

Making 360° Video Watchable in 2D: Learning Videography for Click Free Viewing

Yu-Chuan Su Kristen Grauman
The University of Texas at Austin

Abstract

360° video requires human viewers to actively control “where” to look while watching the video. Although it provides a more immersive experience of the visual content, it also introduces additional burden for viewers; awkward interfaces to navigate the video lead to suboptimal viewing experiences. Virtual cinematography is an appealing direction to remedy these problems, but conventional methods are limited to virtual environments or rely on hand-crafted heuristics. We propose a new algorithm for virtual cinematography that automatically controls a virtual camera within a 360° video. Compared to the state of the art, our algorithm allows more general camera control, avoids redundant outputs, and extracts its output videos substantially more efficiently. Experimental results on over 7 hours of real “in the wild” video show that our generalized camera control is crucial for viewing 360° video, while the proposed efficient algorithm is essential for making the generalized control computationally tractable.

1. Introduction

360° cameras are becoming very popular, thanks to emerging Virtual Reality (VR) technologies and applications. More than a dozen new 360° cameras were released in 2016 [1], and the market is expected to grow by more than 100% per year in the next few years [2]. Watching panoramic photos and videos is becoming a common experience on content distribution sites like YouTube and Facebook, and many content creators are adopting the new medium. For example, BBC News distributes news in 360° videos online [3]. Whereas traditional cameras are restricted to a field of view (FOV) even narrower than human perception, a 360° camera captures the entire visual world from its optical center.

This broad capture offers new freedom to videographers and video consumers alike. A videographer no longer has to determine which direction to capture in the scene, freeing her to experience the moment rather than the act of recording a video. Meanwhile, a human video consumer has the freedom to explore the visual content based on her inter-




Figure 1: Our goal is to control the direction and field of view of a virtual camera within a 360° video in order to record a video that looks as if it were captured by a human videographer.

est, without being severely restricted by choices made by the videographer. For example, a news correspondent can traverse a war zone without consciously considering how to portray the scene, and subsequent viewers will still have an immersive experience about the tragedy and witness events in more detail than the videographer may even be able to attend to in the moment. Similarly, a parent at his child’s bustling birthday party can passively record 360° memories to be perused more deliberately later.

On the other hand, the medium also introduces new challenges. The most common interface for watching 360° videos is to display a small portion of the video as a 2D normal field of view (NFOV) video captured by a virtual camera¹. The video viewer now has to decide “where and what” to look at by controlling the direction of the virtual camera throughout the full duration of the video. The display can be a normal screen, a mobile device, or a VR headset, and the control signal will come from the mouse, the pose of the device, or the head movement respectively. These choices determine the content seen by the viewer and thus the user experience. Because the viewer has no information

¹For example, try clicking around on <https://www.youtube.com/watch?v=aTTzKwLPqFw>

about the content beyond the current FOV, it may be difficult to determine where to look in real time, e.g., a 360° video viewer can easily fail to notice that there is something approaching the camera from the opposite direction. In fact, the viewer may have to watch the video multiple times in order to find a proper way to control the virtual camera that navigates through the content of interest. While 360° videos may alternatively be displayed in their entirety using equirectangular projection, the unfamiliar format and distortion make such video hard to watch.

In light of this challenge, 360° video is a particularly appealing domain to invoke *automatic videography* techniques, which aim to convert unedited materials into an effective video presentation that conveys events [7, 9, 10, 18, 19, 31, 36, 37, 39]. While automatic videography in prior work has largely dealt with virtual environments and hand-crafted heuristics [7, 9, 18, 31], our recent work shows the potential of *learning* how to extract informative portions of 360° video as a presentable NFOV video [37]. In [37], we introduce the *Pano2Vid* problem, which takes a 360° video as input and as output generates NFOV videos that look like they could have been captured by a human observer equipped with a real NFOV camera. Our AUTOCAM algorithm learns videography tendencies directly from human-captured Web videos. By controlling the pose of a virtual NFOV camera within the input video, it removes the burden of deciding “where” to look when watching 360° videos.

We propose a new algorithm for the Pano2Vid problem that broadens the scope of camera control to generate more realistic videos. First, we generalize the task of Pano2Vid to allow not only spatial selections within the 360° video but also changes in the FOV. This allows the virtual camera control to fully mimic human videographer tendencies, because changing the FOV, i.e., zooming, is a common technique in both professional and amateur videography. Second, toward a more computationally efficient algorithm, we present a coarse-to-fine search approach that iteratively refines the camera control while reducing the search space in each iteration. The new approach makes the generalized task encompassing zoom computationally tractable. Finally, to account for the fact that valid Pano2Vid solutions are often multimodal, we explore how to generate a *diverse* set of plausible output NFOV videos, overcoming the redundancy that deters a straightforward optimization approach.

We experiment with more than 7 hours of real-world 360° video and capture 12 hours of manually edited 360° data in order to characterize results both quantitatively and qualitatively. We demonstrate that the proposed generalization of the Pano2Vid problem has dramatic effects. Compared to both the existing solution as well as strong baselines driven by saliency or center biases, our method’s auto-zoom improves results by up to 43.4%. In addition, we achieve a significant advantage in computational cost, reducing run-time by more than 84%.

