



HAL
open science

Gabor Noise Revisited

Vincent Tavernier, Fabrice Neyret, Romain Vergne, Joëlle Thollot

► **To cite this version:**

Vincent Tavernier, Fabrice Neyret, Romain Vergne, Joëlle Thollot. Gabor Noise Revisited. j•FIG 2018
- Journées Françaises d'Informatique Graphique, Nov 2018, Poitiers, France. pp.1-4. hal-01926451

HAL Id: hal-01926451

<https://hal.science/hal-01926451v1>

Submitted on 20 Nov 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Gabor Noise Revisited

V. Tavernier, F. Neyret, R. Vergne and J. Thollot

Univ. Grenoble Alpes, CNRS, Grenoble INP (Institute of Engineering), LJK, INRIA, 38000 Grenoble, France

Abstract

Gabor noise ingredients — points distribution, weights, kernel — can be changed. We show that minor implementation changes allow for huge 17 – 24× speed-up with same or better quality.

Keywords: textures, procedural, noise

1. Introduction

Procedural noise texturing is an important element of the CG toolbox. Gabor noise [LLDD09] and its numerous follow-ups [BLV*10, LLD11, GLLD12, GSV*14, GLM17] proposed a way to procedurally generate stochastic noise instances with controlled spectral properties. It belongs to *sparse convolution* methods [Lew89, vW91], and as such, it combines 3 ingredients: one or several splat kernels, a large set of random splat positions, and a random weighting of splats. The result is an instance of a *Gaussian texture* (*i.e.* having random uncorrelated Fourier phases), which process reproduces the Power Spectrum Distribution (*PSD*) of the kernel and is normalized to unitary variance.

This procedural method is especially powerful: the look can be controlled by interactive GUI-based PSD design [LLDD09] or example-based automatic PSD reproduction [GLLD12]; it applies to curved surfaces [LLDD09] as well as to solid texturing [LLD11]; and, as a procedural noise, is able to generate infinite spans of zoomable textures.

However, this approach is too costly for being used in interactive applications, as it requires 30 to 100 splats per pixel depending on the target quality. Another problem arises concerning resulting texture properties. If these are known and controlled for the process itself, the quality and normalization of the resulting *instances* are not ensured. Indeed even subtle statistic variations may lead to drastic changes in the final appearance, as noises are usually non-linearly post-treated in real-world shaders.

In this paper, we revisit each ingredient of the Gabor noise method – kernel, point distribution, weighting. For each, we envision the alternatives and discuss their impact on performances and quality based on objective criteria.

2. Gabor noise ingredients

Gabor noise is defined as the convolution of a Poisson point distribution with the real part of a Gabor kernel mul-

tiplied by uniform random weights. The seminal Gabor noise [LLDD09] is defined in practice as follows for one base kernel $g(\mathbf{x})$ at location \mathbf{x} :

$$n(\mathbf{x}) = \sum_{\{\mathbf{x}_i\}} w_i g(\mathbf{x} - \mathbf{x}_i)$$

$$g(\mathbf{x}) = K e^{-\frac{\pi}{r^2} \|\mathbf{x}\|^2} \cos(2\pi F_0 \mathbf{x} \cdot \boldsymbol{\omega}_0) \quad \text{if } \|\mathbf{x}\| \leq r,$$

where $\{\mathbf{x}_i\}$ are Poisson point process impulses, and w_i are random weights (*i.i.d.* $W \sim \mathcal{U}(-1, 1)$). The user parameters control the magnitude $K \in \mathbb{R}^+$, the spatial frequency $F_0 \in \mathbb{R}^+$, the main orientation $\boldsymbol{\omega}_0 \in \mathbb{R}^2$, the kernel radius $r \in \mathbb{R}^+$, and the splat density (*i.e.* splats per pixel, or impulses per splat) λ . In practice, complex textures are obtained by combining several $n_j(\mathbf{x})$ functions (with their own set of parameters). For efficiency, implementations rely on virtual cells of size r and impulses are generated cell-wise, so that only cells that may contribute to the current pixel are evaluated. In the following we thus consider the number N of impulses per cell rather than λ ($\lambda = \pi N$).

