to maximizing the spectral content of $\|\tilde{\hat{\phi}}\|^2$ within the wedge of revolution.

We can derive an upper bound on the worst-case spectral content of any camera motions. The amount of photon energy collected by a camera within a fixed exposure time T is bounded. Therefore, by Parseval's theorem, the norm of every ω_{x_0,y_0} slice of $\tilde{\hat{\phi}}$ (i.e. $\tilde{\hat{\phi}}(\omega_{x_0}, \omega_{y_0}, \omega_t)$) is bounded [49]:

$$\int \|\tilde{\hat{\phi}}(\omega_{x_0}, \omega_{y_0}, \omega_t)\|^2 d\omega_t \leq T \tag{1.11}$$

Every ω_{x_0,y_0}-slice intersects the wedge of revolution for a segment of length $2S_{obj}\|\omega_{x_0,y_0}\|$. To maximize the worst-case spectral power, the optimal camera would spread the captured energy uniformly in this intersection. Therefore, we can derive an upper bound on the worst-case spectral power by dividing the captured energy by the segment length:

$$\min_{\omega_t} \|\tilde{\hat{\phi}}(\omega_{x_0}, \omega_{y_0}, \omega_t)\|^2 \leq \frac{T}{2S_{obj}\|\omega_{x_0,y_0}\|} . \tag{1.12}$$

Since the PSFs spectra $\hat{k}^j_{s_{x,y}}$ are slices through $\hat{\phi}^j$, this bound also applies for the PSFs' spectral power:

$$\min_{s_{x,y}} \|\tilde{\hat{k}}_{s_{x,y}}(\omega_{x_0}, \omega_{y_0})\|^2 \leq \frac{T}{2S_{obj}\|\omega_{x_0,y_0}\|} . \tag{1.13}$$

The optimal bound Eq. 1.12 applies to any types of integration kernel ϕ regardless of the number of shots taken during the time budget T.

■ 1.1.2 Orthogonal parabolic motions

We seek a motion path whose spectrum covers the wedge of revolution and approaches the bound in Eq. 1.12. We also seek to cover the spectrum with the fewest images, because as we take more images within the time budget, the delay between subsequent shots reduces the effective time budget, degrading the spectral performance.

We could compute the optimal camera motion by inverting the Fourier transform of the bound in Eq. 1.12. However, the inverse Fourier transform of this bound is not a physically valid motion of the form in Eq. 1.7. To illustrate this, we invert the bound for 1D motions [49] in Fig. 1.3. We can see that the corresponding optimal motion in the spatial domain is not a realizable motion. If we invert the bound for 2D motions in Eq. 1.12, we observe the same phenomenon in 3D: the optimal path is not a realizable motion.

Our solution captures two images of a scene with two orthogonal parabolic motions. We show analytically that the orthogonal parabolic motions capture the wedge of revolution with

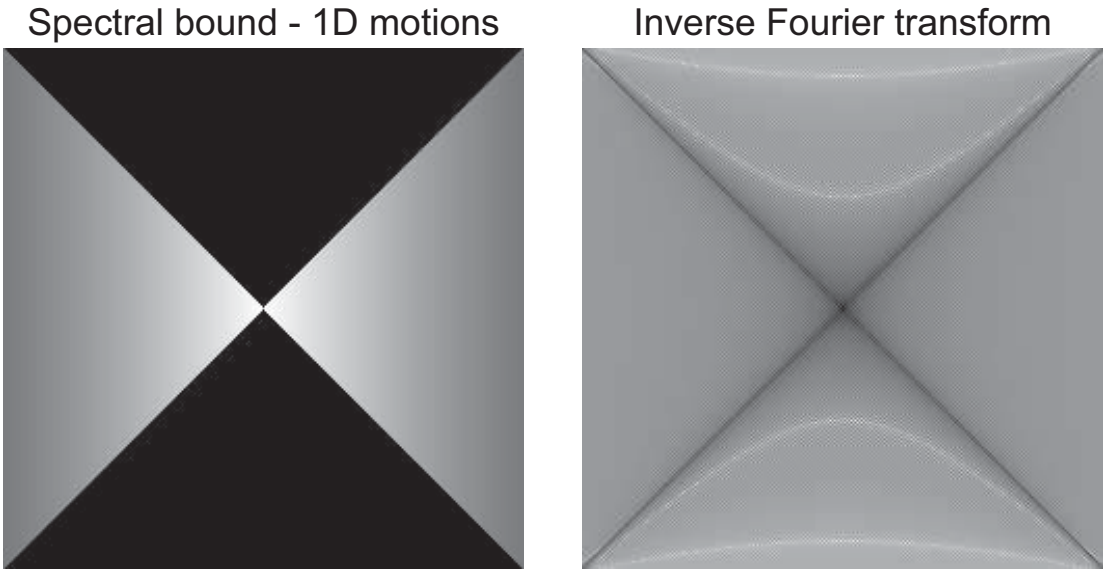

Figure 1.3: *We explicitly invert the spectral bound for 1D motions to illustrate that the explicit inversion of the spectral bound in Eq. 1.12 does not result in a physically realizable motion of the form in Eq. 1.7.*

the worst-case spectral power greater than $2^{-1.5}$ of the upper bound.

