Density Measure for Line-Drawing Simplification.

Stéphane Grabli¹

Frédo Durand²

François X. Sillion¹

¹ ARTIS*/GRAVIR-IMAG - INRIA ² MIT CSAIL

Figure 1. Simplified illustrations produced by our system using an indication strategy. The initial and simplified versions are respectively on the left and on the right. The objective is to keep a few complex regions at the borders of visually-dense regions to suggest their overall complexity.

Abstract

We present an approach for clutter control in NPR line drawing where measures of view and drawing complexity drive the simplification or omission of lines. We define two types of density information: the a-priori density and the causal density, and use them to control which parts of a drawing need simplification. The a-priori density is a measure of the visual complexity of the potential drawing and is computed on the complete arrangement of lines from the view. This measure affords a systematic approach for characterizing the structure of cluttered regions in terms of geometry, scale, and directionality. The causal density measures the spatial complexity of the current state of the drawing as strokes are added, allowing for clutter control through line omission or stylization. We show how these density measures permit a variety of pictorial simplification styles where complexity is reduced either uniformly, or in a spatially-varying manner through indication.

1. Introduction

Line drawing can produce legible renditions of complex scenes with a remarkable economy of means. In re-

cent years, the field of Non-Photorealistic Rendering has proposed a variety of techniques to produce line drawings from 2D and 3D inputs. However, when the scene complexity grows, the resulting images may suffer from clutter as too many lines are drawn on a small area. This problem is raised by complex structures, such as brick walls, tiled roofs, trees, and becomes more pronounced as these structures are viewed at grazing angles. In contrast to photography, drawings afford omission of details or abstraction, and artists have developed a number of pictorial techniques to prevent clutter while preserving shape and information. For example, they omit structures that are too small, exploit repetition in the scene, and omit texture detail. They carefully control the local amount of strokes, or *density*, in order to avoid clutter, focus attention, and create dynamism.

In particular, repetitive or semi-repetitive structures such as texture, vegetation, or clusters of similar objects raise intriguing cognitive and pictorial issues because of the high clutter they generate and because similarity might be emphasized or exploited. We identify two pictorial strategies used by artists to address clutter in line drawing of repetitive or near-regular structures. They differ in their focus (emphasize vs. exploit repetition) and visual style (uniform vs. spatially-varying drawing complexity).

Uniform pruning ensures low complexity by omitting lines homogeneously. This leads to a picture of uniform density where the original view complexity is com-

* ARTIS is a research project in the GRAVIR/IMAG laboratory, a joint unit of CNRS, INPG, INRIA and UJF

pletely hidden. Regularity is emphasized in that the depiction of the regular structure is regular as well. In practice, uniform pruning can be achieved through level of detail, where each pattern is simplified by omitting secondary strokes; or through sub-sampling of the patterns where entire patterns are omitted, for example drawing every other line in a grid.

Indication exploits repetitive structures and relies on non-uniform simplification to lower the overall complexity but suggest the full complexity in small regions [16]. The artist draws in full detail only a few parts of a repetitive structure so as to suggest its overall complexity, for example, a few tiles on a roof. The objective is to preserve enough pattern structures to convey information in all its complexity.

These strategies can be combined with semantic information to emphasize important parts of the drawing through selective simplification.

We believe that careful control of clutter is fundamental for compelling computer depiction. For this, it is crucial to devise simplification approaches, but also systematic tools to estimate complexity in the view and in the drawing. Previous approaches have presented powerful methods to control tone and complexity in a line drawing. However, they often rely on manual specification of omission strategies and density thresholds. Our work focuses on the automatic determination of complexity, repetition, and simplification strategies for clutter control in line drawing.

