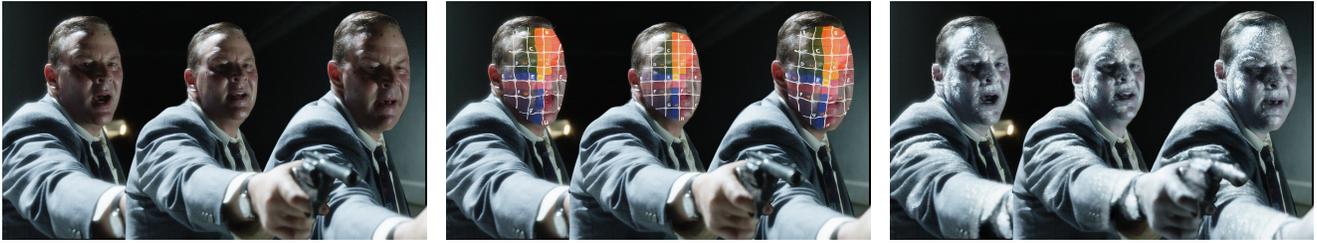


Warping with Accumulated Motion Vectors

Stephen A. Willey Gregory Keech
Double Negative*



Using MagicWarper to cover a character in ice by providing a single frame DMP

©2016 Marvel TV

Abstract

From facial lacerations that follow slight stretching of the skin, to deformations of hard bodied objects or subtle nuances of water simulations, our tool can automatically warp an image across an entire sequence from any arbitrary frame to any other using accumulated motion vectors, minimising the keyframes needed to produce high quality results.

Keywords: motion vectors, compositing, nuke, magic

Concepts: •Computing methodologies → Image processing;

1 Introduction

We present MagicWarper, which was developed to allow artists at Double Negative to quickly propagate textures or paint throughout an image sequence by providing a single arbitrary frame as the input.

This tool creates and utilises accumulated motion vectors across a sequence in order to automatically warp an input into the correct position for each frame. By pre-accumulating motion vectors for a variety of frame intervals and using a hierarchical temporal downsampling technique, we are able to quickly calculate any given frame jump.

We demonstrate the use of our tool by painting a single frame and propagating the changes backwards and forwards through a close-up shot of complex facial movements. This tool is effective regardless of image resolution – it is only limited by the quality of the motion vectors.

2 Our Approach

We accumulate motion vectors from an initial calculated pass, rather than calculate them on for the fly for every frame, due to the overall time saving this approach offers. Motion vectors are only computed for consecutive frames. These are then accumulated and stored on disk at intervals increasing with the power of 2 such that for frame n we have motion vectors on disk mapping to frame $n + 1$, $n + 2$, $n + 4$ and so on. We can then construct the required motion from

*e-mail: {saw,grk}@dneg.com

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). © 2016 Copyright held by the owner/author(s).

SIGGRAPH '16 Talks, July 24-28, 2016, Anaheim, CA,

ISBN: 978-1-4503-4282-7/16/07

DOI: <http://dx.doi.org/10.1145/2897839.2927456>

frame n to any other frame in the sequence with minimal disk access and CPU load. Although a binary tree approach would appear to offer greater efficiency still, it was found that such a method was inadequate in providing coherent forwards and backwards motion estimation to achieve the desired effect.

Our biggest concern when developing this tool was speed and accuracy for the artists. We allow for the use of an arbitrary keyframe that can easily be changed at any point, as well as allowing more to be added if necessary. If, for example, the tracked area is occluded at any point, the artist is able to simply create a new keyframe to account for this as occlusion detection is not currently implemented within this tool. Vector accumulation is lossless and if the initial singleton motion vectors are of high quality that level will be maintained throughout even over large ranges. In the cases where high quality motion vectors are not available, artists can add keyframes to re-align the warped area of the input image and mitigate drift.

The process of adding two vectors together is ostensibly very simple. However, it is not just a case of overlaying one frame over the other and adding the values together. We must first use the forward vectors from the current frame in order to locate the correct pixel values in the next frame to add. We created a template script in Nuke that allows us to perform this action once in order to produce the 2-frame accumulations, and then feed these back in to produce the 4-frame accumulations, and so on.

By pre-calculating power-of-2 accumulations we are able to minimise the number of reads when using the MagicWarper tool. This means we can quickly display results regardless of the frame jump being performed. If a user wishes to jump 19 frames, for example, then we simply need to perform the standard accumulation algorithm three times using the pre-calculated 16-, 2-, and 1-frame jump vectors, rather than reading nineteen single frames or attempting to calculate a single motion pass which can introduce sub-optimal results. We find this to provide a good balance between speed of calculation and storage space required. We could, alternatively, go for a brute force approach and pre-calculate all of the vectors required to get from each frame to every other frame, but the storage requirements for such a method would be prohibitive.

This tool has been used for many TV projects, including Hyde and Agent Carter, as well as films such as Spectre and the upcoming blockbuster Geostorm.

References

SEYMOUR, M., 2006. Art of optical flow. https://www.fxguide.com/featured/art_of_optical_flow/, 28 Feb.