Camera motion Let ϕ_1, ϕ_2 be the 3D integration kernels of x and y parabolic camera motions. The kernels are defined by the integration curves f_1, f_2:

$$\begin{aligned} f_1(t) &= [a_x(t+T/4)^2, 0, t], \quad t = [-T/2...0] \\ f_2(t) &= [0, a_y(t-T/4)^2, t], \quad t = [0...T/2] \end{aligned} \tag{1.14}$$

At time t, the derivative of the x-parabolic camera motion $f_1(t)$ is $2a_x(t-T/4)$, therefore the camera essentially tracks an object moving with velocity $2a_x(t-T/4)$ along x axis. During exposure, the x-parabolic camera tracks *once* every moving object with velocity within the range $[-2a_xT/4...2a_xT/4]$. Similarly, the y-parabolic camera covers the velocity range $[-2a_yT/4...2a_yT/4]$. For the reason that will be clarified below, we set

$$a_x = a_y = \frac{2\sqrt{2}S_{obj}}{T} \tag{1.15}$$

The maximal velocity of the sensor becomes $S_{sens} = \sqrt{2}S_{obj}$. That is, the velocity range covered by these parabolas is $[-S_{sens}...S_{sens}]$.

the power spectrum is constant along ω_x and ω_y:

$$\|\hat{\phi}^{static}(\omega_x,\omega_y,\omega_t)\|^2 = T^2 sinc^2(\omega_t T) \tag{1.20}$$

The Fourier transform of the PSF is a slice of the 3D spectrum $\hat{\phi}$, and is a sinc whose width depends on the object velocity $\|\hat{k}^{static}_{s_{x,y}}\|^2 = T^2 sinc^2((s_x\omega_x + s_y\omega_y)T)$. For fast object motions this sinc highly attenuates high frequencies. In fact, if the object motion is fast it is better to reduce the exposure time (this increases the width of the sinc) despite reducing the total amount of energy collected.

A flutter shutter camera In a flutter shutter camera [63], the integration curve of a static camera is temporally modulated (Fig. 1.1(b)). Therefore, the spectrum of the integration curve $\phi^{flutter}$ is constant along ω_x, ω_y and is modulated along ω_t:

$$\|\hat{\phi}^{flutter}(\omega_x,\omega_y,\omega_t)\|^2 = \|\hat{m}(\omega_t)\|^2 \tag{1.21}$$

where $\hat{m}$ is the Fourier transform of the shutter code. This code can be designed to be more broadband than the sinc function in a static camera. However, the spectrum is constant along ω_x, ω_y. Therefore, the worst-case spectral performance is bounded as follows:

$$\min_s \|\hat{k}^{flutter}_s(\omega_x,\omega_y)\|^2 = T/(2S_{obj}\Omega) \tag{1.22}$$

for all (ω_x, ω_y) [49], where Ω is the spatial bandwidth of the camera. As a result, the flutter shutter poorly captures the low frequency image contents.

A linearly moving camera If the camera is moving at a constant velocity (Fig. 1.1 third column), the integration curve is a slanted straight line $\phi^{linear}(t) = [s_x t, s_y t, t]$ (Fig. 1.1(c)). By linearly moving the camera, we can track the object that moves at the camera's speed, but we still suffer from a sinc fall-off for objects whose velocities are different from the camera's velocity.

A camera with parametric motions We design other cameras with parametric motions and analyze their performance numerically. Although we cannot analytically compute the spectral performance of each camera, we observe in the Fourier domain that the blur kernels contain zeros. We define four cameras: (i) a camera with a raised-cosine motion, (ii) a camera with a circular motion with a constant angular speed, (iii) a camera with a circular motion with a