Contributions

We present a general approach to control line-drawing simplification based on line omission and stylization. We introduce measures of density to quantify visual complexity in a drawing. We first define an *a-priori* density information that captures how visually complex the drawing will be if all the input lines from the current view are drawn. It is computed at multiple scales and orientations and can be exploited to analyze local structure in order to exploit and prevent cluttered repetitive patterns. Next, during rendering, we estimate the current drawing complexity through a so-called *causal* density. This density information is updated each time a stroke is added in a way similar to Salisbury et al. [10]. Omission or stylistic decisions can then be taken from this information to finely control the final image complexity. We also show that rendering may need to be preceded by a stroke prioritization since the order in which strokes are drawn matters. This relates to prioritized stroke textures [16, 10] and tonal art maps [14, 7] where the ordering is defined manually. We show how these tools can be used to achieve different pictorial clutter control. In particular, we demonstrate both uniform pruning and indication.

We emphasize that both *a-priori* and *causal* densities are necessary for appropriate clutter control and that it is their

combination that provides fine simplification. The *a-priori* density permits the planning of simplification by analyzing the pattern of potential strokes in a region of the picture. However, during rendering, the *a-priori* density does not provide information about the current drawing. In contrast, the *causal* density provides up-to-date information about the current drawing but does not have the ability to look ahead and exploit structure as finely. Their combination provides comprehensive information and affords powerful pictorial clutter control.

2. Related Work

Line-art illustration has received much attention in NPR [5, 13]. The related clutter issues have been addressed by early papers such as the work by Winkenbach et al. [16] who introduce the notion of *indication*, where complex textures are drawn fully only at a few locations to suggest the complexity of the pattern but reduce clutter. They introduce and leverage the notion of prioritized strokes [16, 10], which can be seen as a pen-and-ink textural half-toning patterns. Strokes that form a texture are manually sorted by order of importance and a given tone is obtained by drawing strokes in order of importance until the right intensity is reached. Prioritization ensures that the most salient features of the texture are drawn first. Our approach builds on this work and extends it from textures to general scenes.

The notions of clutter, density and tone are quite related and other NPR techniques have built on half-toning to produce a tone with stroke primitives, e.g. [9].

Deussen et al. [3] propose a simplification approach dedicated to vegetation and trees. Based on the tree hierarchy, complex groups of objects, such as leaves, are replaced by simpler primitives. In addition, a threshold on the z-buffer allows them to render only edges with large depth discontinuities. This permits powerful simplification but heavily relies on the hierarchical representation of vegetation and on the z-buffer.

The work closest to ours is by Wilson et al. [15]. They use line omission to generate lighter drawings of complex scenes. They also use an estimation of the potential drawing density and rely on strokes prioritization to drive omission. We introduce more advanced density information and local structure analysis, and show how the combination of *a-priori* and *causal* density provides comprehensive control.

2.1. Overview of the approach

Our approach to clutter control works in the context of NPR line drawing. The method takes as input a large set of line primitives that is a superset of the final drawing. These lines can come from a 2D or a 3D source (silhouettes, etc.). We will see that the analysis of this input set of lines allows

us to plan simplification and extract local patterns. Our simplification schemes rely only on line omission and the modification of line attributes. We do not address here the topological modification of strokes.

We assume a rendering method where lines are processed sequentially: Each line is stylized into a stroke and rendered, and we proceed to the next one. This sequentiality matches the process of drawing where the artist can see the current state of the drawing before making decisions about the next stroke. We use a causal notion of density where the current local complexity can be evaluated and affect how and whether to render the next stroke. This is the role of the *causal density*. It is causal in that it is updated after each stroke is rendered, and a given stroke will cause influence only onto subsequent ones. For example, causal density can be used to omit a given stroke if the current local density is too high. Causal density is described in Section 4.

Causality makes important the order in which the lines are drawn. As a result, it can be crucial to treat strokes in an appropriate order so that, e.g. the “important” strokes are rendered first and are less likely to be omitted.

This causal density information is essential to control the actual drawing complexity. However, it is limited in its ability to plan ahead and exploit repetitive structures. For example, the indication strategy suggests global complexity by drawing only some regions in full detail, usually at the boundary of the pattern. Therefore we introduce the *a-priori* density information that measures the visual complexity of the “potential” drawing, made of the entire input set of lines. The a-priori density is evaluated at the beginning of the drawing process, before the sequential rendering occurs. An informal interpretation is that it corresponds to the preliminary knowledge that an artist has of the scene. The a-priori density information and the corresponding pattern analysis are described in Section 3.

