Comfort-driven Disparity Adjustment for Stereoscopic Video

Miao Wang[1] (✉), Xi-Jin Zhang[1], Jun-Bang Liang[1], Song-Hai Zhang[1], and Ralph R. Martin[2]

Abstract Pixel disparity—the offset of corresponding pixels between left and right views—is a crucial parameter in stereoscopic three-dimensional (S3D) video, as it determines the depth perceived by the human visual system. Unsuitable pixel disparity distribution throughout an S3D video may lead to visual discomfort. We present a unified and extensible stereoscopic video disparity adjustment framework which improves the viewing experience for an S3D video: keep the perceived 3D appearance as unchanged as possible while minimizing discomfort. We first analyse disparity and motion attributes of S3D video in general, then derive a wide-ranging visual *discomfort metric* from existing perceptual comfort models. An objective function based on this metric is used as the basis of a hierarchical optimisation method to find a disparity mapping function for each input video frame. Warping-based disparity manipulation is then applied to the input video to generate the output video, using the desired disparity mappings as constraints. Our comfort metric takes into account *disparity range*, *motion* and *stereoscopic window violation*; the framework could easily be extended to introduce further visual comfort models. We demonstrate the power of our approach using both animated cartoon and real S3D videos.

Keywords Stereoscopic Video Editing, Video Enhancement, Perceptual Visual Computing, Video Manipulation.

1 Department of Computer Science and Technology, Tsinghua University, Beijing, 100084, P.R.China. E-mail: wangmiao11@mails.tsinghua.edu.cn, zhangxijin91@gmail.com, williamm2006@126.com, shz@tsinghua.edu.cn

2 Cardiff University, Cardiff, CF243AA, United Kingdom. E-mail: ralph@cs.cf.ac.uk.

1 Introduction

With the recent worldwide increase in stereoscopic display hardware, there has been great interest in both academia and industry in stereoscopic three-dimensional (S3D) movie production, for instance, glasses-free multi-view display technology [22, 38] and perceptual disparity models [5, 6]. Viewing the 3D world through a display screen differs from natural viewing—it introduces vergence-accommodation conflicts [10, 11]. As a result, poor scene design in S3D movies can lead to visual fatigue. In addition to vergence-accommodation conflict, other factors such as motion and luminance also affect the human visual system (HVS), and may make the viewer feel uncomfortable. Most of these factors have a close relationship to *binocular disparity*—the difference in an object's location on the left and right retinas [30]. The brain uses binocular disparity to extract depth information via a process of stereopsis.

The goal of making a movie stereoscopic is to add realism by providing a feeling of depth, but care must be taken to avoid visual discomfort. It is a tedious process to accordingly tune the perceptual depth of S3D videos during shooting, even for professionals with years of experience [26]. Existing S3D video post-processing technology [14, 19] helps to manipulate original disparity of S3D images and videos. Given the desired disparity mapping, these methods manipulate the original disparity to meet the requirements. Unfortunately, such approaches require manually input disparity targets or manipulation operators for guidance. A general, content-driven, solution for ensuring the comfort of S3D video is still lacking.

In this paper, we provide an automatic solution to the disparity tuning problem using a unified and extensible comfort-driven framework. Unlike previous works that focus on user-guided S3D video disparity retargeting [14, 19], we automatically manipulate the
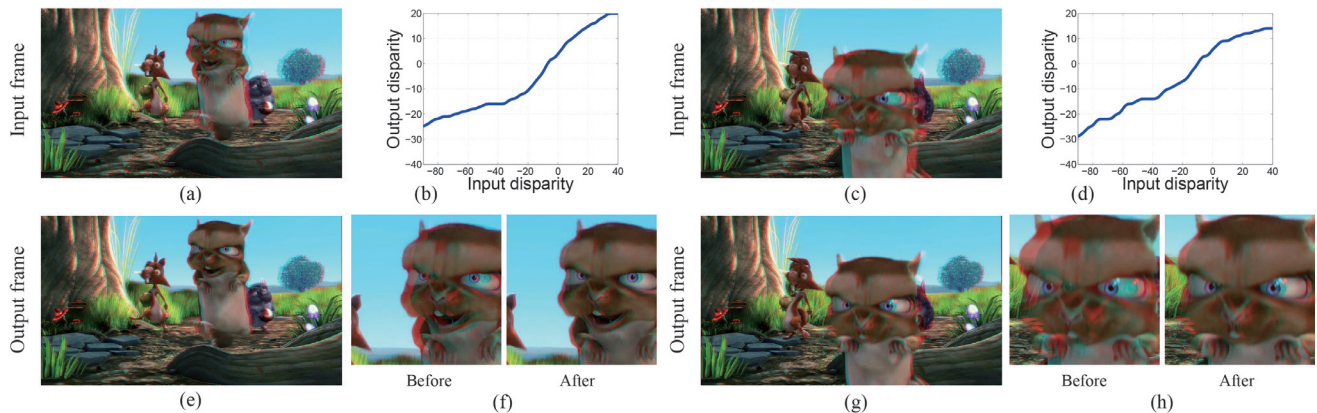
Fig. 1   Inputs and outputs: given an input stereoscopic 3D video (sample frames (a) and (c)), our framework automatically determines a comfort-driven disparity mapping (b) and (d) for every frame. Output video frames (e) and (g) are produced by applying these mappings to the input video frames, improving visual comfort. (f) and (h) show close-ups of frames before and after manipulation (©Blender Foundation).

disparity of an original S3D video, to improve visual comfort while maintaining satisfactory parts of the original content whenever possible. The challenge of this problem is to build a bridge between S3D visual comfort and the automatic manipulation of video content. By taking advantage of existing S3D visual comfort models, we derive a general *discomfort* metric which we use to evaluate and predict the discomfort level. We build on this metric to define an objective function for use in optimising disparity mapping functions. Our metric may be further extended if needed, to introduce further visual comfort models. We optimise the mappings over the whole video, using a hierarchical solution based on a genetic algorithm. The output video is generated by applying the disparity mappings to the original video using a warping-based technology. To our knowledge, our framework is the first system which can automatically improve visual comfort by means of comfort-driven disparity adjustment.

The major contributions of our work are thus:

- A unified S3D video post-processing framework that automatically reduces visual discomfort by disparity adjustment.
- A discomfort metric that combines several key visual comfort models; it could easily be extended to incorporate others too if desired. It provides a basis for an objective function used to optimise disparity.
- A hierarchical optimisation method for computing a disparity mapping for each video frame.

2   Related work

Causes of visual discomfort experienced when watching S3D movies have been investigated, with a view to improving such movies. Mendiburu [26] qualitatively determined various factors such as excessive depth and discontinuous depth changes that contribute to visual fatigue. Liu et al. [20] summarized several principles, and applied them to photo slideshows and video stabilization.

Various *mathematical models* have also been proposed to quantitatively evaluate discomfort experienced by the HVS. Besides viewing configurations such as viewing distance [31], time [4] and display screen type, effects particular to stereoscopic content have also been widely investigated [5–7, 13, 17, 27, 34], which we now consider in turn.

