EXERCISES 9

Exercise 9.1. Simulate data from the following model

$$x_t | x_{t-1} \sim \mathcal{N}(x_t; a x_{t-1}, \sigma_x^2),$$
$$y_t | x_t \sim \mathcal{N}(y_t; x_t, \sigma_y^2).$$

For reproducibility, choose $x_0 = 1$ (instead of sampling from a prior μ_0). Choose $a = 0.9$, $\sigma_x = 0.01$ and $\sigma_y = 0.1$. Simulate this system $T = 100$ time steps. Plot your x and y in the same graph with different colors and give some examples what the model output can model in the real world.

Exercise 9.2. Imagine you are asked to develop a *volatility* model. Intuitively, volatility is a hidden quantity (that nobody knows about) which controls the variance of the observed stock (or option) prices. If the volatility is high, the stock prices (or returns) will have high variance, if the volatility is low, the prices (or returns) will have lower variance. Develop a volatility model. What you need to do essentially is:

- To define a Markov transition kernel to model the volatility variable, defined it as x_t (similar to above). Note that a naive Gaussian kernel would give you negative values and is therefore not suitable. You can use a Gaussian kernel but then you need to figure out how to use it modelling variance.
- Define your likelihood – this is more clear. If higher volatility (x_t) means higher variance in observed returns (y_t), how can you model it? Since we model returns (not prices directly), negative values in y_t are allowed.
- Simulate data and show your results meaningfully model volatility.

As a sanity check, your model should exhibit the following behaviour if you choose x as a decaying or growing system. However, this is **not** the only output. The model is more

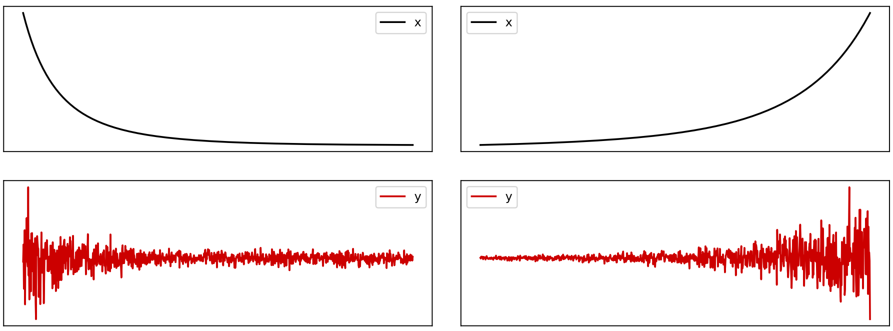

Two examples of a simulated data from a volatility model. You can see that as x decays, the variance of y decays – or if x grows, the variance y grows (it becomes more erratic).

important. Hopefully you will be able to generate some realistic data from your model, such as the figure in the next page. You are free to search the literature for volatility models and use any model you wish in published literature (only state-space models, not more general ones).

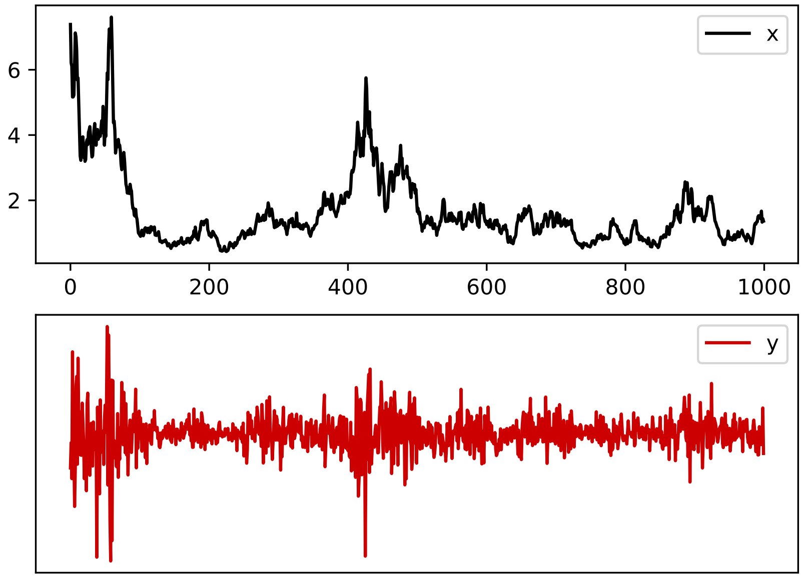

Exercise 9.3. Let us consider the following modification of our model:

$$\begin{aligned}
\theta &\sim p(\theta), \\
x_0|\theta &\sim \mu(x_0|\theta), \\
x_t|x_{t-1}, \theta &\sim f(x_t|x_{t-1}, \theta), \\
y_t|x_t, \theta &\sim g(y_t|x_t, \theta).
\end{aligned}$$

1. The main difference between this model and a regular state space model is the existence of the parameter θ which also has its own prior. Let us assume that we would like to sample from $p(\theta|y_{1:T})$, i.e., the parameter posterior. Assuming that we can sample from $p(x_{0:T}|y_{1:T}, \theta)$ for a given θ and that we can sample from $p(\theta|x_{0:T}, y_{1:T})$ for a given $x_{0:T}$, describe the Gibbs sampler as an idealised algorithm to sample from $p(\theta|y_{1:T})$. Suggest a method to simulate from $p(x_{0:T}|y_{1:T}, \theta)$ as a subroutine of this Gibbs sampler.

2. Simulating $p(x_{0:T}|y_{1:T}, \theta)$ at once as a block sampler might be too expensive. Describe a method simulating each X_t variable independently, conditioning on $X_{-t}, \theta, y_{1:T}$. Follow simplifications due to conditional independence and write down the algorithm.