It is important to note that these quantities can be queried at arbitrary locations, scales, and orientation in the case of the a-priori density. Simplification and stylization decisions can be based on a simple density criterion at the stroke location, but it can also involve a more complex analysis based on density queries at different scales, locations or orientations. In this paper, we show how simple yet powerful analysis tools can be defined using such queries.

Finally, the philosophy of our approach is to separate stylistic decisions from technical ones. We provide density *information*, and we leave it to the user to decide exactly how to exploit this information, in the context of programmable NPR styles where the user can define *shaders* that drive stylization and omission of lines [6]. The various strategies proposed in this paper to query and exploit density are meant as illustrations of the power of this information, not as hard-coded drawing styles. In this work, we perform clutter control using line omission or by control-

ling the attributes of a stroke, but more advanced line simplification could exploit our density measures, e.g. [12].

3. A-Priori Density

While the notion of drawing complexity (or density) is intuitive, its precise definition requires care; Notions of normalization and scale must be carefully treated. In addition, because of the one-dimensional nature of lines and strokes, their orientation plays a major role. For example, if we have one vertical line in the middle of a large set of horizontal lines, we might want to treat these differently. In Figure 9 for instance, we used directionality to distinguish the horizontal and the vertical bars of the grid, in order to avoid removing horizontal lines, due to the high number of vertical crossing lines. We define an estimator for line density from these considerations.

To address these requirements, we borrow inspiration from image decompositions such as steerable pyramids [4], and represent the notion of density at different scales and orientations.

3.1. A Line Density Estimator

Intuitively, we define density at a given point and at scale s as the sum of the length of the lines included in a circle of radius s normalized by the area of the circle. This normalization is important to ensure scale independence. This density is measuring a length by unit surface. To understand the effect of the scale s , consider a regular pattern of vertical lines. When the radius s grows, the length of each line inside the disk grows linearly, and the number of lines that intersect the disk also grows linearly. The resulting quadratic growth in length is compensated for by the normalization by the disk area, and density is mostly independent of s .

In practice, we use a spatially-weighted average with a circular Gaussian kernel of variance σ . We also decouple information about different orientations and define density for a given direction $\vec{u}$ using a falloff w_o on the direction of the line. The density of a set of lines $\mathcal{L}$ at a point Q , for scale σ and orientation $\vec{u}$ is then:

$$d(Q, \sigma, \vec{u}) = \int_{P \in \mathcal{L}} w_d(P, Q, \sigma) w_o(P, \vec{u}) dl \quad (1)$$

where w_d is the normalized circular Gaussian function of standard deviation σ :

$$w_d(P, Q, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{\|\vec{PQ}\|^2}{2\sigma^2}} \quad (2)$$

and w_o the orientation weighting function that depends on $\theta(P)$, the angle between the line tangent at point P and $\vec{u}$:

$$w_o(P, \vec{u}) = \begin{cases} |\cos(\frac{n\theta(P)}{2})|, & \text{if } \theta(P) \in [-\frac{\pi}{n}, \frac{\pi}{n}], \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Figure 2. A-Priori Density Map Pyramids for the house model (drawn in negative with a higher gamma for viewing convenience). Each column corresponds to a different direction and shows three levels (the base, $\sigma = 2$ and 5 pixels) of the corresponding Gaussian pyramid. The fifth column shows the a-priori density map with respect to all four directions.

n controls the range of angles a given point of $\mathcal{L}$ contributes to. This falloff function is used to ensure proper normalization when the directions are discretized. It is related to the steerable interpolation weights [4].

As discussed above for the simple case of the disk and vertical lines, we use a normalized version of the Gaussian function for this estimator, so as to ensure that the mean lines density value for the drawing is the same whatever σ was used to define the scale of each punctual estimation.

3.2. A-Priori Density Maps

For efficient computation and access, we follow the pyramid [1] and steerable pyramid [4] approaches and store the density measure for dyadic scales and a number of orientation bands. The illustrations of this paper were done using four directions, whose angles with the horizontal are $(0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4})$. This has proved sufficient for characterizing most structural aspects. We also set $n = 4$ in formula 3, equal to the number of directions; this way a segment contributes at most to two consecutive directional maps. Therefore, each map is associated to a specific direction and scale, and stores the line density measured at pixel positions, with respect to this direction and this scale.

