Detection and Removal of Rain from Videos*

Kshitiz Garg and Shree K. Nayar
Department of Computer Science, Columbia University New York, 10027
Email: {kshitiz,nayar}@cs.columbia.edu

Abstract

The visual effects of rain are complex. Rain consists of spatially distributed drops falling at high velocities. Each drop refracts and reflects the environment, producing sharp intensity changes in an image. A group of such falling drops creates a complex time varying signal in images and videos. In addition, due to the finite exposure time of the camera, intensities due to rain are motion blurred and hence depend on the background intensities. Thus, the visual manifestations of rain are a combination of both the dynamics of rain and the photometry of the environment. In this paper, we present the first comprehensive analysis of the visual effects of rain on an imaging system. We develop a correlation model that captures the dynamics of rain and a physics-based motion blur model that explains the photometry of rain. Based on these models, we develop efficient algorithms for detecting and removing rain from videos. The effectiveness of our algorithms is demonstrated using experiments on videos of complex scenes with moving objects and time-varying textures. The techniques described in this paper can be used in a wide range of applications including video surveillance, vision based navigation, video/movie editing and video indexing/retrieval.

1 Steady and Dynamic Weather Conditions

Outdoor vision systems are used for various purposes such as tracking, recognition and navigation. Despite their widespread use, current systems do not account for common weather conditions such as rain, snow, fog and mist. In order to develop vision systems that perform under all weather conditions, it is essential to model the visual effects of the various weather conditions and develop algorithms to remove them.

Weather conditions vary widely in their physical properties and in the visual effects they produce in images. Based on their differences, weather conditions can be broadly classified as *steady* (fog, mist and haze) or *dynamic* (rain, snow and hail). In the case of steady weather, individual droplets are too small ($1-10\,\mu$m) to be visible to a camera, and the intensity produced at a pixel is due to the aggregate effect of a large number of droplets within the pixel's solid angle (see Figure 1(a)). Hence, volumetric scattering models such as attenuation and airlight [9] can be used to adequately describe the effects of steady weather. Algorithms [10] have been recently developed to remove the effects of steady weather from images.

On the other hand, the constituent particles of dynamic weather conditions such as rain, snow and hail are larger ($0.1-10$ mm) and individual particles are visible in the image. An example is shown in Figure 1(b), the streaks of rain are caused by individual drops. Here, aggregate scattering models previously used for steady conditions are not applicable . The analysis of dynamic weather conditions requires the development of stochastic models that capture the spatial and temporal effects of a large number of particles moving at high speeds (as in rain) and with possibly complex trajectories (as in snow).

(a) A scene in fog (b) A scene in rain

(c) Rain and moving objects (d) Rain and ripples

Figure 1: The visual appearances of steady (fog) and dynamic (rain) weather conditions differ widely. (a) An image of a scene taken under foggy conditions. The intensity at each pixel is due to the aggregate effect of a large number of droplets within the pixel's solid angle. (b) An image of a scene taken on a rainy day showing streaks caused by the motion blur of individual drops. An algorithm that detects and removes rain must be robust to complex scene and camera motions as in (c) and at the same time insensitive to other time-varying textures such as water ripples as in (d).

In this work, we focus on the problem of rain. Rain consists of a distribution of a large number of drops of various sizes, falling at high velocities. Each drop behaves like a transparent sphere, refracting and reflecting light from the environment towards the camera. An ensemble of such drops falling at high velocities results in time varying intensity fluctuations in images and videos. In addition, due to the finite exposure time of the camera, intensities due to rain are motion blurred and therefore depend on the background. Thus, the visual manifestations of rain are a combined effect of the *dynamics* of rain and the *photometry* of the environment.

Rain has been studied extensively in the fields of atmospheric sciences, signal communication and remote sensing [8, 13, 6]. Most of these studies use active illumination sources (lasers) and specialized detectors (photo-cells) to examine the effects of rain on a transmitted signal. However, the effects of rain on a camera viewing a scene in a natural environment are very different and remain unexplored.

*This research was conducted at the Columbia Vision and Graphics Center in the Computer Science Department at Columbia University. It was funded in parts by a DARPA HID Contract (N00014-00-1-0916) and an NSF Grant (IIS-99-87979).

In computer graphics, rain has been rendered using particle systems [11] or heuristic models [12]. However, these rendering methods are not based on the physical properties of rain and fail to represent the complex visual effects of rain. One possible approach to analyzing the appearance of rain is to learn a model from real rain videos (for instance, dynamic textures [2]). However, as mentioned before, the appearance of rain depends on several factors including the physical properties of rain, the environment and the camera settings. Hence, learning a general model for the appearance of rain in an arbitrary scene and camera setting is hard.

In this paper, we present the first comprehensive analysis of the visual effects of rain on imaging systems. We begin by briefly summarizing the physical properties of rain such as the spatial distribution, shapes, sizes and velocities of drops. Then, we develop two separate models that capture the dynamics and the photometry of rain. Based on these models, we develop efficient algorithms for detecting and removing rain from videos. We demonstrate the effectiveness of our algorithms through experiments on videos of complex scenes with moving objects and time-varying textures. Our results can be used in a wide range of applications such as video surveillance, vision based navigation, movie editing and video indexing/retrieval.

2 Physical Properties of Rain

Rain is a collection of randomly distributed water droplets of different shapes and sizes that move at high velocities. The physical properties of rain have been extensively studied in atmospheric sciences [4, 1, 7, 6, 13]. Here, we briefly summarize these properties and make observations that are relevant to our goal of modeling the appearance of rain.

The size of a raindrop typically varies from 0.1 mm to 3.5 mm. The distribution of drop sizes in rain is given by the Marshall-Palmer distribution [7]. Figure 2(a) shows the distribution for a typical rainfall. Note that the density of drops decreases exponentially with the drop size.

The shape of a drop can be expressed as a function of its size [1]. Figure 2(b) shows the shapes of raindrops of various sizes. Smaller raindrops are generally spherical in shape while larger drops resemble oblate spheroids. In a typical rainfall, most of the drops are less than 1 mm in size, as seen in Figure 2(a). Hence, most raindrops are spherical and we will use this approximation in our work.