2. Related Work

Video saliency Saliency studies visual content that attracts viewers’ attention [17, 20, 21, 30, 35, 43], where attention is usually measured by gaze fixations under free viewing settings. Although the research originated in static images, there is increasing work that studies video saliency [21, 35, 43]. Both video saliency and Pano2Vid try to predict spatial locations in videos. However, saliency targets locations that are eye-catching in 2D image coordinates whereas Pano2Vid predicts directions in spherical coordinates that videographers would try to capture with cameras. Also, saliency usually depends on local image content whereas Pano2Vid depends on the content and composition of the entire FOV.

Video retargeting Video retargeting adapts a source video by cropping and scaling to better fit the target display while minimizing the information loss [5, 22, 23, 26, 29, 34]. Both retargeting and our algorithm select portions of the original video to display to the user, but retargeting takes an already well-edited video as input and tries to generate a new version that conveys the same information. In contrast, our input 360° video is not pre-edited, and our goal is to generate multiple outputs that convey different information. Pano2Vid also entails a more severe reduction in spatial extent, e.g., compared to retargeting a 2D video from the Web to display nicely on a mobile device.

Virtual cinematography Most existing work on virtual cinematography studies virtual camera control in virtual (computer graphics) environments [7, 9, 18, 31] or else a specialized domain such as lecture videos [10, 36, 39]. Aside from camera control, some prior works also study automatic editing of raw materials like videos or photos [4, 12, 13, 19]. The goal is to generate an effective video presentation automatically to reduce human labor filming or editing. Existing approaches typically rely on heuristics that encode popular cinematographic rules. Our problem differs from the above in that we take unrestricted real 360° videos as input. Furthermore, we learn the cinematography tendencies directly from Web videos. Our work is most related to the AUTOCAM method [37], which also treats 360° video and takes a data-driven approach. In contrast to [37], we generalize the problem by allowing more degrees of freedom in the camera control (i.e. zoom) and considering diverse hypotheses. In addition, we propose a more efficient algorithm that makes the generalized problem computationally tractable.

Video summarization Video summarization aims to generate a concise representation for a video by removing temporal redundancy while preserving the important events [8, 14–16, 24, 27, 32, 33, 38, 42]. The goal is different from ours in the sense that video summarization selects content temporally whereas we select content spatially. Also, the outputs of our algorithm are continuous videos that look as if they were captured by a hand-held camera in the scene,

whereas the output of a video summarization algorithm is usually keyframes or concatenated disjoint video clips.

Diverse solution search Generating a diverse set of candidate solutions has been widely discussed in different domains [6, 11, 25, 28, 41]. A common approach is to find the solutions iteratively and encourage the solution of the current iteration to be different from that of previous iterations by penalizing the similarity between them. This approach has been applied to segmentation and pose estimation in a Markov Random Field [6], as well as machine translation in diverse beam search [11]. Our approach differs in that we formulate an exact solution search using dynamic programming instead of a probabilistic model or an approximate search. Also, our approach guarantees the minimum distance between solutions obtained in different iterations, whereas existing approaches rely on a predefined penalty term which fails to provide the same guarantee.

3. Preliminaries: AUTOCAM

First we provide background on our existing AUTOCAM solution [37]. AUTOCAM starts by sampling spatio-temporal glimpses (ST-glimpses) from the 360° video. An ST-glimpse is a five-second NFOV video clip with a 65.5° FOV² recorded from the 360° video by directing the camera to a fixed direction in the 360° camera axes. The algorithm samples candidate ST-glimpses at 18 azimuthal angles and 11 polar angles every five seconds:

$$\begin{aligned} \theta \in \Theta &= \{0, \pm 10, \pm 20, \pm 30, \pm 45, \pm 75\}, \\ \phi \in \Phi &= \{0, 20, \dots, 340\}, \\ t \in T &= \{0s, 5s, \dots, L - 5s\}, \end{aligned} \quad (1)$$

where L is the video length. Each candidate ST-glimpse is defined by the camera principal axis (θ, ϕ) direction and time in the video t :

$$\Omega_{t,\theta,\phi} \equiv (\theta_t, \phi_t) \in \Theta \times \Phi. \quad (2)$$

AUTOCAM then learns the “capture-worthiness” scores—the likelihood of a ST-glimpse appearing in human captured NFOV videos. Based on the assumption that 1) the content in human captured NFOV videos are mostly capture-worthy and 2) most ST-glimpses are not capture-worthy, it learns a classifier that predicts whether a video clip is a NFOV video or a ST-glimpse. The posteriors on test ST-glimpses are their capture-worthiness scores. The NFOV videos are crawled from YouTube and convolutional 3D features (C3D) [40] are used as the video representation.

After obtaining the capture-worthiness score of each candidate ST-glimpse, AUTOCAM constructs a camera trajectory (i.e. camera direction over time) by finding a path over the ST-glimpses that maximizes the *aggregate* capture-worthiness score, while simultaneously producing human-like smooth camera motions. It realizes the smooth camera

²Note FOV refers to the horizontal field of view throughout the paper.




Figure 2: Examples showing the importance of zooming. Our zoom in/out result is on the right for each pair.

motion by restricting the trajectory from choosing an ST-glimpse that is displaced from the previous one by more than a threshold ϵ , i.e.

$$|\Delta\Omega|_{\theta} = |\theta_t - \theta_{t-1}| \leq \epsilon_{\theta}, \quad |\Delta\Omega|_{\phi} = |\phi_t - \phi_{t-1}| \leq \epsilon_{\phi}. \quad (3)$$

In practice, we restrict the new ST-glimpse to be within the 8-adjacency of the previous ST-glimpse in the spherical coordinates. The problem can be reduced to a shortest path problem and solved efficiently using dynamic programming. The algorithm generates K NFOV outputs from each 360° input by 1) computing the best trajectory ending at each ST-glimpse location (of $18 \times 11 = 198$ possible), and 2) picking the top K of these. The algorithm assumes the input consists of pixel values on the unit sphere around the camera optical center and does not assume a specific 360° camera model. Our results include video frames from at least four common 360° camera models.