Vergence-accommodation conflict is widely accepted to be a key factor in visual discomfort. These ideas may be used to quantitatively determine a *comfort zone* within which little discomfort arises [31]. Stereoscopic fusion disparity range is modeled in [12], based on viewing distance and display sizes. Didyk et al. [5] model perceptual disparity based on experiments with sinusoidal stimuli; the ideas can be used to produce backward-compatible stereo and personalized stereo. This work was later extended to incorporate the influence of luminance contrast [6]. Our metric includes a disparity range term, based on the *comfort zone* model in [31]. It allows us to decide whether the disparity of a given point lies within the comfort zone.
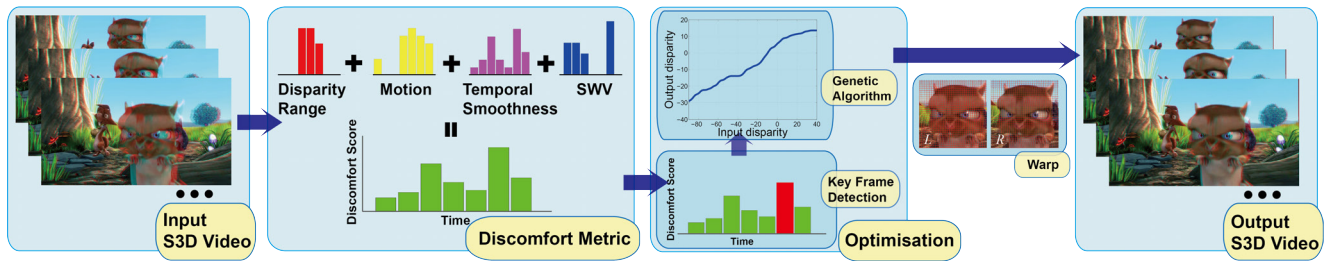
Fig. 2 Pipeline. The input of our system is a stereoscopic 3D video. Discomfort level of every frame is evaluated using the proposed metric. Discomfort intervals and key frames inside each interval are determined. A disparity mapping for every frame is optimised, based on the key frames, using a hierarchical optimisation method. Finally, the output video is generated by applying the mappings to the original video by warping.

Motion is another important factor in perceptual discomfort [7, 13, 37]. In [37], the contribution of the velocity of moving objects to visual discomfort is considered. Jung et al. [13] give a visual comfort metric based on salient object motion. Cho and Kang [4] conducted experiments with various combinations of disparity, viewing time and motion-in-depth, measuring the visual discomfort. Du et al. [7] proposed a comfort metric for motion which takes into account the interaction of motion components in multiple directions, and depths. Such literature shows that visual comfort is improved when objects move at lower velocities or lie closer to the screen plane. Movements perpendicular to the screen (along the z-axis) plays a more powerful role in comfort than movements parallel to the screen plane (the x-y plane).

Abrupt depth changes at scene discontinuities may also induce discomfort: for instance, a sudden jump from a shot focused in the distance to an extreme close-up can be disorienting. Disparity-response time models [27, 34] have been determined by a series of user-experience experiments. To reduce discomfort caused by depth changes, depths in shots should change smoothly.

Stereoscopic window violation describes a situation in which any object with negative disparity (in front of the screen) touches the left or right screen boundary. Part of the object may be seen by one eye but hidden from the other eye, leading to confusion by the viewer as to the object's actual position; this too causes fatigue [39]. Yet further factors are discussed in a recent survey [25]. As our approach provides a post-processing tool, we consider factors related to scene layout rather than camera parameters. These factors are disparity range, motion, stereoscopic window violation and depth continuity; they are meant to cover the major causes of discomfort, but our approach could be extended to include others too.

Use of post-processing technology has increased in recent years, helping amateurs to create S3D content and directors to improve S3D movie appearance. Lo et al. [21] show how to perform copy & paste for S3D, to create new stereoscopic photos from old ones; constraints must be carefully chosen. Later, Tong et al. [35] extend this work to allow pasting of 2D images into stereoscopic images. Kim et al. [16] provide a method to create S3D line drawings from 3D shapes. Niu et al. [28] give a warping-based method for stereoscopic image retargeting. Lang et al. [19] provide a disparity manipulation framework which applies desired disparity mapping operators to the original video using image warping. Kellnhofer et al. [14] optimise the depth trajectories of objects in an S3D video, providing smoother motion. Kim et al. [15] propose to compute multi-perspective stereoscopic images from a light field, meeting users' artistic control requirements. Masia et al. [24] propose a light field retargeting method that preserves perceptual depth on a variety of display types. Koppal et al. [18] provide an editor for interactively tuning camera and viewing parameters. Manually tuned parameters of cameras are applied to video; the results are then immediately fed back to the user. However, there is presently a gap between mathematical comfort models and post-processing applications—few technologies automatically work in a comfort-driven manner.

In a similar vein to our work, the OSCAM system[29] automatically optimises the camera convergence and interaxial separation to ensure that 3D scene contents are within the comfortable depth range. However this work is limited to process virtual scenes with known camera settings. Tseng et al. [36]

automatically optimise parameters of S3D cameras, taking into account the depth range and stereoscopic window violation. The major differences between their work and ours are, firstly, they optimise the camera separation and convergence, while our system automatically generates an output video with a better viewing experience. Secondly, their objective functions are derived from either a simple depth range or few general principles while ours rely on mathematical models. We build upon existing S3D post-processing approaches, especially warping-based ones, to build a bridge between comfort models and a practical tool.

3   Overview

In this section, we explain our notation, and then sketch our proposed framework.

We adapt the measure of *binocular disparity* from [7], expressed as *angular disparity*. Assuming the viewer focuses on the display screen with a vergence angle θ', the angular disparity at a 3D point P with a vergence angle θ is measured as the difference of vergence angles $\theta' - \theta$ (see Figure 3(a)). We also use the concept of *pixel disparity* in the metric and disparity mapping optimisation. The *pixel disparity* of a feature point f_L in the left view L is defined as an *integer* offset $f_R - f_L$ where f_R is the corresponding feature location in the right view R (see Figure 3(b)). Given these definitions, both the angular disparity and pixel disparity are negative when the 3D point P is in front of the screen and positive when it is behind the screen. A disparity mapping is a function $\phi(d)$ that given an input disparity value d, returns a new output disparity value d'. In this paper, ϕ is presented in discrete form: given a set of τ different disparity values $D_{\text{in}} = \{d_{\min}, \ldots, d_\tau\}$, and a corresponding set of output disparity values $D_{\text{out}} = \{d'_{\min}, \ldots, d'_\tau\}$, we regard $\phi : D_{in} \to D_{out}$ as a disparity mapping, where $d'_i = \phi(d_i)$.
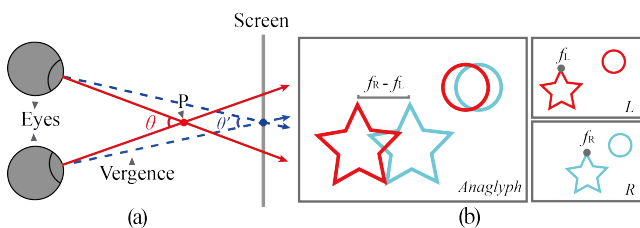
Fig. 3 Definitions: (a) angular disparity and (b) pixel disparity.

As explained in the Introduction, our comfort-driven disparity mapping framework automatically adjusts the disparity in an S3D video to improve visual comfort. Given an input S3D video to be optimised, we first

Fig. 4   Comfort zone. Left: anaglyph 3D images. Right: disparities beyond the comfort zone shown in blue.

evaluate the discomfort level of every frame, using the proposed metric, then determine intervals which cause discomfort and key frames inside each interval throughout the video (see Section 4). Next, based on the key frames, we optimise a disparity mapping ϕ for every frame using a hierarchical optimisation method (see Section 5), using an objective function derived from the discomfort metric. Finally, the mappings are applied to the original video by video warping. The pipeline is illustrated in Figure 2.

4   Discomfort Metric

An objective function measuring discomfort level is essential for automatic S3D video comfort optimisation. In this section, we present a general discomfort metric which is used to determine the objective function for disparity mapping optimisation. The metric takes into account *disparity range*, *motion*, *stereoscopic window violation* and *temporal smoothness*, all of which have been shown to have a major impact on the HVS. Each factor is formulated as a cost function. The temporal smoothness term relates pairs of successive frames (so is a binary term) while others are only dependent on one frame (so are unary terms). The wide-ranging nature of this metric enables us to evaluate the discomfort level in the round. The disparity range term measures the severity of vergence-accommodation conflict. The motion term evaluates discomfort brought about by eye movements. Retinal rivalry arises from inconsistent screen boundary occlusions, and is assessed by the stereoscopic window violation term. Flickering resulting from temporal inconsistency is evaluated by the temporal smoothness term. We now discuss each term individually and then explain how they are combined.