As a drop falls through the atmosphere, it reaches a constant *terminal velocity*. The terminal velocity v of a drop is also related to its size a and is given by [4]

$$v = 200\sqrt{a}\ , \tag{1}$$

where a is in meters and v is in meters/s.

The individual raindrops are distributed randomly in 3D space. This distribution is usually assumed to be uniform [6, 13]. Moreover, it can be assumed that the statistical properties of the distribution remains constant over time [6]. These assumptions are applicable in most computer vision scenarios.

3 Appearance Model for Rain

In this section, we analyze image formation through rain. The complex spatial and temporal intensity fluctuations in images produced by rain depend on several factors: (a) drop distribution and velocities; (b) environment illumination and background scene; and (c) the intrinsic parameters of the camera. We first develop a correlation model that captures the dynamics of rain based on the distribution and velocities of raindrops. Then, we develop a physics-based motion blur model that describes the brightnesses produced by streaks of rain.

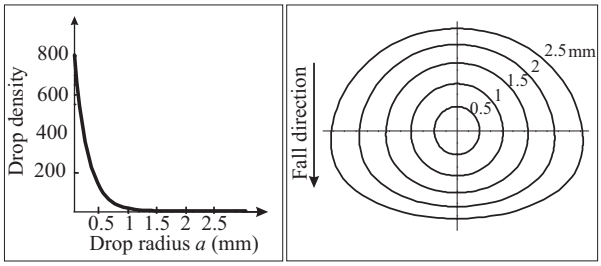

(a) Drop size distribution (b) Shapes of raindrops

Figure 2: (a) The Marshall-Palmer distribution [7] gives the number density of raindrops as a function of drop size. Note that the density of drops decreases exponentially with drop size. (b) The shapes of raindrops of various sizes ($0.5 - 2.5$ mm). Due to air pressure, the bases of the drops are flattened in the direction of fall. From (a) we note that a large fraction of rain drops are small and hence spherical.

3.1 Dynamics of Rain

Consider a camera observing a volume of rain. Drops are randomly distributed in this volume and fall with high velocities. The projection of these drops onto the image plane produces a time-varying random field in image space which represents the dynamics of rain. For now, we consider only the image projections of the drops and not their intensities. Thus, the dynamics of rain may be represented by a binary field

$$b(\vec{r},t) = \begin{cases} 1, & \text{if drop projects to location } \vec{r} \text{ at time } t; \\ 0, & \text{otherwise,} \end{cases} \tag{2}$$

where $\vec{r}$ represents the spatial coordinates in the image and t is time. Initially, we consider both the space and time parameters, $\vec{r}$ and t, to be continuous.

As mentioned in Section 2, we assume that the distribution of drops in the volume is uniform over space and time. Under this condition, the binary random field $b(\vec{r},t)$ is *wide sense stationary* in space and time [6]. This implies that the correlation function $R_b(\vec{r_1},t_1;\vec{r_2},t_2)$ depends only on differences in the image coordinates ($\Delta\vec{r} = \vec{r_1} - \vec{r_2}$) and the difference in time ($\Delta t = t_1 - t_2$). That is:

$$\begin{aligned} R_b(\vec{r_1},t_1;\vec{r_2},t_2) &\equiv \frac{1}{L}\int_0^L b(\vec{r_1},t_1+t)\, b(\vec{r_2},t_2+t)\, dt \\ &= R_b(\Delta\vec{r},\Delta t)\,, \end{aligned} \tag{3}$$

where, the correlation R_b is computed over a large time period $[0, L]$. $R_b(\Delta\vec{r},\Delta t)$ can be computed by measuring the temporal correlation with time lag Δt between the values of the binary field at points $\vec{r}$ and $\vec{r} + \Delta\vec{r}$. An important constraint arises due to the straight line motion of the drops. Consider a drop that falls with image velocity $\vec{v}_i$. After time Δt, the displacement of this drop is $\vec{v}_i\Delta t$. Hence, the binary field at time instants t and $t + \Delta t$ are related as

$$b(\vec{r} + \vec{v}_i\Delta t, t + \Delta t) = b(\vec{r},t)\,. \tag{4}$$

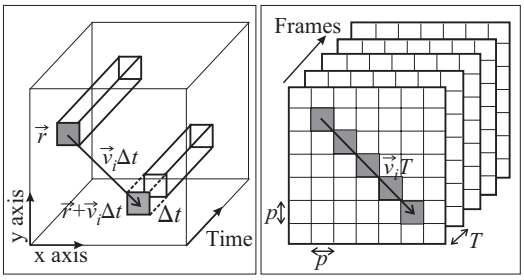

(a) (b)

Figure 3: (a) Spatio-temporal correlation in a continuous space-time volume. Two image points $\vec{r}$ and $\vec{r_1} = \vec{r} + \vec{v}_i\Delta t$ are temporally correlated due to rain ($R_b(\vec{v}_i\Delta t, \Delta t)$ is high). (b) In discrete domain, computing the correlation $R_b(\vec{m}p, 0)$ is equivalent to computing $R_b(\vec{v}_i\Delta t, \Delta t)$ over the entire duration $[0 \leq \Delta t \leq T]$. Hence, $R_b(\vec{m}p, 0)$ is high for pixels (shown shaded) separated by distance $[0 \leq \vec{m}p \leq \vec{v}_i T]$.

As a result, the correlation $R_b(\vec{r}, t; \vec{r} + \vec{v}_i\Delta t, t + \Delta t)$ is high. From equation (3), we write

$$R_b(\vec{r}, t; \vec{r} + \vec{v}_i\Delta t, t + \Delta t) = R_b(\vec{v}_i\Delta t, \Delta t)\,. \qquad (5)$$

This implies that the value of the binary field b at *any* two image coordinates, separated by $\vec{v}_i\,\Delta t$ in space are correlated with time lag Δt. This is illustrated in Figure 3(a).