Figure 2 shows a subset example of such maps for the house model, whose initial arrangement of lines can be seen in Figure 1. In addition to the four directional pyramids of images, we store an extra pyramid for the omni-directional

a-priori density that encodes the contributions of all lines regardless of their orientation.

The rendering of the pyramids basis maps is made using OpenGL: for each of the four directions, each line segment is given a color reflecting its directional weight w_o with respect to the considered direction (see section 3.1) and rendered using additive blending. We then build Gaussian pyramids from each of the four basis maps.

The most direct use of this information is the identification of dense areas of a drawing, with respect to a given scale and orientation, available through simple queries. Such areas correspond to dark regions in Figure 2.

A-priori density could already be used to drive line omission, for example by assigning each stroke an “omission probability” inversely proportional to its a-priori density, as illustrated by Figure 3. This technique allows for controlling the final mean density, but does not provide any fine control over strokes placement, as afforded by the causal density (see section 4)

3.3. Density variations across space, scale and orientation

We now show how the oriented multi-scale information contained in our a-priori density pyramids can be used to analyze structure and enable simplification strategies such as indication. For this, we borrow inspiration from image and pattern analysis. As seen in formula 1, the a-priori den-

Figure 3. A-priori density is used to drive line-omission: the probability for each line to be dismissed is inversely proportional to its density. This provides coarse control over the mean final density but no fine control over the placement of the strokes, leading to this “random” look. The causal density (section 4) is dedicated to this objective, as demonstrated by Figure 8 for instance.

sity is a 3-parameterized function of space (x,y) , scale (σ) and direction $(\vec{u})$. Important structural aspects of drawing complexity can be revealed by studying the variation of density along each parameter dimension. To gain insight, we fix two parameters out of three and study the density response when the third parameter varies.

Spatial boundaries In many simplification schemes, the boundaries of dense areas receive special treatments. In particular, indications tend to be located in these regions (see section 5). Boundaries can be identified by studying the spatial derivative of density $d_{\sigma_0, \vec{u}_0}(P)$, at a given scale σ_0 and for a given direction $\vec{u}_0$. Similar to edge detection, this can be done by computing the gradient image of the map of scale σ_0 and direction $\vec{u}_0$. More advanced detection schemes inspired from image processing could also be used.

Directional distribution As mentioned above, different pictorial strategies might be used depending on whether the complexity of a dense region is due to lines directed along a few preferred directions or whether no principal direction stands out. We say that density is *anisotropic*, when one or few directions dominate, *isotropic* otherwise. The degree of isotropy can be characterized by studying $d_{\sigma_0, P_0}(\vec{u})$, at a given scale σ_0 and for a given point P_0 . In Fig. 4 we show how the profile of the function $d_{\sigma_0, P_0}(\vec{u})$ varies depending on the arrangement of the surrounding lines, which permits the determination of the isotropy or anisotropy of the density. In particular we emphasize the differences between two anisotropic profiles measured for pixels of the house model’s roof, and put them in contrast with an isotropic profile observed on a pixel of the tree model.

Scale When dealing with repetitive structure, it is often important to characterize the scale of a pattern. This is also

Figure 4. Scale and directional profiles of the a-priori density. Top: four pixels chosen at particular positions on the roof and tree images. Middle: The scale profile of the a-priori density for pixels 1,2,3. Although having very different densities at low scales, the three curves converge toward a plateau at $\sigma = 4$. Bottom: Directional profiles for pixels 1,3,4. The two first profiles (from the roof image) characterize a high anisotropic density respectively along the horizontal and $\pi/4$ directions. The last profile (from the tree image) corresponds to a high isotropic density.

crucial in our case because lines are essentially Dirac distributions, and the scale at which density is queried can have a strong influence on the result: at the smallest scale, the image is essentially a binary function that corresponds to the presence of lines. We show that a fine scale analysis can be carried out, which permits precise selection of the lines to omit. For example, we can draw a roof with one tile out of two, or with a few groups of n tiles. For this, it is essential to have a systematic way to extract the characteristic scale in order to make the simplification criteria independent of the scene.