4. Approach

In this section, we present our algorithm. First we generalize the original Pano2Vid problem to enable zooming (Sec. 4.1). Next, we describe our coarse-to-fine approach to reduce the computational cost of camera trajectory selections (Sec. 4.2). Finally, we introduce an iterative approach to generate a diverse set of output trajectories (Sec. 4.3).

4.1. Zoom Lens Pano2Vid

We first generalize the Pano2Vid problem to Zoom Lens Pano2Vid, which allows zooming in the virtual camera control. Zooming is the technique of changing the focal length (f) of the lens. This is equivalent to changing the FOV of the camera because they are related by

$$\text{FOV} = 2 \arctan\left(\frac{d}{2f}\right), \quad (4)$$

where d is the horizontal sensor size and is a constant for the camera. The technique is widely used in videography: almost every digital camera nowadays provides the functionality, and most of us have the experience of adjusting the camera FOV by zooming to obtain the desirable framing. On the other hand, many existing manual interfaces for viewing 360° videos do not provide the control for FOV. In-




Figure 3: Generalized ST-glimpses.³

stead of trying to comply with these interfaces, we believe a proper definition of Pano2Vid must follow the full experience of videography, i.e. generating a video as if it is captured by a human videographer in the scene. Fig. 2 shows examples where zooming helps to improve the framing.

To achieve the effect of zooming, we sample ST-glimpses not only with different principal axis directions but also with different focal lengths. Assume the focal length for the 65.5° FOV is f^0 , we sample ST-glimpses with three different focal lengths

$$f \in F = \{0.5f^0, f^0, 1.5f^0\}, \quad (5)$$

which results in $\text{FOV} \in \{104.3^\circ, 65.5^\circ, 46.4^\circ\}$ respectively. The 104.3° FOV corresponds to an ultra wide angle lens and is the largest FOV commonly used in photography. The 65.5° and 46.4° FOV cover the range of standard lenses. The sampling for the principal axis direction is kept the same as Eq. 1. The new ST-glimpses are therefore defined by:

$$\Omega_{t,\theta,\phi,f} \equiv (\theta_t, \phi_t, f_t) \in \Theta \times \Phi \times F. \quad (6)$$

See Fig. 3.

When constructing the camera trajectories, we allow the algorithm to select ST-glimpses with different focal lengths. Therefore, the problem becomes finding a path over a series of 3D grids ($\Theta \times \Phi \times F$). In addition to the smoothness constraints in Eq. 3, we restrict the change in focal length between consecutive ST-glimpses:

$$|\Delta\Omega|_f = |f_t - f_{t-1}| \leq 0.5f^0, \quad (7)$$

which is a prior saying that human videographers tend to use gradual changes in zoom. While a basic dynamic programming problem similar to that used in AUTOCAM can find the trajectories in our formulation, the computational cost grows linearly with respect to the number of available focal lengths and becomes prohibitive. To make the algorithm for solving the Zoom Lens Pano2Vid problem practical, we next introduce a new computationally efficient approach for optimizing the trajectories.




Figure 4: Coarse-to-fine camera trajectory search. We first construct the trajectory on a coarse sample of ST-glimpses and then refine it on a dense sample of ST-glimpses around the trajectory. It reduces computational cost by avoiding processing all candidate ST-glimpses.

4.2. Coarse-to-Fine Camera Trajectory Search

The computational bottleneck for the algorithm is estimating the capture-worthiness score of each ST-glimpse. The algorithm has to first render the 360° ST-glimpse into NFOV video and then extract the C3D feature, both of which are computationally intensive. A basic dynamic programming approach like AUTOCAM requires the capture-worthiness scores for all candidate ST-glimpses. Even if we assume we can render the NFOV video and extract C3D features in real time, the processing time would be orders of magnitude longer than the input video length due to the large number of candidate ST-glimpses, i.e. several hours to process even a 1 minute 360° video. Therefore, we aim to reduce the number of capture-worthiness scores required. The most straightforward way would be to downsample the number of candidate ST-glimpses. However, this would lead to coarser camera control and a degradation in the quality of the output videos. In other words, the computational overhead not only makes the algorithm slow but it also restricts the granularity of virtual camera control.

Instead, we propose a coarse-to-fine approach that preserves the effective number of candidate ST-glimpses while reducing the number of ST-glimpses that require computation of the capture-worthiness score. The basic idea is to first construct the trajectory over coarsely sampled ST-glimpses and then refine the solution over densely sampled ST-glimpses centered around the initial trajectory. See Fig. 4. Furthermore, to keep the total cost sublinear in the number of available focal lengths, we construct the coarse trajectory with a single focal length and enable zooming only when refining the solution. Because we process ST-glimpses densely only in a small portion of the video, the total number of capture-worthiness scores required by the algorithm decreases. The proposed coarse-to-fine approach is based on the observation that the capture-worthiness scores of neighbor ST-glimpses are positively correlated, and the optimal trajectory in densely sampled ST-glimpses leads to a candidate solution in coarsely sampled ST-glimpses. Although the solution is not guaranteed to be the same as that of dynamic programming over all candidate ST-glimpses, empirical results verify that it perform well.

³Figures are best viewed in color.