4.1   Individual Terms

Disparity Range. Excessive disparity leads to strong adverse reactions in the HVS due to vergence-accommodation conflict [10, 11]. To reduce the

resulting discomfort, one intuitive approach is to compress the disparity range, but severe compression makes an S3D video appear flat, and ultimately imperceptibly different from 2D. Instead, we evaluate how far each 3D point is from the screen plane and penalize points that fall outside the *comfort zone*. In [31], the *far* and *near* comfort zone boundaries B_{far} and B_{near} are introduced. In terms of *angular disparity*, these may be written:

$$B_{\mathrm{far}} = 2\tan^{-1}[\frac{d_a}{2m_{\mathrm{far}}d_f}(1 - T_{\mathrm{far}}d_f)] - 2\tan^{-1}[\frac{d_a}{2d_f}] \quad (1)$$

$$B_{\mathrm{near}} = 2\tan^{-1}[\frac{d_a}{2m_{\mathrm{near}}d_f}(1 - T_{\mathrm{near}}d_f)] - 2\tan^{-1}[\frac{d_a}{2d_f}],$$

where the constants in their model are $m_{\mathrm{far}} = 1.129$, $m_{\mathrm{near}} = 1.035$, $T_{\mathrm{far}} = 0.442$ and $T_{\mathrm{near}} = -0.626$. d_a is the angular disparity (in degrees) of a pixel and d_f is the viewing distance (in metres), which, in our viewing configuration, is set to 0.55 m.

In this formulation, the angular disparity $d_a(p)$ of a pixel p is within the comfort zone range is determined by:

$$\delta(p) = \begin{cases} 1 & \text{if} \quad B_{\mathrm{near}} < d_a(p) < B_{\mathrm{far}} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

The fraction of pixels in frame f whose disparity is outside the comfort zone is computed, and used to define the disparity range penalty term $E_d(f)$ for frame f:

$$E_d(f) = 1 - \frac{1}{N}\sum_{p \in f}\delta(p), \quad (3)$$

where N is the number of pixels in frame f. Figure 4 shows examples where disparities of certain pixels lie beyond the comfort zone.

Motion is an important source of visual discomfort [13, 37]. In [7], a novel visual comfort metric for S3D motion is proposed. This metric is a function of both the combination of velocity and depth, and luminance frequency. It returns a comfort value from 1 to 5 (the higher, the more comfortable). We adopt this model in our metric and assign to every video frame a motion *discomfort* value. We first assign a motion discomfort value $V_c(p) = \omega_n(5 - M_p(p))$ for every pixel p, where ω_n is a coefficient normalising $V_c(p)$ to $[0,1)$, set to 0.25. $M_p(p)$ is the pixel-wise motion comfort value calculated as in [7]:

$$M_p(p) = \sum_{k=0}^{n} C(v_{xy}, v_z, d, \frac{f_{l_0}}{2^k}) \times \frac{L_k(p)}{\sum_k L_k(p)},$$

where $C(v_{xy}, v_z, d, \frac{f_{l_0}}{2^k})$ is a model of motion comfort based on planar velocity v_{xy}, spatial velocity v_z, angular disparity d and luminance frequency $\frac{f_{l_0}}{2^k}$. $L_k(p)$

(a)   (b)
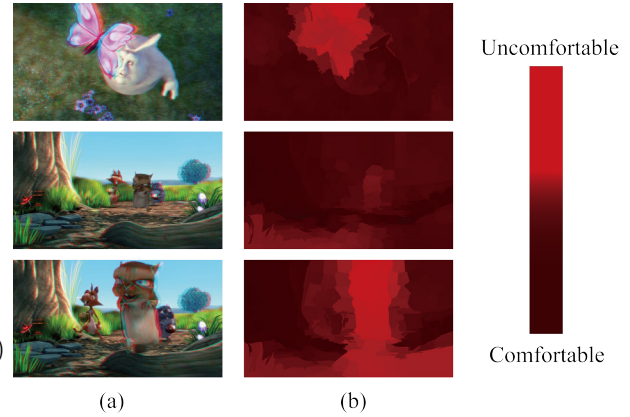
Fig. 5 Motion discomfort estimation: (a) anaglyph 3D frames (©Blender Foundation); (b) estimated discomfort caused by motion.

is the contrast value of the $(2^{k+1} + 1)$-neighborhood at p at the k-th Laplacian level of the Laplacian pyramid of the luminance; see [7] for further details.

After computing a discomfort value for every pixel, we determine the motion discomfort for the whole frame. In [7], average motion comfort values are calculated for individual saliency-based segments [3], assigning an importance value to every segment. The segments are obtained by graph-based segmentation [8]. They assume that the most uncomfortable region in a frame dictates the discomfort of the whole frame. However, we find that calculating the most salient and uncomfortable region in separate images without considering temporal coherence can lead to motion comfort instability. Instead, we modify their approach to perform SLIC superpixel segmentation [1], consider multiple discomfort-causing segments, and regard every segment as having the same importance. We extract an average motion comfort value for the top-K (K=20 by default) segment discomfort values as a motion penalty. The motion discomfort $E_m(f)$ for the whole frame f is:

$$E_m(f) = \frac{1}{K}\sum_{k=1}^{K} T(V_s(s_k)), \quad (4)$$

where $V_s(s) = \sum_{p \in s} V_c(p)/m$ is the average motion discomfort value for a segment s having m pixels. $T(\cdot)$ is the set of segment motion discomfort values, in descending order, as computed in [7]. Figure 5 shows example S3D frames with segment-wise discomfort maps according to motion.

Stereoscopic Window Violation occurs when an object is virtually located in front of the screen (i.e. have negative disparity) but is occluded by the screen boundary. This is confusing as a nearer object appears to be occluded by a further one, causing retinal
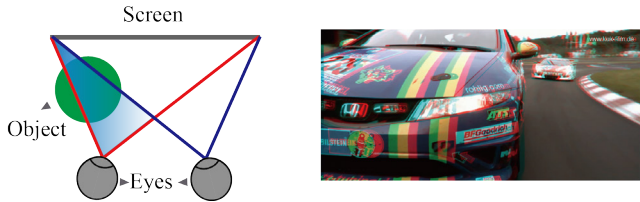
Fig. 6 Stereoscopic window violation (SWV). Left: a toy example illustrating SWV. Part of the object in green falling in the light blue region can only be seen by the left eye. Right: a real S3D photo showing SWV (©KUK Filmproduktion GmbH). There is inconsistent content in the leftmost part of the photo, leading to viewer discomfort.

Fig. 7 Typical frames (©Blender Foundation) and discomfort scores. (a) discomfort scores for frames in an S3D video clip. The discomfort interval is marked in blue. Key frames selected by our algorithm are highlighted in red. (b) shows three frames and corresponding discomfort scores from (a).

rivalry [26]. If this happens, only one eye can see part of the object, leading to visual inconsistency and hence discomfort (see Figure 6). One practical way to alleviate this is to trim off the offending part.

