INFINITE PHOTOGRAPHY: NEW MATHEMATICAL MODEL FOR HIGH-RESOLUTION IMAGES

Tapio Helin

Department of Mathematics and System Analysis
Helsinki University of Technology
P.O. Box 1100 (Otakaari 1 M), FI-02015 TKK, Finland

Matti Lassas, Samuli Siltanen

Department of Mathematics and Statistics
University of Helsinki
P.O. Box 68 (Gustaf Hallstromin katu 2b) FI-00014, Finland

Abstract. A new mathematical image model is introduced using the photographic process as the starting point. Images are represented as infinite sequences of photons allowing analysis at arbitrarily high resolution and leading to novel computational methods for processing, representation, transmission and storage of images. The resulting space of infinite photographs is proved to have a metric structure and to be intimately connected with bounded Borel measures. Theorems are proved indicating that the imaging power of the photographic process exceeds function spaces in the high resolution limit; this implies in particular that natural photographic images need to be modelled by generalized functions. Furthermore, computational results are presented to illustrate the novel algorithms based on photon sequences. The algorithms include stochastic halftoning, representation of cartoon images with outlines, anti-aliasing, blurring and singularity extraction.

Contents

1. Introduction 1
2. The metric space of infinite photographs 3
2.1. Resolutions and pixels 3
2.2. Pixel images from exposures 4
2.3. The snapshot metric 6
2.4. The space P of infinite photographs 7
3. Infinite photographs and Borel measures 8
3.1. Mapping exposures to measures 8
3.2. Mapping measures to exposures 10
3.3. The set of regular exposures 11
3.4. Embedding finite photographs to P 13
4. Decomposition and fractal analysis of images 15
4.1. Singularity indication and local dimension 16
4.2. Photon-based edge detection 18
5. Computational methods 19
5.1. Monte Carlo simulation of the illumination map 19
5.2. Halftoning with black photons 20
5.3. Modelling cartoon images including outlines 22
5.4. Photon-based Gaussian blurring 22
5.5. Singularity extraction 22
6. Acknowledgements 24
References 24

1. Introduction

We introduce a new mathematical image model based on the physical process of photography. Images are represented as infinite sequences of photons, allowing analysis at arbitrarily high resolution and leading to novel computational methods for the analysis, processing, representation, transmission and storage of images. Moreover, the model provides new understanding of the structure of photographs: they are finite approximations of Borel measures. This result is in line with the statistical evidence presented by Mumford and Gidas [28] suggesting that digital photographs of natural scenes need to be modelled by more general objects than functions.

The starting point of our approach is a simplified model of a monochromatic digital camera comprising a light-sensitive surface called *sensor*, an optical arrangement, and a shutter for preventing light from passing through the optics. Photographs are taken by opening the shutter for a suitable period of time, allowing a large number of photons to arrive at the sensor which is divided into square-shaped disjoint subsets called *pixels*. Each pixel then measures the wavelength-dependent energy delivered by the photons, and the gray level of the pixel is determined by the ratio between energy detected at that pixel and the total energy received by the whole sensor.

We model the sensor surface as a unit square. An *exposure* is an infinite ordered sequence $e = (e_1, e_2, e_3, \dots)$, where each $e_j = (z_j, \lambda_j)$ models a photon with wavelength $\lambda_j > 0$ arriving at the sensor at the location z_j. When a digital photograph is taken, the first J photons in an exposure e reach the sensor. The resulting gray level of a pixel D is

$$\frac{\mathcal{E}_g(e, J, D)}{J|D|}, \tag{1.1}$$

where $|D|$ denotes the area of the pixel and

$$\mathcal{E}_g(e, J, D) := \sum_{\substack{z_j \in D \\ 1 \le j \le J}} g(\lambda_j) \tag{1.2}$$

is the energy detected at the pixel D. Here $g : \mathbb{R} \to \mathbb{R}$ is the *spectral sensitivity function* of the sensor satisfying $0 \le g(\lambda) \le M$ with some sensor-dependent constant $M < \infty$.

The above is a mathematical model of digital black-and-white photography based on finite exposures. Let us make two observations clarifying why we need to consider *infinite* sequences of photons.

First, images formed by practical sensors always contain some error. Typical sensor noise can be modeled as a combination of quantum mechanical effects in the photon counting process and normally distributed additive random error arising from sensor electronics. Assume for the sake of the argument that the latter is negligible. Then the photon count at each pixel is accurately modeled by a Poisson distributed random variable with expectation σ and standard deviation $\sqrt{\sigma}$, and the signal-to-noise ratio $\sigma/\sqrt{\sigma} = \sqrt{\sigma}$ tends to infinity as σ grows. Hence, the ideal limit image formed by infinitely many photons is free of noise and therefore an interesting object to study.

Second, consider taking a correctly exposed photograph of a natural scene with a digital camera. Suppose we want a higher-resolution picture of the same scene, and to that end we use another camera with identical optics and same-size sensor divided into more pixels. Then, according to the first observation above, we need to extend the exposure time to get an image with the same signal-to-noise ratio as the first one. Therefore, allowing infinite sequences of photons is necessary for the analysis of photographs at arbitrarily high resolutions.

The gray level resulting from an infinite number of photons is

$$\lim_{J\to\infty} \frac{\mathcal{E}_g(e, J, D)}{J|D|}. \tag{1.3}$$

We construct a sequence of successively finer pixelizations of the sensor and define digital images formed by e at any resolution using formula (1.3) in each pixel. A concept of distance between two exposures is defined by comparing their pixel images at all resolutions simultaneously, leading to the *metric space of infinite photographs* consisting of equivalence classes of exposures: two exposures belong to the same class if and only if their pixel images coincide at all resolutions.

Our infinite photography model is inspired by the physics of counting photons with digital detectors. However, while the mathematical model applies to arbitrarily high resolutions, it does not properly describe quantum mechanical effects at sub-nanometer scale. The value of our model is in (i) deeper understanding of the structure of digital photographs provided by asymptotic analysis, and in (ii) new image processing methods that can be applied consistently at any given finite resolution.

Concerning (i), we prove that a set of uniformly bounded Borel measures can be identified with a subset of infinite photographs. In particular, our result implies that photographs need to be modelled by generalized functions in the high resolution limit. Also, our mathematical construction using infinite sequences of photons comes with a price:

it turns out that there exist irregular photon sequences that produce unexpected limit images. However, excluding such anomalies leads to a set of *regular exposures* that can be used to define an embedding of finitely exposed photographs as a subset of infinite photographs.

Concerning (ii), we describe in Section 5 practical algorithms based on infinite photography. Pixel images and more general Borel measures can be mapped to exposures using a Markov chain Monte Carlo algorithm, leading to a scalable image representation. Further, applying this method to a negative image and interpreting photons as spots of ink constitutes a novel and flexible stochastic halftoning method applicable in graphical printing technology. Moreover, the familiar techniques of anti-aliasing and blurring allow photon-based implementations, and the Radon-Nikodym decomposition of Borel measures can likewise be turned into a new method for extracting singularities.

Although not discussed further in this paper, our Monte Carlo algorithm can be used for creating useful distributions of photons for virtual illumination in computer graphics, see [15]. Also, infinite photography may prove useful for the development of active detectors based on controlling the movement of photons, see for instance [5].

Our image model provides a new contribution to the tradition of modeling photographs with more general objects than functions. In addition to [28] mentioned above, extensive analysis and some numerical algorithms are available for Meyer's image model based on dual spaces [25, 13, 4]. Also, our work is connected via the Radon-Nikodym decomposition to $u + v$ models dividing images into two parts such as geometry and texture [3, 33, 39, 21, 11].

2. The metric space of infinite photographs

We present a mathematically rigorous definition of the metric space of infinite photographs. The construction involves adding some technical details to the streamlined explanation in the introduction.

2.1. **Resolutions and pixels.**

We take $\mathbb{D}$ to be the two-dimensional torus constructed by gluing together sides of the square $[0, 1]^2$ appropriately. More specifically, let us periodize the plane $\mathbb{R}^2$ by identifying all points of the form $(x + m, y + n)$ with arbitrary $m \in \mathbb{Z}$ and $n \in \mathbb{Z}$. We will identify points $z \in \mathbb{D}$ with planar points $(x, y) \in \mathbb{R}^2$ satisfying $0 \le x < 1$ and $0 \le y < 1$. Considering the torus instead of the square simplifies the computational treatment of sensor boundaries.

Let us introduce a series of pixelizations of $\mathbb{D}$ at various resolutions. For each resolution parameter $n = 0, 1, 2, \dots$ define pixels

$$D_n(k,\ell) := \left\{ (x,y) \;\middle|\; \begin{array}{rcccl} (\ell-1)2^{-n} & \le & x & < & \ell\, 2^{-n}, \\ 1-k\,2^{-n} & \le & y & < & 1-(k-1)2^{-n} \end{array} \right\}$$

for $1 \le k \le 2^n$ and $1 \le \ell \le 2^n$. For fixed $n \ge 0$ the pixels are disjoint:

$$D_n(k,\ell) \cap D_n(k',\ell') = \emptyset \quad \text{if } k \neq k' \text{ or } \ell \neq \ell',$$

and they cover the image domain completely:

$$\mathbb{D} = \bigcup_{k=1}^{2^n} \bigcup_{\ell=1}^{2^n} D_n(k,\ell).$$

In particular we have $D_0(1,1) = \mathbb{D}$. Moreover, any pixel $D_n(k,\ell)$ is a union of exactly four pixels of the form $D_{n+1}(k',\ell')$. See Figure 1.

2.2. **Pixel images from exposures.** Define the *set of exposures* as

$$E := (\mathbb{D} \times \mathbb{R})^{\mathbb{N}},$$

where $\mathbb{N} = \{1, 2, 3, \dots\}$. More explicitly, exposures $e \in E$ are of the form $e = (e_j)_{j=1}^{\infty} = (e_1, e_2, e_3, \dots)$, where the photons e_j can be written as $e_j = (z_j, \lambda_j) = ((x_j, y_j), \lambda_j)$ with $0 \le x_j < 1$ and $0 \le y_j < 1$.

Throughout the paper we consider spectral sensitivity functions with the following properties.

Definition 1. *We call a continuous function $g : \mathbb{R} \to \mathbb{R}$ a* spectral sensitivity function *if $0 \le g(\lambda) \le M$ for some $0 < M < \infty$ and there exist $\lambda^{(0)} \in \mathbb{R}$ and $\lambda^{(M)} \in \mathbb{R}$ so that $g(\lambda^{(0)}) = 0$ and $g(\lambda^{(M)}) = M$.*

Given a resolution parameter $n \ge 0$, the pixel image $\mathcal{I}_n(e, J)$ produced by the J first photons of e is the following $2^n \times 2^n$ matrix:

$$\mathcal{I}_n(e,J) := \begin{bmatrix} \mathcal{I}_n(e,J,1,1) & \cdots & \mathcal{I}_n(e,J,1,2^n) \\ \vdots & \ddots & \vdots \\ \mathcal{I}_n(e,J,2^n,1) & \cdots & \mathcal{I}_n(e,J,2^n,2^n) \end{bmatrix}, \tag{2.1}$$

where the matrix elements, or pixel values, are defined by

$$\mathcal{I}_n(e,J,k,\ell) := \frac{\mathcal{E}_g(e,J,D_n(k,\ell))}{2^{-2n}J} = \frac{1}{2^{-2n}J} \sum_{\substack{z_j \in D_n(k,\ell) \\ 1 \le j \le J}} g(\lambda_j), \tag{2.2}$$

where $1 \le k \le 2^n$ is row index and $1 \le \ell \le 2^n$ is column index and $\mathcal{E}_g$ is defined in (1.2). Note that (2.2) is consistent with (1.1) since $|D_n(k,\ell)| = 2^{-2n}$. The assumption $0 \le g(\lambda) \le M$ implies the estimates

$$0 \le \mathcal{I}_n(e; J, k, \ell) \le 2^{2n} M \tag{2.3}$$

and

$$0 \le \sum_{1\le k,\ell\le 2^n} \mathcal{I}_n(e;J,k,\ell) = \frac{1}{2^{-2n}J} \sum_{\substack{z_j\in\mathbb{D}\\ 1\le j\le J}} g(\lambda_j) \le 2^{2n}M. \tag{2.4}$$

Definition 2. *The pixel image $\mathcal{I}_n(e,J)$ defined in (2.1) is called the* J-snapshot *of $e \in E$ at resolution n.*

Next we construct ideal images, called *snapshots*, created by exposures with infinitely many photons. For this we need the concept of *Banach limit* B-lim, which is a linear functional on the space of sequences $(c_j)_{j=1}^\infty$ of real numbers. The Banach limit has the following property:

$$\liminf_{j\to\infty} c_j \le B\text{-}\lim_{j\to\infty} c_j \le \limsup_{j\to\infty} c_j. \tag{2.5}$$

Moreover, if the sequence $(c_j)_{j=1}^\infty$ is convergent, then

$$B\text{-}\lim_{j\to\infty} c_j = \lim_{j\to\infty} c_j,$$

but the Banach limit is well-defined even for divergent sequences [31].

Definition 3. *The snapshot of $e \in E$ at resolution n is*

$$\mathcal{I}_n(e) := \begin{bmatrix} \mathcal{I}_n(e,1,1) & \cdots & \mathcal{I}_n(e,1,2^n) \\ \vdots & \ddots & \vdots \\ \mathcal{I}_n(e,2^n,1) & \cdots & \mathcal{I}_n(e,2^n,2^n) \end{bmatrix}, \tag{2.6}$$

where the pixel values are defined by

$$\mathcal{I}_n(e,k,\ell) := B\text{-}\lim_{J\to\infty} \mathcal{I}_n(e,J,k,\ell). \tag{2.7}$$

Note that a combination of (2.3), (2.5) and (2.7) yields the estimate

$$0 \le \mathcal{I}_n(e;k,\ell) \le 2^{2n}M. \tag{2.8}$$

The reason for using B-lim in formula (2.7) instead of more explicitly defined operations, such as lim inf or lim sup, is the linearity of B-lim with respect to the sequence. We will use the linearity property in several proofs below.

The next example shows that in general we cannot replace B-lim in (2.7) by lim since the sequence $(\mathcal{I}_n(e,J,k,\ell))_{J=1}^\infty$ may diverge.