The above correlation is analyzed in continuous domain. However, imaging systems have a finite pixel size p and a finite integration time T. In a discrete domain, let us denote the correlation by $R_b(\vec{m}p, nT)$, where $\vec{m}$ is the displacement in integer image coordinates and n is the time lag in number of frames. The discrete binary field at any frame is obtained by integrating the continuous binary field over the time duration T. Hence, computing the correlation $R_b(\vec{m}p, nT)$ is equivalent to computing $R_b(\vec{v}_i\Delta t, \Delta t)$ over the entire time interval $[nT \leq \Delta t \leq (n+1)T]$. As a result, $R_b(\vec{m}p, nT)$ is high for all pixels separated by the distance $[\vec{v}_i nT \leq \vec{m}p \leq \vec{v}_i(n+1)T]$. Figure 3(b) shows the pixels for which the zeroth time lag correlation $R_b(\vec{m}p, 0)$ is high, where $[0 \leq \vec{m}p \leq \vec{v}_i T]$.

Note that different drops may have different (unknown) image velocity magnitudes $|\vec{v}_i|$ depending on their sizes and distances from the camera. However, within a local region, drops fall more or less in the same direction $\vec{v}_i/|\vec{v}_i|$. Hence, irrespective of the drop velocity magnitudes, the correlation R_b remains high in the direction $\vec{v}_i/|\vec{v}_i|$ of rain and low in all other directions. In summary, the binary field b produced by rain exhibits the following important properties :

- Since drop distribution is uniform over space and time, the binary field b due to rain is wide sense stationary.
- The temporal correlation between pixels in any neighborhood is high in the direction of rain and can be used to detect rain and compute its direction.

3.2 Photometry of Rain

We now model the intensities produced in an image by falling raindrops. In previous work [3], we studied the appearance of a *stationary drop* and its relation to the radiance distribution of the environment. We first summarize the main points of this study. Then, we develop a model for the motion-blurred intensities produced by a *falling drop*.

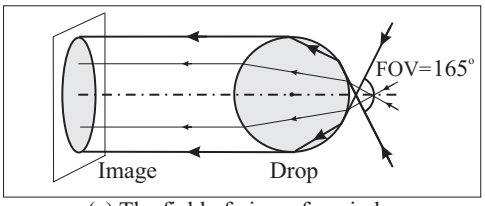

(a) The field of view of a raindrop

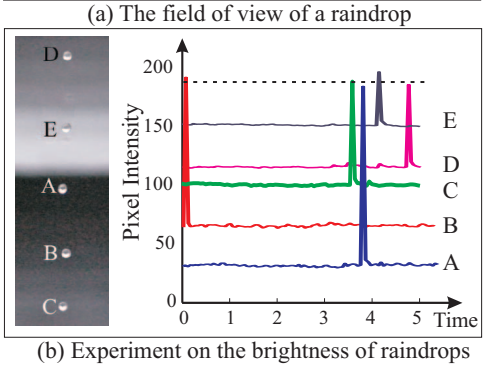

(b) Experiment on the brightness of raindrops

Figure 4: Geometry and brightness of a raindrop. (a) The field of view of a raindrop is approximately 165°. (b) Experiment verifying the average brightnesses of raindrops. The background is a plane with 5 horizontal stripes of different shades of gray. The average intensities produced in the drop-sized regions A through E are plotted as a function of time. Note that the drops (spikes) are brighter than their backgrounds. Further, *in the absence of motion-blur*, the brightnesses of the drops are roughly the same and independent of the background.

3.2.1 Brightness of a Stationary Raindrop

Raindrops behave like lenses refracting and reflecting (both specularly and internally) scene radiances towards the camera. We have developed detailed geometric and photometric models [3] for the refraction through and reflection from a spherical raindrop. These models show that raindrops have a large field of view of approximately 165° (see Figure 4(a)) and the incident light that is refracted towards the camera is attenuated by only 6%. Based on these optical properties of a drop, we make the following observations:

- Raindrops refract light from a large solid angle of the environment (including the sky) towards the camera. Specular and internal reflections further add to the brightness of the drop. Thus, a drop tends to be much brighter than its background (the portion of the scene it occludes).
- The solid angle of the background occluded by a drop is far less than the total field of view of the drop itself. Thus, in spite of being transparent, the average brightness within a stationary drop (*without motion-blur*) does not depend strongly on its background.

We verified these observations using a video of drops falling under an overcast sky. The video was captured with low exposure time to prevent motion blur of drops. The background scene consisted of horizontal stripes of different brightnesses, as shown in Figure 4(b). The average intensities of drop-sized regions marked A through E are plotted as a function of time. Each spike corresponds to a sudden increase in brightness when a drop passes through the marked region. Note that the brightnesses of the drops (peak values of the spikes) are much higher than the corresponding background intensities. Also, these peak values are approximately the same even though the

background intensities are very different (see the dashed line in Figure 4(b)).

3.2.2 Photometry of Rain Streaks: Motion Blur

Falling raindrops produce motion-blurred intensities due to the finite integration time of a camera. These intensities are seen as *streaks* of rain. Unlike a stationary drop, the intensities of a rain streak depend on the brightness of the (stationary) drop as well as the background scene radiances and integration time of the camera. Let us now analyze these dependencies.

Consider a video camera with a linear radiometric response and exposure (integration) time T, observing a scene with rain. To determine the intensity I_d produced at a pixel effected by a raindrop, we need to examine the irradiance of the pixel over the time duration T. Figure 5(a) shows a raindrop passing through a pixel within the time interval $[t_n, t_n + T]$. In the Appendix, we show that the time τ that a drop projects onto a pixel is far less than T. Thus, the intensity I_d is a linear combination of the irradiance E_{bg} due to the background of the drop and the irradiance E_d due to the drop itself:

$$I_d(\vec{r}) = \int_0^{\tau} E_d \, dt + \int_{\tau}^{T} E_{bg} \, dt. \tag{6}$$

