

Following recent work by Weissman *et al.*, in the spirit of results on universal compression, we compare the performance of universal denoisers on discrete memoryless channels to that of the best k -th order omniscient denoiser, namely one that is tuned to the transmitted noiseless sequence. We show that the additional loss incurred in the worst case by any denoiser on a length- n sequence grows like $\Omega(c^k/\sqrt{n})$, where $c > 1$ is a constant depending on the channel parameters and the loss function. This shows that for fixed k the additional loss incurred by the Discrete Universal Denoiser (DUDE) is no larger than a constant multiplicative factor of the best possible. We also extend this result to the compound decision problem.