The minimal N to be used in order to get a given target quality is let to user trial-and-error, and the alternatives for the different components (weights, kernels, and point distribution) that could decrease the cost for a given target quality have not been studied. In this paper we propose a study of alternate Gabor noise ingredients tested with statistical measures on process and instances. Our test set is defined bellow.

Point distribution: The Gabor noise method is costly firstly because it requires many splats: a pixel must receive enough contributions for expected statistical properties to arise. The problem with a Poisson point distribution is that it tends to produce clusters and voids, which wastes samples. However there is no reason to stick to Poisson. Many alternative point distributions could be used instead, with better coverage uniformity over the texture space. In particular, blue noises and low discrepancy distributions provide more regular coverage. Among all possible point distributions we therefore study and compare the original Poisson process to a stratified Poisson process, a jittered grid process (for both rectangular




Figure 1: Alternative point distributions studied in this work, shown for $N = 16$. (a): Seminal Poisson point process. (b): Stratified Poisson – i.e. constant number of impulses per cell. (c),(d): Jittered grid distribution, from N sub-cells along a square (respectively, hexagonal) regular grid. (e),(f): Sobol and scrambled Sobol sequences generating point coordinates in each cell.

and hexagonal grids), and point distributions sampled using low-discrepancy Sobol sequences (see fig. 1).

Weights: Splats are modulated by uniform random weights on $[-1, 1]$ inherited from the shot noise model in electrical engineering¹. But the weighting role, effect and usefulness have not been studied. Splats with a low weight count for little in the sum, which seems like a waste of computation. We therefore compare the original uniform weighting on $[-1, 1]$ to the minimalist Bernoulli distribution on $\{-1, 1\}$.

Kernel: The true Gabor kernel [Gab46] is a Gaussian times $\exp(i\omega x)$. The practical version used in Gabor noise seminal paper is the real part of it, i.e. a cosine multiplied by a Gaussian truncated to $\approx 5\%$ value. The choice of cosine rather than a sine has never been discussed, while the later would enforce a strict zero mean. We thus compare these two kernels for all our test cases.

3. Evaluation

Whatever the variant Gabor noise algorithm, the seminal properties must be respected: producing an instance of a Gaussian texture process (i.e. random uncorrelated phases) reproducing the target PSD with a normalized variance. We therefore measure and compare the quality of the process for each alternative and function of the number of impulses per cell N . As measuring the randomness of Fourier phases is particularly difficult (see a survey in [Lec15]), we evaluate an equivalent property that is the *Gaussianity* of the process.

However, the statistics of the process are not enough for a Computer Graphics application where the user wants to produce one image instance out of the process while still hoping for “nice” textural properties. Indeed not only the overall statistics or PSD are important but also the appearance of the generated textures. More specifically, the *stationarity* of textures properties has to be ensured. This includes process stationarity (i.e. no spatial bias) and windowed texture properties in instances: a texture is about spatial variations of values, but above a given scale the statistics should not vary.

This leads to three statistical tests: the Gaussianity of the process, its spatial stationarity along the texture space, and the large-scale stationarity inside an image instance. We define precisely these tests and apply them to the alternative ingredients in the following sections.

¹ https://en.wikipedia.org/wiki/Shot_noise

Moreover, the average and contrast of an image texture instance could differ to the one of the process, biasing the a priori normalization. We should thus ensure that this bias is kept low. But the measure of windowed large-scale stationarity already provides a stronger criterion so no extra test is required.

3.1. Global process Gaussianity

High-intensity shot noise, the underlying model of Gabor noise, converges as a process to a Gaussian distribution. To choose the number N of splats sufficient for an acceptable “convergence” to a Gaussian texture, previous work only estimated visually the quality of obtained instances. Here, convergence is not in terms of asymptotic values: an instance will never converge while adding more splats, the pattern will keep evolving (just more and more slowly). Convergence is in terms of probability law, so we evaluate it as the Gaussianity of the noise process as N grows. We measure this as the *Cramer von Mises criterion*² (CvM) of the noise process compared to the Gaussian probability distribution using 7