Here, we have dropped the parameters $(\vec{r}, t)$ on the right hand side for brevity. If the motion of the background is slow, E_{bg} can be assumed to be constant over the exposure time T. Then, the above equation simplifies to

$$I_d = \tau \bar{E}_d + (T - \tau) \, E_{bg} \ , \quad \bar{E}_d = \frac{1}{\tau} \int_0^{\tau} E_d \, dt \, , \tag{7}$$

where, $\bar{E}_d$ is the time-averaged irradiance due to the drop. For a pixel that does not observe a drop, we have $I_{bg} = E_{bg} \, T$. Thus, the change in intensity ΔI at a pixel due to a drop is

$$\Delta I = I_d - I_{bg} = \tau \, (\bar{E}_d - E_{bg}) \, . \tag{8}$$

Recall from Section 3.2.1 that raindrops are much brighter than their backgrounds. Thus, $\bar{E}_d > E_{bg}$ and ΔI is positive. By substituting $I_{bg} = E_{bg} \, T$ in equation (8), we obtain a relation between ΔI and I_{bg} as

$$\Delta I = -\beta \, I_{bg} + \alpha \, , \ \ \beta = \frac{\tau}{T} \ , \ \ \alpha = \tau \, \bar{E}_d. \tag{9}$$

In the Appendix, we derive the time τ for which a drop remains within a pixel as a function of the physical properties of the drop (size and velocity). Also, we show that τ and hence β are constant for all pixels within a streak. In addition, since the brightness of the (stationary) drop is weakly effected by the background intensity, the average irradiance $\bar{E}_d$ can be assumed to be constant for pixels that lie on the same streak (see Section 3.2.1). Thus, the change in intensities ΔI observed at all pixels along a streak are *linearly* related to the background intensities I_{bg} occluded by the streak.

In the Appendix, numerical bounds are also derived for the parameters β and τ. We show that the maximum value of τ is approximately 1.18 ms, which is much less than the typical exposure time $T \approx 30$ ms of a video camera. As a result, the slope β is shown to to lie within the range $0 < \beta < 0.039$. Based on these bounds, we make the following observations:

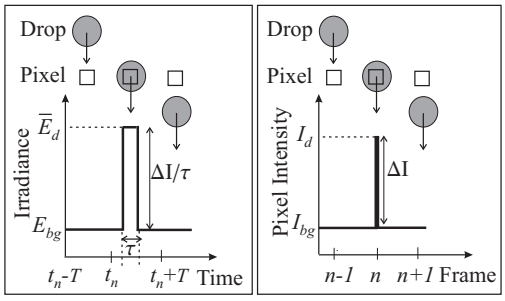

(a) (b)

Figure 5: The intensity change at a pixel due to a falling raindrop. (a) The average irradiance at the pixel due to the rain drop is $\bar{E}_d$ and that due to the background scene is E_{bg}. Note that $\bar{E}_d > E_{bg}$. The drop projects onto a pixel for time $\tau < 1.18$ ms, which is far less than the typical exposure time T of a camera. (b) Intensities of a pixel in three frames. A drop stays over the pixel in only a single frame and produces a positive intensity fluctuation of unit frame width.

- The time a drop stays at a pixel is less than the integration time of a typical video camera. Thus, a drop produces a positive intensity change ($\Delta I > 0$) of unit frame width at a pixel as illustrated in Figure 5(b).
- The change in intensities observed at all pixels along a rain streak are linearly related to the background intensities I_{bg} occluded by the streak. The slope β of this linear relation depends only on the physical properties of the raindrop. This can be used to detect rain streaks.

4 Detection of Rain in Videos

Based on the dynamics and photometric models of rain, we now develop a robust algorithm to detect (segment) regions of rain in videos. Although our models do not explicitly take into account scene motions, we will show that they provide strong constraints which are sufficient to disambiguate rain from other forms of scene motions.

4.1 Applying Constraints of Photometric Model

Consider a video of a scene captured in rain such as the one shown in Figure 6. We apply constraints derived using the photometric model to detect *candidate* pixels effected by rain in each frame of the video. In Section 3.2.2, it was shown that a drop produces a positive intensity fluctuation of unit frame duration. Hence, to find candidate rain pixels in the n^{th} frame, we need to only consider intensities I_{n-1}, I_n and I_{n+1} at each pixel corresponding to the 3 frames $n-1$, n and $n+1$, respectively (see Figure 5(b)). If the background remains stationary in these three frames[1], then the intensities I_{n-1} and I_{n+1} must be equal and the change in intensity ΔI due to the raindrop in the n^{th} frame must satisfy the constraint

$$\Delta I = I_n - I_{n-1} = I_n - I_{n+1} \geq c \, , \tag{10}$$

where c is a threshold that represents the minimum change in intensity due to a drop that is detectable in the presence of noise. The result of applying this constraint with $c = 3$ gray levels is shown in Figure 6(a). The selected pixels (white) include almost all the pixels effected by rain.

[1]When the background is not stationary over the three frames, we may miss some streaks. However, since we detect rain using a sufficient number of frames (say, 30 frames) missed rain pixels may be detected in other frames.

