While working on the XDoG Compendium paper we ran into a fairly strong negative result. Essentially, while most DoGs aren't actually very good approximations to an LoG, the space of sharpening filters that can be created using a DoG with $k \leq 1.6$ is more or less identical to the space of LoG sharpeners. In other words, while changing the k value certainly will change the output of our filters, in almost all cases, the new output is one that could have been achieved using the original k along with a different set of p and σ values.

We considered adding an appendix on the significance of k values to the paper, but ultimately decided against it. What follows is the draft writeup.

1 Choice of k

In their seminal paper on the theory of Edge Detection, Marr and Hildreth noted that as k approaches 1, the difference of Gaussians approaches the Laplacian of Guassians (LoG) up to a factor of scale. More exactly, given

$$\begin{aligned} g_\sigma(\mathbf{x}) &:= \frac{1}{2\pi\sigma^2} e^{\frac{||\mathbf{x}||^2}{2\sigma^2}}, \\ d_\sigma(\mathbf{x}, k) &:= g_\sigma(\mathbf{x}) - g_\sigma(k\mathbf{x}). \end{aligned}$$

We can show that,

$$\lim_{k \to 1} \frac{k}{1-k} d_\sigma(\mathbf{x}, k) = \sigma^2 \nabla^2 g_\sigma(\mathbf{x}).$$

Unfortunately, the scale factor $\frac{k}{1-k}$ approaches negative infinity as k approaches 1. For this reason the DoG responses at more accurate k values require increasingly large scalings, resulting in a tradeoff between the accuracy of the approximation and the precision that can be achieved. Marr and Hildreth suggested $k = 1.6$ as a precision/accuracy tradeoff likely to be appropriate in most electrical engineering applications.

However, it is not clear that $k = 1.6$ is necessarily the best choice when the DoG is being used in the context of an edge stylization filter. Might a smaller k, one that provides a more accurate LoG approximation, produce superior results? Might using an exact LoG calculation, rather than an approximation, lead to a higher quality filter?

The answer to both these questions is a strong "no". To see why, first recall that the XDoG is defined in terms of a sharpening filter, which may be written as follows,

$$s(\mathbf{x}, \sigma, p, k) := g_\sigma(\mathbf{x}) + p d_\sigma(\mathbf{x}, k). \tag{1}$$

We can define a similar sharpening filter using the scale invariant Laplacian,

$$s'(\mathbf{x}, \sigma, p) := g_\sigma(\mathbf{x}) - p\sigma^2 \nabla^2 g_\sigma(\mathbf{x}). \tag{2}$$

Changing the value of k used in Eq. (1) from k to k' is only useful if it somehow improves the range of possible sharpening filters. More formally, it is only useful if

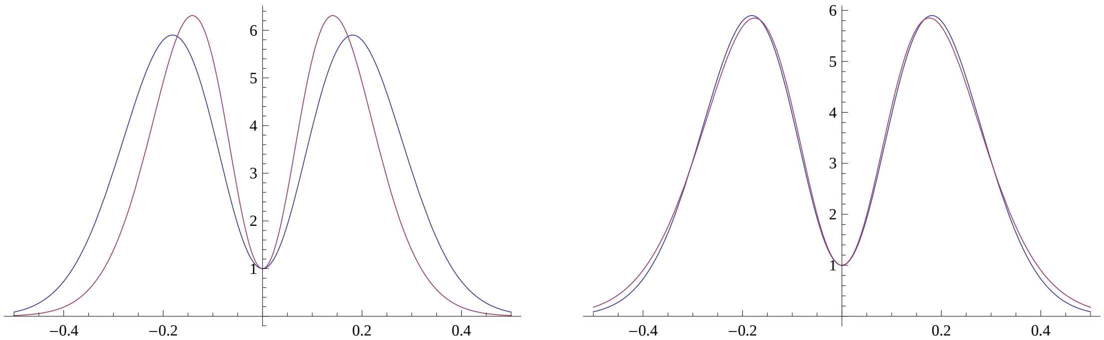

Figure 1: The left graph shows the frequency response of the Laplacian sharpening filter, $S'(\omega, 1.2, 7.5)$, in blue. The optimally approximating DoG sharpener of the form $S(\omega, 1.2, p, 1.6)$ occurs when $p \approx 17.24$. It is shown in purple. Clearly, there are significant differences between the two filters. These differences exist because the DoG calculated with $k = 1.6$ can provide only a crude approximation for the Laplacian when the two filters are applied at identical scales, i.e., when $\sigma' = \sigma$. However, if we allow the DoG to be applied at a different scale than the LoG, then it becomes possible to find a nearly exact match. Such a match is shown by the purple curve in the right hand graph, which plots $S(\omega, .96, 15.86, 1.6)$.

there are values p' and σ' for which $s(\mathbf{x}, \sigma', p', k')$ produces interesting edge images, and for which it is impossible to find values p and σ such that $s(\mathbf{x}, \sigma, p, k)$ implies a similar sharpening filter.

Analyzing the degree to which changing k may influence the range of edge images thus requires a means of measuring the similarity between two filters. There are many ways this might be done. We choose a simple Fourier space metric.

We start by comparing DoG sharpening to Laplacian sharpening. The filters in equations (1) and (2) are rotationally symmetric, and thus they can be reduced to one dimensional functions by the substitution $||\mathbf{x}||^2 \to r^2$. Transforming those functions to Fourier space yields,

$$\begin{aligned} S(\omega, \sigma, p, k) &= e^{-2\pi^2\sigma^2\omega^2}(1+p) - e^{-2\pi^2\sigma^2k^2\omega^2}, \\ S'(\omega, \sigma, p) &= e^{-2\pi^2\sigma^2\omega^2}(1+4p\pi^2\omega^2). \end{aligned}$$

We now define the error in any approximation of the Laplacian sharpening filter $s'(\mathbf{x}, \sigma', p')$ by a DoG sharpening filter $s(\mathbf{x}, \sigma, p, k)$ as the squared sum of differences between their responses to frequencies above the Nyquist rate,

$$E(\sigma, p, k, \sigma', p') := \int_{-\frac{1}{2}}^{\frac{1}{2}} \left(S(\omega, \sigma, p, k) - S'(\omega, \sigma', p')\right)^2 d\omega.$$

Let us assume that we will use σ values large enough that our filters will have little response to frequencies beneath the Nyquist rate, so $\sigma' > 1$. Numerical optimization

Figure 2: Result images generated using a sharpener based on the LoG (left) compared with that generated via a DoG (right). The parameters for each filter are those given for the right hand graphs in Figure 1. Note that there are no noticeable differences between the two results.

can be used to find values σ and p such that the approximation error $E(\sigma, p, k, \sigma', p')$ is minimized.

When both p and σ are allowed to vary as needed to provide an optimal match, experimentation reveals that for all p' and $\sigma' > 1$, any changes introduced by switching from a LoG to a DoG can be compensated for by making adjustments to both p and σ. Figure 1 provides some example matches.

Testing on real images shows that even in a case selected to maximize the error in the optimal DoG match, there are no noticeable differences between the edge image implied by an LoG sharpener and that of its optimally approximating DoG sharpener. Figure 2 compares two filtered images resulting from a relatively high error case.

Thus, there is no advantage to changing from the k=1.6 DoG to an exact LoG, as the space of possible edge effects remains the same. Reducing k below 1.6 results in sharpening filters that more accurately approximate the space of Laplacian sharpening filters, but, as that filter space is already effectively captured given $k = 1.6$, such a change can have no appreciable effect on the range of edge images that may be created using the XDoG filter.