We start the algorithm by sampling ST-glimpses at

$$\begin{aligned}
 \theta &\in \Theta' = \{\pm 10, \pm 30, \pm 75\}, \\
 \phi &\in \Phi' = \{0, 40, \dots, 320\}, \\
 t &\in T' = \{0s, 10s, \dots, L - 10s\}, \\
 f &\in F' = \{0.5f^0\}.
 \end{aligned} \tag{8}$$

We use the focal length $f = 0.5f^0$ which corresponds to the largest FOV so the visual content of these initial ST-glimpses cover that of other focal lengths at the same direction. Eq. 8 downsamples the candidate ST-glimpses by a factor of two from Eq. 1, so the number of capture-worthiness scores that need to be computed is only

$$\frac{|\Theta' \times \Phi' \times T'|}{|\Theta \times \Phi \times T|} \times \frac{1}{|F|} \approx 4.5\% \tag{9}$$

the number of candidate ST-glimpses. We solve for the trajectory using dynamic programming, but set the smoothness constraint to 2ϵ to account for the coarser samples

$$|\Delta\Omega|_\theta = |\theta_t - \theta_{t-1}| \leq 2\epsilon, \quad |\Delta\Omega|_\phi = |\phi_t - \phi_{t-1}| \leq 2\epsilon. \tag{10}$$

Denoting the ST-glimpses selected by the trajectory as

$$\Omega_{t,\theta,\phi}^0 \equiv (\theta_t^0, \phi_t^0) \tag{11}$$

for $t \in T'$, we then interpolate the ST-glimpses for $t \in T \setminus T' = \{5s, 15s, \dots\}$ to obtain the full trajectory.

To refine the trajectory, we sample ST-glimpses

$$\Omega_{t,\theta,\phi,f}^1 = (\theta_t^1, \phi_t^1, f_t^1) \tag{12}$$

that are adjacent to the original trajectory in direction

$$|\theta_t^1 - \theta_t^0| \leq \epsilon_\theta, \quad |\phi_t^1 - \phi_t^0| \leq \epsilon_\phi \tag{13}$$

following Eq. 1 and 5. We then solve the same trajectory search problem over the sampled ST-glimpses using dynamic programming with the smoothness constraints in Eq. 3 and Eq. 7. The number of candidate ST-glimpses is greatly reduced by the adjacency constraint in Eq. 13 and is only 5% of all candidate ST-glimpses.

4.3. Diverse Camera Trajectory Search

So far we have 1) presented our approach for the generalized Pano2Vid problem allowing variable FOV, and 2) devised a fast algorithm for optimizing it to produce the best hypothesis NFOV video output. Next, we wish to expand that single best solution to a set of *diverse* plausible outputs.

The motivation to generate a diverse set of output videos stems from the fact that, by definition, there may be multiple valid Pano2Vid solutions for each 360° video. For example, one might capture a soccer game by tracking the ball or by focusing on a particular player. Both of them will lead to a plausible presentation for the game and should be a valid output. In fact, when we ask human editors to manually extract NFOV videos from 360° data, the outputs for any two editors on the same source video have only about 47% over-

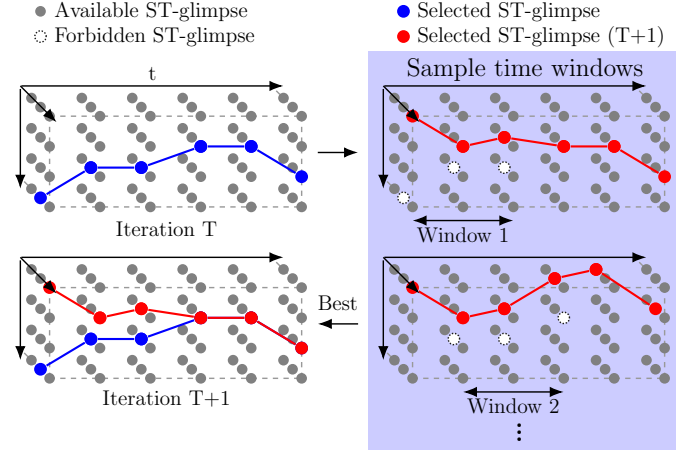


Figure 5: Diverse trajectory search generates trajectories iteratively. In each iteration, we construct multiple trajectory search problems by sampling time windows and removing previously selected ST-glimpses in the window from the search space. We solve all the problems and take the best solution as the output trajectory.

lap on average. In addition, many applications of Pano2Vid would prefer a set of candidate solutions instead of a single output. An editing aid system would be more useful if the editor has the freedom to choose from different reasonable algorithm-provided initializations. Similarly, a 360° video player that allows the viewers to choose from different NFOV video presentations is likely to achieve a better user experience because the viewers can decide what to see based on their preferences.

It is difficult to encourage diversity in a single pass of dynamic programming, because all the potential solutions are constructed concurrently and the distances between them are hard to control. In prior work, the AUTOCAM algorithm generates multiple trajectories by requiring them to end at different spatial locations in the video, i.e., with different ST-glimpses. However, this requirement does not guarantee the distance between different trajectories. In the worst case, they can be exactly the same except in the last ST-glimpses, leading to poor diversity in the outputs.

Instead, we generate trajectories iteratively and encourage diversity by imposing the minimum distance constraint between trajectories generated in different iterations. To realize the constraint, we sample a time window and forbid the trajectories of the current iteration from selecting the same ST-glimpses as the solution of previous iterations in the window. Therefore, the length of the time window determines the minimum distance between the solutions of different iterations. We sample the time window at multiple temporal locations and construct the optimal trajectory for each window. We take the best trajectory among them in terms of accumulated capture-worthiness score as the solution of the current iteration. This avoids critical glimpses being excluded from the solution space even if it is selected by previous trajectories. See Fig. 5.