To measure stereoscopic window violation (SWV), we use a term $E_v(f)$ for frame f. We first detect violations by checking pixels near left and right boundaries: if pixels touching the frame boundaries have negative disparity, they violate the stereoscopic window. The SWV penalty for frame f is then defined by counting the number of pixels included in violating objects:
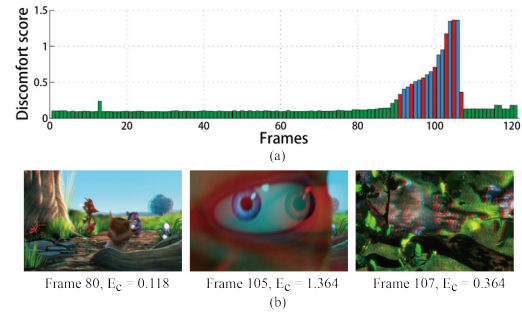
$$E_v(f) = \frac{1}{N} \sum_{s \in R_b} n(s) \tag{5}$$

where s stands for image segments extracted as before, and R_b is an approximation of violating objects in the form of segments; every segment in R_b has a negative average disparity. R_b is initially set as boundary segments with a negative average disparity, and is then iteratively augmented by adding new neighbouring segments with negative average disparity until no new segments with negative average disparity are found. $n(s)$ is the number of pixels in segment s and N is the number of pixels in frame f.

Temporal Smoothness. To avoid sudden depth changes, the disparity should vary smoothly and slightly, as needed. In [20], the importance of temporal smoothness is emphasised in 3D cinematography; they suggest that the disparity range of successive frames should vary smoothly. Following the definition of *disparity map similarity* in [20], we define the similarity of disparity between neighbouring frames f and f' using Jensen-Shannon divergence [23]:

$$E_s(f, f') = H\left(\frac{\Psi(f) + \Psi(f')}{2}\right) - \frac{H(\Psi(f)) + H(\Psi(f'))}{2}, \tag{6}$$

where $\Psi(f)$ is a pixel disparity histogram for frame f with $d_{max} - d_{min} + 1$ bins; d_{max} is the largest integer pixel disparity value in f and d_{min} is the smallest integer pixel disparity value in f. $H(X)$ is the Shannon entropy for distribution X. Intuitively, the more unlike the disparity histograms are, the higher the value of E_s.

4.2 Discomfort Metric

Our general discomfort metric for a set of successive frames $\widehat{F}$ in an S3D video is formulated as a linear combination of the above terms in Equations 3–6, summed over the frames:

$$E_c(\widehat{F}) = \sum_{f \in \widehat{F}} E_c(f), \tag{7}$$

$$E_c(f) = \lambda_d E_d(f) + \lambda_m E_m(f) + \lambda_v E_v(f) + \lambda_s E_s(f, f'),$$

where f' is the successor frame to f in $\widehat{F}$. λ_d, λ_m, λ_v and λ_s are weights balancing the penalty terms, set to 1, 0.4, 10, and 0.1 respectively. The weights are determined via experiments. We did a small scale perceptual study on 10 subjects with 10 input videos: we enumerated every weight from 0 to 20 with step 0.1, generating 1.6×10^9 possible combinations of weights. After calculated the corresponding general metrics for input videos based on each group of weights, we let 5 subjects view each video and evaluate its comfort level by assigning integer scores from 1 to 5. We finally selected the group of weights, under which the metric score best reflects subjects comfort feelings. The weights was further validated by the other 5 subjects' evaluation. This metric can be used to predict the overall discomfort level of part or all of an S3D video. An S3D video frame is predicted as visually comfortable if the discomfort value $E_c < 0.3$. Figure 7(b) shows exemplar frames with their corresponding discomfort values.

The metric has a general form, with a default set of weights balancing the penalty terms. If considered unimportant, certain terms can be ignored, by setting their corresponding weights to 0. Alternatively,

additional terms of a similar kind could also be included with proper weight configuration (as an example, we present a variation of the metric, driving perceptual depth enhancement, by adding another unary term to each frame in the video (see Section 6)). We intentionally do not include all factors that cause visual fatigue—there are many such factors. Instead, we claim that the above metric includes many of the most significant factors, and the way we have formulated it allows ready extension to include other comfort models using additional penalty terms. The ideas in the rest of the paper do not depend on the precise form of this metric, only that such a metric can be formulated. We next show how to use this metric to define the objective function used to optimise pixel disparity mapping.

5 Optimisation of Pixel Disparity Mapping

Based on the above visual discomfort metric, we next derive the objective function used for disparity mapping optimisation. A genetic algorithm is used in a hierarchical approach to optimise disparity mapping: given a set of input disparity values, we compute a corresponding target output disparity for each value.

5.1 Objective Function

The visual discomfort metric E_c measures the discomfort level of S3D video frames. However, directly using it as an objective function in an optimisation process leads to unsatisfactory results: clearly, mapping all disparity values to zero would minimise E_c, making it equal to zero at all times. Also, making large changes to the disparity without scaling the sizes of objects leads to a change in the perceived size of the original content. We thus add an additional unary term $E_n(\phi, f)$ to every frame f with the intent that optimisation should change the original disparities as little as possible. $E_n(\phi, f)$ measures differences between new and original disparities:

$$E_n(\phi, f) = \frac{1}{N} \sum_{d \in [d_{\min}, d_{\max}]} \Psi_d(f) |\phi(d, f) - d|, \quad (8)$$

where N is the number of pixels in frame f, d is the integer pixel disparity value and $\Psi_d(f)$ is the disparity histogram count for disparity d in frame f, as in Equation 6. $\phi(d, f)$ is disparity mapping for disparity d in frame f. This formulation gives a cost for the mapping ϕ, punishing large changes from the original disparity distribution. This additional term allows us to find a suitable disparity mapping for each video frame that improves visual comfort while also preserving the original appearance.

We denote the objective function for optimising a sequence of mappings $\widehat{\Phi}$ of a sequence of frames $\widehat{F}$ in a S3D video as $E(\widehat{\Phi})$; it is defined as:

$$\begin{aligned} E(\widehat{\Phi}) &= \sum_{f \in \widehat{F}} (\lambda_d E_d(\Gamma \phi_f(f)) + \lambda_m E_m(\Gamma \phi_f(f)) \\ &+ \lambda_v E_v(\Gamma \phi_f(f)) + \lambda_n E_n(\phi_f, f) \\ &+ \lambda_s E_s(\Gamma \phi_f(f), \Gamma \phi_{f'}(f'))), \quad (9) \end{aligned}$$

where f' is the successor frame to f, and $\Gamma_{\phi_f}(f)$ is a function which applies the mapping operator ϕ_f to frame f to produce a new frame with the desired new pixel disparities. λ_n is a further weight, set to 0.01 by default.

5.2 Hierarchical Optimisation

The objective function in Equation 9 is complex; we use an efficient hierarchical approach to optimise it in a coarse-to-fine manner along the time-line. We observed that in S3D movies, frames causing discomfort usually appear together, forming discomfort intervals. Thus, we firstly extract *discomfort intervals* for the whole video: we manipulate the disparity only for frames which cause discomfort, and leave the others unchanged. The discomfort intervals are determined using Equation 7: a discomfort interval is a set of continuous frames from starting frame f_s to ending frame f_e, within which the discomfort metric $E_c(\{f, f'\})$ is above a threshold $\alpha = 0.3$, where f and f' are consecutive frames inside the interval. During optimisation, at coarser levels, inside every discomfort interval we determine *key frames* where the disparity changes drastically or there is a local maximum of discomfort. Frames at discomfort interval boundaries are also taken as key frames having a fixed identity pixel disparity map ($\phi(d) = d$). Next, we use a genetic algorithm to optimise pixel disparity mappings of the key frames, treating the key frames as neighbours. After optimising the key frames at this hierarchy level, we fix the disparity mappings of the current key frames, and continue to seek new key frames for finer intervals between any two successive key frames at the current level. The mappings of the current key frames are used as boundary conditions for the next level. This process is recursively performed until fewer than ten frames exist between each neighbouring pair of key frames. Finally, the disparity mapping for remaining frames between these key frames is interpolated. We now give further details of key steps.