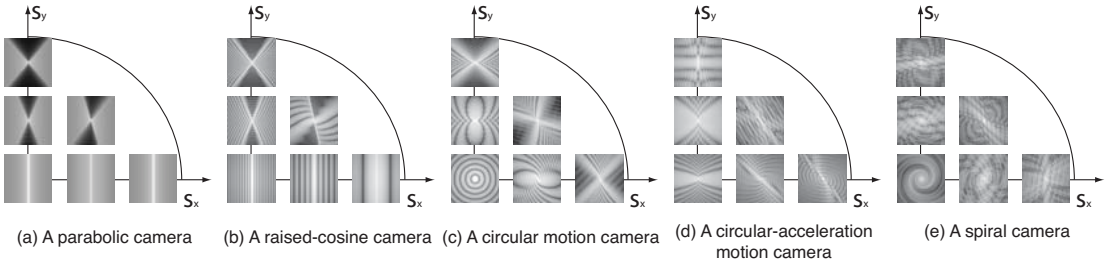

Figure 1.7: *We numerically analyze the spectral performance of five different cameras: (i) a camera with a raised-cosine motion, (ii) a camera with a circular motion with a constant angular speed, (iii) a camera with a circular motion with a constant angular acceleration, and (iv) a camera with a spiral motion. This figure shows that even a two-image solution of each of these cameras cannot capture all frequencies without zeros for all object motions.*

constant angular acceleration, and (iv) a camera with a spiral motion. Each camera moves in the defined camera path ϕ during the exposure $t \in [-T/2...T/2]$.

(i) A camera with a raised-cosine motion: The parametric motion that defines this camera is as follows:

$$\{\phi : [x,y,t] = [\alpha\,(1+cos(\omega t))\,,0,t]\} \tag{1.23}$$

where $\omega = \pi/(T/2)$ and $\alpha = S_{obj}/\omega$. This camera moves in 1D and covers each velocity *twice* as opposed to once in a parabolic camera. Fig. 1.7(b) shows the blur kernel spectra for different object motions. The zero pattern is quite similar to that of a parabolic camera in Fig. 1.7(a), but the zeros also appear in frequencies that are well covered by a parabolic camera. From this, we argue that even a two-image solution of a raised-cosine camera cannot cover all frequencies without zeros, as in an orthogonal parabolic camera pair.

(ii) A camera with a circular motion with a constant angular speed: We can define the motion of this camera as

$$\{\phi : [x,y,t] = [\alpha\cos(\omega t), \alpha\sin(\omega t), t]\} \tag{1.24}$$

where $\omega = \pi/(T/2)$ and $\alpha = S_{obj}/(2\omega)$. The camera sensor moves along a circular path at a constant speed ($S_{obj}/2$), and therefore the camera essentially tracks each object motion with this speed in all possible angles. However, this camera fails to capture other motions, and therefore we observe many zeros in blur kernel spectra, as shown in Fig. 1.7(c).

(iii) A camera with a circular motion with a constant angular acceleration: We can modify the above camera to track each object speed twice, in two different orientations. The

idea is to move the camera circularly but at a constant angular acceleration:

$$\left\{\phi : [x,y,t] = [\alpha\cos(\omega t^2), \alpha\sin(\omega t^2), t]\right\} \tag{1.25}$$

where $\omega = \pi/(T/2)^2$ and $\alpha = S_{obj}/(2\omega)$. While this camera performs well for many object velocities, the blur spectra still contain many zeros, as shown in Fig. 1.7(d).

(iv) A camera with a spiral motion:

$$\{\phi : [x,y,t] = [\alpha t\cos(\omega t), \alpha t\sin(\omega t), t]\} \tag{1.26}$$

where $\omega = k\pi/(T/2)$ and $\alpha = S_{obj}/\sqrt{(1+\omega^2T^2/4)}$. This camera tracks each speed once during exposure, but not in all directions. Fig. 1.7(e) shows that blur kernels for different object velocities contain substantial amount of zeros.

Two shots Taking two images with cameras defined above can simplify the kernel estimation task, but it does not substantially enhance the spectral coverage of these cameras. Optimizing the exposure lengths of each shot [3], and in the case of a flutter shutter camera also optimizing the random codes in each shot, do not eliminate their fundamental limitations: their power spectra are constant along $\omega_{x,y}$ and hence they spend the energy budget outside the wedge of revolution. Previous two-image solutions to deblurring, such as [13, 15, 64, 90], fall into the category of taking two images with a static or a linearly moving camera. These methods can correctly find the motion kernels, but the image reconstruction quality is limited since the spectrum coverage is low.

A parabolic camera with a single exposure Blur kernels from a single-exposure parabolic camera are invariant to 1D constant-velocity motions, and can be shown to approach the optimal bound for a set of 1D linear motions with bounded speed [49]. A single parabolic camera, however, is neither motion invariant nor optimal for 2D motions. When an object moves in a direction orthogonal to the camera movement axis (i.e. a x-parabolic camera imaging an object moving in the y direction), the spectral coverage along the orthogonal frequencies (i.e. ω_y) is poor. We have shown in Fig. 1.1 (n-o) several Fourier slices in which the captured spectra contain zeros.

Synthetic simulation We compare the deblurring performance of (i) a pair of static cameras, (ii) a pair of flutter shutter cameras, (iii) a single parabolic camera and (iv) an orthogonal parabolic camera through simulations (Fig. 1.8). For all cameras, we fix the total exposure time