Example 1. *Let g be as in Definition 1 with $M > 0$. Fix $n \ge 0$ and let k,ℓ be the unique indices satisfying $(\frac{1}{2},\frac{1}{2}) \in D_n(k,\ell)$. Set $J_1 = 1$ and $e_1 = ((\frac{1}{2},\frac{1}{2}),\lambda^{(M)})$; then by (1.2) we have $\mathcal{E}_g(e,J_1,D_n(k,\ell)) = g(\lambda^{(M)}) = M$. We proceed by induction. Assume that J_m first photons of e have been constructed with $m \ge 1$. Then repeat these two steps* ad infinitum*:*

- *If* $\lambda_{J_m} = \lambda^{(M)}$, *then set* $e_j := ((\frac{1}{2}, \frac{1}{2}), \lambda^{(0)})$ *for* $j = J_m + 1, \dots, J_{m+1}$ *with so large* J_{m+1} *that*

$$\mathcal{E}_g(e, J_{m+1}, D_n(k,\ell)) = \frac{1}{J_{m+1}} \sum_{j=1}^{J_{m+1}} g(\lambda_j) < \frac{M}{4}.$$

- *If* $\lambda_{J_m} = \lambda^{(0)}$, *then set* $e_j := ((\frac{1}{2}, \frac{1}{2}), \lambda^{(M)})$ *for* $j = J_m + 1, \dots, J_{m+1}$ *with so large* J_{m+1} *that*

$$\mathcal{E}_g(e, J_{m+1}, D_n(k,\ell)) = \frac{1}{J_{m+1}} \sum_{j=1}^{J_{m+1}} g(\lambda_j) > \frac{M}{2}.$$

Now the sequence $(\mathcal{I}_n(e, J, k, \ell))_{J=1}^{\infty}$ *diverges since*

$$\begin{aligned} \mathcal{I}_n(e; J_m, k, \ell) &> M 2^{2n-1} \qquad \text{for odd } m, \\ \mathcal{I}_n(e; J_m, k, \ell) &< M 2^{2n-2} \qquad \text{for even } m. \end{aligned}$$

2.3. **The snapshot metric.** For each $n \geq 0$ we introduce a pseudometric d_n in the set E by the formula

$$d_n(e, e') := c_n \|\mathcal{I}_n(e) - \mathcal{I}_n(e')\|_2, \tag{2.9}$$

where $c_n = \sqrt{2} \cdot 2^{2n-1}$ and $\|\cdot\|_2$ denotes the Euclidean distance between matrices interpreted as elements in $\mathbb{R}^{2^n \times 2^n}$. Let us check that d_n really is a pseudometric. First the triangle inequality:

$$\begin{aligned} d_n(e, e'') &= c_n \|\mathcal{I}_n(e) - \mathcal{I}_n(e'')\|_2 \\ &\leq c_n \|\mathcal{I}_n(e) - \mathcal{I}_n(e')\|_2 + c_n \|\mathcal{I}_n(e') - \mathcal{I}_n(e'')\|_2 \\ &= d_n(e, e') + d_n(e', e''). \end{aligned}$$

Then symmetry: $d_n(e, e') = c_n \|\mathcal{I}_n(e) - \mathcal{I}_n(e')\|_2 = c_n \|\mathcal{I}_n(e') - \mathcal{I}_n(e)\|_2 = d_n(e, e')$, and finally, $d_n(e, e) = c_n \|\mathcal{I}_n(e) - \mathcal{I}_n(e)\|_2 = 0$.

The estimates (2.8) and (2.4) imply $\|\mathcal{I}_n(e)\|_2 \leq 2^{2n} M$ and thus we get a bound on $d_n(e, e')$. Namely,

$$\begin{aligned} d_n(e, e') &= c_n \Big(\sum_{k=1}^{2^n} \sum_{\ell=1}^{2^n} |\mathcal{I}_n(e; k, \ell) - \mathcal{I}_n(e'; k, \ell)|^2 \Big)^{1/2} \\ &\leq c_n \Big(\|\mathcal{I}_n(e)\|_2^2 + \|\mathcal{I}_n(e')\|_2^2 \Big)^{1/2} \\ &\leq \sqrt{2} \cdot 2^{2n-1} ((2^{2n} M)^2 + (2^{2n} M)^2)^{1/2} \\ &= M \end{aligned} \tag{2.10}$$

where we have used the fact that intensities are always positive. We use the sequence $d_0, d_1, d_2, \dots$ to construct a pseudometric involving

all resolutions simultaneously. Define

$$d(e,e') := \max_{n\geq 0} \frac{d_n(e,e')}{n+1}. \tag{2.11}$$

Note that (2.10) implies that $d(e,e') \leq M$.

The following result defines examples with $d(e,e') = 0$ but $e \neq e'$.

Lemma 1. *Let $e = (e_1, e_2, e_3 \dots) \in E$ and $K > 0$. Define another exposure by $e' := (e_{K+1}, e_{K+2}, e_{K+3}, \dots) \in E$. Then for any $n \geq 0$ we have $\mathcal{I}_n(e) = \mathcal{I}_n(e')$.*

Proof. Fix $n \geq 0$ and $1 \leq k \leq 2^n$ and $1 \leq \ell \leq 2^n$. We need to show

$$B\text{-}\lim_{J\to\infty} \mathcal{I}_n(e'; J, k, \ell) = B\text{-}\lim_{J\to\infty} \mathcal{I}_n(e; J, k, \ell). \tag{2.12}$$

Denote $D_n = D_n(k,\ell)$ and compute

$$\begin{aligned}
\mathcal{E}_g(e', J, D_n) &= \sum_{\substack{z_{j+K}\in D_n \\ 1\leq j\leq J}} g(\lambda_j) \\
&= \sum_{\substack{z_j\in D_n \\ 1\leq j\leq J+K}} g(\lambda_j) - \sum_{\substack{z_j\in D_n \\ 1\leq j\leq K}} g(\lambda_j) \\
&= \mathcal{E}_g(e, J+K, D_n) - G_K, \qquad (2.13)
\end{aligned}$$

where $|G_K| < \infty$. By (2.13) and the linearity of the Banach limit we have

$$\begin{aligned}
B\text{-}\lim_{J\to\infty} \frac{\mathcal{E}_g(e', J, D_n)}{J} &= B\text{-}\lim_{J\to\infty} \frac{\mathcal{E}_g(e, J+K, D_n)}{J} - G_K \left(B\text{-}\lim_{J\to\infty} J^{-1} \right) \\
&= B\text{-}\lim_{J\to\infty} \frac{\mathcal{E}_g(e, J+K, D_n)}{(J+K)-K} \\
&= \left(B\text{-}\lim_{J\to\infty} \frac{\mathcal{E}_g(e, J+K, D_n)}{J+K} \right) \left(\lim_{J\to\infty} \frac{1}{1-\frac{K}{J+K}} \right) \\
&= B\text{-}\lim_{J\to\infty} \frac{\mathcal{E}_g(e, J+K, D_n)}{J+K},
\end{aligned}$$

and (2.12) follows from (2.2). □

2.4. **The space P of infinite photographs.** Define an equivalence relation $\sim$ between exposures $e \in E$ and $e' \in E$ by

$$e \sim e' \quad \text{if and only if} \quad d(e,e') = 0,$$

where d is the pseudometric defined in (2.11). Denote the equivalence class of e by $[e] = \{e' \in E : e' \sim e\}$.

Recall that the difference between a pseudometric and a metric is that a metric needs to additionally satisfy that if the distance between e and e' is zero then $e = e'$; see [16]. It is easy to see that the pseudometric d of E extends to a metric on the quotient space $E/\!\sim = \{[e] \,:\, e \in E\}$. Namely, assume that $d(e, e') = 0$ so that $[e] = [e']$. Then by (2.9) and (2.11) we can estimate for any fixed $n \geq 0$

$$\|\mathcal{I}_n(e) - \mathcal{I}_n(e')\|_2 = \frac{1}{c_n}\, d_n(e, e') \leq \frac{n+1}{c_n} d(e, e') = 0,$$

so $\mathcal{I}_n(e) = \mathcal{I}_n(e')$. Thus we can set

$$d_n([e], [e']) := d_n(e, e') \tag{2.14}$$

with arbitrary choice of representatives of equivalence classes $e \in [e]$ and $e' \in [e']$, and

$$d([e], [e']) := \max_{n \geq 0} \frac{d_n([e], [e'])}{n+1}. \tag{2.15}$$

Definition 4. *The set $P := E/\!\sim$ is called* the space of infinite photographs. *The metric d of P is called the* snapshot metric.

Note that the snapshot map $\mathcal{I}_n : E \to \mathbb{R}^{2^n \times 2^n}$ defined in (2.6) extends naturally to a map $\mathcal{I}_n : P \to \mathbb{R}^{2^n \times 2^n}$.

3. Infinite photographs and Borel measures

3.1. **Mapping exposures to measures.** In this section we show how an exposure $e \in E$, together with a spectral sensitivity function g as in Definition 1, induces a Borel measure μ^e on $\mathbb{D}$. Define

$$\mu^*(D) = \inf \left\{ \sum_{\nu=1}^{\infty} 2^{-2n_\nu} \mathcal{I}_{n_\nu}(e, k_\nu, \ell_\nu) \;:\; D \subset \bigcup_{\nu=1}^{\infty} D_{n_\nu}(k_\nu, \ell_\nu) \right\} \tag{3.1}$$

for any subset $D \subset \mathbb{D}$. Note that $0 \leq \mu^*(D) \leq M$. Due to [27, Prop. 2.4.] the mapping μ^* is an outer measure.

Now μ^* satisfies the Caratheodory condition: if $D_1 \subset \mathbb{D}$ and $D_2 \subset \mathbb{D}$ are such that

$$0 < \inf \{|x_1 - x_2| \;:\; x_1 \in D_1, x_2 \in D_2\},$$

then

$$\mu^*(D_1 \cup D_2) = \mu^*(D_1) + \mu^*(D_2).$$

This implies that Borel sets are μ^*-measurable [10, Thm. 4.12], and we can define a Borel measure μ^e as the restriction of μ^* to Borel sets.

Since the definition (3.1) only depends on the snapshot map $\mathcal{I}_n$, we see that $\mu^e = \mu^{e'}$ whenever $e \sim e'$. Thus the following mapping is well-defined.

Definition 5. *Denote by $\mathcal{B}_M(\mathbb{D})$ the set of all positive Borel measures μ on $\mathbb{D}$ satisfying $\mu(\mathbb{D}) \leq M$. The* limit image map $\mathcal{M} : P \to \mathcal{B}_M(\mathbb{D})$ *is defined by*

$$\mathcal{M}([e]) := \mu^e,$$

where μ^e is the restriction of the outer measure μ^* to Borel sets.

The snapshot metric space for exposures was constructed so that converging sequence must convergence at every resolution. Similarly we can ask whether the snapshots approximate the induced measure somehow. Let $e \in E$ be an exposure and $\mu = \mathcal{M}([e])$ with a given spectral sensitivity function g. Now we can define a sequence $\mu_0, \mu_1, \mu_2, \ldots$ of simple measures by

$$\mu_n = \sum_{k=1}^{2^n} \sum_{\ell=1}^{2^n} \mathcal{I}_n(e,k,\ell)\chi_{D_n(k,\ell)}(z)dz, \tag{3.2}$$

where dz is Lebesgue measure on $\mathbb{D}$. Then we have the following approximation result.

Theorem 1. *The sequence $\mu_1, \mu_2, \ldots$ converges to μ in the weak topology of Borel measures.*

Proof. We need to show that

$$\lim_{n\to\infty} \int_{\mathbb{D}} \phi \, d\mu_n = \int_{\mathbb{D}} \phi \, d\mu$$

for an arbitrary continuous function $\phi \in C(\mathbb{D})$. The measures μ_n and μ are regular Borel measures with finite variation:

$$|\mu_n|(\mathbb{D}) \leq M \quad \text{and} \quad |\mu|(\mathbb{D}) \leq M.$$

Let now $\phi \in C(\mathbb{D})$ and note that due to compactness of the torus $\mathbb{D}$ the function ϕ is absolutely continuous. Let $\epsilon > 0$. Let $m = m(\phi)$ be such that we can find a piecewise constant function

$$\phi_m(x) = \sum_{k=1}^{2^m} \sum_{\ell=1}^{2^m} a(k,\ell)\chi_{D_m(k,\ell)}(x)$$

for which $\|\phi - \phi_m\|_\infty < \epsilon/(2M)$. When $j \geq m$ we have $\mu_j(D_m(k,\ell)) = \mu(D_m(k,\ell))$ for all $1 \leq k \leq 2^m$ and $1 \leq \ell \leq 2^m$ and thus

$$\begin{aligned}\left|\int_{\mathbb{D}} \phi d\mu - \int_{\mathbb{D}} \phi d\mu_j\right| &\leq \left|\int_{\mathbb{D}} (\phi - \phi_m) d\mu_j\right| + \left|\int_{\mathbb{D}} (\phi - \phi_m) d\mu\right| \\ &\quad + \left|\int_{\mathbb{D}} \phi_m (d\mu - d\mu_j)\right| \\ &< \epsilon/2 + \epsilon/2 + 0 = \epsilon,\end{aligned}$$

which concludes the proof. □

Example 2. *Assume that $g(\lambda) = 1$ for some $\lambda \in \mathbb{R}$ and define an exposure $\delta = (\delta_j)_{j=1}^{\infty} \in E$ by $\delta_j = (z_0, \lambda)$, where $z_0 = (\frac{1}{2}, \frac{1}{2})$. Because all photons arrive at the same point, this exposure models photographing a bright and distant star in otherwise black night sky.*

Let us determine the J-snapshots of δ. For a fixed $n \geq 0$ let k and ℓ be the unique indices for which $z_0 \in D_n(k, \ell)$. Then

$$\mathcal{I}_n(\delta, J, k', \ell') = \frac{2^{2n}}{J} \sum_{\substack{z_j \in D_n(k', \ell') \\ 1 \leq j \leq J}} g(\lambda) = \begin{cases} 2^{2n} \text{ if } k' = k \text{ and } \ell' = \ell, \\ \\ 0 \text{ otherwise,} \end{cases} \tag{3.3}$$

so all pixel values are constant with respect to J. See Figures 2 and 3 for illustration of J-snapshots of δ with varying values of J.

Consider now the measure μ^δ induced by δ. For every Borel set D containing z_0 it is easy to see that $\mu^\delta(D) = 1$ and $\mu^\delta(\mathbb{D} \setminus D) = 0$. This follows since any cover $D \subset \bigcup_{\nu=1}^{\infty} D_{n_\nu}(k_\nu, \ell_\nu)$ includes a pixel that contains z_0. Thus we conclude that $\mu^\delta = \delta_{z_0}$, a point mass at z_0.

3.2. **Mapping measures to exposures.** Our starting point here is a positive Borel measure $\mu \in \mathcal{B}_M(\mathbb{D})$ and a spectral sensitivity function g as in Definition 1 with $M \geq \mu(\mathbb{D}) \geq 0$.

Let $(\Omega, \Sigma, \mathbb{P})$ be a complete probability space and let $Z^{(j)} : \Omega \to \mathbb{D}$ with $j = 0, 1, 2, \ldots$ be a sequence of independent random variables. If $\mu(\mathbb{D}) > 0$ assume that each $Z^{(j)}$ has probability distribution $\mu(\mathbb{D})^{-1}\mu$, and in the case $\mu(\mathbb{D}) = 0$ assume that each $Z^{(j)}$ is uniformly distributed on $\mathbb{D}$.