5.2.1 Key Frame Determination

Key frame determination is a crucial step in the hierarchical disparity mapping optimisation. Since the
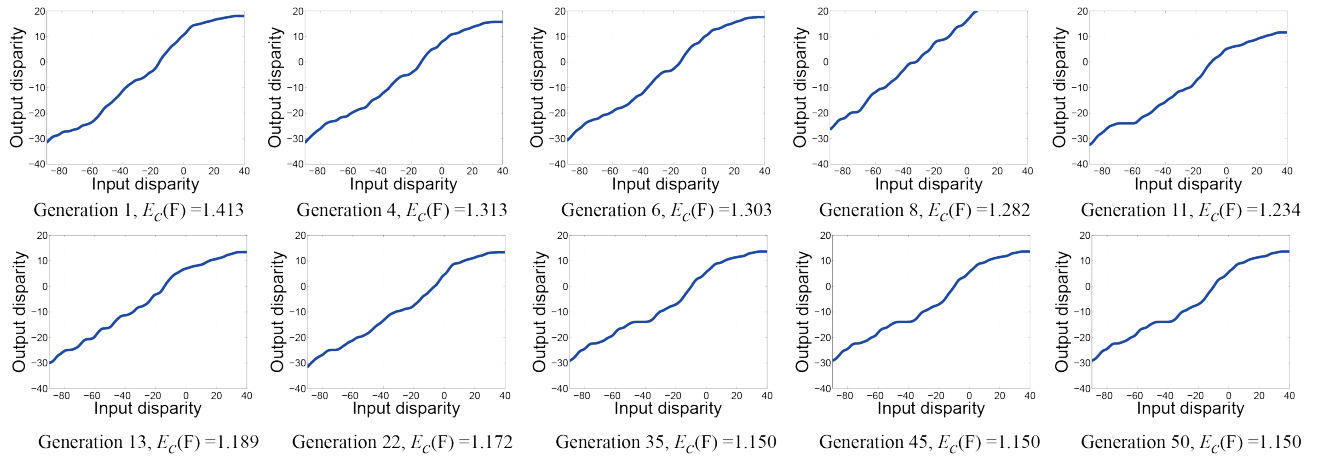
Fig. 8 Best disparity mapping solutions for improving comfort level of the frame shown in Figure 1(c), at various generations of the genetic algorithm. The frame set F contains four key frames. During optimisation, the discomfort cost E_c of F is reduced.

optimisation is performed in a coarse-to-fine manner, at coarser levels, key frames should provide a story line overview of frames at finer levels, especially in terms of disparity. Motivated by this requirement, inside each discomfort interval we mark a frame as a key frame when there is a sudden depth change or the discomfort metric reaches a local maximum within a window of Υ_l frames for each level l. By default, Υ_l at level l is set to a quarter of the interval length between the two boundary key frames. Specifically, we use the inequality $E_s(f, f') > \beta$ to determine whether frame f is a key frame at a drastic depth change; by default $\beta = 0.5$. After optimising key frames at level l, new key frames at level $l + 1$ are recursively determined, by seeking new key frames at level $l + 1$ between every adjacent pair of key frames at level l. We stop when fewer than ten frames exist between each neighbouring pair of key frames.

5.2.2 Heuristic Optimisation Using Genetic Algorithm

After finding key frame sets F at level l, we use a heuristic algorithm to optimise disparity mappings of these key frames. Without loss of generality, assume we are optimising a discomfort interval with t detected key frames $F = \{f_1, \ldots, f_t\}$. Including the additional key frames at the discomfort interval boundaries, the augmented key frame set becomes $\underline{F} = \{f_s, f_1, \ldots, f_t, f_e\}$, with fixed identity disparity mappings for f_s and f_e as boundary conditions. We

regard every successive pair of frames in $\underline{F}$ along the time-line as neighbours in a coarse view. We optimise the key frame mappings $\Phi = \{\phi_1, \ldots, \phi_t\}$ at coarser levels using the objective function adapted from Equation 9:

$$E(\Phi) = \tag{10}$$
$$\sum_{f \in F} (\lambda_d E_d(\Gamma\phi_f(f)) + \lambda_m E_m(\Gamma\phi_f(f)) + \lambda_v E_v(\Gamma\phi_f(f))$$
$$+ \lambda_n E_n(\phi_f, f) + \lambda_s E_s(\Gamma\phi_f(f), \Gamma\phi_{f'}(f'))),$$

where f' is the successor frame to f in $\underline{F}$. This objective function is used as *fitness assessment* in genetic algorithm.

A genetic algorithm is used to optimise the disparity mapping ϕ for each frame f using this objective function as a fitness function. We use the GALib implementation of a *steady-state* genetic algorithm [33]; 50% of the population is replaced on each generation. The genome for each individual is a vector of real numbers, which is used to store target disparity mapping values (with a conversion between integer and real numbers). Uniform crossover [32] is used with Gaussian mutation [9], which adds a random value from a Gaussian distribution to each element of an individual's state vector, to create offspring.

Genome of Individuals. The genome representation needs to be carefully designed; poor choice can lead to GA divergence. The target output disparity mapping values $D_{\phi_f} = \{\phi(d_{\min}), \ldots, \phi(d_{\max})\}$ of the mapping function ϕ for

a frame f is an *elementary unit* in each individual's genome. The disparity mapping $\phi(x)$ should be a non-decreasing function, i.e. if $x_1 < x_2$, then $\phi(x_1) \le \phi(x_2)$, to avoid depth reversal artifacts in the output video. We enforce this requirement by using an *increment*-based representation. We represent the mapping values $D_{\phi_f} = \{\phi(d_{\min}), \dots, \phi(d_{\max})\}$ as $\widetilde{D_{\phi_f}} = \{\phi(d_{\min}), \Delta_1, \dots, \Delta_{p-1}\}$ where d ranges over all integer pixel disparity values between $d_{\min}$ and $d_{\max}$, and $\Delta_i = \phi(d_{i+1}) - \phi(d_i)$ is a non-negative mapping value increment). Obviously, we can recover D_{ϕ_f} from the relationship $D_{\phi_f}[i] = \widetilde{D_{\phi_f}}[i-1] + \widetilde{D_{\phi_f}}[i]$. The non-negativity of Δ_i is guaranteed by additional bound $\overline{b_i}$ and lower bound $\underline{b_i}$ on each integer element of $\widetilde{D_{\phi_f}}$:

$$\overline{b_i} = \begin{cases} d_{min} + (d_{\max} - d_{\min})/2 & \text{if } i = 1 \\ 5 & \text{otherwise,} \end{cases}$$

and

$$\underline{b_i} = \begin{cases} d_{min} - (d_{\max} - d_{\min})/2 & \text{if } i = 1 \\ 0 & \text{otherwise.} \end{cases}$$

These upper upper and lower bounds also prevent the mappings from making over-large increments Δ_i. This constraint is supported by the *steady-state* genetic algorithm. The full genome of each individual is a vector of integers which concatenates the mapping values $\widetilde{D_\Phi} = \{\widetilde{D_{\phi_1}}, \dots, \widetilde{D_{\phi_t}}\}$ for the t key frames in F.

Evolution. The state of every individual in the first generation is initialised using random mappings. The objective function in Equation 10 is used for individual fitness assessment. The *uniform crossover* probability is $p_c = 0.7$ and the *Gaussian mutation* probability is $p_m = 0.05$. The population size n_p is set to 100 and the GA terminates after $n_g = 50$ generations. As a steady-state GA is used, the last generation includes the best solution found for the desired mappings Φ'. Figure 8 illustrates the mappings corresponding to the best solution in different generations.