The scale of dense repetitive areas can be characterized by studying $d_{P_0, \vec{u}_0}(\sigma)$ at a given point P_0 and for a given direction $\vec{u}_0$. In the case of a complex region made of small repetitive patterns, this function contains a plateau: since

our density estimator is normalized by the area of its support region (section 3.1), it gives the same result at multiple scales for periodic structures, as soon as the scale is higher than the pattern size and smaller than the distance to the region boundary. Thus, the pattern size can be inferred from the σ value at which this plateau begins (just as the frequency of the line distribution in the area). Likewise, the distance to the structure’s boundary can be obtained from the σ value at which the plateau ends. Figure 4 shows how the scale profiles for three different points of the house model’s roof, converge toward such a plateau, though being very distinct at small scales. In Figure 8, we use this property to retrieve the size of a tile, in order to drive line omission at two different levels. More details about this illustration are given in section 5.

So far, we have discussed density and variations at a single spatial point. For 1D primitives, this information can be aggregated over a line using a variety of simple statistical tools such as mean, min, max, or variance, depending on the application. It is then possible to extract the same spatial, directional, and scale characteristics for a line, i.e. to tell for each line in the drawing whether it belongs to a dense region, whether this region is dense isotropically or anisotropically, etc. In this way, lines can be precisely characterized, allowing for advanced omission and stylization.

4. Causal Density

The a-priori density discussed thus far permits the analysis of the potential drawing if all the lines were drawn. It provides powerful analysis tools, but unfortunately does not allow for fine control of the final drawing’s appearance, as illustrated in Figure 3. In particular, it cannot guarantee that the density of the actual drawing does not exceed a given threshold, nor that no pair of strokes is too close in the image. This is why we also use a *causal* density that complements the a-priori density and reflects the current state of the drawing. It is updated after each stroke is drawn and can be used to stylize or decide to omit subsequent strokes.

4.1. A Stroke density estimator

The causal density estimator works on the arrangement of strokes rendered in the current drawing. We chose to use the standard normalized Gaussian function of standard deviation σ , convolved with the luminance image I of the drawing as our estimator. Similar to the definition of the a-priori density in eq. 1, the estimation of the stroke density in image I at a point Q can then be written as:

$$d(Q, \sigma) = \int_{P \in I} w_d(P, Q, \sigma) I(P) dP \quad (4)$$

where w_d is the Gaussian function, defined by formula 2. It indicates how “dark” the drawing locally is with respect to the given scale defined by σ . Each added stroke contributes to “darken” the image an amount that depends on its color, size, thickness and on the scale σ .

As for the line density estimator, we use a normalized version of the Gaussian function so that the mean stroke density of the drawing does not depend on the scale at which the queries are made.

In our approach, the causal density can be queried at multiple scales. However, because causal density is refreshed after every stroke, we have not found it beneficial to store it in a pyramid. Queries are implemented by actually integrating the information around an area and queries at a large scale are more costly.

For performance reasons, we have chosen not to encode causal density at multiple orientations. In contrast to the line density estimator, it does not take directionality into account. Nonetheless, a directionality dependency might be taken into account using the directional information afforded by the a-priori density information.

4.2. Causality and Stroke Ordering

As pointed out by Winkenbach et al. [16] and Salisbury et al. [10], stroke prioritization is crucial when stroke omission is used to control tone or density in the final image. Because our second measure of density is causal, the first strokes have more influence on the rest of the drawing and it can be important to make sure that the important ones are drawn first. For example, when using line omission, the most important lines have to be drawn first to minimize their chance of being dismissed.

This ordering depends on stylistic choices. In our approach, it can be defined using any type of information on the line, 3D or 2D, including the a-priori density estimation. Figure 5 shows two simplified versions of a sunflower obtained using causal density. In the first one, the ordering gives the highest priority to strokes with high depth discontinuities. In contrast, in the second picture, strokes are drawn in an arbitrary order. Notice in particular that, in the ordered version, the limit of the flower’s center appears more pronounced and the shape of the seeds is better suggested.