For each pixel $D_n(k, \ell)$ with $n \geq 0$ and $1 \leq k \leq 2^n$ and $1 \leq \ell \leq 2^n$ consider the sequence $(\chi_{D_n(k,\ell)}(Z^{(j)}))_{j=1}^{\infty}$ of random variables. The strong law of large numbers implies almost surely

$$\lim_{J \to \infty} \frac{1}{J} \sum_{j=1}^{J} \chi_{D_n(k,\ell)}(Z^{(j)}(\omega)) = \begin{cases} \dfrac{\mu(D_n(k, \ell))}{\mu(\mathbb{D})}, & \text{if } \mu(\mathbb{D}) > 0, \\ \\ |D_n(k, \ell)|, & \text{if } \mu(\mathbb{D}) = 0. \end{cases} \tag{3.4}$$

Let $\Omega_{n,k,\ell}$ denote the subset of Ω in which (3.4) fails. The countable union over all sets $\Omega_{n,k,\ell}$ has zero measure:

$$\mathbb{P}\left(\bigcup_{\substack{n \geq 0 \\ 1 \leq k, \ell \leq 2^n}} \Omega_{n,k,\ell} \right) = 0.$$

Thus for every $\mu \in \mathcal{B}_M(\mathbb{D})$ one can always choose a (not necessarily unique) $\omega_\mu \in \Omega$ for which (3.4) holds for any choice of n, k and ℓ.

Definition 6. *The* illumination map $\mathcal{P} : \mathcal{B}_M(\mathbb{D}) \to P$ *is defined by*

$$\mathcal{P}(\mu) := [\,((Z^{(j)}(\omega_\mu), \lambda))_{j=1}^{\infty}\,],$$

where λ *satisfies* $g(\lambda) = \mu(\mathbb{D})$.

It is straightforward to show that $\mathcal{P}(\mu)$ is well-defined. We note that due to (3.4) any exposure $e \in \mathcal{P}(\mu)$ satisfies

$$\mathcal{I}_n(e, k, \ell) = 2^{2n}\mu(D_n(k, \ell)) \tag{3.5}$$

for any $n \in \mathbb{N}$ and $1 \leq k \leq 2^n$ and $1 \leq \ell \leq 2^n$.

We remark that the map $\mathcal{P}$ is deterministic although its construction includes a random procedure.

In Example 2 we constructed a point measure using infinite sequences of photons. Measures supported on one-dimensional curves can be represented by photon sequences as well, leading to a scalable anti-aliasing technique:

Example 3. *Let* γ *be a continuous curve* $\gamma : [0, 1] \to \mathbb{D}$ *and let us consider the measure* μ *satisfying*

$$\mu(E) = \int_0^1 \chi_E(\gamma(t))dt$$

for any measurable $E \subset \mathbb{D}$. *Assume that* $g(\lambda) = 1$ *for some* $\lambda \in \mathbb{R}$. *We construct an exposure* $c \in E$ *by choosing a uniformly distributed random sequence of points* $s_1, s_2, \ldots$ *on the interval* $[0, 1]$ *and defining photons* c_j *by*

$$c_j := (\gamma(s_j), \lambda), \qquad j = 1, 2, 3, \ldots$$

It is now easy to see that $\mathcal{P}(\mu) = [c]$. *See Figure 4 for a simple sine curve example where, up to scaling the coordinates,* $\gamma(s) = (s, \sin(2\pi s))$.

3.3. **The set of regular exposures.** In this subsection we discuss some important properties of the limit image map $\mathcal{M}$ and the illumination map $\mathcal{P}$. Let us first study two examples illustrating some counter-intuitive behaviour.

Example 4. *Define exposures* $\delta^+ = (\delta_j^+)_{j=1}^{\infty}$ *and* $\delta^- = (\delta_j^-)_{j=1}^{\infty}$ *by*

$$\delta_j^{\pm} = (z_0 \pm (\frac{3}{10j}, \frac{3}{10j}), \lambda).$$

See Figures 2 and 3 for J-snapshots of δ^+ *and* δ^-.

Let δ be as in Example 2. It is easy to see that $\lim_{J\to\infty} \mathcal{I}_n(\delta^+, J, k, \ell) = \mathcal{I}_n(\delta, k, \ell)$ *for all* n, k, ℓ*, showing that* $d(\delta, \delta^+) = 0$*. Furthermore,*

$$\mathcal{M}(\delta) = \delta_{z_0} = \mathcal{M}(\delta^+),$$

where δ_{z_0} denotes a point measure at z_0.

However, for any $n \geq 1$ and $j \in \mathbb{N}$ we have $z_0 - (\frac{3}{10j}, \frac{3}{10j}) \notin D_n(k, \ell)$ if $z_0 \in D_n(k, \ell)$. This means that

$$\mathcal{I}_n(\delta^-, k, \ell) = \lim_{J\to\infty} \mathcal{I}_n(\delta^-, J, k, \ell) = 0$$

for all pixels $D_n(k, \ell)$ with $n \geq 1$ and containing z_0. For the distance of δ and δ^- this implies that $d(\delta, \delta^-) \neq 0$. Moreover, we can show that

$$\mathcal{M}(\delta^-) \equiv 0.$$

Namely, denote by $t_\nu = \sum_{j=1}^{\nu} \frac{1}{2^{j+1}}$, $\nu \geq 1$, a sequence of points converging to $\frac{1}{2}$ as ν increases. Let $D_{n_\nu}(k_\nu, \ell_\nu)$ be the pixel with lower left corner at (t_ν, t_ν) and $n_\nu = \nu + 2$. We notice that $\mathcal{I}_{n_\nu}(\delta^-, k_\nu, \ell_\nu) = 0$. Consider now the definition of $\mu^{\delta^-}(\mathbb{D})$ in (3.1). By writing $\mathbb{D}$ as a union of disjoint pixels including $D_{n_\nu}(k_\nu, \ell_\nu)$ one sees that $\mu^{\delta^-}(\mathbb{D})$ must be zero.

Example 5. *Take $0 \leq x < 1$ and $0 \leq y < 1$ to be irrational numbers and set $z_0 = (x, y)$. Consider the exposures δ^+ and δ^- defined in Example 4. Then due to our dyadic pixel construction we have $\mathcal{M}(\delta^+) = \delta_{z_0} = \mathcal{M}(\delta^-)$.*

Example 4 shows that there exists no such positive Borel measure μ on $\mathbb{D}$ that $\mathcal{P}(\mu) = [\delta^-]$ since due equation (3.5) we have $\mu(D_n(k, \ell)) = 2^{-2n}\mathcal{I}_n(\delta^-, k, \ell)$ and this leads to a contradiction with the assumption that μ is countably additive.

It turns out that we can characterize the subset of P that is natural for mappings $\mathcal{M}$ and $\mathcal{P}$. Let us define regular exposures and regular photographs as follows.

Definition 7. *An element $e \in E$ is a* regular exposure *if the equality $\mu^e(D_n(k, \ell)) = 2^{-2n}\mathcal{I}_n(e, k, \ell)$ holds and the limit $\lim_{J\to\infty} \mathcal{I}_n(e; J, k, \ell)$ exists for any $n \geq 0$ and for all indices $1 \leq k \leq 2^n$ and $1 \leq \ell \leq 2^n$. The subset of regular exposures is denoted by $E_0 \subset E$.*

Further, let P_0 be the subset of P such that every equivalence class $[e] \in P_0$ has a regular exposure representative. This P_0 is called the set of regular infinite photographs*.*

Let us study in the following how regular exposures are related to mappings $\mathcal{M}$ and $\mathcal{P}$.

Theorem 2. *The maps $\mathcal{M} : P \to \mathcal{B}_M(\mathbb{D})$ and $\mathcal{P} : \mathcal{B}_M(\mathbb{D}) \to P$ have the following properties:*

(i) *$\mathcal{M} \circ \mathcal{P}$ is the identity map on $\mathcal{B}_M(\mathbb{D})$,*
(ii) *$\mathcal{P} \circ \mathcal{M} : P \to P$ is a projection onto P_0,*
(iii) *$\mathcal{P} \circ \mathcal{M}$ is the identity map on P_0.*

Proof. Let us prove (i)-(iii) via five separate claims: we show that

(a) $\mathcal{M} : P_0 \to \mathcal{B}_M(\mathbb{D})$ is injective,
(b) $\mathcal{P} : \mathcal{B}_M(\mathbb{D}) \to P$ is injective,
(c) $\mathcal{P}(\mathcal{M}([e])) = [e]$ for any $[e] \in P_0$ and
(d) $\mathcal{M}(\mathcal{P}(\mu)) = \mu$ for any $\mu \in \mathcal{B}_M(\mathbb{D})$.

Note that two positive Borel measures μ and μ' coincide on $\mathbb{D}$ if for any pixel $D_n(k,\ell)$ with $n \in \mathbb{N}$ and $1 \leq k \leq 2^n$ and $1 \leq \ell \leq 2^n$ we have

$$\mu(D_n(k,\ell)) = \mu'(D_n(k,\ell)).$$

This is due to the fact that pixels $D_n(k,\ell)$ span the Borel σ-algebra on $\mathbb{D}$, see [32].

To prove (a) let $[e], [e'] \in P_0$ and assume $[e] \neq [e']$. Then there exists n, k and ℓ such that $\mathcal{I}_n(e,k,\ell) \neq \mathcal{I}_n(e',k,\ell)$. Since e and e' are regular we have $\mu^e(D_n(k,\ell)) \neq \mu^{e'}(D_n(k,\ell))$. This implies that $\mathcal{M}([e]) \neq \mathcal{M}([e'])$.

Now for problem (b) let $\mu, \mu' \in \mathcal{B}_M(\mathbb{D})$ and $\mu \neq \mu'$. Then there exists n, k and ℓ such that $\mu(D_n(k,\ell)) \neq \mu'(D_n(k,\ell))$. Let $e \in \mathcal{P}(\mu)$ and $e' \in \mathcal{P}(\mu')$. According to equation (3.5) we have $\mathcal{I}_n(e,k,\ell) \neq \mathcal{I}_n(e',k,\ell)$ and furthermore $\mathcal{P}(\mu) \neq \mathcal{P}(\mu')$.

(c) Let $[e] \in P_0$, $f \in \mathcal{P}(\mathcal{M}([e]))$ and denote $\mu = \mathcal{M}([e])$. We see that by definition $\mathrm{Ran}(\mathcal{P}) \subset P_0$ and thus $[f] \in P_0$. From this notion and equation (3.5) we have that

$$\mathcal{I}_n(f,k,\ell) = 2^{2n}\mu(D_n(k,\ell)) = \mathcal{I}_n(e,k,\ell). \tag{3.6}$$

for any $n \geq 0$ and $1 \leq k \leq 2^n$ and $1 \leq \ell \leq 2^n$. This proves that $[e] = [f] = \mathcal{P}(\mathcal{M}([e]))$.

(d) Let $\mu \in \mathcal{B}_M(\mathbb{D})$ and denote $[e] = \mathcal{P}(\mu)$ and $\mu' = \mathcal{M}(\mathcal{P}(\mu))$. For any pixel $D_n(k,\ell)$ we have

$$\mu'(D_n(k,\ell)) = 2^{-2n}\mathcal{I}_n(e,k,\ell) = \mu(D_n(k,\ell))$$

since $e \in E_0$ and second equality due to equation (3.5). □

3.4. **Embedding finite photographs to P.** In practice, every photograph is formed by a finite number of photons. For our infinite photography model to be useful, we need to find an appropriate way to embed finite exposures into the space P. There are two main properties we wish from the embedding.

Photographic experiments show that detecting more photons leads to an increase in image quality. This can be easily observed using any digital camera: take a picture of the same scene with two different sensitivity settings, for example ISO 800 and ISO 100, and the latter photograph will be less noisy. This suggests that finite photographs of natural scenes are created by truncating an infinite exposure for which ordinary limit exists in (2.7).

Our way of embedding practical photographs into the infinite photography model is to represent them as *repeated* finite sequences.

Definition 8. *Define an infinite sequencing operation $S_N : E \to E$ for $N > 0$ as follows. For any $e = (e_1, e_2, \dots) \in E$ set*

$$S_N e = (e_1, e_2, \dots, e_N, e_1, e_2, \dots, e_N, e_1, e_2, \dots, e_N, \dots).$$

The following result shows the convergence of photographs at any resolution as the exposure time increases.

Theorem 3. *Let g be as in Definition 1 and take let $e \in E_0$ to be a regular exposure. Then*

$$\lim_{N\to\infty} d(S_N e, e) = 0. \tag{3.7}$$

Proof. Let $n \geq 0$ be arbitrary and calculate using the ℓ^1-norm

$$\begin{aligned}
&\|\mathcal{I}_n(S_N e) - \mathcal{I}_n(e)\|_1 \\
&= \sum_{k=1}^{2^n}\sum_{\ell=1}^{2^n} \left| \lim_{J\to\infty} \frac{\mathcal{E}_g(S_N e, J, D_n(k,\ell))}{2^{-2n}J} - \lim_{J\to\infty} \frac{\mathcal{E}_g(e, J, D_n(k,\ell))}{2^{-2n}J} \right| \\
&= 2^{2n} \sum_{k=1}^{2^n}\sum_{\ell=1}^{2^n} \left| \frac{1}{N} \sum_{\substack{z_j \in D_n(k,\ell) \\ 1\leq j\leq N}} g(\lambda_j) - \lim_{J\to\infty} \frac{1}{J} \sum_{\substack{z_j \in D_n(k,\ell) \\ 1\leq j\leq J}} g(\lambda_j) \right|. \qquad (3.8)
\end{aligned}$$

We see that (3.8) implies

$$\begin{aligned}
\lim_{N\to\infty} d_n(S_N e, e) &= \lim_{N\to\infty} \|\mathcal{I}_n(S_N e) - \mathcal{I}_n(e)\|_2 \\
&\leq \lim_{N\to\infty} \|\mathcal{I}_n(S_N e) - \mathcal{I}_n(e)\|_1 = 0,
\end{aligned}$$

where the last inequality follows from the 2-norm being always smaller than 1-norm on $\mathbb{R}^{2^{2n}}$. So we have proved formula (3.7) with d replaced by d_n.