² https://en.wikipedia.org/wiki/Cramer%3%A9r%E2%80%9393von_Mises_criterion

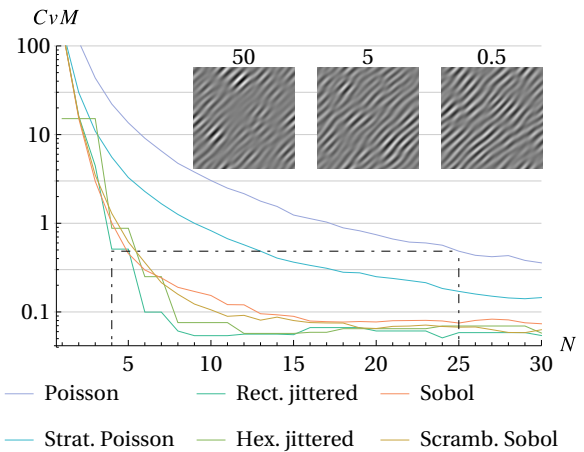


Figure 2: Compared Gaussianity-wise convergence. *Cramer von Mises criterion* value for the alternative point distributions for increasing N . The inset images show instances of seminal Gabor noise for CvM values of 50, 5 and 0.5. Alternative point distributions converge faster; for example, the rectangular jittered grid reaches the same CvM value for $N = 4$ as the Poisson process for $N = 25$ (mixed, black).

million samples. Zero means perfect Gaussianity. As a calibration of this metric in terms of corresponding visual quality, we use the seminal Gabor noise process as a reference. We can see that the CvM should be less than 0.5 for a reasonable visual quality, coherently corresponding to the estimation of $N = 30$ in [LLDD09] for the seminal Gabor noise.

Result. The seminal process based on the Poisson point distribution converges to a Gaussian distribution as N increases, albeit very slowly. As fig. 2 shows, alternative point distributions (stratified Poisson and jittered grids) greatly speed up the process convergence, achieving the same Gaussianity criterion as the seminal high quality Gabor noise with up to 6 times fewer splats and as few as $N = 4$.

In addition, fig. 3 shows that the use of Bernoulli weights (solid curves) results in even faster convergence of the noise process for Poisson and Stratified Poisson. This makes the stratified Poisson point distribution using Bernoulli weighting the best alternative, reaching a constant 15× convergence speed-up compared to the seminal Gabor noise process.

Our experiments also revealed that changing the kernel from cosine to sine does not change the process convergence.

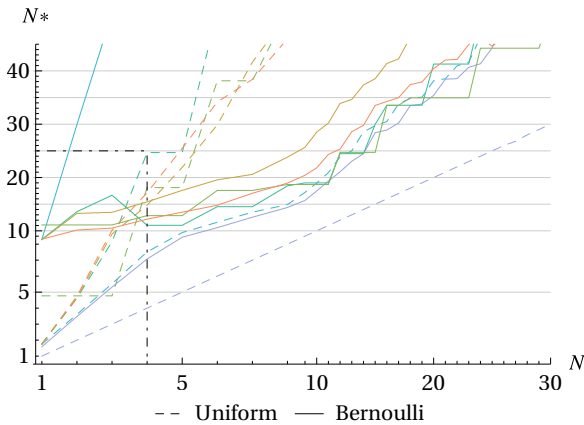


Figure 3: Convergence gain. Same as fig. 2, but plotting the seminal Gabor impulse count N^* required to get the same CvM quality as each tested alternative using an impulse count of N . In addition to the uniform weights (dashed), we also plot the gains using Bernoulli weights (solid curves).

3.2. Process at a given location

The seminal Gabor noise process is stationary by construction since it draws on Poisson distribution, which is translation-invariant and thus has no location bias, but the alternatives point distributions we bench are not. We must then ensure that such bias does not show up in the final texture. The global process statistics presented above are blind to these local effects, so we must also explicitly study stationarity. For this, we measure the variation of the mean and variance of the process at a given location within a cell. In the absence of bias, these statistics should be the same as the global process and be constant over the considered texture space. Our experiments show that this is the case for the proposed alternatives with a relative interdecile range under 5% of the considered statistics, showing that these methods preserve the translation invariance of the original Gabor noise