Defining pertinent criteria for feature-line ordering is an exciting avenue of future work.

4.3. Stroke Omission, Density and Regularity

A large class of simplification strategies are based on line omission. This can be achieved using causal density, by testing, for each stroke, the measured density against a thresh-

Figure 5. It is essential to order the lines prior to using causal density. (a) all the visible lines. (b) use of causal density for line omission when lines are ordered using depth discontinuity, and (c), without line ordering.

old τ to decide to draw it or not. The corresponding naïve algorithm is:

```

for each stroke  $s$ 
   $d = \text{density of the drawing under } s \text{ at scale } \sigma$ 
  if  $d < \tau$ 
    draw( $s$ )
  
```

Two parameters control this algorithm: the scale σ and the threshold τ .

The resulting drawing can be qualified in two terms, the average *final density* (intuitively the total number of lines drawn) and the *regularity* of the distribution of lines.

The objective of simplification is to get a smaller density than initially. The final set of strokes can be seen as a subsample of the input lines. As with any sampling approach, the notion of regularity is important to characterize how spatially-uniform the final distribution is. For a target density threshold τ tested at a given scale σ , different drawings can be obtained from the same arrangement of lines, depending on the order in which lines are processed. A drawing where lines appear more evenly-spaced is said to be more regular than one displaying clusters of tight lines. Figure 6 illustrates this property. This notion of regularity is related to the Fourier spectrum of non-uniform sampling patterns used in antialiasing [8] and to the discrepancy used to study Monte-Carlo integration [11], but we will leave the discussion at an informal level.

The effects of the two parameters σ and τ on the result, i.e. the final average density and the regularity of the drawing are important and deserve a discussion. Consider a situation where a simple threshold τ is used on the causal density queried at a given scale σ . Table 1 summarizes the dependencies between these different values by showing how regularity and the final average density vary when either the threshold τ or the scale σ are increased.

We informally discuss the results of the last row of table 1, which are obtained when σ increases while τ stays

Figure 6. Despite the difference in the distribution of the lines observed in these two images, the density (computed at the scale of the gray circle at its center) is the same. The lines distribution on the right is said to be more regular than the one on the left.

parameter	test: $d < \tau$		test: $d < \frac{\tau}{\sigma^2}$	
	regularity	density	regularity	density
$\tau \nearrow$	$\searrow$	$\nearrow$	$\searrow$	$\nearrow$
$\sigma \nearrow$	$\searrow$	—	—	$\searrow$

Table 1. Effects of parameters scale σ and threshold τ , on regularity and final average density.

constant. It is first important to notice that, since we use a normalized definition of density, the density estimation does not change with σ when the scale is above the characteristic scale of a repetitive pattern, (see section 3.3). Then, we observe that when the scale is increased, the precise location of lines is less constrained, leading to more irregular sets of lines. Indeed, The threshold τ is now enforced on larger neighborhoods as defined by σ . This effect is illustrated in Figure 7.

Figure 7. Effect of increasing the scale σ on the regularity of the lines distribution. Right: $\sigma = 2, \tau = 0.2$. Left: $\sigma = 7, \tau = 0.2$. The final density is the same in both illustrations, but only the image on the right contains pairs of close lines.

On the other hand, if the threshold is inversely proportional to σ^2 , which is equivalent to working with a non-normalized version of the Gaussian function, then increasing σ decreases the final drawing density, as the number of strokes allowed on a given area is the same whatever the size of this area. In this case, the regularity remains glob-

ally unchanged as lowering the number of strokes allowed on a given area decreases the chances of clutter for this area.

The desired result of a simplification based on line omission, often consists in having at least a certain amount of free area surrounding each line, which is obtained by setting a scale corresponding to the minimum spacing between two strokes and a low threshold.

5. Applications and Results

We now present results demonstrating how these two density definitions complement each other and permit the implementation of simplification strategies.

