Supplementary Material

Warp interpolation: linear and bidirectional interpolation

Linear interpolation picks the closest image:

$$I_{interpolated}(a,b) = \begin{cases} I_1(x,y) & if\,\alpha < 0.5, \\ I_2(x,y) & otherwise. \end{cases} \tag{1}$$

where a and b are the moved pixels between the images I_1 and I_2 at position $\alpha \in [0,1]$:

$$\begin{aligned} a &= \begin{cases} x+\alpha F_{x1} & if\,\alpha < 0.5, \\ x+\alpha F_{x2} & otherwise. \end{cases} \\ b &= \begin{cases} y+\alpha F_{y1} & if\,\alpha < 0.5, \\ y+\alpha F_{y2} & otherwise. \end{cases} \end{aligned} \tag{2}$$

where $F_{\{x,y\}1}$ is the motion field from I_1 to I_2 and $F_{\{x,y\}2}$ is the motion field from I_2 to I_1. In bidirectional interpolation two images using both flows separately are blended together:

$$I_{interpolated}(x,y) = (1-\alpha)I_{b1} + \alpha I_{b2} \tag{3}$$

$$I_{b\{1,2\}}(c,d) = I_{\{1,2\}}(x,y) \tag{4}$$

$$\begin{aligned} c &= \alpha F_{x\{1,2\}} \\ d &= \alpha F_{y\{1,2\}} \end{aligned} \tag{5}$$

There are some differences in the visual appearance between the two methods and the choice should depend on what the artist desires. By using bidirectional interpolation, we achieve a smooth interpolation with some ghosting. The amount of this ‘artefact’ depends on how different the images are (i.e., the performance of the motion field estimation). Since we are only using one of the input images in linear interpolation we do not receive any ghosting. However, instead we see a seamless transition between the two images. That is, we get the impression that the input is a video rather than a sparse set of images. The linear approach is more sensitive to the performance of the flow estimation, and has more noticeable popping on dynamic objects and significant frame-to-frame changes. Ultimately, the choice is left to the artist as both methods are visually distinct.