N. Bertschinger
M. Kaschube
V. Ramesh

Given a binary image, in which pixels x_i can take on values ± 1, $J_{n,m}$ is the pair-potential between nodes n and m (which is typically greater than zero if n and m are direct neighbors, else zero), and let $H \geq 0$ be the external field. The energy of this system is then:

$$E(x) = E(x; J, H) = -\left(\frac{1}{2}\Sigma_{n,m} J_{n,m} x_m x_n + \Sigma_n H x_n\right) \tag{1}$$

We denote $\beta = \frac{1}{T}$ as the inverse temperature, and the probability of a state x is:

$$P(x) = P(x|J, H, \beta) \propto e^{-\beta E(x)} \tag{2}$$

The normalizing factor can be denoted as $Z = \Sigma_x e^{-\beta E(x)}$. In Gibbs sampling, we select a pixel x_n and consider its probability given all of its direct neighbors x_m : $P(x_i = 1|\{x_m\})$:

$$P(x_n = 1|\{x_m\}) = \frac{P(x_n = 1, \{x_m\})}{P(x_n = 1, \{x_m\}) + P(x_n = -1, \{x_m\})} \tag{3}$$

By fixing x_n, we can calculate the probability using the energy-function from above:

$$\begin{aligned} P(x_n = 1|\{x_m\}) &= \frac{exp(\beta \Sigma_m J_{n,m} x_m + H)}{exp(\beta \Sigma_m J_{n,m} x_m + H) + exp(-\beta \Sigma_m J_{n,m} x_m - H)} = \\ &= \frac{1}{1 + exp(-\beta \Sigma_m J_{n,m} x_m + H)} \end{aligned} \tag{4}$$

A sampling step consists in calculating the above equation, and setting $x_n = 1$ with probability $P(x_n = 1|\{x_m\})$, and $x_n = -1$ otherwise.

Note: parameters J,H and β and the relative probability is invariant under scaling (i.e. normalization by Z), so J may be fixed to 1.

Note: a live-simulation of Gibbs-sampling on an Ising-model can be found here: http://cs.stanford.edu/people/karpathy/visml/ising_example.html

Ex. 10.1: Gibbs and Ising Punkte: __ / 10

(a) Implement the above model, generate a random image with pixel values beeing either -1 or +1. Set parameters $J = 1$, $H = 0$, $\beta = 2$. Sample 100.000 times, visualize the image after every 10.000 steps (it should look like Figure 1).

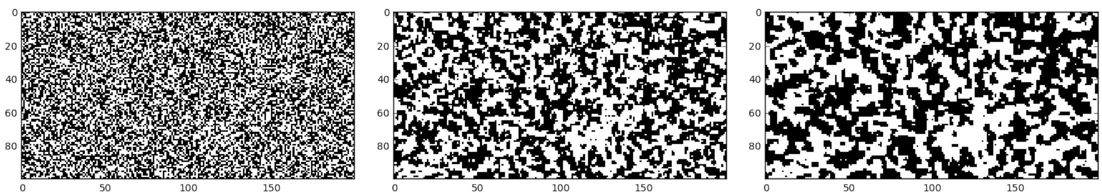

Figure 1: Draws from the Ising-model using Gibbs sampling after a) random initialization, b) 50.000 samples, c) 100.000 samples with parameters $J = 1$, $H = 0$,$\beta = 2$

(b) Download *image_small.png* from the course website. With a probability of 10% for every pixel, flip it's value. Integrate this noisy image as the external field in the above model. For a burn-in-period of 50.000 samples, sample but do not save the results. After that, sample 500.000 times, aggregate the results and average them. Re-binarize the image. (set $\beta = 2$ as above). Your final results should look like Figure 2

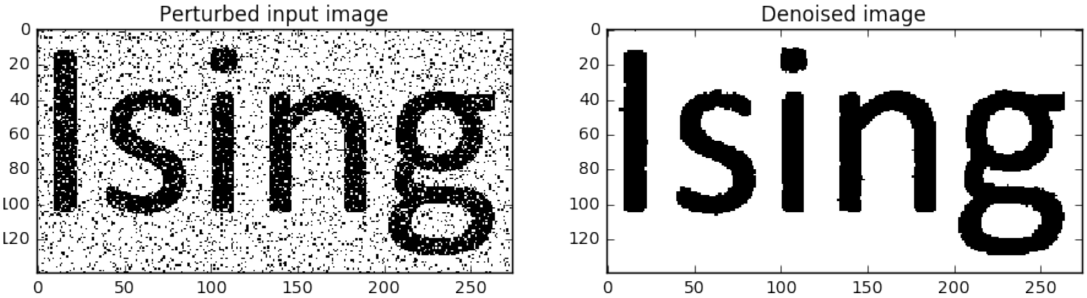

Figure 2: Gibbs/Ising denoising, left: perturbed image, right: denoised image after 500.000 samples.