It remains to prove convergence in (3.7) with the actual pseudometric d. Let $\epsilon > 0$. Choose an integer q so large that $M/(q+1) < \epsilon$. Further, choose $Q = Q(q)$ so that $d_n(S_N e, e) < \epsilon$ with every $N > Q$

and $0 \leq n \leq q$. When $N > Q$ we have by (2.10)

$$\max_{n\geq 0} \frac{d_n(S_N e, e)}{n+1} = \max\left\{ \max_{0\leq n\leq q} \frac{d_n(S_N e, e)}{n+1}, \max_{n>q} \frac{d_n(S_N e, e)}{n+1} \right\} < \epsilon,$$

and the proof of (3.7) is complete. □

4. Decomposition and fractal analysis of images

The well-known decomposition theorem by Lebesgue states that a Borel measure μ on $\mathbb{D}$ can be uniquely divided into two parts

$$\mu = \phi + \lambda,$$

where ϕ is absolutely continuous with respect to Lebesgue measure m and λ is singular in the sense that it vanishes outside a set of Lebesgue measure zero. In the following we denote $a(\mu) = \phi$ and $s(\mu) = \lambda$. Furthermore, the Radon-Nikodym derivative $f_\mu = \frac{d(a(\mu))}{dm}$ belongs to $L^1(\mathbb{D})$ and we can write

$$\mu = f_\mu dm + s(\mu). \tag{4.1}$$

In our image model, an exposure $e \in E$ induces a Borel measure $\mu^e = \mathcal{M}([e])$, which we can write in the form

$$\mu^e = f_{\mu^e} dm + s(\mu^e). \tag{4.2}$$

Now the question is how to construct exposures e^a and e^s such that

$$\mathcal{M}([e^a]) = f_{\mu^e} dm \quad \text{and} \quad \mathcal{M}([e^s]) = s(\mu^e). \tag{4.3}$$

In Section 4.1 we develop theoretical tools that can be used for locating the support of the singular part of a measure and for analysing the pointwise degree of singularity using *local dimension*. Further, in Section 4.2 we outline a photon-based edge detection method.

Our approach is connected by formula (4.2) to so-called $u + v$ image models [25, 33, 39, 21], where typically u represents a regular part of the image, while v is a singular part (texture or noise). In some models, including ours, the singular part v cannot be represented as a function, leading to serious difficulties in computational approaches. Some promising methods have been introduced with the problem of rapid oscillations [25, 13, 4], and in Section 5.5 below we present a novel photon-based algorithm.

4.1. **Singularity indication and local dimension.** In this subsection we construct tools for studying the support and singularity of $s(\mu)$ in (4.1). First let us show how to estimate the location of the support. Notice that the proof regarding Lebesgue's decomposition of measures (given e.g. in [32]) is constructive in nature. Namely, the singular part $s(\mu)$ with respect to Lebesgue measure in (4.1) is obtained by restricting μ to a specific measurable set $A \subset \mathbb{D}$ with $m(A) = 0$ so that

$$(s(\mu))(E) = \mu(A \cap E) \tag{4.4}$$

for any measurable $E \subset \mathbb{D}$. Naturally instead of A any set $\tilde{A} = A \cup B$ with $m(B) = 0$ can be used in (4.4). Next we explain how A can be recovered uniquely up to a zero m-measurable set.

Assume μ is a positive and bounded Borel measure on $\mathbb{D}$ and define using the Lebesgue measure m an auxiliary measure $\rho = \mu + m$. We see that the mapping

$$f \mapsto \int_{\mathbb{D}} f d\mu$$

is bounded linear functional in $L^2(\mathbb{D}, \rho)$, i.e., the L^2 integrable functions on $\mathbb{D}$ with respect to the measure ρ. By the Riesz representation theorem there exists $\psi \in L^2(\mathbb{D}, \rho)$ such that

$$\int_{\mathbb{D}} f d\mu = \int_{\mathbb{D}} f \psi d\rho \tag{4.5}$$

for all $f \in L^2(\mathbb{D}, \rho)$. It can be shown that $0 \leq \psi(z) \leq 1$ in ρ-a.e. Now by setting

$$A = \{z \in \mathbb{D} \mid \psi(z) \text{ is defined and } \psi(z) = 1\} \tag{4.6}$$

one obtains $s(\mu)$ by defining $(s(\mu))(E) = \mu(A \cap E)$ (for proof see [32, Thm. 6.10.]). Let us denote

$$\psi_r(z) = \frac{\mu(B(z,r))}{\mu(B(z,r)) + m(B(z,r))}$$

for any $z \in \mathbb{D}$ and $r > 0$. The function ψ_r is a fundamental tool for us since computation of $\mu(B(z,r))$ can be done efficiently using exposure data.

Lemma 2. *For ρ-almost all $z \in \mathbb{D}$ it holds that*

$$\psi(z) = \lim_{r \to 0} \psi_r(z). \tag{4.7}$$

Proof. By [23, Cor. 2.14] ρ-almost every point $z \in \mathbb{D}$ is a Lebesgue point of ψ and at these points

$$\psi(z) = \lim_{r \to 0} \frac{1}{\rho(B(z,r))} \int_{B(z,r)} \psi(x) d\rho(x).$$

Therefore by substituting $f = \chi_{B(z,r)}$ into (4.5) we get

$$\psi(z) = \lim_{r\to 0} \frac{\mu(B(z,r))}{\rho(B(z,r))} = \lim_{r\to 0} \frac{\mu(B(z,r))}{\mu(B(z,r)) + m(B(z,r))} \tag{4.8}$$

for ρ-almost all $z \in \mathbb{D}$. □

Hence by estimating the limit $\lim_{r\to 0} \psi_r(z)$ we have means to evaluate whether $z \in A$. We call ψ the *singularity indication function*.

Measure theory also contains ways to classify the local behaviour of a measure. One of them is by computing so-called *(lower) local dimension* or *local Hölder exponent* [17] of μ at $z \in \mathbb{D}$. It is defined by

$$\dim_{loc} \mu(z) = \liminf_{r\to 0} \frac{\log(\mu(B(z,r)))}{\log r} \tag{4.9}$$

when there exists no $r > 0$ such that $\mu(B(z,r)) = 0$. When such an r exists we define $\dim_{loc} \mu(z) = \infty$. It is easy to see that $0 \leq \dim_{loc} \mu(z) \leq \infty$. For further information about local dimension see [17].

Later in Section 5.5 we show how to use exposure data for computing $\psi(z)$ and $\dim_{loc} \mu(z)$. Before that let us now discuss some connections between these concepts. As was discussed earlier we want to locate the points $z \in A$. It is then interesting to study what the local dimension at these points is. It turns out that the convergence speed of $\psi_r(z)$ plays a role here.

Lemma 3. *Let us assume that $z \in \mathbb{D}$ satisfies $\psi(z) > 0$. Then we have that*

$$\dim_{loc} \mu(z) = 2 - \limsup_{r\to 0} \frac{\log(1 - \psi_r(z))}{\log r} \leq 2.$$

Proof. Let $r > 0$ be fixed. The definition of $\psi_r(z)$ implies that

$$\mu(B(z,r)) = \frac{\psi_r(z)}{1 - \psi_r(z)} m(B(z,r)).$$

The logarithm of the right-hand-side simplifies

$$\log\left(\frac{\psi_r(z)}{1 - \psi_r(z)} m(B(z,r))\right) = \log(\psi_r(z)) + \log(m(B(z,r))) - \log(1 - \psi_r(z)).$$

Since $\lim_{r\to 0} \frac{\log(\psi_r(z)))}{\log r} = 0$ and $\lim_{r\to 0} \frac{\log(m(B(z,r)))}{\log(r)} = 2$ we have that

$$\dim_{loc} \mu(z) = 2 + \liminf_{r\to 0} \frac{-\log(1 - \psi_r(z))}{\log(r)}$$

which proves the equality. Here $\dim_{loc}\mu(z) \leq 2$ since $\frac{\log(1-\psi_r(z))}{\log r}$ is positive for $r < 1$. □

Corollary 1. *Assume that $z \in \mathbb{D}$ satisfies equation* (4.7) *and $\dim_{loc}\mu(z) < 2$. Then we have $\lim_{r\to 0}\psi_r(z) = 1$.*

Proof. There exists a sequence $\{r_j\}_{j=1}^{\infty}$, $r_j > 0$, such that

$$\dim_{loc}\mu(z) = \lim_{j\to\infty}\frac{\log(\mu(B(z,r_j)))}{\log(r_j)}$$

and

$$\frac{\log(\mu(B(z,r_j)))}{\log(r_j)} \leq 2-\epsilon \tag{4.10}$$

for some $\epsilon > 0$. Inequality (4.10) yields that

$$\mu(B(z,r_j)) \geq r_j^{2-\epsilon}$$

and by assumption we have

$$\psi(z) = \lim_{j\to\infty}\psi_{r_j}(z) \geq \lim_{j\to\infty}\frac{r_j^{2-\epsilon}}{r_j^{2-\epsilon}+\pi r_j^2} = 1.$$

This yields the claim. □

Note that conditions $\lim_{r\to 0}\psi_r(z) = 1$ and $\dim_{loc}\mu(z) < 2$ are not equivalent for a fixed $z \in \mathbb{D}$. It is possible to construct a bounded measure μ (e.g. with suitably weighted sum of point masses) in such a way that at some point $z \in \mathbb{D}$ both $\lim_{r\to 0}\psi_r(z) = 1$ and $\dim_{loc}\mu(z) = 2$ hold. We do not discuss this discrepancy further in this paper. Instead, in Section 5.5 we simply approximate the set A by sets

$$A_d = \{z \in \mathbb{D} \mid \dim_{loc}\mu(z) < d\}$$

for some $0 < d < 2$.

4.2. **Photon-based edge detection.** We discuss a gradient-based approach for detecting edges in infinite photographs. Furthermore, we study how to differentiate images given in exposure form.

Let us assume that e is a regular exposure and $\mu = \mathcal{M}([e])$. Furthermore, assume that L is a linear first order differential operator (e.g. $L = \frac{\partial}{\partial x}$) and $L\mu$ is a bounded measure. Notice that this assumption is not implicated by our model. Hence we are considering only a subset of all possible images here.

Now we would like to locate the singular part of the total variation of $L\mu$ using the method proposed above. In practise, to be able to fully use the exposure data e we have to relax this objective a bit.

Denote by $\phi_r \in C^\infty(\mathbb{D})$ a sequence of positive functions such that $\lim_{r\to 0}\langle \phi_r, \phi\rangle_{L^2(\mathbb{D})} = \phi(0)$ for any $\phi \in C^\infty(\mathbb{D})$. Furthermore, in the following $\mu * \phi_\epsilon$ denotes the convolution on torus $\mathbb{D} = \mathbb{T}^2$.

Now $|\phi_\epsilon * L\mu|$-measure of a ball $B(z,r)$ can be computed as follows:

$$\begin{aligned} |\phi_\epsilon * L\mu|(B(z,r)) &= \langle |(L\phi_\epsilon) * \mu|, \chi_{B(z,r)}\rangle \\ &= \sup_{\substack{h\in C(\mathbb{D})\\ |h|\le 1}} \langle (L\phi_\epsilon) * \mu, \chi_{B(z,r)} h\rangle \\ &= \sup_{\substack{h\in C(\mathbb{D})\\ |h|\le 1}} \langle \mu, (L\phi_\epsilon) * (\chi_{B(z,r)} h)\rangle \qquad (4.11) \end{aligned}$$

Now if our data is $S_N e$ the duality pairing reduces to

$$\langle \mu^{S_N e}, (L\phi_\epsilon) * (\chi_{B(z,r)} h)\rangle = \frac{1}{N}\sum_{1\le j\le N} g(\lambda_j)(L\phi_\epsilon) * (\chi_{B(z,r)} h)(z_j)$$

The supremum in (4.11) can be approximated by using a sufficiently large number of randomly chosen functions $h_1, h_2, ..., h_T \in C(\mathbb{D})$ and taking a maximum. Thus we have presented a method for approximating the value of $|\phi_\epsilon * L\mu|(B(z,r))$. We postpone the implementation of this method to a further study.

5. Computational methods

5.1. **Monte Carlo simulation of the illumination map.** The construction of the illumination map in Definition 6 was based on a random construction and the strong law of large numbers. In the same way we are able to simulate it stochastically.

Assume that the image has the form $\mu = f_\mu dm$ with a piecewise continuous and bounded function f_μ satisfying $\int_{\mathbb{D}} f_\mu dm = 1$. Further, assume that the point values $f_\mu(x,y)$ are available for any $(x,y) \in \mathbb{D}$. Let us describe a simple approach for simulating the J first photons based on the Metropolis-Hastings algorithm [14, 24].

(1) Pick an initial photon location (x_1, y_1) arbitrarily. Set $\ell = 1$.
(2) Choose a candidate location (x,y) by drawing x and y from the uniform probability density on the interval $[0,1)$.
(3) We need to decide whether to accept the candidate as the next photon in the sequence or to repeat the previous photon. The decision is based on the number

$$\alpha = \min\left\{1, \frac{f_\mu(x,y)}{f_\mu(x_\ell, y_\ell)}\right\}.$$

Draw a random number r from the uniform probability distribution on the interval $[0,1]$. If $\alpha \geq r$ then accept the candidate:

$$(x_{\ell+1}, y_{\ell+1}) = (x, y),$$

otherwise reject the candidate and repeat the previous photon:

$$(x_{\ell+1}, y_{\ell+1}) = (x_\ell, y_\ell).$$

(4) If $\ell < J$ then set $\ell \leftarrow \ell + 1$ and go to 2. Otherwise stop.

The exposure data $e = ((x_j, y_j), \lambda^{(1)})_{j=1}^{\infty}$ produced by Metropolis-Hastings algorithm satisfies $\mathcal{M}(e) = \mu$. Results concerning properties of the chain can be found e.g. in [26].

Let us test the above algorithm in the simple cases of a constant measure $f_\mu \equiv 1$, a completely white image, and modify the result to construct a middle gray constant image as well.

Example 6. *We construct an exposure $w \in E$ by drawing points z_j from the uniform probability distribution on $\mathbb{D}$ and setting $w_j := (z_j, \lambda^{(1)})$ where $g(\lambda^{(1)}) = 1$. See Figure 5 for the resulting approximate pixel images $\mathcal{I}_n(w, J)$.*

We modify w to cover a constant gray value as well. Assume that $g(\lambda^{(0)}) = 0$ and set $\lambda_j = \lambda^{(1)}$ for every odd value of j and $\lambda_j = \lambda^{(0)}$ for every even value of j. Then we construct an exposure $w' \in E$ by $w'_j = (z_j, \lambda_j)$. Then every other photon in w' delivers unit energy and every other delivers zero energy. The result represents a constant image with gray level, as seen in Figure 6.

Example 7. *Let us turn to a nontrivial example. Figure 7(a) shows a monochrome 512×512 pixel digital image of a strawberry. The above algorithm is then used to construct a finite exposure with $J = 1500000$ photons. Approximate pixel images $\mathcal{I}_n(e, J)$ are shown in top right of Figure 7 for $n = 6$ and in top right of Figure 8 for $n = 8$.*

5.2. **Halftoning with black photons.** Given a grayscale pixel image, the goal of halftoning is to produce a set of black dots on a white surface in such a way that to the human eye the pixel image and the set of dots appear as similar as possible. Halftoning is routinely used in graphical printing industry, and there are plenty of quick and simple methods available for halftoning a given image so that the local average gray levels of the pixel image and the halftoned rendition match very closely. However, a typical problem is that if the black dots are distributed too systematically, then disturbing patterns are perceived by human observers. There are several ways around this in the literature [19], including stochastic halftoning methods.