5.3 Warping-based Manipulation

After optimising pixel disparity mappings for each frame of the video, we have to adjust the input video using these mappings. In [19], a warping-based method is given to adjust disparity to match desired disparity mappings. Their approach first extracts sparse stereo correspondences, followed by warping of left and right frames respectively with constraints applied to the vertices of a mesh grid placed over each frame. The output is thus a deformed mesh as well as the warped frame. We use this technology to generate the output
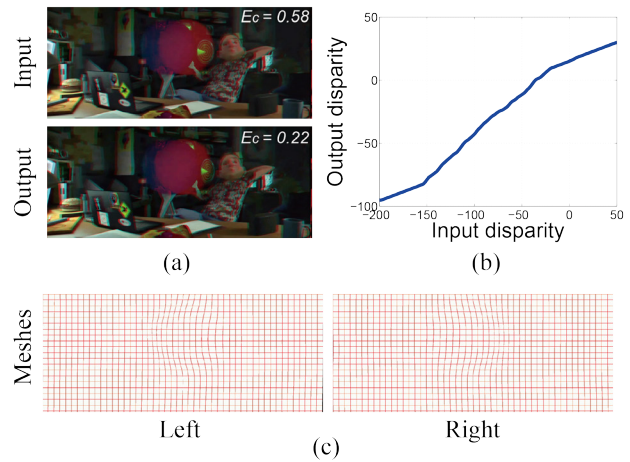
Fig. 10 Disparity Mapping. (a) an input S3D frame and corresponding output frame. The 'ball' outside the comfort zone is pushed back towards the screen. (b) disparity mapping generated by our algorithm. (c) deformed meshes for left and right view, indicating the warping effect.

video.

6 Results

The experiments were carried out on a computer with an Intel Core i7-4790K CPU with 32GB RAM. All videos were uniformly scaled to fit the screen size (1920×1080 pixels) to the extent possible before calculation. We calculate dense pixel correspondence between the left view and right view to estimate the pixel disparity in S3D videos using optical flow [2]. Motion in the x-y plane is also estimated using this method, between consecutive frames in left view. Calculating the discomfort metric for one S3D video frame of size 1920×1080 takes less than 0.2 second. The most time-consuming part is hierarchical optimisation, but the time taken is variable. It is dominated by the key frame determination step; it takes up to 15 minutes to optimise ten key frames together in our implementation, using a single core.

We have tested our approach on S3D video clips whose lengths are less than one minute. As explained in [7], the proposed motion comfort metric was derived from experiments on short videos. All of the results were obtained using default parameters. With a lot of experiments, we found that our system is insensitive to parameters.

Our method provides smooth scene transitions between successive shots. Representative frames of a
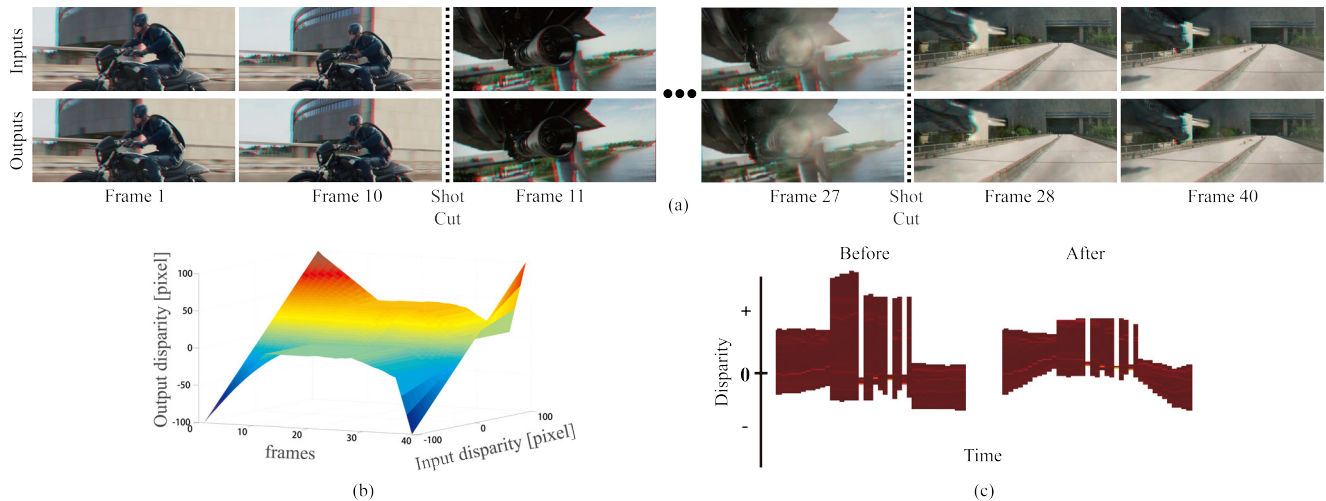
Fig. 9 Representative anaglyph frames of our results, with a fluent depth storyboard. (a) sample input and output frames (frame 1 and frame 40 are fixed to their original disparities). (b) pixel disparity mappings along the time-line (colour encodes output disparity value). (c) depth storyboard before and after the manipulation, with colour encoding frequency of the occurrence of disparity values.

video clip with shot cuts are shown in Figure 9(a). Boundary frames 1 and 40 do not cause discomfort, so are fixed to retain their original disparities. Our algorithm detects drastic disparity changes between these boundary frames and automatically adjusts disparities to provide smoother depth transitions by finding suitable disparity mappings (see Figure 9(b)). In this example, frames where shot cuts occur are detected as key frames. This is because the values of motion term and temporal smoothness term reach a local maximum with in a window. As can be seen in Figure 9(c), after manipulating the video, the depth storyboard suffers less from sudden jumps in disparity. While the last part of the video initially has a constant disparity range, which after processing becomes a slowly increasing disparity range, this does not lead to any perceptual artifacts: (i) slow transitions in disparity are often used to control disparity at shot cuts, (ii) the rate of disparity change is small, and (iii) the warping provides a smooth solution.

Figure 10 gives an example of automatic correction of excessive disparity range. The ball popping out towards the viewer in the center of the frame makes it difficult for the viewer to comfortably perceive the depth. Our correction pushes the ball a little closer to the screen. Pushing the ball back into the screen too far would change the content too much, in disagreement with the film maker's intent. The deformed meshes of the left and right views used for the warping-based disparity manipulation are also shown. Discomfort scores in our metric before and after manipulation are 0.58 and 0.22 respectively.
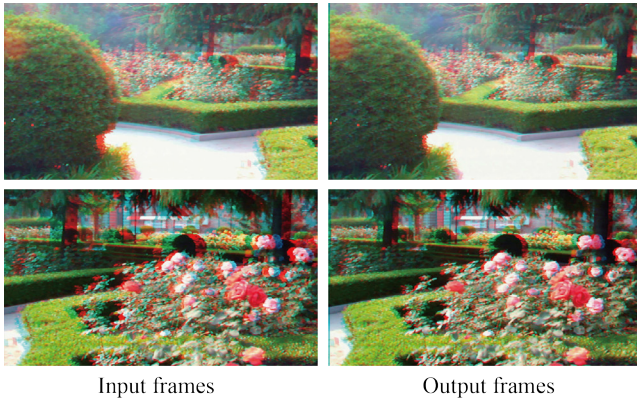

Input frame   Output frame

Fig. 11 Eliminating stereoscopic window violation. Left: input frames with SWV (©KUK Filmproduktion GmbH). Right: in the manipulation result, the popped out parts are pushed back towards the screen.

Figure 11 gives an example of eliminating stereoscopic window violation. In the original input frame, the front of the car appears in front of the screen plane, but is occluded by the picture boundary. This causes the leftmost part of the image to be seen only by the left eye. Such inconsistent content gives an unpleasant viewing experience. Our approach detects such violation automatically detected and eliminates it by pushing the car behind the screen.

We further tested our framework using videos from a consumer stereoscopic camera (Fuji FinePix REAL 3D W3). Typical frames from one video are shown in Figure 12. The perceptual depth range is excessive, making it hard to fuse the scene. In the result, the depth of the building in the far distance is reduced, while the disparity values of the flowers are subtly changed.