The computation times required to generate the illustrations of this paper vary between a few seconds and a few minutes: Using a-priori density is fast thanks to the precomputation of the density maps, on the other hand, the causal density is evaluated upon request each time a stroke is to be drawn and is therefore the most time-consuming operation (in particular at large scales).

Uniform pruning We describe first a simplification strategy where we emphasize regularity and omit lines uniformly. We decide to consider uniformity in object space in order to preserve the geometric regularity of the scene.

Figure 8 demonstrates two levels of uniform pruning on the tiled roof of a house. In both cases, the lines belonging to large regions of high visual complexity are first identified using the omni-directional map of the a-priori density, for scale $\sigma = 4$. Next, from the scale profile (section 3.3) of the a-priori density we set the kernel’s size of the causal density estimator to nearly match twice the size of a tile, for the left image, and twice as much for the middle one. Thus, the images are simplified uniformly, keeping respectively all or half of the tiles.

Figure 9 demonstrates the high control over line omission by simulating the 3D perspective effect through multi-scale queries of the causal density; Our goal is to draw a simplified version (top) of a grid (bottom) in the most natural way: each bar is represented with a single line instead of two, and only half of the vertical bars are drawn. The bars are uniformly distributed along the grid, with a decreasing spacing as we are moving away from the viewpoint, so as to respect the perspective impression. This was done by subjecting the lines to the causal density, the size of the Gaussian kernel depending on the depth of the processed line (the kernel’s size is approximately twice the spacing between two lines in the simplified image). Furthermore, we relied on the directional information provided by the a-priori density to set higher thresholds when the a-priori density is low in the direction of the current stroke. This is required to avoid systematic omission of horizontal lines, whose causal density is obviously higher, due to the potentially high number of vertical crossing lines.

Figure 9. A uniformly simplified illustration (top) of a grid (bottom) obtained using causal density. Each bar is drawn using a single line and exactly half of the bars were homogeneously dismissed. The dismissal thresholds are furthermore direction-dependent so as to avoid the systematic omission of horizontal lines of higher causal density.

Fig. 10 illustrates the use of the various density measures to control clutter but preserve the impression of visual mass. The complex object includes many near-parallel lines (not necessarily straight) and we make a heavy use of directional a-priori density information to treat lines according to the presence of near-parallel neighbors. We first select the lines that are surrounded by nearly-parallel lines, i.e. having a highly anisotropic density, using the a-priori density directional profile (section 3.3). These lines are then ordered by decreasing length, and drawn using causal density for line omission (bottom left). In a second pass, we take care of lines surrounded by a more isotropic density (no preferred direction in the set of neighboring lines) and remove the ones that are in a region of high overall a-priori density at a fairly large scale. Notice how the drawing is successfully simplified (the image on the bottom right shows all the lines that are totally omitted) while retaining a proper meaning. We furthermore use a programmable line shader [6] to thicken the resulting orphan lines.

Indication This simplification strategy exploits repetition by suggesting the overall complexity through few well-positioned detailed parts.

We notice that these indications are most useful at the boundaries of dense regions and use this property to produce the simplified images of the house and tree, figure 1 at the beginning of the paper.

In the house illustration, we first select all lines from large regions of high density by querying the omni-directional map of the a-priori density. The re-

Figure 8. Uniform pruning of the tiled roof. From the scale profile (section 3.3) of the *a-priori* density, the proper kernel size is inferred for the causal density, so as to draw a simplified roof with all the tiles (left) or half of them (middle). The right column shows close-ups on (top) the initial set of lines (middle) the left image, (bottom) the middle image. Notice how, although drawing all tiles, the middle image yet shows less complexity than the initial one.

Figure 11. Intermediate maps used for the house illustration of Figure 1. Left: The omni-directional density map. Right: The gradient image computed on this map. This image is not actually stored as gradient computations are made on the fly.

maintaining is similar for both the house and tree illustrations. The lines are subjected to the causal density with an omission threshold proportional to the value of the gradient computed on the omni-directional *a-priori* density map at a fairly low scale. This way, lines are preferably kept in areas of high gradient, i.e. next to the boundaries of high-density regions. Figure 11 shows the *a-priori* density map that was used to produce the house illustration and the corresponding gradient image.