Our image model based on sequences of photons can be applied to halftoning by using the Metropolis-Hastings algorithm to the inverted pixel image, and interpreting photons as black dots of ink. The result is a novel halftoning method.

Let us describe the stochastic halftoning algorithm. Assume given a grayscale image G'' with $N \times N$ pixels and nonnegative pixel values $G''(i,j) \geq 0$. Invert the picture by defining

$$G' := \max_{1 \leq i,j \leq N} G''(i,j) - G''.$$

Further, introduce a normalized image by defining

$$G := \frac{G'}{\frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} G'(i,j)}.$$

Then we can interpret G as a piecewise constant simple function defined on $\mathbb{D}$ and satisfying

$$\int_{\mathbb{D}} G(z) dz = 1,$$

where dz denotes Lebesgue measure on $\mathbb{D}$. Then we can use the Metropolis-Hastings algorithm of Section 5.1 for constructing a sequence of J photons distributed according to the absolutely continuous measure G.

Once the J photons have been chosen using the Metropolis-Hastings algorithm we need to print the black dots. However, there are a couple of problems. First, what should be the area of the black dot printed in place of a photon? Second, if we print a constant size dot at the location of each photon, the repeated photons appear as the same as just one photon, which is not the correct result.

We solve the above problems approximately as follows. Let us compute the average gray level of the original image G'' by setting

$$G''' := \frac{G''}{\max_{1 \leq i,j \leq N} G''(i,j)}$$

and calculating

$$a := \int_{\mathbb{D}} G'''(z) dz = \frac{1}{N^2} \sum_{i=1}^{N} \sum_{j=1}^{N} G'''(i,j).$$

We interpret the number a as follows: a bitmap image on $\mathbb{D}$ has the same average gray level than G'' if the area of black color is a and the area of white color is $1 - a$. The term *bitmap image* means here that all pixel values of the image are either zero (interpreted as black) or one (interpreted as white).

The outcome of our halftoning algorithm is a bitmap image H of size $\widetilde{N} \times \widetilde{N}$ pixels. For the halftoning idea to make sense we must take $\widetilde{N} > N$, for example $\widetilde{N} = 4N$ would do.

In the locations of photons appearing only once we draw a black dot of area a/J. In the location any photon repeated $\nu > 1$ times we draw a black dot of area $\nu a/J$. The result is a halftoned version of G''. Let us still clarify what we mean by "dot" here. In the context of the $\widetilde{N} \times \widetilde{N}$ bitmap image H we reproduce dots as square areas of black pixels. Of course, the fact that a pixel in the image H has area $\widetilde{N}^{-2}$ implies that we cannot reproduce dots of exactly the right size, but the larger $\widetilde{N}$ we take the better we can approximate the correct area.

See Figure 10 for a halftoned gradient and Figure 11 for a halftoned image of a key.

5.3. **Modelling cartoon images including outlines.** It is possible to combine photon-based representations of continuous images and curves in the same exposure. This way we can introduce a novel representation of cartoon images consisting of piecewise continuous grayscale images and outlines of shapes. The outlines will be printed as anti-aliased curves as narrow as possible in the chosen resolution.

Figures 7 and 8 show an image of a strawberry together with outline at low and high resolutions, respectively.

5.4. **Photon-based Gaussian blurring.** Let us close this section by illustrating a simple photon-based image processing operation, namely Gaussian blurring. We construct an exposure e' by randomly translating each photon in the above exposure $e = ((x_j, y_j), \lambda_j)$ computed from the strawberry image as follows:

$$e' = ((x_j + \frac{1}{100}\texttt{randn}, y_j + \frac{1}{100}\texttt{randn}), \lambda_j),$$

where `randn` stands for an independent realization of a normally distributed random variable with mean zero and unit variance.

See Figure 9 for snapshots of e'.

5.5. **Singularity extraction.** Let us now tackle the question about how to decompose exposure data in the sense of Lebesgue discussed in Section 4. Assume that e is a regular exposure. Next we formulate a heuristic method for obtaining two exposures e_n^a and e_n^s such that their induced measures approximate with respect to n the absolutely continuous and singular part of $\mathcal{M}([e])$. In practise we never have the full data of e, but instead a truncated data $S_N e$ for some finite N.

Therefore the two key parameters of our algorithm are n and N and during the following scheme we assume

(1) that $N = N(n)$ is always large with respect to n in the sense of Theorem 3.7 so that
$$d_n(e, S_N e) \approx 0$$
(2) and that we can approximate
$$\mu^e(B(x,r)) \approx \mu^{S_N e}(B(x,r)) = \frac{1}{N} \sum_{\substack{z_j \in B(x,r) \\ 1 \le j \le N}} g(\lambda_j).$$

Let us now assume that n is given. Denote by $z^n_{k,\ell}$ the center point of pixel $D_n(k,\ell)$ for each $1 \le k \le 2^n$ and $1 \le \ell \le 2^n$ and recall $\lambda^{(0)}$ was chosen so in Definition 1 that $g(\lambda^{(0)}) = 0$.

Remark 1. *We noted earlier before Lemma 3 that the convergence speed of $\psi_r(z)$ plays a role when computing $\dim_{loc}\mu(z)$ for a given point $z \in \mathbb{D}$. In fact, it is straightforward to formulate Lemma 3 in the following form using an auxiliary function $h(r)$ that satisfies $h : \mathbb{R}_+ \to [0,1]$ and $\lim_{r\to 0} h(r) = 0$. Assume that function $\psi_r(z)$ satisfies*

$$\psi_r(z) \ge 1 - h(r). \tag{5.1}$$

Then it follows that

$$\dim_{loc}\mu(z) \le 2 - \limsup_{r\to 0} \frac{\log(h(r))}{\log r}. \tag{5.2}$$

For example,

(i) *if $h(r) = r^{2-a}$, $a < 2$, condition (5.1) implies that $\dim_{loc}\mu(z) \le a$ and*
(ii) *if $h(r) = -(\log r)^{-1}$, then for point z which has a neighbourhood not intersecting $\operatorname{supp}(s(\mu))$ the condition (5.1) is not valid.*

Using the observations above we formulate the algorithm to decompose $S_n e$ to e^a_n and e^s_n:

(1) Compute value of
$$\psi^N_{r_n}(z^n_{k,\ell}) = \frac{\mu^{S_N e}(B(z^n_{k,\ell}, r_n))}{\mu^{S_N e}(B(z^n_{k,\ell}, r_n)) + m(B(z^n_{k,\ell}, r_n))}$$
for all $1 \le k \le 2^n$ and $1 \le \ell \le 2^n$, where $r_n = 2^{-n}$.
(2) Define $\psi^N_{r_n}$ everywhere by linearly interpolating the value at z from three closest neighbours $z^n_{k,\ell}$.
(3) For all photons $e_j = (z_j, \lambda_j) \in e$, $j = 1, ..., N$ do the following
   (a) Compute $t = \psi^N_{r_n}(z_j)$.

(b) If $t > 1-h(r_n)$ then set $e^a_{n,j} := (z_j, \lambda^{(0)})$ and $e^s_{n,j} := (z_j, \lambda_j)$.
(c) If $t \leq 1-h(r_n)$ then set $e^a_{n,j} := (z_j, \lambda_j)$ and $e^s_{n,j} := (z_j, \lambda^{(0)})$

(4) Set $e^a_n = (e^a_{n,j})_{j=1}^N$ and $e^s_n = (e^s_{n,j})_{j=1}^N$.

It is straightforward to see now that e^a_n and e^s_n are reconstructed in such a way that $\mathcal{M}([S_N e]) = \mathcal{M}([e^a_n]) + \mathcal{M}([e^s_n])$. By this construction we do not necessarily approximate the whole singular part of $\mathcal{M}([e])$. In fact, as can be seen from Remark 1 we do not expect to see singularities for which the convergence speed of ψ_r is slower as h.

We have implemented the algorithm above in Figures 12 and 13 using convergence speed $h(r_n) = -(\log r_n)^{-1}$. Moreover, after separating the absolutely continuous and singular part we have computed the local dimension of e^s. In Figure 12 the original image has been a line measure on a circle and gray background. A snapshot at resolution $n = 10$ is given on the top. The decomposition is shown in the middle. Furthermore, the inverse of local dimension $\dim_{loc} \mu^e(x)$ is plotted on the bottom. Here we have used $N = 2 \cdot 10^6$ photons. In Figure 13 the original image a photo of sunlight arriving from behind a tree. A snapshot taken at resolution $n = 10$, and the decomposition and the inverse local dimension are plotted as in Figure 12. For the sunlight picture we have used $N = 10^7$ photons. We have used the inverse value of local dimension in order to visualize the absolutely continuous background as black and the singular points as bright spots.

6. Acknowledgements

The authors thank Eero Saksman for valuable comments, especially for pointing out the benefits of using the Banach limit. During part of the preparation of this work, SS worked as professor at the Department of Mathematics of Tampere University of Technology. TH was supported by Emil Aaltonen foundation and Graduate School of Inverse problems (Academy of Finland). The research work of all three authors was funded in part by the Finnish Centre of Excellence in Inverse Problems Research (Academy of Finland CoE-project 213476).

References

[1] Adler, R.L., Kitchens, B.P., Martens, M., Tresser, C.P., Wu, C.W.: The mathematics of halftoning. IBM J. Res. Develop. **47**(1), 5–15 (2003)
[2] Angenent, S., Pichon, E., Tannenbaum, A.: Mathematical methods in medical image processing. Bull. Amer. Math. Soc. (N.S.) **43**(3), 365–396 (2006)
[3] Aubert, G., Kornprobst, P.: Mathematical Problems in Image Processing, Partial Differential Equations and the Calculus of Variations, 2nd edn. Springer, New York (2006)

[4] Aujol, J-F., Aubert, G., Blanc-F'eraud, L., Chambolle, A.: Image Decomposition into a Bounded Variation Component. Journal of Mathematical Imaging and Vision **22**(1), 71–88 (2005)
[5] Boiko, D.L., Gunther, N.J., Brauer, N., Sergio, M., Niclass, C., Beretta, G.B., Charbon, E.: A quantum imager for intensity correlated photons, New J. Phys. (2009) doi:10.1088/1367-2630/11/1/013001
[6] Choi, H., Baraniuk, R.G.: Wavelet statistical models and Besov spaces, Nonlinear estimation and classification, Lecture Notes in Statist., **171**, Springer, New York, 9–29 (2003)
[7] Daubechies, I.: Ten lectures on Wavelets, CBMS-NSF Series in Appl. Math. Montperlier, VT: Capital City. (1992)
[8] Falconer, K.: Fractal Geometry, John Wiley and Sons, Chichester (1990)
[9] Fraysse, A.: Generic validity of the multifractal formalism, SIAM J. Math. Anal. **39**(2), 593–607 (2007)
[10] Gariepy, R., Ziemer, W.: Modern Real Analysis, PWS Publishing Company (1995)
[11] Garnett, J.B., Le, T.M., Meyer, Y., Vese, L.A.: Image decompositions using bounded variation and generalized homogeneous Besov spaces, Appl. Comput. Harmon. Anal. **23**(1), 25–56 (2007)
[12] Gousseau, Y., Morel, J.M.: Are natural images of bounded variation?, SIAM J. Math. Anal. **33**(3), 634–648 (2001)
[13] Haddad, A., Meyer, Y.: An improvement of Rudin-Osher-Fatemi model, Appl. Comput. Harmon. Anal. **22**(3), 319–334 (2007)
[14] Hastings, W.K.: Monte Carlo sampling methods using Markov chains and their applications, Biometrika (1970) doi:10.1093/biomet/57.1.97
[15] Jensen, H.W.: Realistic image synthesis using photon mapping, A K Peters, Natick, (2001)
[16] Kelley, J.L.: General Topology, Van Nostrand Company, Inc., Toronto-New York-London (1955)
[17] Handbook of Measure Theory. Vol II. In: Pap, E. (eds.), Elsevier, Amsterdam (2002)
[18] Harte, D.: Multifractals. Theory and applications, Chapman Hall/CRC, Boca Raton (2001)
[19] Lau, D.L., Arce, G.R.: Modern Digital Halftoning, CRC Press (2001)
[20] Li, Y., Dai, C.: Multifractal dimension inequalities in a probability space, Chaos Solitons Fractals **34**(2), 213–223 (2007)
[21] Lieu, L.H., Vese, L.A.: Image restoration and decomposition via bounded total variation and negative Hilbert-Sobolev spaces, Appl. Math. Optim. **58**(2), 167–193 (2008)
[22] Mallat, S., Zhong, S. et al.: Wavelet transform maxima and multiscale edges, in Wavelets and Their Applications, M. B. Ruskai et al., Eds. Boston, MA: Jones and Bartlett. (1991)
[23] Mattila, P.: Geometry of Sets ans Measures in Euclidian Spaces: Fractals and rectifiability, Cambridge University Press (1995)
[24] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller and E. Teller. Equations of state calculations by fast computing machine. J. Chem. Phys., **21**:1087–1091 (1953)

[25] Meyer, Y.: Oscillating patterns in image processing and in some nonlinear evolution equations, The fifteenth Dean Jacqueline B. Lewis memorial lectures. University Lecture Series, **22**. American Mathematical Society, (2001)
[26] Nummelin, E.: MC's for 'MCMC'ists, International Statistical Review **70**(2), 215–240 (2002)
[27] Mukherjea, A., Pothoven, K.: Real and Functional Analysis. Part A: Real Analysis. Plenum Press, New York (1984)
[28] Mumford, D., Gidas, B.: : Stochastic models for generic images, Quarterly of applied mathematics **59** (2001)
[29] Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. Comm. Pure Appl. Math. **42**(5), 577–685 (1989)
[30] Olsen, L.: Multifractal geometry. Fractal geometry and stochastics, II (Greifswald/Koserow, 1998), 3–37, Progr. Probab., 46, Birkhuser, Basel (2000)
[31] Rao, K.C.: Functional Analysis. Alpha Science Intl Ltd, (2002)
[32] Rudin, M.: Real and Complex Analysis, McGraw-Hill Book Company (1987)
[33] Rudin, L., Osher, S. Fatemi, E.: Nonlinear Total Variation based noise removal algorithms, Physica D **60**, 259–268 (1992)
[34] Simoncelli, E.P.: Statistical Modeling of Photographic Images, in Handbook of Video and Image Processing, ed. Alan Bovik, Academic Press. (2005)
[35] Srivastava, A.; Lee, A. B.; Simoncelli, E. P.; Zhu, S.-C. On advances in statistical modeling of natural images. Special issue on imaging science (Boston, MA, 2002). J. Math. Imaging Vision **18**(1), 17–33 (2003)
[36] Starck, J.-L., Elad, M., Donoho, D.L.: Image decomposition via the combination of sparse representations and a variational approach. IEEE Trans. Image Process. **14**(10), 1570–1582 (2005)
[37] Turiel, A., Parga, N.: The Multifractal Structure of Contrast Changes in Natural Images: From Sharp Edges to Textures, Neural Comput. **12**, 763–793 (2000)
[38] Turiel, A., Del Pozo, A.: Reconstructing Images From Their Most Singular Fractal Manifold, IEEE Transactions on Image Processing, **11**(4), 345–350 (2002)
[39] Vese, L., Osher, S.: Modeling Textures with Total Variation Minimization and Oscillating Patterns in Image Processing, J. Sci. Comput. **19**, 553–572 (2003)

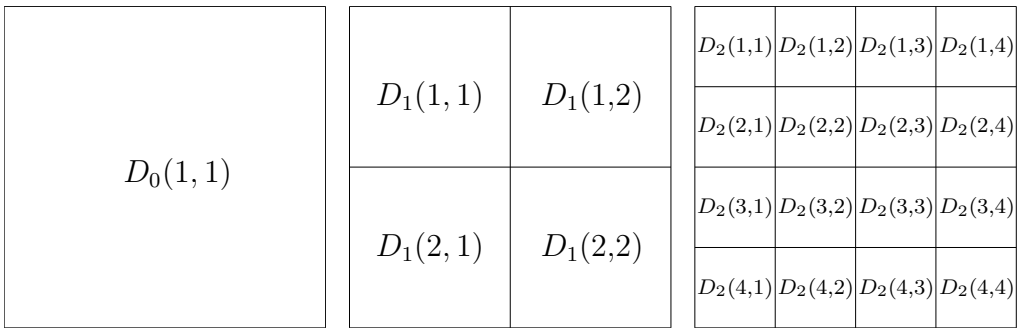

FIGURE 1. Canonical pixelizations for resolutions $n = 0, 1, 2$.