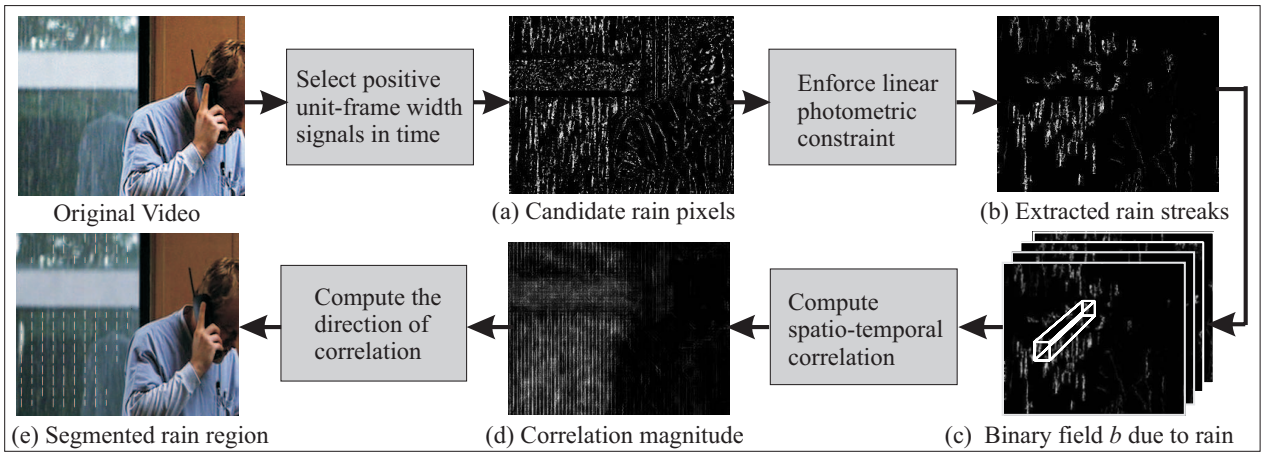

Figure 6: The rain detection algorithm applied to a video. (a) Using the photometric model to identify pixels that are effected by rain. Candidate rain pixels are shown in white. Note that there are several false positives (non-rain pixels). (b) Applying the linear photometric constraint to streak intensities (equation (9)). This reduces false positives significantly to yield an estimate of the binary rain field b. (c) The computation of spatio-temporal correlation in the binary field b. (d) A correlation map computed using 11×11 neighborhoods over 30 frames. Pixels with high intensity values represent rain pixels while the non-rain pixels have low intensity values. (e) Needle map showing the segmentation of the scene into regions with and without rain. The needle map is kept sparse for clarity.

In the presence of object motions in the scene, the above constraint also detects several false positives. Some of the false positives can be seen in and around the moving person in Figure 6(a). To reduce such false positives, we apply the photometric constraint in equation (9) as follows. For each individual streak[2] in frame n, we verify whether the intensity changes ΔI along the streak are linearly related to the background intensities I_{n-1}, using equation (9). The slope β of the linear fit is estimated. Then, streaks that do not satisfy the linearity constraint, or whose slopes lie outside the acceptable range of $\beta \in [0 - 0.039]$, are rejected. Figure 6(b) shows a significant decrease in false positives after applying this constraint. By applying these constraints to all the frames, an estimate of the binary rain field b is obtained (see Figure 6(c)).

4.2 Applying Constraints of Dynamics Model

Although a significant reduction in false positives is achieved using the photometric constraint, some false positives will remain. In this step, we further reduce the false positives using the dynamics model. In Section 3.1, we showed that in a binary field produced by rain strong temporal correlation exists between neighboring pixels in the direction of rain. Using the estimated binary field b, we compute the zeroth order temporal correlation R_b of a pixel with each of its neighbors in a local $(l \times l)$ neighborhood, over a set of frames $\{n, n-1, \ldots, n-f\}$. Figure 6(d) shows the correlation values obtained for all (11×11) neighborhoods in frame n, computed using the previous $f = 30$ frames. Bright regions indicate strong correlation. The direction and strength of correlation is computed for each neighborhood center which is depicted in Figure 6(e) as a needle map. The direction of the needle indicates the direction of correlation (*direction of the rainfall*) and its length denotes the strength of correlation (*strength of the rainfall*). The needle map is kept sparse for clarity. Weak and non-directional correlations occur at pixels with no rain[3] and hence are rejected. Thus, constraints of the photometric and dynamics models can be used to effectively segment the scene into regions with and without rain, even in the presence of complex scene motions.

5 Removal of Rain from Videos

Once the video is segmented into rain and non-rain regions, we apply the following simple method to remove rain from each frame of the video. For each pixel with rain in the n^{th} frame, we replace its intensity I_n with an estimate of the background obtained as $(I_{n-1} + I_{n+1})/2$ (see Figure 5(b)). This step removes most of the rain in the frame. However, since drop velocities are high compared to the exposure time of the camera, the same pixel may see different drops in consecutive frames. Such cases are not accounted for by our detection algorithm[4]. Fortunately, the probability of raindrops effecting a pixel in more than three consecutive frames is negligible. In the case of a pixel being effected by raindrops in 2 or 3 consecutive frames, we remove rain by assigning the average of intensities in the two neighboring pixel (on either side) that are not effected by raindrops. Our results show that this additional step can be very effective for rain removal.

Note that we only remove streaks that can be detected. Severely defocused streaks and streaks on bright backgrounds produce very small changes in intensities that are difficult to detect in the presence of noise. Hence, we are unable to remove such streaks. In addition, we do not handle the steady effects of rain in this work. Raindrops far from the camera are much smaller than a pixel. Hence, the intensity at a pixel is due to a large number of drops in the pixel's field of view. These aggregate effects are similar to the effects of fog. Hence, defogging algorithms [10] can be used to remove the steady effects of rain.

[2] A linear time sequential labeling algorithm (Matlab function "bwlabel") [5] is used to segment streaks in the binary video frames.

[3] Typically, the false positive streaks due to camera and object motions vary in direction over time and thus do not exhibit directional correlation.

[4] In detecting rain, we did not consider this case in order to minimize false detection due to object/camera motions.

6 Experimental Results

We conducted experiments on several videos with rain to demonstrate the effectiveness of our algorithms. In all experiments, we chose the photometric threshold $c = 3$ and the spatio-temporal correlation was computed using (11×11) neighborhoods over $f = 30$ frames. The rain segmentation has a temporal lag since 30 frames were used to compute spatio-temporal correlation.

Figure 7(a) shows 3 frames from the movie "Magnolia", where a person is moving and speaking over the phone. Rain is visible through the window. The camera moves and zooms in on the person. The detection task is challenging due to the fast moving textures (shirt creases and folds on the arm). Despite these complexities, our algorithm robustly detects only pixels with rain (see Figure 7(b)). Note that we are unable to detect rain in pixels with a bright background (white wall) because the changes in intensities produced by rain are very low. Derained frames are shown in Figure 7(c) and the differences between the derained and original frames (scaled by a constant) are shown in Figure 7(d). Similar results are shown for a clip from the movie "Seven" in Figure 8(I). Despite the fast back and forth motion of the person's head and the book, our algorithm detects only pixels effected by rain.

Figure 8(II) shows results for a scene with raindrops falling and forming a pool of water. The ripples of water may be viewed as a temporal texture with frequencies similar to those produced by rain. Even in this case, the algorithm detects only pixels with rain (see Figure 8(II)(b)). The derained frames are shown in Figure 8(II)(c). These examples demonstrate that our algorithm is effective for scenes with complex motions and at the same time is insensitive to time-varying textures that have temporal frequencies similar to those due to rain.