In practice, we set the length of the time window to 10% the original video length and sample 20 different windows distributed evenly over time. Once the length and location of the time window is specified, the optimal trajectory can be found using dynamic programming on a modified shortest path problem where the ST-glimpses selected by previous solutions in the window are removed from the search space. To improve computational efficiency, we divide the unit sphere into 6 regions (by 3 azimuthal angles and 2 polar angles) and generate an output per region by finding the best trajectory ending in the region. This leads to 6 trajectories per iteration. The iteration ends after K -trajectories have been generated.

5. Experiments

Next we validate our method on challenging videos. See project webpage for video examples and comparisons.

5.1. Dataset

We use the Pano2Vid dataset introduced in [37].⁴ It consists of 86 360° videos crawled from YouTube using four keywords: “Soccer,” “Mountain Climbing,” “Parade,” and “Hiking”, for a total length of 7.3 hours. The dataset also provides “HumanCam” data: URLs for 343 hours of NFOV YouTube videos. We use half of that data to train logistic regression capture-worthiness classifiers and half for evaluation (details below). Following [37], we train the classifier using leave-one-360°-video-out strategy for each keyword.

Some evaluation metrics below compare our outputs to human-selected camera trajectories from 360° videos (“HumanEdit”). To collect these trajectories, we ask human subjects to edit the videos. The editors control the angle and FOV of a virtual camera with their mouse, overlaid on the 360° video’s equirectangular projection so the user can see all the visual content at once. Please see supp. file for details and examples. We collect HumanEdit data for 40 videos, each of them annotated by 3 editors. Overall, we collect 480 trajectories totaling 717.2 minutes of video, and roughly 18 hours of annotation time.

5.2. Baselines

We compare the following methods in the experiments.

- **CENTER** — random trajectories biased toward the “center” of the 360° video axes ($\theta = \phi = 0^\circ$). The bias accounts for the fact that user-generated 360° videos often contain interesting content close to the center, possibly because the 360° camera design allows the users to use it as if it were a NFOV camera. We sample the camera direction for the next time-step from a Gaussian centered around the current direction, which starts from the center.
- **EYE-LEVEL** — static trajectories that place the virtual camera on the equator ($\theta = 0^\circ$). The equator usually

corresponds to eye-level in 360° video where most interesting events happen. We sample the azimuthal angle ϕ every 20° for 18 different camera directions.

- **SALIENCY** — replace the capture-worthiness scores in AUTOCAM by saliency scores.⁵ The saliency is computed by a popular method [17] over the 360° video frame in equirectangular projection.
- **AUTOCAM** — to our knowledge, the only prior work tackling this problem [37].

5.3. Evaluation Metrics

We adopt the metrics proposed in [37], generalizing them as need to account for Zoom Lens Pano2Vid.

HumanCam-based metrics This group of metrics evaluates whether the output videos look like human-captured NFOV video (HumanCam). The more indistinguishable the algorithm outputs are from HumanCam, the better the algorithm. There are three metrics:

- **Distinguishability** quantifies if it is possible to tell the algorithm output apart from HumanCam. It is measured by 5-fold cross validation error rate of discriminative classifiers trained with HumanCam as positives and algorithm outputs as negatives. *Higher* error is better.
- **HumanCam-likeness** measures the relative distance from algorithm outputs to HumanCam data in a semantic feature space. It trains classifiers with HumanCam video as positives and all algorithm-generated videos as negatives using leave-one-360°-video-out strategy. An algorithm-generated video is ranked based on its distance to the decision boundary, and we compute the normalized mean rank; lower is better.
- **Transferability** measures how well semantic classifiers trained on HumanCam videos transfer to algorithm-generated videos, and vice versa. The more transferable the classifiers are, the more similar the HumanCam and algorithm outputs are. We take the four search keywords as labels to learn a multi-class classifier on one domain and measure transferability using the test error on the other domain.

We use logistic regression classifiers in all metrics.

HumanEdit-based metric Whereas the above metrics score output videos by their resemblance to human-captured videos in general, the HumanEdit metric measures the similarity between algorithm-generated camera trajectories and the manually created trajectories in the same 360° video. The more similar they are, the better the algorithm. This metric captures the subjective preference of human editors but is easily reproducible, in contrast to one-off user studies. In particular, we compute the **overlap** of the camera FOV in each frame. The overlap is approximated by

⁵We also considered a saliency baseline that permits zoom like our method, but it fared worse than all others.

⁴<http://vision.cs.utexas.edu/projects/Pano2Vid>

Table 1: Pano2Vid performance: HumanCam-based and HumanEdit-based metrics. The arrows in column 3 indicate whether lower scores are better (\Downarrow), or higher scores (\Uparrow). Our full method (OURS) significantly outperforms the baselines, and the relative improvements over the best performing baseline are up to 43.4%.