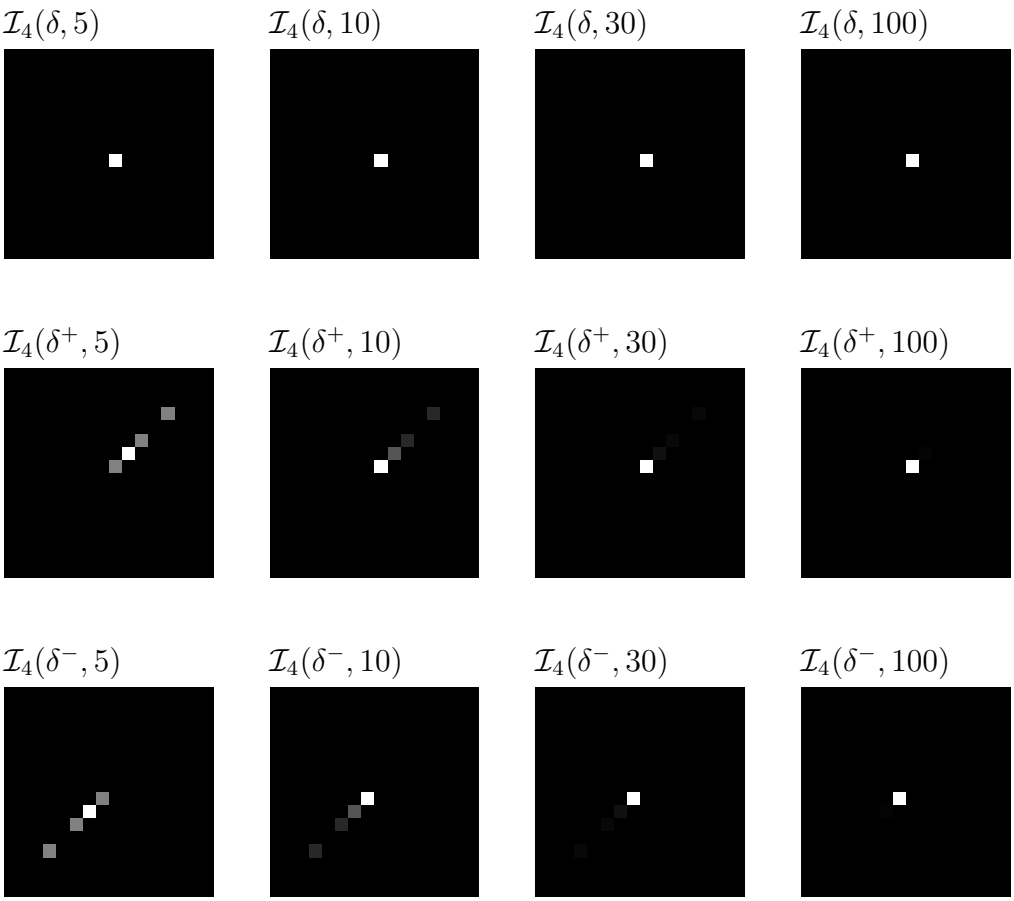

FIGURE 2. J-snapshots of the exposures δ and δ^+ and δ^- at resolution $n = 4$ (size 16×16 pixels). See Example 2 on page 10 for the definition of δ, and Example 4 on page 11 for the definitions of δ^+ and δ^-. Note that the images related to δ in the top row are actually independent of J, as shown by formula (3.3). According to the theory, the images in the middle row converge to the images in the top row as the number of photon increases. The scale is the same in all the images: black is zero and white is $2^{2n} = 2^8 = 256$.

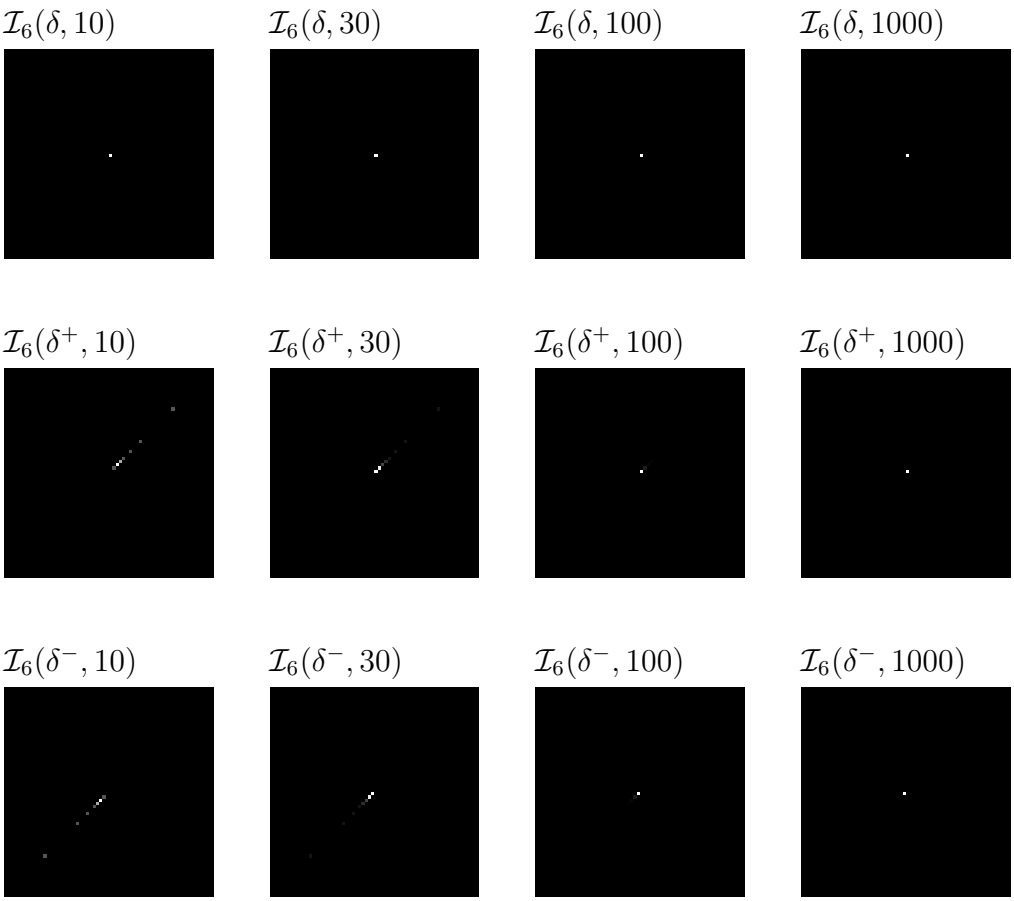

FIGURE 3. J-snapshots of the exposures δ and δ^+ and δ^- at resolution $n = 6$ (size 64×64 pixels). See Example 2 on page 10 for the definition of δ, and Example 4 on page 11 for the definitions of δ^+ and δ^-. Note that the pixel images related to δ in the top row are actually independent of J, as shown by formula (3.3). According to the theory, the images in the middle row converge to the images in the top row as the number of photon increases. Comparison to Figure 2 shows that the convergence rate is slower in the higher-resolution case. The scale is the same in all the images: black is zero and white is $2^{2n} = 2^{12} = 4096$.

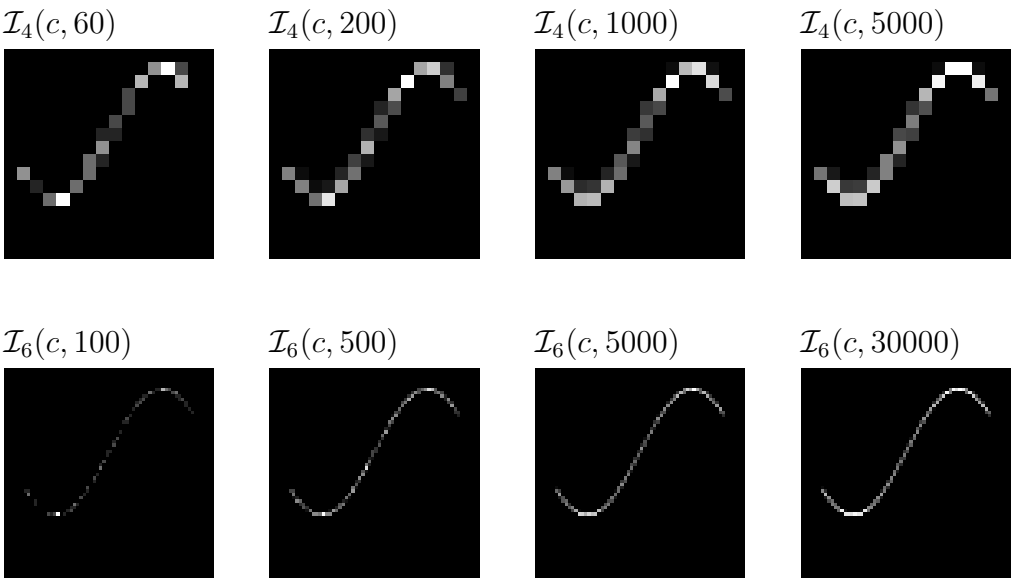

FIGURE 4. Pixel images produces by the curve exposure $c \in E$ constructed in Example 3. The images in the top row are at resolution $n = 4$ and have size 16×16. The images in the bottom row are at resolution $n = 6$ and have size 64×64. Note the adaptive anti-aliasing (smoothing of corners by gray pixel values instead of having only black and white pixels) provided by the photon-based approach. Note also how more photons are needed at the higher resolution to achieve good image quality. The scales are different in all the eight plots; black is always zero but each image has been scaled so that maximal pixel value is white.

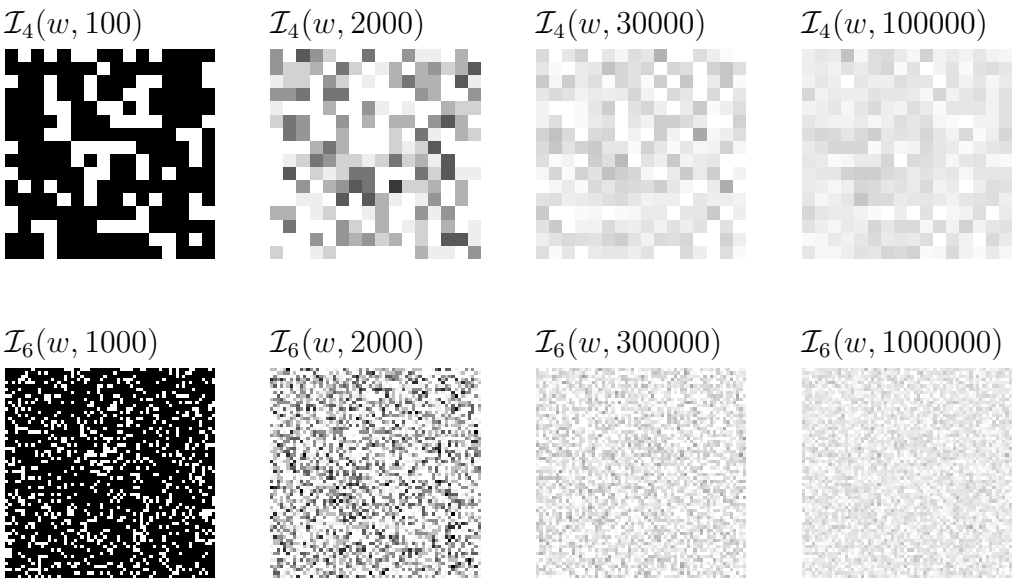

FIGURE 5. Pixel images formed by the exposure $w \in E$ defined in Example 6. The images in the top row are at resolution $n = 4$ and have size 16×16. The images in the bottom row are at resolution $n = 6$ and have size 64×64. The scale is the same in all eight images: black corresponds to zero and white corresponds to pixel value 1.2 (or greater). Note that convergence is slower when resolution parameter n is greater.

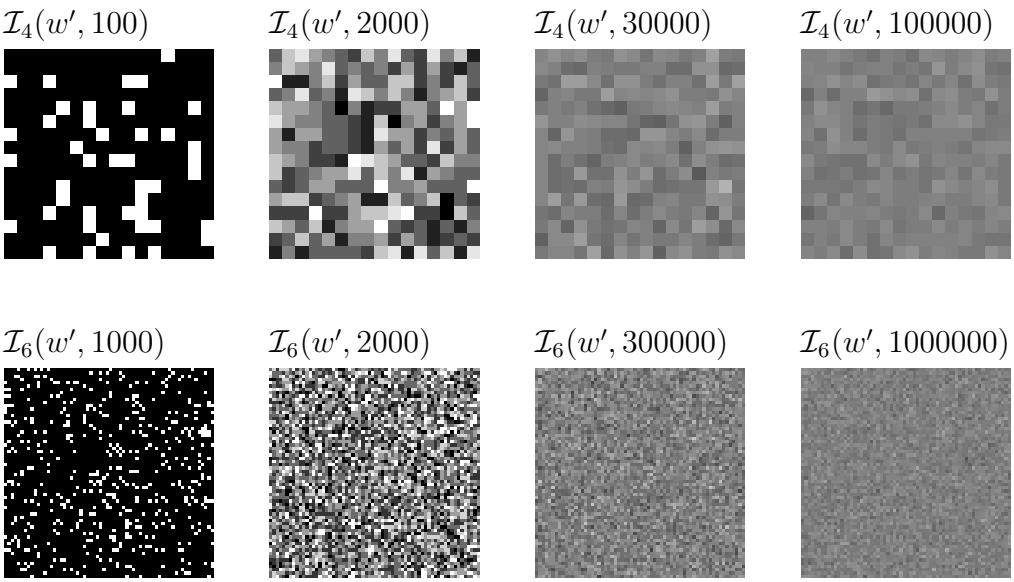

FIGURE 6. Pixel images formed by the exposure $w' \in E$ defined in Example 6. The images in the top row are at resolution $n = 4$ and have size 16×16. The images in the bottom row are at resolution $n = 6$ and have size 64×64. The scale is the same in all eight images: black corresponds to zero and white corresponds to pixel value 1 (or greater). Note that convergence is slower when resolution parameter n is greater. Compare to Figure 5.