7 Conclusion

We have developed a comprehensive model for the visual appearance of rain. Based on this model, we presented efficient algorithms for the detection and removal of rain from videos. Note that simple temporal filtering methods are not effective in removing rain since they are spatially invariant and hence degrade the quality of the image in regions without rain. In contrast, our method explicitly detects pixels effected by rain and removes the contribution of rain only from those pixels, preserving the temporal frequencies due to object and camera motions. The proposed algorithm can be used in a wide range of applications including video surveillance, video/movie editing, and video indexing/retrieval. Our models also have implications for efficient and realistic rendering of rain. Currently, we do not handle the steady effects of rain and we do not remove severely defocused rain streaks. In future work, we wish to address these issues and also extend our analysis to other types of dynamic weather conditions such as snow and hail.

References

[1] K.V. Beard and C. H. Chuang. A New Model for the Equilibrium Shape of Raindrops. *Journal of Atmospheric Science*, 44(11):1509–1524, 1987.

[2] G. Doretto, A. Chiuso, Y.N. Wu, and S. Soatto. Dynamic textures. *IJCV*, 51(2):91–109, 2003.

[3] K. Garg and S. K. Nayar. Photometric Model for Raindrops. *Columbia University Technical Report*, 2003.

[4] R. Gunn and G.D. Kinzer. Terminal Velocity for Water Droplet in Stagnant Air. *Journal of Meterology*, 6:243–248, 1949.

[5] B.K.P. Horn. *Robot Vision*. The MIT Press, 1986.

[6] R.M Manning. *Stochastic Electromagnetic Image Propogation*. McGraw-Hill, Inc, 1993.

[7] J.S Marshall and W.M.K Palmer. The Distribution of Raindrops with Sizes. *Journal of Meterology*, 5:165–166, 1948.

[8] B.J Mason. *Clouds, Rain and Rainmaking*. Cambridge Press, 1975.

[9] E.J. McCartney. *Optics of the Atmosphere: Scattering by molecules and particles*. John Wiley and Sons, 1975.

[10] S.G. Narasimhan and S.K. Nayar. Vision and the Atmosphere. *IJCV*, 48(3):233–254, August 2002.

[11] W. T. Reeves. Particle System- a techinque for modeling a class of fuzzy objects. *Computer Graphics (SIGGRAPH 83 Conference Proceedings)*, 17, 1983.

[12] K. Starik and M. Werman. Simulation of Rain in Videos. *Texture Workshop, ICCV*, 2002.

[13] T. Wang and R.S Clifford. Use of Rainfall-Induced Optical Scintillations to Measure Path-Averaged Rain Parameters. *JOSA*, 8:927–937, 1975.

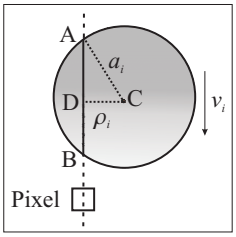

Figure 9: Schematic to compute the time τ that a drop stays over a pixel

A Bounds for Photometric Model Parameters

We derive an expression for the time τ that a drop stays over a pixel. Figure 9 shows a drop passing through a pixel. Consider the right angled triangle $\triangle ADC$. Then, τ is given by,

$$\tau = \frac{AB+1}{v_i} = \frac{2\sqrt{a_i^2 - \rho_i^2}+1}{v_i} < \frac{2a_i+1}{v_i}, \qquad (11)$$

since $0 < \rho_i \leq a_i$. A streak is visible for a drop that projects to a region larger than a pixel i.e $2\,a_i > 1$. Then, we obtain a conservative upper bound $\tau < 4\,a_i/v_i$. If f is the effective focal length of the camera and z is the distance of the drop from the camera, then τ can be related to the physical properties of the drop (size a and fall velocity v) by substituting $a_i = \frac{f\,a}{z}$, and $v_i = \frac{f\,v}{z}$, to get the range $0 < \tau < 4\,a/\,v$. Substituting v from equation (1), we obtain range of possible values for τ (and hence for $\beta = \tau/T$) in terms of drop size a as:

$$0 < \tau < \sqrt{a}/50 \;\; sec \;\;, \;\; 0 < \beta < \sqrt{a}/(50\,T) \;, \qquad (12)$$

where, T is the exposure time, typically $1/30$ sec for a video camera. Since the maximum value of $a = 3.5 \times 10^{-3}$ m, the possible ranges for τ and β can be obtained as $0 < \tau < 1.18$ ms and $0 < \beta < 0.039$. Now let us consider the dependence of τ on ρ_i for all pixels within a streak. From Figure 9, we see that ρ_i does not change along the direction of the streak. Further, since the width of the streak (image size of the drop) is only a few pixels, the change in ρ_i across the streak is also negligible. Thus, the dependence of τ on ρ_i is weak and can be neglected. In other words, τ and β are constant for all pixels on a streak.