		CENTER	EYE-LEVEL	SALIENCY	AUTOCAM [37]	OURS w/o DIVERSITY	OURS
Distinguishability	Error rate (%)	\Uparrow	1.93	4.03	7.70	12.05	17.41
HumanCam-Likeness	Mean Rank	\Downarrow	0.659	0.707	0.612	0.522	0.279 0.267
Transferability	Human \rightarrow Auto	\Uparrow	0.582	0.607	0.597	0.517	0.590 0.591
	Auto \rightarrow Human		0.526	0.552	0.549	0.584	0.618 0.617
Overlap	Trajectory	\Uparrow	0.271	0.335	0.359	0.343	0.436 0.442
	Frame		0.498	0.555	0.580	0.530	0.629 0.630

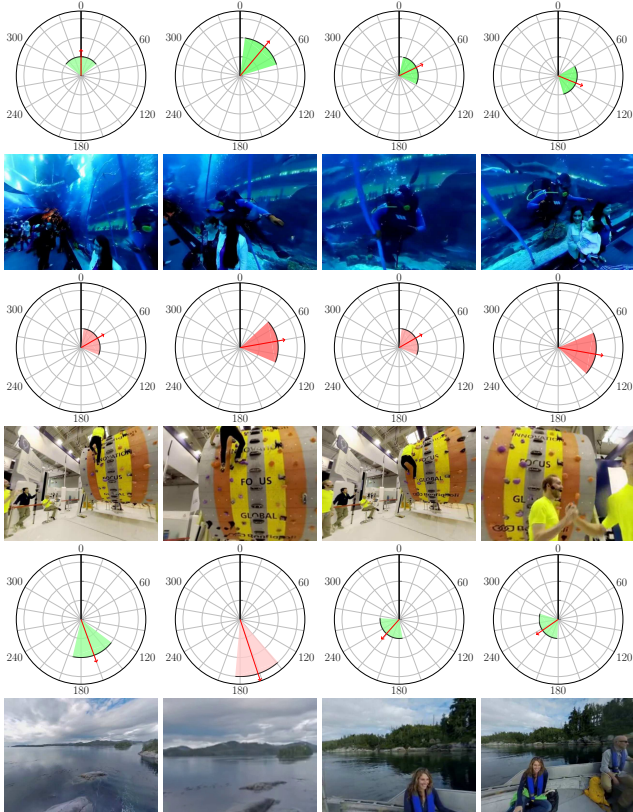


Figure 6: Example frames of our algorithm outputs and the corresponding camera poses. The circular sector shows the camera FOV and azimuthal angle, and the color shows the polar angle. Red/green indicates the angle is greater/smaller than 0, and more saturated color indicates larger value. In the first example, the camera tracks the diver. In the second example, it zooms in to capture special moments, e.g., when the person climbs to the top or high fives with another person. In the third example, the camera first captures distant scene with long focal length and close objects with short focal length. Also see videos on the project webpage.

$\max(1 - \frac{2\Delta\Omega}{FOV_H + FOV_A}, 0)$, where FOV_H and FOV_A correspond to the FOV of algorithm and human controlled camera respectively. We report overlap results under two pooling strategies, **trajectory**, which rewards outputs similar to at least one HumanEdit trajectory as a whole, and **frame**, which rewards outputs similar to any HumanEdit trajectory at each frame. See supp. for details.

Diversity We gauge diversity by the number of distinct groups of trajectories among the outputs, where trajectories

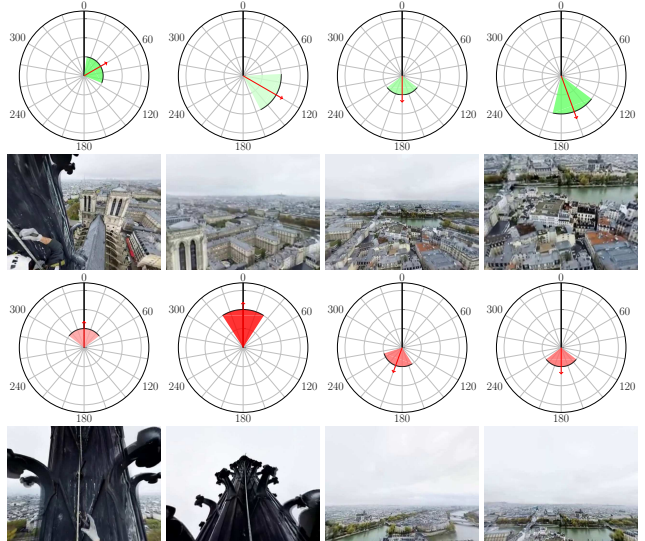


Figure 7: Two trajectories extracted from the same 360° video. Our algorithm presents the same scene in different manners.

differing in angle or FOV for fewer than 10% of the video frames are considered to be the same group.

Computational cost We measure computational cost by the average processing time for 1 minute of 360° video. The time is measured on a machine equipped with one Intel Xeon E5-2697 v2 processor (24 cores) and one GeForce GTX Titan Black GPU (including I/O).

5.4. Output Quality

First we quantify the quality of the algorithm-generated videos using the HumanCam-/HumanEdit-based metrics. We obtain the trajectories for AUTOCAM from the authors of [37]. We take the top $K=20$ outputs, following [37].

Table 1 shows the results. Our full method (OURS) significantly outperforms all other methods. Compared to the best performing baseline (AUTOCAM), our method improves Distinguishability by 43.4% and ranks 25.5% better on average in the HumanCam-Likeness. We also see a 23% improvement in the Trajectory Overlap metric. The superior performance clearly shows the importance of the proposed zoom lens. Figures 6 and 7 show examples. Interestingly, the camera zooms out (i.e. $FOV > 65.5^\circ$) more often than it zooms in. In fact, 76% of the ST-glimpses selected by OURS have 104.3° FOV, and editors select 104.3° FOV in 55% of the HumanEdit data. These results suggest that a




Figure 8: Computational cost vs. output quality for our method and AUTOCAM [37]. Computational cost is measured by the average processing time per 1 minute of input 360° video. Quality is measured by the Distinguishability and Trajectory Overlap; higher is better (\uparrow) for both metrics.

larger FOV is preferable when viewing 360° videos, possibly because the object of interest is usually closer to the camera. Please see project webpage for videos.