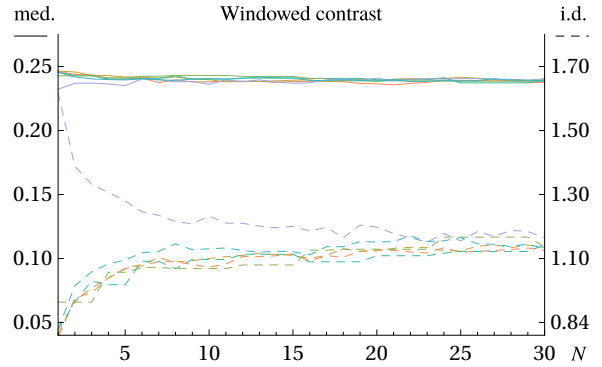


Figure 4: Measure of image stationarity as variations in windowed contrast. *Windowed contrast is the spatial variance in a given image window. We measure its stationarity along the image as the median (solid, left axis) and the interdecile normalized by the median (dashed, right axis) of its variations for each alternate point distribution, relative to N (more robust than mean and variance since these variance samples are not normally distributed). Here we show only the case of Bernoulli weights. The uniform case is quite similar, only with worse interdecile ranges for small N values. Note that the expected median of contrast is 0.25, since we scale and offset the noise so as the $[0,1]$ value range includes 95% of the histogram.*

process despite all the alternate point processes themselves being more or less spatially biased.

3.3. Instance stationarity and normalization

Satisfactory global statistics can be obtained with well spread local artifacts, while users expect *spatial stationarity* of the local texture properties above some macroscopic “pattern” scale. For instance the Poisson distribution tends to have clustered and void areas despite the process being homogeneous. More generally, unbiased processes can still cause macro-fluctuations along the instances.

For Gabor noise the natural texture scale on which considering instance large-scale stationarity is the diameter of a kernel, since pixels are decorrelated by construction above this size. We thus consider instance image statistics windowed at this scale. Thanks to the ergodicity yielded by this decorrelation, we can characterize statistical fluctuations from non-overlapping instance sub-images within a large image instance by studying statistical fluctuations of small instances spanning 3×3 cells (the minimum cell span required to observe one full kernel).

Variations in image statistics are expected in stochastic procedural texture instances, but if we can ensure that image statistics are close to process statistics, then we can obtain an acceptable texture normalization a priori, without requiring a second pass (which would be costly, complicated or impractical in many real-world use cases of procedural noise).

Result. Our experiments show (see fig. 4) that for all distributions the median of contrast variations conforms to the process variance (*i.e.* the instance normalization is not biased),

and that the benched alternative point distributions result in significantly lower contrast variations over instances than the seminal Gabor noise, especially when using Bernoulli weighting. It also unintuitively shows that the contrast relative interdecile range doesn't tend to zero: some contrast variations are unavoidable. It stabilizes around 1.15 pretty early, which means that it's no use increasing N in the hope of getting better stationarity.

Oppositely, we measured that exotic processes with strong location bias yield instances which can sometimes appear stationary but significantly deviate from the process statistics, depending in complex ways on N , F_0 and ω_0 (or even particular instance seeds). This then makes *a priori* normalization impossible.

We also studied the variation of the mean of instance images (of 3×3 cells). Although its variation is low (less than 0.5% of the noise value range around the analytical expected value 0), using sine for the harmonic part for the kernel instead of cosine actually ensures that the image mean is perfectly 0 over all instances.

4. Conclusion and future work

Our study results in the following recommendation: replacing the Poisson point distribution with stratified Poisson (*i.e.* fixed number of splats per cells), and replacing uniform weighting with Bernoulli $\{-1, 1\}$. The recommended CvM quality of 0.5 corresponding to $N = 30$ with the seminal method is now obtain with $N = 2$ (corresponding to 6.3 splats per pixel). In addition to the accelerated convergence, implementation performance is also improved by the fact that seminal Poisson point process requires the costly generation of Poisson random numbers to determine the varying N for each different cell. All this provides a total performance gain of 17 folds (or up to 24 for higher N , see fig. 5) with minimal modifications to the seminal algorithm. This is achieved without introducing spatial bias in instances, and indeed improving the usability by lowering the variability of instance image statistics. Once applied to a real case combining several kernels (see fig. 6), this gain is confirmed.