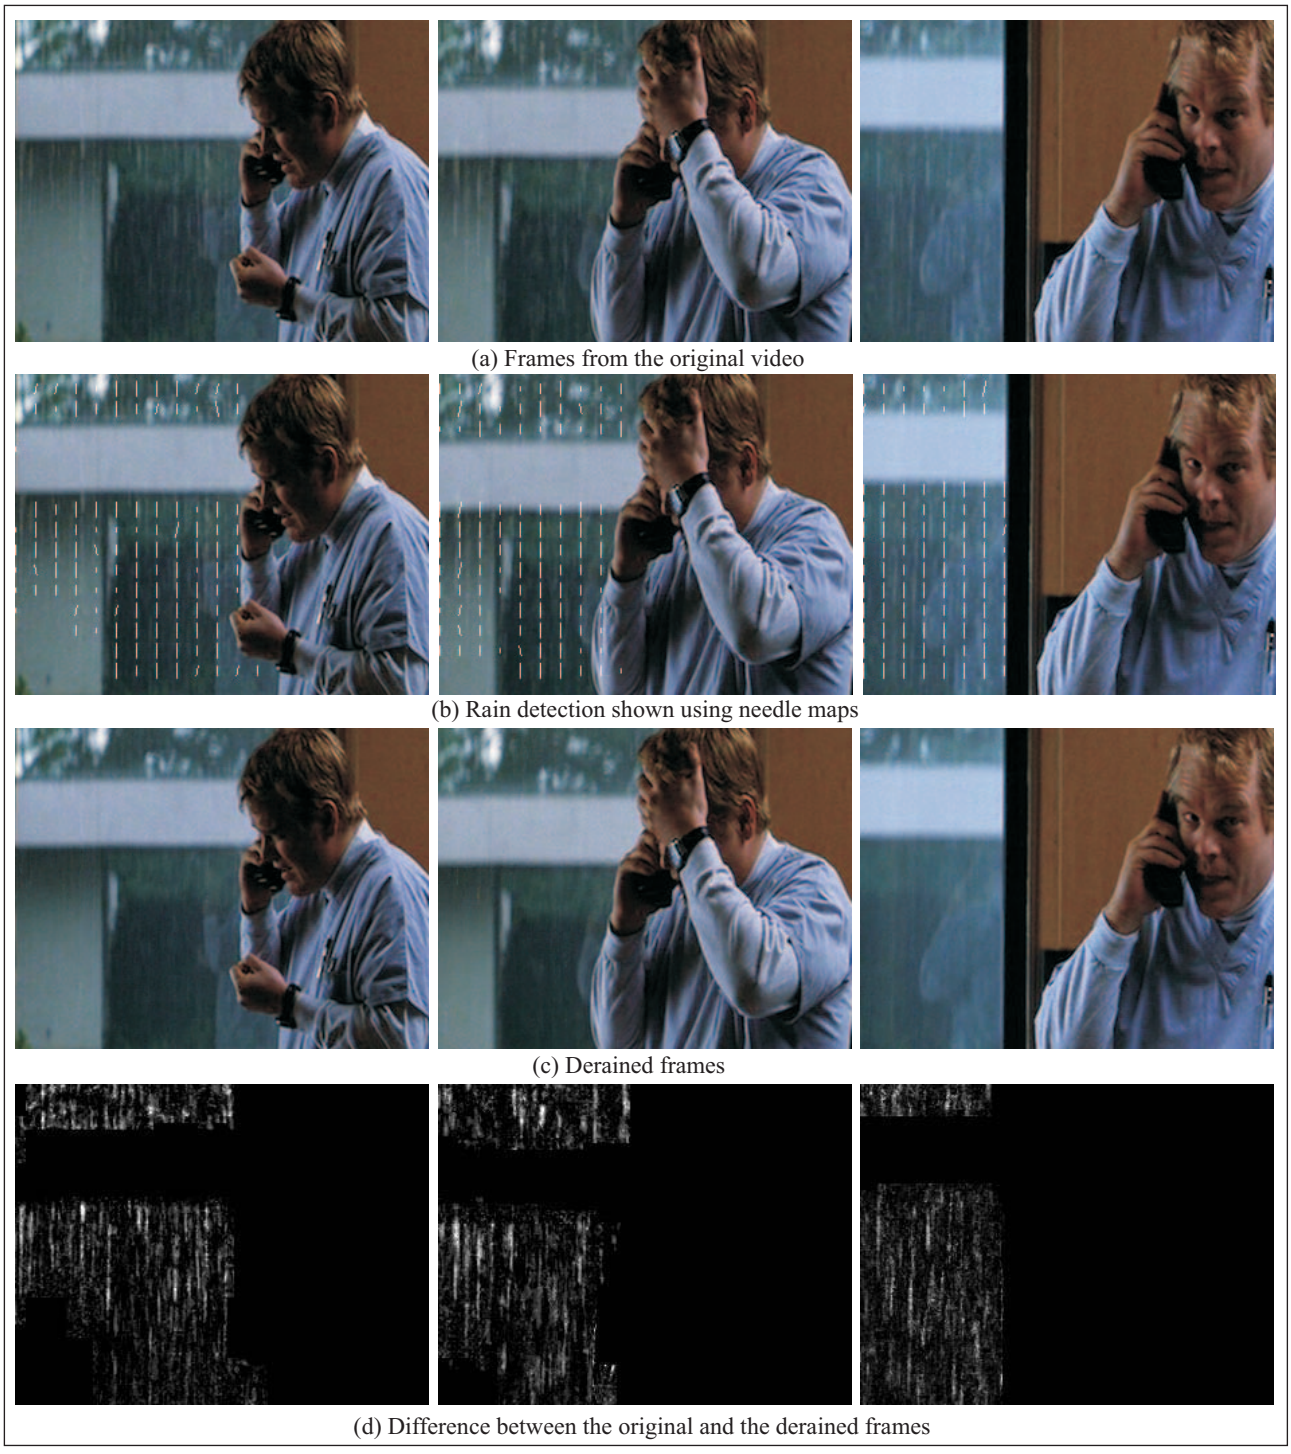

(a) Frames from the original video

(b) Rain detection shown using needle maps

(c) Derained frames

(d) Difference between the original and the derained frames

Figure 7: Detection and removal of rain for a clip of the movie "Magnolia" (Original images courtesy of New Line Productions-©1999 New Line Productions, Inc.). (a) Three frames from the original video. The scene consists of a person moving and speaking on the phone and rain visible through the glass window. The camera pans and zooms in on the person. Fast moving textures (shirt creases and folds on the moving arm) make the detection task challenging. (b) Detected regions with rain represented as a needle map. The needle map is kept sparse for clarity. The direction of the needle at a pixel represents the direction of rain and the length of needle represents the strength of the rain. (c) Derained frames. (d) Difference between the original frames and the derained frames. For illustration purposes, the difference images have been scaled by a constant.