Table 1 also shows that the CENTER and EYE-LEVEL baselines perform poorly, indicating that hand coded heuristics based on prior knowledge are not enough and a content-dependent method is necessary. EYE-LEVEL performs better than CENTER, which reflects the fact that EYE-LEVEL is a generic prior while the CENTER prior only holds when the 360° camera is asymmetric and the user uses it as an ordinary camera. Although SALIENCY is content-dependent, it captures content that attracts gaze, which appears to be a poor proxy for the Pano2Vid task. It underperforms OURS in all metrics except the Human \rightarrow Auto transferability.

Although our method outperforms all baselines, we notice that the learned capture-worthiness is unable to capture preferences induced by context. For example, the algorithm fails to concentrate on family members in a family video. Also, the smoothness constraint may be too strong in some scenes where the camera is unable to adjust to rapidly changing content. See project page for failure cases.

5.5. Computational Cost

Fig. 8 shows the computational cost versus output quality, measured by HumanCam Distinguishability and HumanEdit Trajectory Overlap. The computational cost for OURS *without* coarse-to-fine processing is 2.88 times that of AUTOCAM, which shows the need for a more computationally efficient algorithm. OURS *w/* FAST reduces the computational cost by 84% compared to OURS while performing similarly in Trajectory Overlap and only 6% worse in Distinguishability. Comparing the trajectories generated by the two methods, the coarse-to-fine approach tends to favor the largest FOV and ignore trajectories that select the smallest FOV throughout the video, because the initial trajectories are constructed with the largest FOV. This may cause the degradation in Distinguishability due to the fact that the 104.3° FOV introduces distortion in output frames. OURS *w/* FAST takes less than 50% of the computation of AUTOCAM yet is more accurate in all metrics. Although further optimization on the implementation may reduce the




Figure 9: The number of distinct trajectories each algorithm generates with respect to the number of algorithm outputs.

processing time, the relative cost will remain the same, as it is linear in the number of ST-glimpses processed.

5.6. Output Diversity

Fig. 9 displays the relative diversity captured by the methods as a function of the number of outputs they generate. Our diverse trajectory search approach generates more distinct trajectories than any other method, with 2 ~ 3 times more distinct outputs than our ablated non-diverse variant. AUTOCAM generates slightly more distinct trajectories than OURS *w/o* DIVERSITY, despite the fact that they both use DP to search for the trajectories. This is because OURS *w/o* DIVERSITY has more degrees of freedom in camera control and can generate closer variants of a given camera trajectory, i.e. it can generate more outputs by varying the choice of the last ST-glimpse than AUTOCAM can. SALIENCY performs poorly in terms of output diversity—fewer than two distinct trajectories on average even among its top 20 outputs. Because saliency scores are computed more densely than the capture-worthiness scores, the correlation of the scores between neighbor directions are stronger, and the algorithm can generate more trajectories that are close variants to the others. Note that the results in Fig. 9 are complementary to that in Table 1. Together, they show that our method can generate diverse yet quality outputs.

6. Conclusion

We explore virtual videography in the context of 360° video. Our system controls a virtual camera in a 360° video to generate videos that look human-captured and properly present the content to a passive human viewer. We generalize the previously proposed Pano2Vid problem by allowing the algorithm to control its field of view dynamically, introduce a coarse-to-fine optimization approach that makes it tractable, and propose a method to encourage diversity and capture the multimodal nature of “good” NFOV videos. Under a battery of metrics—including against human editors—our algorithm outperforms the state of the art, bringing tools one step closer to practical virtual videography.

Acknowledgement. We thank Chia-Chen, Kimberly, Ruohan, Wei-Ju, and YiHsuan for collecting annotation. This research is supported in part by NSF IIS -1514118 and a gift from Intel.