6. Conclusions and Future Work

We have introduced two measures of drawing density to control and prevent clutter. The *a-priori* density is computed once for all, before rendering, on the full set of lines composing the view. It can be queried at different locations, scales and orientation and permits powerful analysis of local structure. In contrast, our *causal density* is updated as

the drawing is created thereby providing information about the current state of the drawing in order to offer finer control on the final result. Taken in isolation, these two measures provide useful information to drive simplification and omission strategies. But it is really when they are exploited together that they offer a powerful and fine control on the final complexity, regularity, and style of the drawing. In particular, we have shown that they facilitate the exploitation or emphasis of regularity in the view using pictorial strategies such as uniform pruning or indication.

Our work opens several avenues of future research. The strategies that we have proposed are only illustrations of the power of appropriate density information. We hope that subsequent work will propose new styles and approaches to clutter analysis and control. In addition, causal density profoundly raises the issue of stroke ordering: Which feature lines are more crucial for the faithful depiction of an object? We believe that geometry, perception, and pattern analysis must be leveraged to yield pertinent estimates of a stroke's importance. This also suggests that more advanced image processing operations on density and other information about the view should be explored. Density information should finally be combined with semantic or cognitive information such as eye-tracking patterns [2].

Acknowledgments We thank Sylvain Paris for useful discussions. This project was supported by MIT France and an *équipe associée* INRIA.

Figure 10. The visual clutter due to near-parallel lines is limited using the directional profile of the a-priori density (section 3.3). Top left: all visible lines. Top right: Final rendering. Bottom left: long lines with high anisotropic density. Bottom center: short lines or lines of high isotropic density. Bottom right: lines that are omitted in the final image.

References

- [1] P. J. Burt and E. H. Adelson. The laplacian pyramid as a compact image code. *IEEE Trans. on Communications*, COM-31,4:532–540, 1983.
- [2] D. DeCarlo and A. Santella. Stylization and abstraction of photographs. In *Proc. SIGGRAPH*, 2002.
- [3] O. Deussen and T. Strothotte. Pen-and-ink illustration of trees. *Proc. SIGGRAPH*, 2000.
- [4] W. T. Freeman and E. H. Adelson. The design and use of steerable filters. *IEEE Trans. PAMI*, 13(9), 1991.
- [5] Gooch and Gooch. *Non-Photorealistic Rendering*. AK-Peters, 2001.
- [6] S. Grabli, E. Turquin, F. Durand, and F. Sillion. Programmable style for npr line drawing. In *Eurographics Symposium on Rendering*, 2004.
- [7] A. Klein, M. Kazhdan, W. Li, W. T. Corra, A. Finkelstein, and T. Funkhouser. Non-photorealistic virtual environments. *Proc. SIGGRAPH*, 2000.
- [8] E. H. W. M. A. Z. Dippe. Antialiasing through stochastic sampling. *Proc. SIGGRAPH*, 1985.
- [9] V. Ostromoukhov. Digital facial engraving. *Proc. SIGGRAPH*, 1999.
- [10] M. Salisbury, M. Wong, J. Hughes, and D. Salesin. Orientable textures for image-based pen-and-ink illustration. *Proc. SIGGRAPH*, 1997.
- [11] P. Shirley. Discrepancy as a quality measure for sample distributions. *Proc. Eurographics*, 1991.
- [12] M. C. Sousa and P. Prusinkiewicz. A few good lines: Suggestive drawing of 3d models. *Proc. Eurographics*, 2003.
- [13] T. Strothotte and S. Schlechtweg. *Non-Photorealistic Computer Graphics. Modeling, Rendering, and Animation*. Morgan Kaufmann, 2002.
- [14] M. Webb, E. Praun, A. Finkelstein, and H. Hoppe. Fine tone control in hardware hatching. In *Proc. NPAR*, June 2002.
- [15] B. Wilson and K.-L. Ma. Rendering complexity in computer-generated pen-and-ink illustrations. In *NPAR*, 2004.
- [16] G. Winkenbach and D. Salesin. Computer-generated pen-and-ink illustration. *Proc. SIGGRAPH*, 1994.