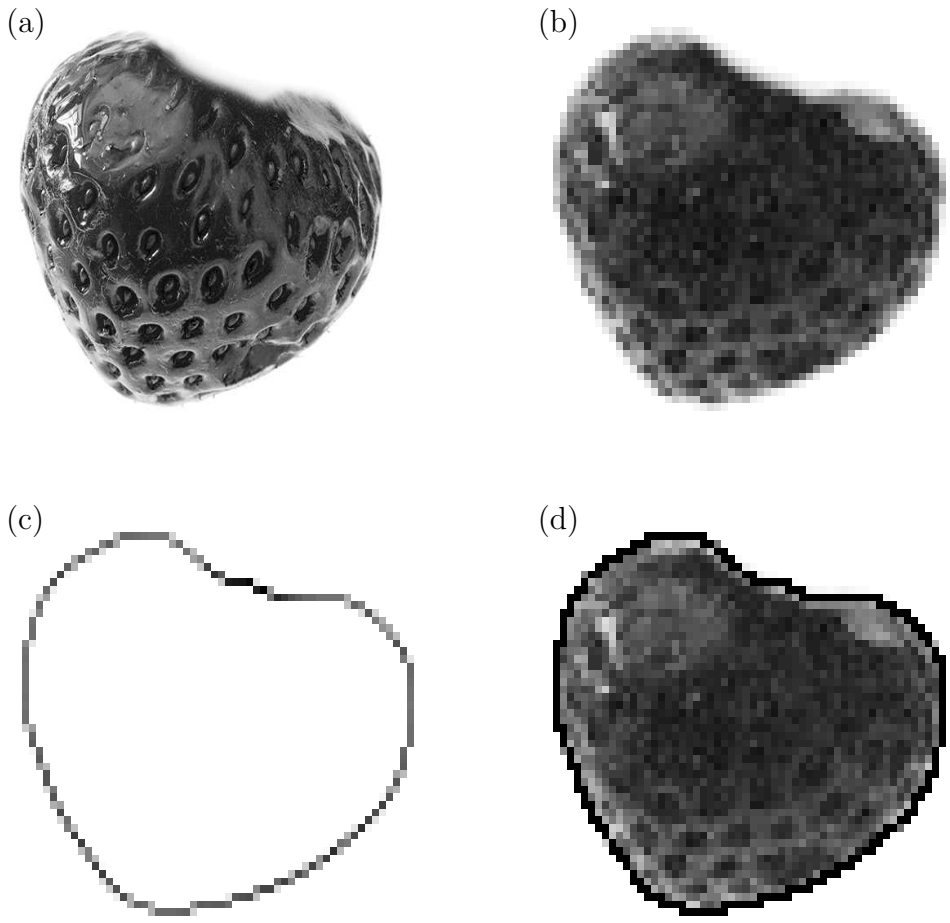

FIGURE 7. **(a)** Monochrome digital photograph of a strawberry. **(b)** Approximate pixel image $\mathcal{I}_6(e, 1500000)$ of size 64×64 computed from random exposure of (a). **(c)** Approximate pixel image $\mathcal{I}_6(c, 1500000)$ of the outline curve. **(d)** Approximate pixel image of an exposure mixing the exposures used to create (b) and (c). This is a low resolution pixel representation of the underlying cartoon image.

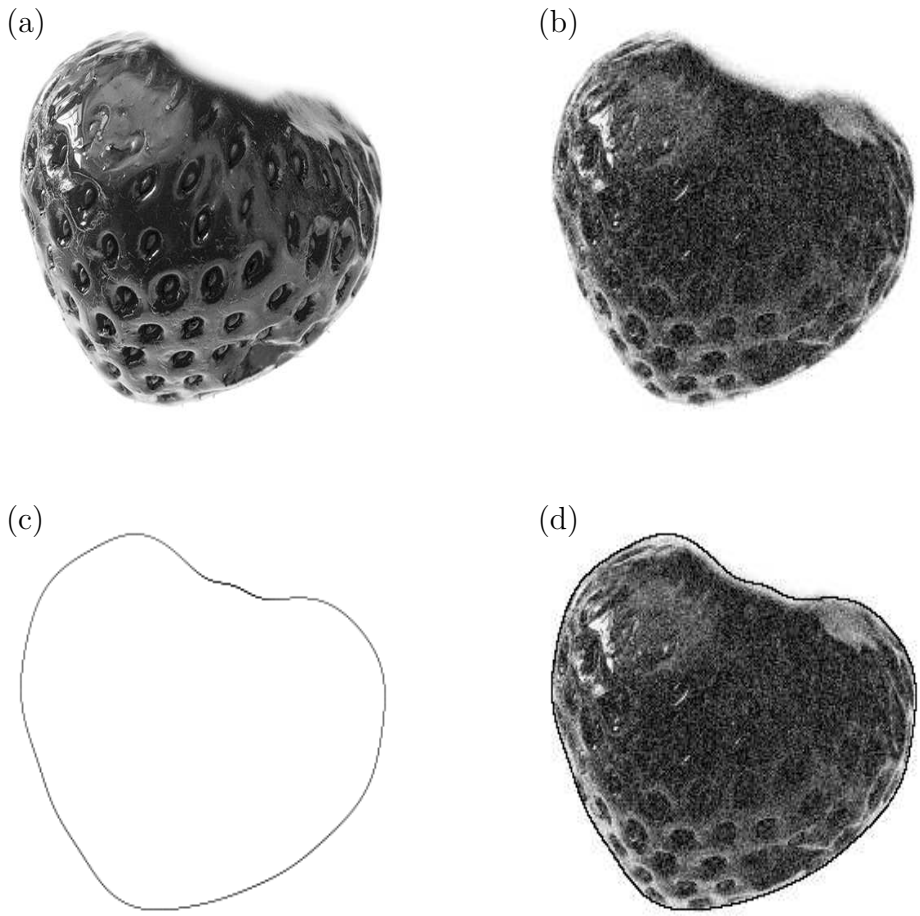

FIGURE 8. **(a)** Monochrome digital photograph of a strawberry. **(b)** Approximate pixel image $\mathcal{I}_8(e, 1500000)$ of size 256×256 computed from random exposure of (a). **(c)** Approximate pixel image $\mathcal{I}_8(c, 1500000)$ of the outline curve. **(d)** Approximate pixel image of an exposure mixing the exposures used to create (b) and (c). This is a higher resolution pixel representation of the underlying cartoon image as compared to Figure 7.

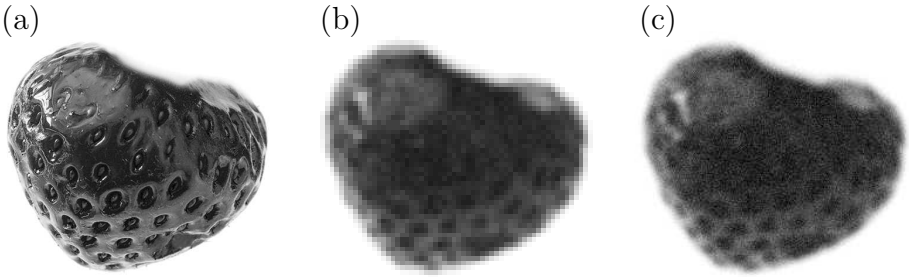

FIGURE 9. **(a)** Monochrome digital photograph of a strawberry. **(b)** Approximate pixel image $\mathcal{I}_6(e', 6000000)$ of size 64×64 computed from the blurred exposure. **(c)** Approximate pixel image $\mathcal{I}_8(e', 6000000)$ of size 256×256 computed from the blurred exposure.

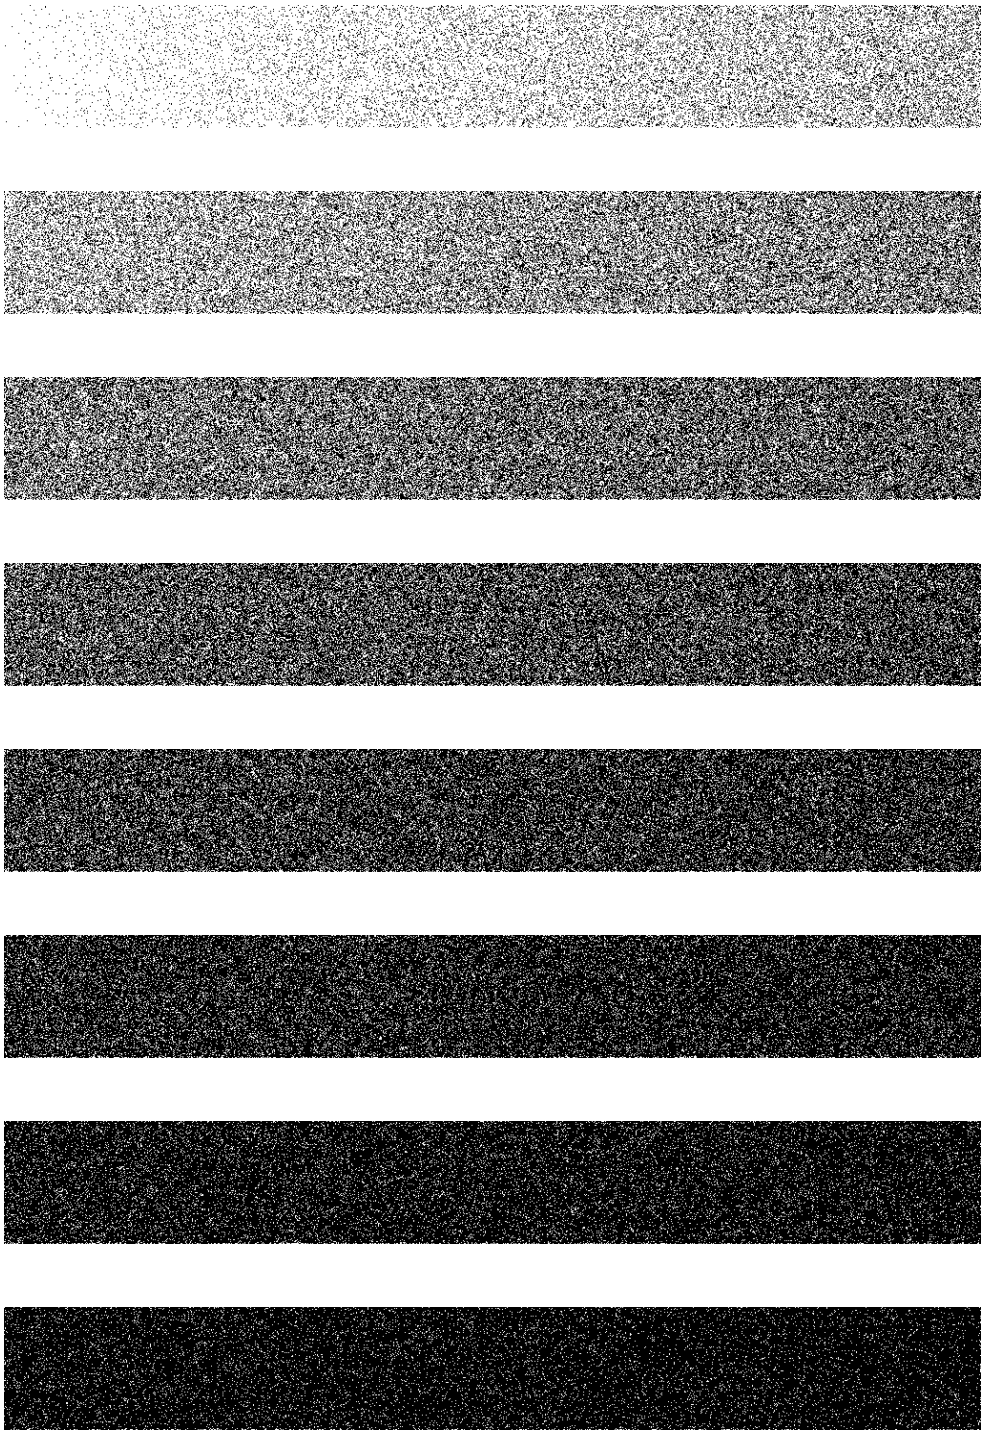

FIGURE 10. Halftoning of a gradient using the stochastic algorithm described in Section 5.2. The size of the image matrix is 128×8192, and 10^7 photons were used.