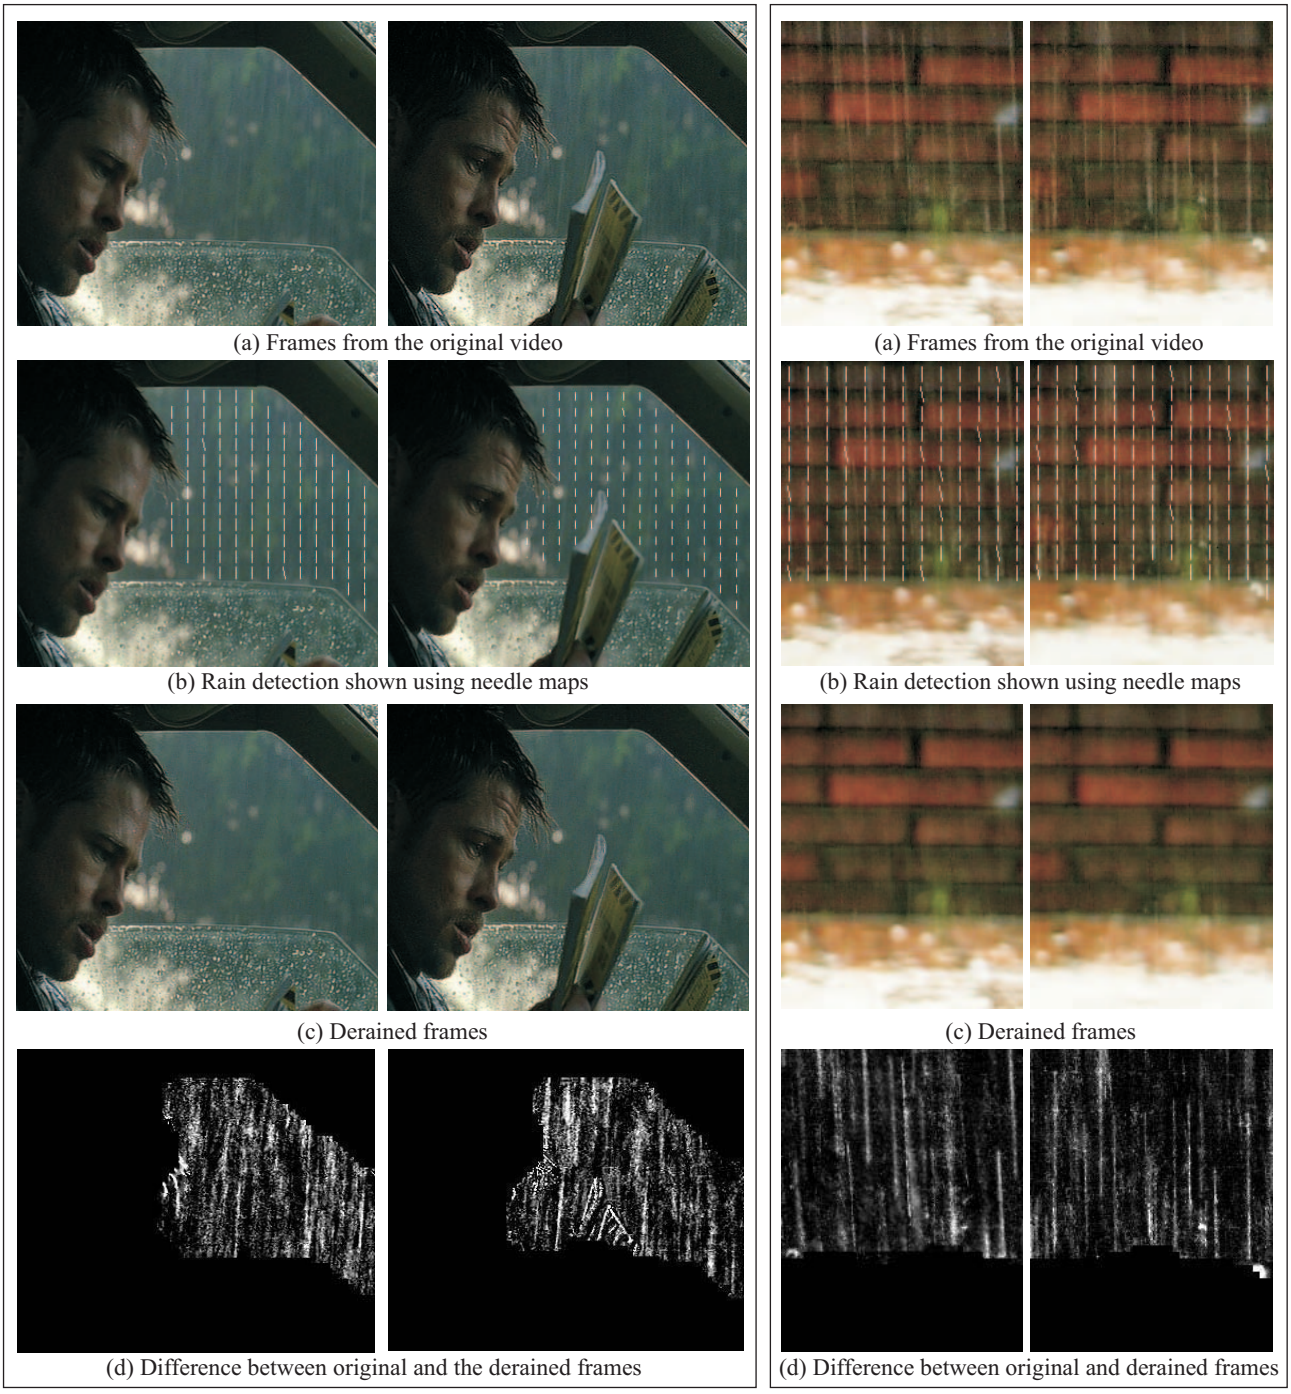

Figure 8: (I) Rain detection and removal for a clip of the movie "Seven" (Original images courtesy of New Line Home Entertainment -©2001 New Line Home Entertainment, Inc.). Despite the fast back and forth motion of the head and the book, the algorithm detects only pixels that have rain. When viewing on a screen, please zoom in to see the subtle rain streaks in the original frames. (II) In this example the rain causes ripples in the pool of water. These ripples may be viewed as a temporal texture with frequencies similar to those produced by rain. However, since the dynamic and photometric properties of rain are very different from those of the ripples, our algorithm detects only pixels with rain. For illustration purposes, the difference frames have been scaled by a constant factor.