Macrocanonical texture synthesis

Valentin De Bortoli

With: Agnès Desolneux, Bruno Galerne, Arthur Leclaire

January 7, 2019

Centre de Mathématiques
et de Leurs Applications

Goal: sample images $x \sim \widetilde{\Pi}(x)$ which look like an input texture x_0 but are not verbatim copies of x_0.

Question: How to combine randomness and structure in an image model?

Macrocanonical model

The probability distribution function $\widetilde{\Pi} \in \mathscr{P}$ is a *macrocanonical model* associated with the exemplar texture $x_0 \in \mathbb{R}^d$, statistics $f : \mathbb{R}^d \to \mathbb{R}^p$ if

$$H(\widetilde{\Pi}) = \max \left\{ H(\Pi),\ \Pi \in \mathscr{P},\ \mathbb{E}_\Pi \left[f(X) \right] := \Pi(f) = f(x_0) \right\} \ .$$

Maximize the entropy (H) under geometrical constraints (f).

Gibbs measure: $\Pi_\theta(x) \propto \exp\left[-\langle \theta, f(x) - f(x_0)\rangle\right]$

Gibbs measures are macrocanonical models

Under mild assumptions there exists $\tilde{\theta} \in \mathbb{R}^d$ such that $\Pi_{\tilde{\theta}}$ is a macrocanonical model associated with the exemplar texture $x_0 \in \mathbb{R}^d$, statistics f.

Two questions remain:

1. how to find the optimal weights $\tilde{\theta}$?
2. how to sample from the model, *i.e.* sample from a Gibbs measure Π_θ?

We denote $V(x,\theta) = \langle \theta, f(x) - f(x_0) \rangle$.

Finding optimal weights

$\tilde{\theta}$ is the minimum of the log-partition function which is a *convex problem*.

Gradient descent dynamics

$$\theta_{n+1} = \theta_n + \delta_{n+1} \Pi_{\theta_n}(\nabla_\theta V(\cdot, \theta_n))$$

Sampling from a Gibbs measure

The potential $x \mapsto V(x,\theta)$ is usually *non-convex* but has *curvature at infinity*.

Langevin dynamics

$$X_{n+1} = X_n - \gamma_{n+1} \nabla_x V(X_n, \theta) + \sqrt{2\gamma_{n+1}} Z_n$$

$\Rightarrow$ **Combining dynamics**

$$X_{k+1}^n = X_k^n - \gamma_n \nabla_x V(X_k^n, \theta_n) + \sqrt{2\gamma_n} Z_{k+1}^n \, , \text{ with } X_0^n = X_{m_{n-1}}^{n-1} \, ,$$

$$\theta_{n+1} = \theta_n - \delta_{n+1} m_n^{-1} \sum_{k=1}^{m_n} \nabla_\theta V(X_k^n, \theta_n) \, ,$$

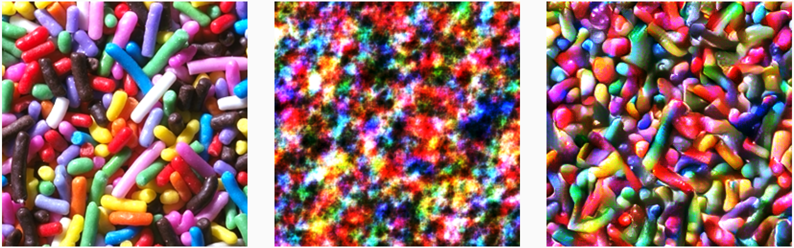

(a) (b) (c)

Figure 1: ***Texture synthesis.*** (a) input texture, (b) is the initialization of the algorithm and (c) the output.