References

- [1] <https://www.cnet.com/news/360-degree-cameras-2016>.
- [2] <https://www.cta.tech/News/Blog/Articles/2016/September/360-Camera-Sales-Set-to-Skyrocket-in-2016.aspx>.
- [3] <https://www.youtube.com/playlist?list=PLS3XGZxi7cBXqnRTtKMU7Anm-R-kyhkyC>.
- [4] I. Arev, H. S. Park, Y. Sheikh, J. Hodgins, and A. Shamir. Automatic editing of footage from multiple social cameras. In *SIGGRAPH*, 2014.
- [5] S. Avidan and A. Shamir. Seam carving for content-aware image resizing. In *ACM TOG*, 2007.
- [6] D. Batra, P. Yadollahpour, A. Guzman-Rivera, and G. Shakhnarovich. Diverse m-best solutions in markov random fields. In *European Conference on Computer Vision*, 2012.
- [7] D. B. Christianson, S. E. Anderson, L.-w. He, D. H. Salesin, D. S. Weld, and M. F. Cohen. Declarative camera control for automatic cinematography. In *AAAI/IAAI, Vol. 1*, 1996.
- [8] K. Dale, E. Shechtman, S. Avidan, and H. Pfister. Multi-video browsing and summarization. In *CVPR*, 2012.
- [9] D. K. Elson and M. O. Riedl. A lightweight intelligent virtual cinematography system for machinima production. In *AIIDE*, 2007.
- [10] J. Foote and D. Kimber. Flycam: Practical panoramic video and automatic camera control. In *ICME*, 2000.
- [11] K. Gimpel, D. Batra, C. Dyer, G. Shakhnarovich, and V. Tech. A systematic exploration of diversity in machine translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, October*, 2013.
- [12] M. Gleicher and J. Masanz. Towards virtual videography. In *Proceedings of the eighth ACM international conference on Multimedia*, 2000.
- [13] M. L. Gleicher, R. M. Heck, and M. N. Wallick. A framework for virtual videography. In *Proceedings of the 2nd international symposium on Smart graphics*, 2002.
- [14] B. Gong, W.-L. Chao, K. Grauman, and F. Sha. Diverse sequential subset selection for supervised video summarization. In *NIPS*, 2014.
- [15] M. Gygli, H. Grabner, H. Riemenschneider, and L. Van Gool. Creating summaries from user videos. In *ECCV*, 2014.
- [16] M. Gygli, H. Grabner, and L. Van Gool. Video summarization by learning submodular mixtures of objectives. In *CVPR*, 2015.
- [17] J. Harel, C. Koch, and P. Perona. Graph-based visual saliency. In *NIPS*, 2006.
- [18] L.-w. He, M. F. Cohen, and D. H. Salesin. The virtual cinematographer: a paradigm for automatic real-time camera control and directing. In *ACM CGI*, 1996.
- [19] R. Heck, M. Wallick, and M. Gleicher. Virtual videography. *ACM Trans. Multimedia Comput. Commun. Appl.*, 3(1), Feb. 2007.
- [20] L. Itti and P. Baldi. Bayesian surprise attracts human attention. *Vision Research*, 49(10):1295–1306, 2009.
- [21] L. Itti, N. Dhavale, and F. Pighin. Realistic avatar eye and head animation using a neurobiological model of visual attention. In *Proc. SPIE 48th Annual International Symposium on Optical Science and Technology*, 2003.
- [22] E. Jain, Y. Sheikh, A. Shamir, and J. Hodgins. Gaze-driven video re-editing. In *SIGGRAPH*, 2015.
- [23] P. Khoenkaw and P. Piamsa-Nga. Automatic pan-and-scan algorithm for heterogeneous displays. In *Springer MTA*, 2015.
- [24] A. Khosla, R. Hamid, C.-J. Lin, and N. Sundaresan. Large-scale video summarization using web-image priors. In *CVPR*, 2013.
- [25] A. Kirillov, B. Savchynskyy, D. Schlesinger, D. Vetrov, and C. Rother. Inferring m-best diverse labelings in a single one. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
- [26] P. Krähenbühl, M. Lang, A. Hornung, and M. Gross. A system for retargeting of streaming video. In *ACM TOG*, 2009.
- [27] Y. J. Lee, J. Ghosh, and K. Grauman. Discovering important people and objects for egocentric video summarization. In *CVPR*, 2012.
- [28] J. Li and D. Jurafsky. Mutual information and diverse decoding improve neural machine translation. *arXiv preprint arXiv:1601.00372*, 2016.
- [29] F. Liu and M. Gleicher. Video retargeting: automating pan and scan. In *ACM MM*, 2006.
- [30] T. Liu, Z. Yuan, J. Sun, J. Wang, N. Zheng, X. Tang, and H.-Y. Shum. Learning to detect a salient object. In *PAMI*, 2011.
- [31] P. Mindek, L. Čmolík, I. Viola, E. Gröller, and S. Bruckner. Automated summarization of multiplayer games. In *ACM CCG*, 2015.
- [32] D. Potapov, M. Douze, Z. Harchaoui, and C. Schmid. Category-specific video summarization. In *ECCV*, 2014.
- [33] A. Rav-Acha, Y. Pritch, and S. Peleg. Making a long video short: Dynamic video synopsis. In *CVPR*, 2006.
- [34] M. Rubinstein, A. Shamir, and S. Avidan. Improved seam carving for video retargeting. In *ACM TOG*, 2008.
- [35] D. Rudoy, D. Goldman, E. Shechtman, and L. Zelnik-Manor. Learning video saliency from human gaze using candidate selection. In *CVPR*, 2013.
- [36] Y. Rui, A. Gupta, J. Grudin, and L. He. Automating lecture capture and broadcast: technology and videography. *Multimedia Systems*, 10(1):3–15, 2004.
- [37] Y.-C. Su, D. Jayaraman, and K. Grauman. Pano2vid: Automatic cinematography for watching 360° videos. In *ACCV*, 2016.
- [38] M. Sun, A. Farhadi, and S. Seitz. Ranking domain-specific highlights by analyzing edited videos. In *ECCV*, 2014.
- [39] X. Sun, J. Foote, D. Kimber, and B. Manjunath. Region of interest extraction and virtual camera control based on panoramic video capturing. In *IEEE TOM*, 2005.
- [40] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [41] A. K. Vijayakumar, M. Cogswell, R. R. Selvaraju, Q. Sun, S. Lee, D. Crandall, and D. Batra. Diverse beam search: Decoding diverse solutions from neural sequence models. *arXiv preprint arXiv:1610.02424*, 2016.
- [42] B. Xiong and K. Grauman. Detecting snap points in egocentric video with a web photo prior. In *ECCV*, 2014.
- [43] Y. Zhai and M. Shah. Visual attention detection in video sequences using spatiotemporal cues. In *ACM MM*, 2006.