Perceptual Depth Enhancement. In Sections 4 and 5, we presented an extensible framework that optimises disparity mappings driven by comfort. A variation of this framework can be used to drive disparity manipulation to provide depth enhancement

Input frames	Output frames

Fig. 12 Processing a video taken by a consumer S3D camera. The depth of the scene is reduced to facilitate stereoscopic fusion.
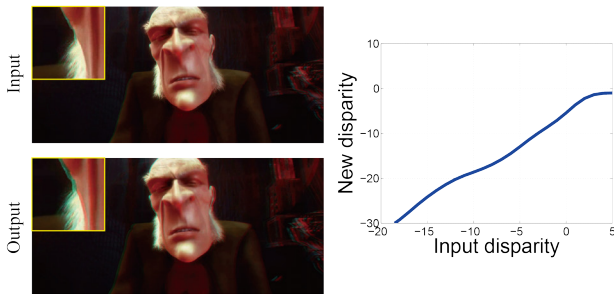
Fig. 13 Enhancing perceptual depth. Left: input and output frames (©Blender Foundation). After enhancement, the head of the man looks more angular. Right: the generated disparity mapping.

(while not introducing visual discomfort), for a greater feeling of depth. This goal can be accomplished by introducing an additional unary term $E_a(\phi, f)$ for each frame with weight $\lambda_a = 1$ to the objective function, with the aim of punishing small disparities after applying the mapping ϕ to the video:

$$E_a(\phi, f) = exp(-\frac{1}{2N} \sum_{d \in [d_{\min}, d_{\max}]} \Psi_d(f) |\phi(d, f)|),$$

(11)

where N is the number of pixels in frame f, d is the integer pixel disparity value and $\Psi_d(f)$ is the disparity histogram count for disparity d in frame f, as in Equation 8. This change allows the perceived depths to be amplified, shown in Figure 13.

6.1 User Study

We conducted two user studies with 20 subjects aged from 18 to 32, to further assess the performance of our proposed comfort-driven disparity adjustment method. The primary aims for the two user studies were to test whether the framework can produce artifact-free results, and its ability to improve visual comfort. Subjects participated by watching S3D videos and filling in questionnaires.

We used a 23-inch interleaved 3D display (1920×1080 pixels, 400 cd/m^2 brightness), with passive polarized glasses. The viewing distance was set to 55 cm, as assumed in the proposed metric. All subjects had normal or corrected-to-normal vision, and were assessed to ensure they had no difficulty in stereoscopic fusion. Videos were displayed at full screen size.

We prepared ten pairs of S3D videos including animated cartoons and real S3D videos. Both videos in a pair had the same content, one being the original and the other being modified by our system. A random order was used for each pair, and displayed three times in succession. Subjects were allowed to pause and carefully examine the content at any time.

In the first user study, we evaluated whether our output video provides greater visual comfort than the original. After watching each video, each subject was asked to rate the comfort level of their viewing experience, in terms of ease of fusing the scene, causing fewer or less severe headaches, and other feelings of discomfort. Five ratings were used, from 1 to 5: *very uncomfortable, uncomfortable, mildly comfortable, comfortable, very comfortable*. In *all* ten pairs of test videos, our results achieved on average a higher comfort score than the original video. The differences in average score in each pair varied from 0.3 to 1.35. A right-tailed paired-sample hypothesis test was conducted, with the null hypothesis H_0: *there was no significant difference between the comfort scores of our outputs and the original videos* and alternate hypothesis H_A: *the comfort scores of our results were significantly higher than those for the original videos* at significance level $\alpha = 0.05$ with $n = 200$ samples. The one-tailed critical value was $t = 1.653$, while the test statistic was $t^* = 9.905$. Since $t^* > t$, the null hypothesis was rejected, indicating that the differences were statistically significant: our approach provides an improved stereoscopic video viewing experience.

The second user study aimed to assess artifacts in our results. Before undertaking the user study, the subject was told to note any disturbing perceptual depth artifacts (e.g. depth reversals or unsuitable depth changes) that caused confusion. After watching each video, the subject was asked to rate both videos for unsuitable perceived depths, which were scored as

follows: $4 = $ *many strong artifacts*, $2 = $ *few strong / many weak artifacts*, $1 = $ *few weak artifacts*, and $0 = $ *no artifacts*. The results showed that 8 out of 20 subjects did not notice artifacts in any video, 2 subjects only saw artifacts in our results and 2 subjects only saw artifacts in the original videos. The other 8 subjects noticed artifacts in both our results and the original videos. The worst score for both our results and the original videos was 2 (few strong / many weak artifacts). To further test whether statistically the two sets of scores have no difference, a two-tailed paired-sample hypothesis test was conducted, with the null hypothesis H_0: *there was no significant difference between the artifact scores of our outputs and the original videos* and alternate hypothesis H_A: *artifact scores of our results and the original videos differ* at significance level $\alpha = 0.05$ with $n = 200$ samples. The two-tailed critical value was $t = 1.972$, while the test statistic was $t^* = 1.236$. This time, the null hypothesis was not rejected, as $|t^*| \leq |t|$. We conclude that there is no significant difference in the perceived level of artifacts in the original videos and our results. Indeed, viewers are fairly insensitive to artifacts in these videos. Full statistics of the user studies are provided in the supplementary material.

Limitations. Our approach has limitations. As optimisation is based on a genetic algorithm, it may only find a local optimum. However, tests in which the genetic algorithm was initialized with differing initial populations led to quite similar output mappings. Secondly, existing individual comfort models work well only for viewers with normal stereoscopic fusion ability, and give an average comfort evaluation. Thus using the discomfort metric with default parameters may not give an accurate comfort evaluation for every individual, especially for those with poor stereoscopic fusion ability. Across individuals, there may well be differences in which aspects of an S3D video cause most discomfort. Moreover, our system cannot predict directors' intention; intentional discomfort shots for artistic visual impact would unfortunately be eliminated by our system.

7 Conclusion

We have suggested a general framework for automatic comfort-driven disparity adjustment, together with a S3D discomfort metric. The metric combines several key factors, and could be of general benefit to S3D movie makers by giving an objective visual comfort evaluation in the round. It underpins our automatic disparity adjustment approach, which is based on disparity mapping optimisation. Our results demonstrate the effectiveness and uses of our approach.

Our work is among the first attempts to tackle this challenging problem, and leaves room for improvement. In our framework, the disparity mapping is automatically determined using a heuristic method, and a closed-form solution for this problem is desirable.

References

[1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Trans. Pat. Anal. Mach. Intell.*, 34(11):2274–2282, 2012.

[2] T. Brox and J. Malik. Large displacement optical flow: descriptor matching in variational motion estimation. *IEEE Trans. Pat. Anal. Mach. Intell.*, 33(3):500–513, 2011.

[3] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. S. Torr, and S.-M. Hu. Global contrast based salient region detection. *IEEE Trans. Pat. Anal. Mach. Intell.*, 37(3):569–582, 2015.

[4] S.-H. Cho and H.-B. Kang. Subjective evaluation of visual discomfort caused from stereoscopic 3d video using perceptual importance map. In *TENCON 2012-2012 IEEE Region 10 Conference*, pages 1–6. IEEE, 2012.

[5] P. Didyk, T. Ritschel, E. Eisemann, K. Myszkowski, and H.-P. Seidel. A perceptual model for disparity. *ACM Trans. Graph.*, 30(4):96:1–96:10, 2011.

[6] P. Didyk, T. Ritschel, E. Eisemann, K. Myszkowski, H.-P. Seidel, and W. Matusik.

A luminance-contrast-aware disparity model and applications. *ACM Trans. Graph.*, 31(6):184:1–184:10, 2012.

[7] S.-P. Du, B. Masia, S.-M. Hu, and D. Gutierrez. A metric of visual comfort for stereoscopic motion. *ACM Trans. Graph.*, 32(6):222:1–222:9, 2013.

