the range data. Artificial planar stimuli, generated in software, are matched to the tilt, slant, distance, and contrast of the natural stimuli. Human observers binocularly viewed natural and artificial surfaces through a small aperture and reported the estimated tilt with a mouse-controlled probe. Human performance in natural scenes is significantly less accurate and more strongly influenced (biased) by the tilt prior than human performance in matched artificial scenes. However, the imprecision of human estimates with natural stimuli is tightly predicted by an ideal observer for the task. The ideal observer reports the Bayes' optimal tilt estimate given three local image cues computed directly from the images. Remarkably, the ideal observer predicts human performance with zero free parameters at all levels of analysis, including trial-by-trial errors. These similarities suggest that the biased, imprecise patterns of human performance are nevertheless lawful, and that they result from the optimal computations on local areas of natural scenes.