

Rendering falling snow using an inverse Fourier transform

Michael S. Langer *

School of Computer Science, McGill University

Linqiao Zhang †

School of Computer Science, McGill University

Keywords: falling snow, non-photorealistic rendering

Methods for rendering falling snow typically use particle systems [Reeves 1983] which requires tens of thousands of particles (snowflakes) and thus can be expensive. Here we present an alternative method for rendering falling snow which does not use particles, and instead uses an inverse Fourier transform.

Our idea is based on a well-known fact that *pure image translation* with velocity (v_x, v_y) pixels/frame produces a plane of energy,

$$\omega_t = -v_x \omega_x - v_y \omega_y$$

in the 3D frequency domain [Watson and Ahumada 1985]. Falling snow differs from pure image translation in that falling snow produces motion parallax: the 2D speed and size of each moving object (snowflake) is determined by its 3D depth. Snowflakes that are further away from the viewer appear smaller in the image plane but also move more slowly in the image plane.

The correlation between size and speed of falling snowflakes can be captured in the frequency domain. The distance d to a snowflake is proportional to spatial frequency $\sqrt{\omega_x^2 + \omega_y^2}$. Distance d is also inversely proportional to the speed ω_t/ω_y , where we assume (temporarily) that the motion is in the y direction. This gives a relation between spatial and temporal frequency:

$$\omega_t = c \frac{\omega_y}{\sqrt{\omega_x^2 + \omega_y^2}} \quad (1)$$

We render falling snow in two steps. The first step is to generate *in the frequency domain* a set of surfaces of the form of Eq. (1), for a small range of constants c . The power spectrum for these surfaces is chosen as follows. For an image of spatial size 512×512 , we limit the power spectrum to three octaves, ranging from 16 to 128 cycles per frame. Spatial frequencies lower than 16 cycles per frame are not used in order to enforce an upper bound on size of moving image structure – that is, snowflakes are small. Spatial frequencies above 128 cycles per frame are not used in order to stay far from the Nyquist limit. For spatial frequencies between 16 and 128 cycles per frame, we assign power proportional to $\frac{1}{\sqrt{\omega_x^2 + \omega_y^2}}$. This puts a constant amount of power in each constant octave band [Field 1987] over the bands used. For each spatial frequency, we then randomize the phase, subject to the conjugacy constraint [Bracewell 1965] for a function $I(x, y, t)$ and its 3D Fourier transform, $\hat{I}(\omega_x, \omega_y, \omega_t)$, namely:

$$\hat{I}(\omega_x, \omega_y, \omega_t) = \overline{\hat{I}(-\omega_x, -\omega_y, -\omega_t)}$$

The second step is to take the inverse 3D Fourier transform of $\hat{I}(\omega_x, \omega_y, \omega_t)$, giving $I(x, y, t)$. We rescale $I(x, y, t)$ to have values in $[0, 1]$, and treat the scaled result as an opacity function $\alpha(x, y, t)$.

We use $\alpha(x, y, t)$ to composite a constant intensity snow layer over a static background image:

$$I(x, y, t) = I_{amb} \alpha(x, y, t) + (1 - \alpha(x, y, t)) I_{bg}(x, y)$$

where I_{amb} is the intensity of the snow. We use $I_{amb} = 250$.

Since the computed opacity function $\alpha(x, y, t)$ is composed of independent sine waves, it is periodic. This allows us to reduce computational cost by rendering a smaller $\alpha(x, y, t)$ – say $128 \times 128 \times 64$ – tiling it to create a larger snowing scene, and looping it to make a longer video.

Fig. 1 shows an example of a background image, along with one frame of the falling snow image sequence. The entire image sequence is shown on the enclosed CD ROM. The sequence should be viewed in loop mode.



Figure 1: Painting *Jimmy's place* (by Gary Johnson) along with one frame of the falling snow video.

Acknowledgements

This research was supported by grants from NSERC and FCAR.

References

- BRACEWELL, R. N. 1965. *The Fourier Transform and Its Applications*. NY:McGraw-Hill.
- FIELD, D. J. 1987. Relations between the statistics of natural images and the response properties of cortical cells. In *Journal of the Optical Society America A*, 4:2370 – 2393.
- REEVES, W. T. 1983. Particle system — a technique for modeling a class of fuzzy objects. In *Proc. of SIGGRAPH '83, Detroit, Michigan*.
- WATSON, A. B., AND AHUMADA, A. J. 1985. Model of human visual motion sensing. In *Journal of the Optical Society of America A*, 2(2):322 – 342.

*e-mail: langer@cim.mcgill.ca

†e-mail: qiao@cim.mcgill.ca