[8] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *International Journal of Computer Vision*, 59(2):167–181, 2004.

[9] N. Higashi and H. Iba. Particle swarm optimization with gaussian mutation. In *Proceedings of Swarm Intelligence Symposium*, pages 72–79. IEEE, 2003.

[10] D. M. Hoffman, A. R. Girshick, K. Akeley, and M. S. Banks. Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of vision*, 8(3):33, 2008.

[11] I. P. Howard. *Perceiving in depth, 2: Stereoscopic Vision*. Oxford University Press, 2012.

[12] E. W. Jin, M. E. Miller, S. Endrikhovski, and C. D. Cerosaletti. Creating a comfortable stereoscopic viewing experience: effects of viewing distance and field of view on fusional range. In *Electronic Imaging 2005*, pages 10–21. International Society for Optics and Photonics, 2005.

[13] Y. J. Jung, S.-i. Lee, H. Sohn, H. W. Park, and Y. M. Ro. Visual comfort assessment metric based on salient object motion information in stereoscopic video. *Journal of Electronic Imaging*, 21(1):011008–1, 2012.

[14] P. Kellnhofer, T. Ritschel, K. Myszkowski, and H.-P. Seidel. Optimizing disparity for motion in depth. In *Computer Graphics Forum*, volume 32, pages 143–152. Wiley Online Library, 2013.

[15] C. Kim, A. Hornung, S. Heinzle, W. Matusik, and M. Gross. Multi-perspective stereoscopy from light fields. *ACM Trans. Graph.*, 30(6):190:1–190:10, December 2011.

[16] Y. Kim, Y. Lee, H. Kang, and S. Lee. Stereoscopic 3d line drawing. *ACM Trans. Graph.*, 32(4):57:1–57:13, July 2013.

[17] F. L. Kooi and A. Toet. Visual comfort of binocular and 3d displays. *Displays*, 25(2):99–108, 2004.

[18] S. Koppal, C. Zitnick, M. Cohen, S. B. Kang, B. Ressler, and A. Colburn. A viewer-centric editor for 3d movies. *Computer Graphics and Applications, IEEE*, 31(1):20–35, Jan 2011.

[19] M. Lang, A. Hornung, O. Wang, S. Poulakos, A. Smolic, and M. Gross. Nonlinear disparity mapping for stereoscopic 3d. *ACM Trans. Graph.*, 29(3):75:1–75:10, 2010.

[20] C.-W. Liu, T.-H. Huang, M.-H. Chang, K.-Y. Lee, C.-K. Liang, and Y.-Y. Chuang. 3d cinematography principles and their applications to stereoscopic media processing. In *Proceedings of the ACM Multimedia*, pages 253–262. ACM, 2011.

[21] W.-Y. Lo, J. van Baar, C. Knaus, M. Zwicker, and M. H. Gross. Stereoscopic 3d copy & paste. *ACM Trans. Graph.*, 29(6):147:1–147:10, 2010.

[22] A. Maimone, G. Wetzstein, D. Lanman, M. Hirsch, R. Raskar, and H. Fuchs. Focus 3D: Compressive Accommodation Display. *ACM Trans. Graph.*, 32(5):153:1–153:13, 2013.

[23] C. D. Manning and H. Schütze. *Foundations of statistical natural language processing*. MIT press, 1999.

[24] B. Masia, G. Wetzstein, C. Aliaga, R. Raskar, and D. Gutierrez. Display adaptive 3d content remapping. *Computers & Graphics*, 37(8):983–996, 2013.

[25] B. Masia, G. Wetzstein, P. Didyk, and D. Gutierrez. A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Computers & Graphics*, 37(8):1012–1038, 2013.

[26] B. Mendiburu. *3D movie making: stereoscopic digital cinema from script to screen*. Taylor & Francis, 2009.

[27] T.-J. Mu, J.-J. Sun, R. R. Martin, and S.-M. Hu. A response time model for abrupt changes in binocular disparity. *The Visual Computer*, pages 1–13, 2014.

[28] Y. Niu, W.-C. Feng, and F. Liu. Enabling warping on stereoscopic images. *ACM Trans. Graph.*, 31(6):183:1–183:7, 2012.

[29] T. Oskam, A. Hornung, H. Bowles, K. Mitchell, and M. Gross. OSCAM - Optimized stereoscopic camera control for interactive 3D. *ACM Trans. Graph.*, 30(6):189:1–189:8, 2011.

[30] S. E. Palmer. *Vision science: Photons to phenomenology*, volume 1. MIT press

Cambridge, MA, 1999.

[31] T. Shibata, J. Kim, D. M. Hoffman, and M. S. Banks. The zone of comfort: Predicting visual discomfort with stereo displays. *Journal of vision*, 11(8):11, 2011.

[32] G. Syswerda. Uniform crossover in genetic algorithms. 1989.

[33] G. Syswerda. A study of reproduction in generational and steady state genetic algorithms. *Foundations of genetic algorithms*, 2:94–101, 1991.

[34] K. Templin, P. Didyk, K. Myszkowski, M. M. Hefeeda, H.-P. Seidel, and W. Matusik. Modeling and optimizing eye vergence response to stereoscopic cuts. *ACM Trans. Graph.*, 33(4):145:1–145:8, 2014.

[35] R.-f. Tong, Y. Zhang, and K.-L. Cheng. Stereopasting: interactive composition in stereoscopic images. *IEEE Trans. Vis. Comp. Graph.*, 19(8):1375–1385, 2013.

[36] K.-L. Tseng, W.-J. Huang, A.-C. Luo, W.-H. Huang, Y.-C. Yeh, and W.-C. Chen. Automatically optimizing stereo camera system based on 3d cinematography principles. In *3DTV-Conference: The True Vision-Capture, Transmission and Display of 3D Video (3DTV-CON), 2012*, pages 1–4. IEEE, 2012.

[37] K. Ukai and P. A. Howarth. Visual fatigue caused by viewing stereoscopic motion images: Background, theories, and observations. *Displays*, 29(2):106–116, 2008.

[38] G. Wetzstein, D. Lanman, W. Heidrich, and R. Raskar. Layered 3d: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Trans. Graph.*, 30(4):95:1–95:12, July 2011.

[39] F. Zilly, M. Müller, P. Eisert, and P. Kauff. The stereoscopic analyzer-an image-based assistance tool for stereo shooting and 3d production. In *ICIP*, pages 4029–4032, 2010.

Miao Wang is currently a PhD candidate at Tsinghua University, Beijing. He received his B.S. degree from Xidian University in 2011. His research interests include computer graphics, image processing and computer vision.

Xi-Jin Zhang Xi-Jin Zhang is currently a Ph.D. student at Tsinghua University, Beijing. He received his BS degree from Xidian University in 2014. His research interests include image and video processing, computer vision and machine learning

Jun-Bang Liang is currently an undergraduate student in Tsinghua University. His research interests include Computer Vision and Computer Graphics.

Song-Hai Zhang Song-Hai Zhang received the PhD degree in 2007 from Tsinghua University. He is currently an associate professor in Department of Computer Science and Technology of Tsinghua University, China. His research interests include image and video processing, geometric computing.

Ralph R. Martin is currently a Professor at Cardiff University. He obtained his PhD degree in 1983 from Cambridge University. He has published more than 200 papers and 12 books, covering such topics as solid and surface modeling, intelligent sketch input, geometric reasoning, reverse engineering, and various aspects of computer graphics. He is a Fellow of: the Learned Society of Wales, the Institute of Mathematics and its Applications, and the British Computer Society. He is on the editorial boards of Computer Aided Design, Computer Aided Geometric Design, Geometric Models, the International Journal of Shape Modeling, CAD and Applications, and the International Journal of CAD/CAM.