In addition, using a sine-based kernel ensures exact zero-mean texture, and using a continuous enve-

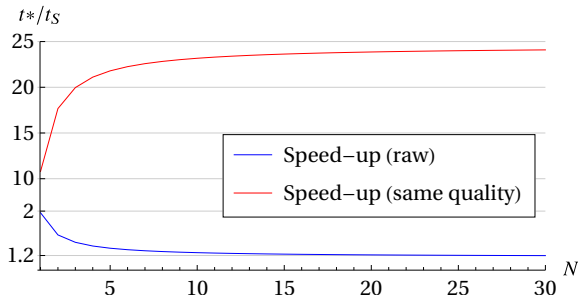


Figure 5: Runtime performance ratio of stratified Poisson with Bernoulli weighting $t_s(N)$ relative to the seminal Gabor noise $t^*(N^*)$ providing the same quality (red curve). Note that for the same splat count we already gain 20% (blue curve) by avoiding the costly random evaluation of N per cell.

lope instead of truncated Gaussian (*e.g.* the C0 alternative $(1 + \exp(-\pi))(\exp(-\pi\|x\|^2) - \exp(-\pi))$, or C1 Kaiser-Bessel window functions) would avoid slight artifacts at splat borders for very low N or if the noise is meant to be differentiated (*e.g.* for bump maps).

For future work, we could also study totally remove weights while using the sine-based kernel, since their main purpose is to ensure the zero-mean which is already guaranteed with this kernel. We could also replace point-location randomness by phase-randomness on regular grids (*i.e.* a scalar random value instead of a random vector). Provided quality is not impeded, these could bring even more performance.

More generally, our approach to optimizing the Gabor noise process mostly relies on properties of sparse convolution noises, and thus could also apply to other derived methods that work on richer power spectra.




Figure 6: Real case with complex power spectrum (3 kernels, cf. inset) and non-linear post-treatment. *Our optimized set of ingredients achieves the same visual quality in 1/17th of the time required by the seminal method. Note that due to the similarities in the 3 kernels, here we can use only $N = 3$ (resp. 45) instead of 3×2 (resp. 3×30).*

References

- [BLV*10] BÉNARD P., LAGAE A., VANGORP P., LEFEBVRE S., DRETTAKIS G., THOLLOT J.: A Dynamic Noise Primitive for Coherent Stylization. *Computer Graphics Forum*. Vol. 29, Num. 4 (juin 2010), 1497–1506.
- [Gab46] GABOR D.: Theory of communication. *Journal of the Institution of Electrical Engineers*. Vol. 93, Part III, Num. 26 (November 1946), 429–457.
- [GLLD12] GALERNE B., LAGAE A., LEFEBVRE S., DRETTAKIS G.: Gabor Noise by Example. *ACM Transactions on Graphics*. Vol. 31, Num. 4 (juillet 2012), Article No. 73.
- [GLM17] GALERNE B., LECLAIRE A., MOISAN L.: Texton Noise. *Computer Graphics Forum* (janvier 2017).
- [GSV*14] GILET G., SAUVAGE B., VANHOEY K., DISCHLER J.-M., GHAZANFARPOUR D.: Local Random-phase Noise for Procedural Texturing. *ACM Trans. Graph.*. Vol. 33, Num. 6 (novembre 2014), 195:1–195:11.
- [Lec15] LECLAIRE A.: *Random phase fields and Gaussian fields for image sharpness assessment and fast texture synthesis*. PhD thesis, Université Sorbonne Paris Cité, juin 2015.
- [Lew89] LEWIS J. P.: Algorithms for Solid Noise Synthesis. SIGGRAPH '89, ACM, pp. 263–270.
- [LLD11] LAGAE A., LEFEBVRE S., DUTRE P.: Improving Gabor Noise. *IEEE Transactions on Graphics and Visualization*. Vol. 17, Num. 8 (août 2011), 1096–1107.
- [LLDD09] LAGAE A., LEFEBVRE S., DRETTAKIS G., DUTRÉ P.: Procedural Noise using Sparse Gabor Convolution. *ACM Transactions on Graphics*. Vol. 28, Num. 3 (2009), 54–64.
- [vW91] VAN WIJK J. J.: Spot Noise Texture Synthesis for Data Visualization. SIGGRAPH '91, ACM, pp. 309–318.