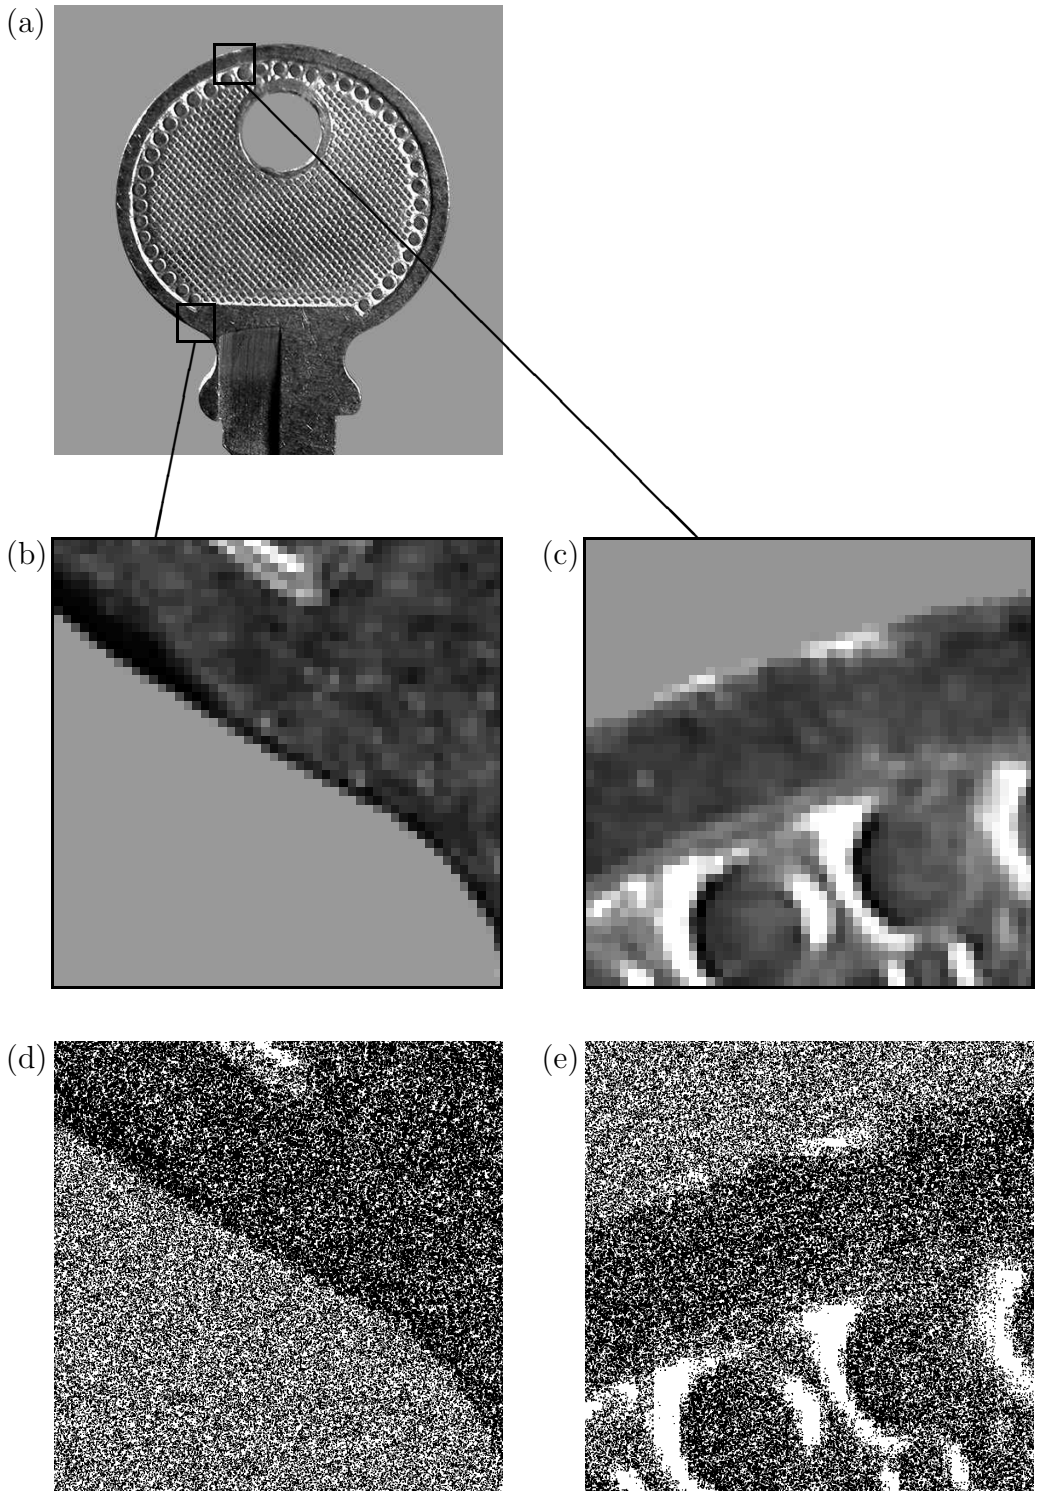

FIGURE 11. Halftoning of a grayscale pixel image using the stochastic algorithm described in Section 5.2. **(a)** The original pixel image of size 512×512. The halftoned binary image (not shown in full size) has size $10^{13} \times 10^{13}$. The number of photons used is 10^7. **(b,c)** Details of the larger pixel image, sizes 52×52. **(d,e)** Parts of the halftoned binary image corresponding to (b,c).

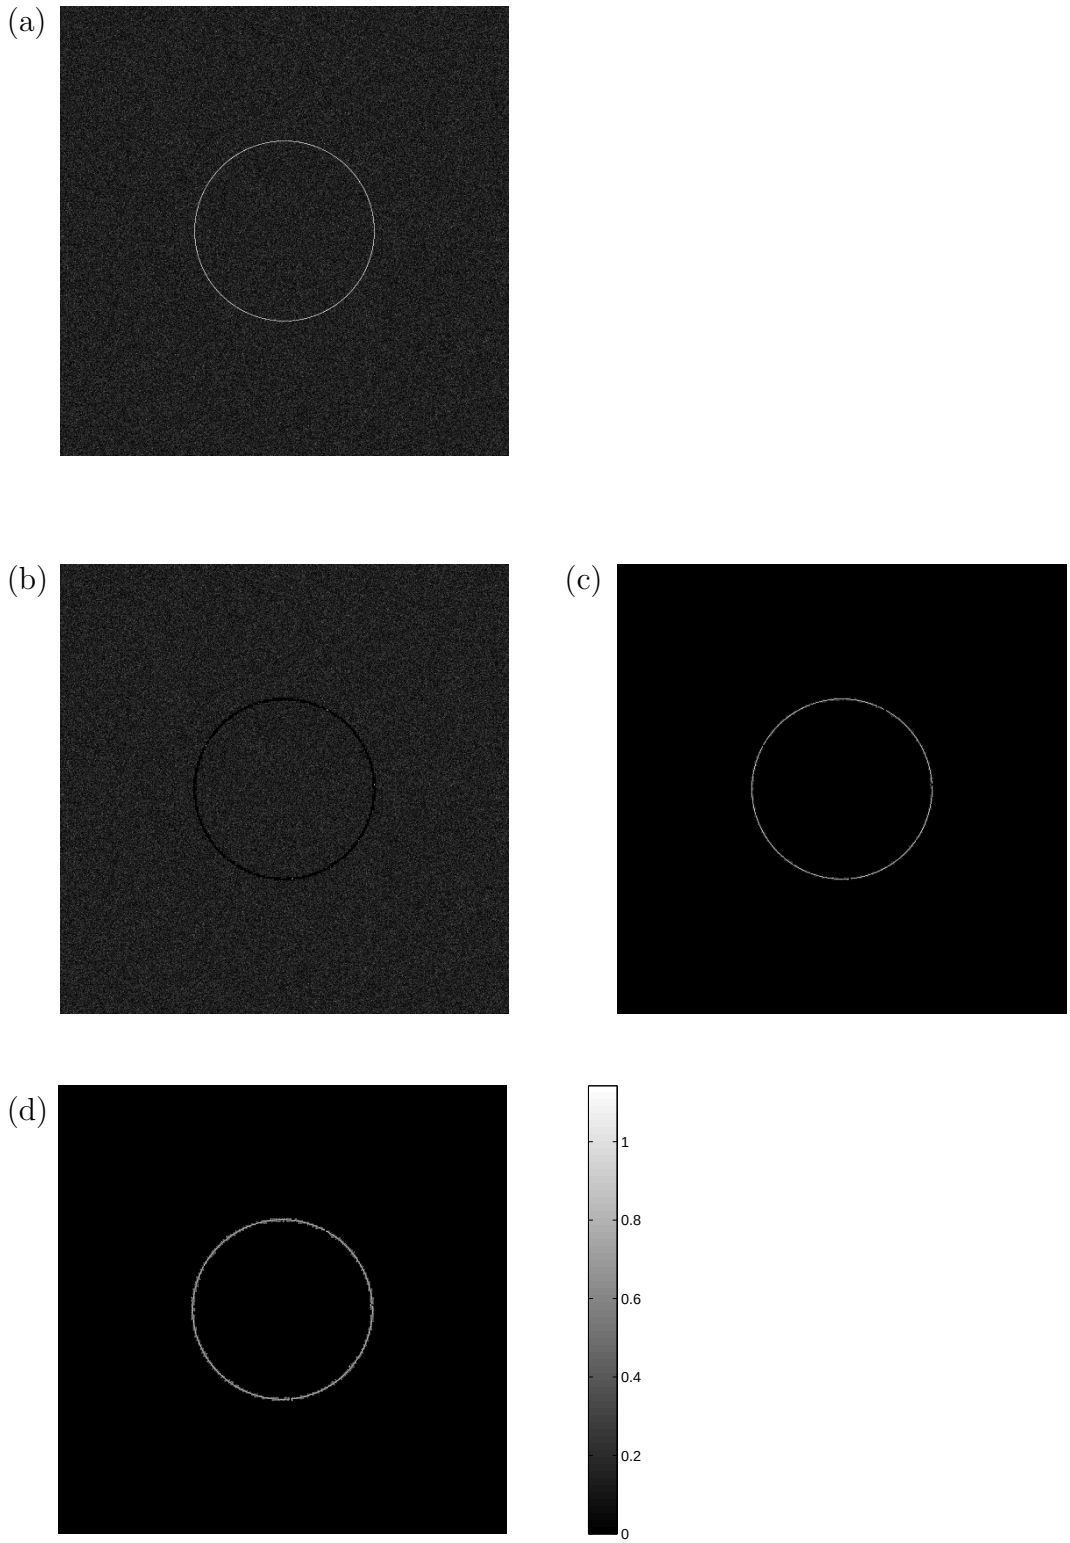

FIGURE 12. A line measure supported on a circle with gray backgroud. **(a)** The original pixel image (J-snapshot). **(b)** The absolutely continuous part of Lebesgue decomposition. **(c)** The singular part of Lebesgue decomposition. **(d)** The inverse value of local dimension computed from singular part (c). The colorbar is related only to (d).

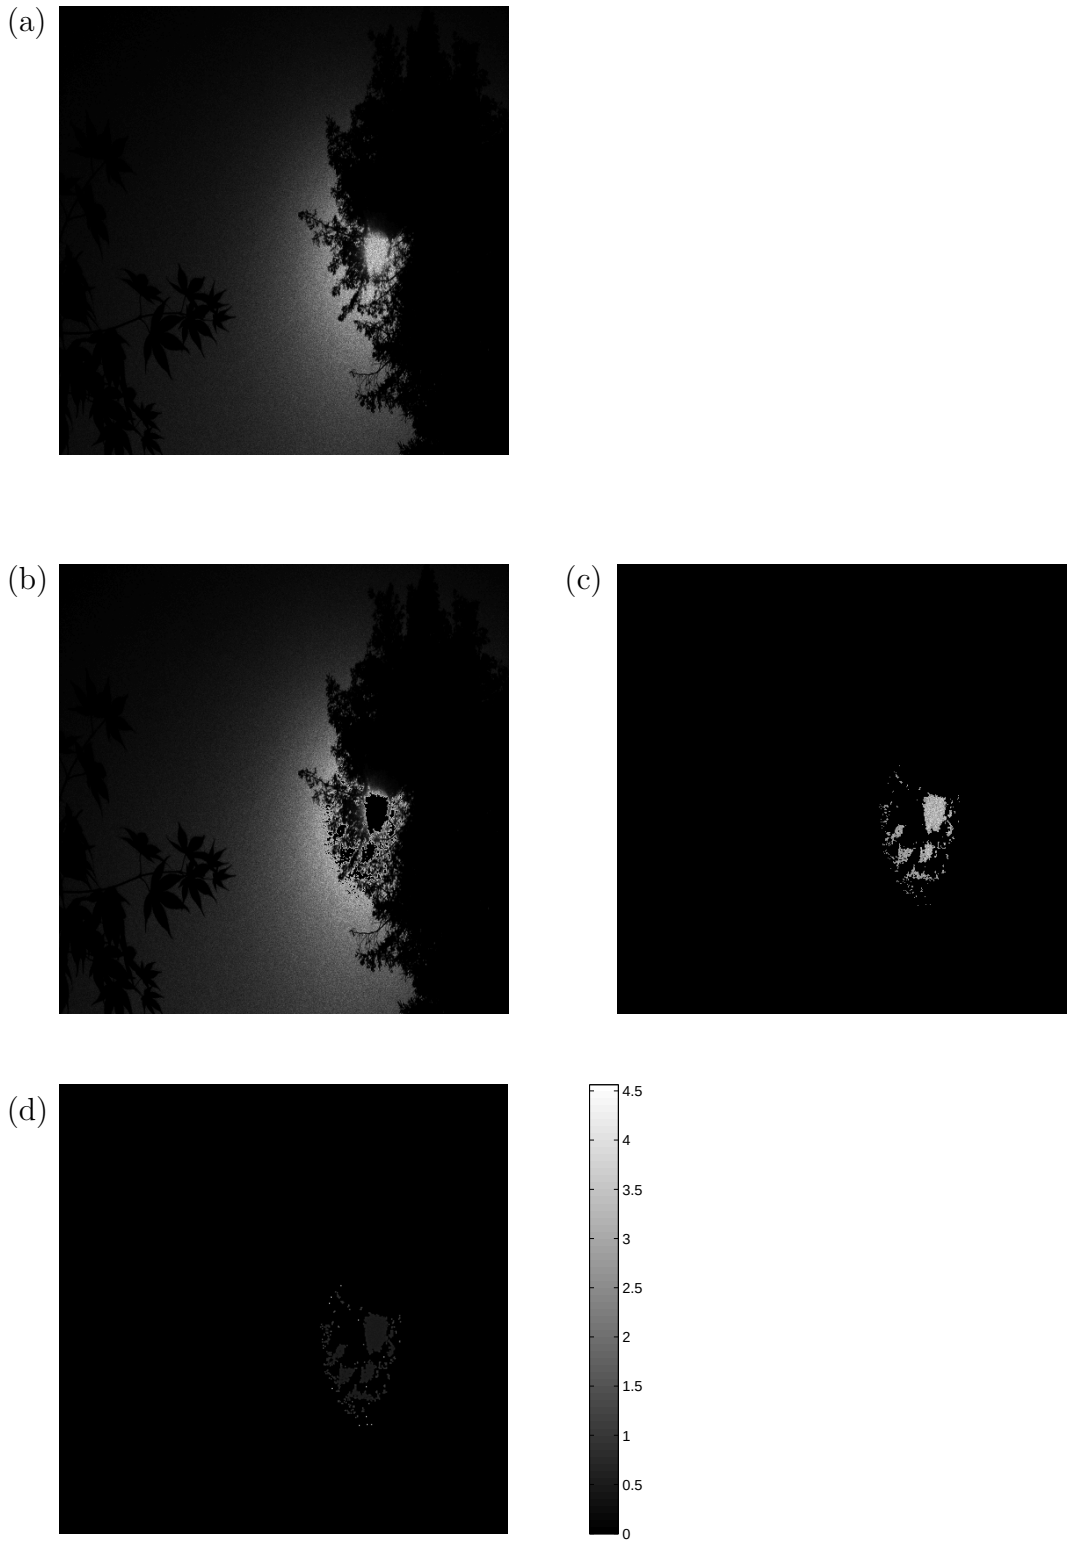

FIGURE 13. Image of sunlight arriving from behind a tree. **(a)** The original pixel image. **(b)** The absolutely continuous part of Lebesgue decomposition. **(c)** The singular part of Lebesgue decomposition. **(d)** The inverse value of local dimension computed from singular part (c). The colorbar is related only to (d).