A differentiable reaction-diffusion model for predicting vegetation dynamics in the Sahel

Jens van der Zee, Diego Marcos, Eric Siero

Background

- The future of vegetation in the Sahel remains uncertain in the face of climate change
- Process-based models may help to understand and predict how vegetation responds to increasing temperatures and changes in precipitation patterns
- These models currently struggle to leverage valuable sources of data such as time series of satellite images
- Recent advances in scientific computing are aimed at developing algorithms that enable process-based models to learn from data to close the gap between simulation and reality

Objective

- Implement a differentiable version of a reaction-diffusion model in PyTorch
- Use satellite observations to learn the parameters of the differentiable reaction-diffusion model

Introduction

- A combination of local self-facilitation and long-range competition for limited water resources leads to the formation of patterned vegetation in the Sahel

- The qualitative behaviour of these ecosystems has been extensively studied using reaction-diffusion models describing the interactions between water and vegetation

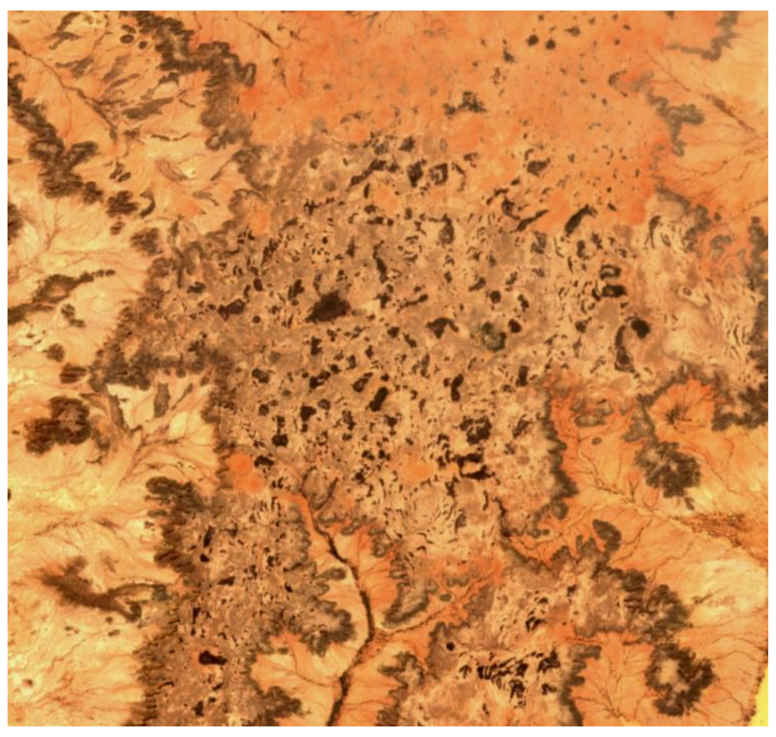

Figure 1. Patterned vegetation near Niamey, Niger

Reaction-diffusion model

$$
\begin{aligned}
H_t &= d_1 \Delta H - l_1 H + p - r_2 HB \\
W_t &= d_2 \Delta W - l_2 W + r_2 HB - r_1 W B(1+\eta B)^q \\
B_t &= d_3 \Delta B - l_3 B + j r_1 W B(1+\eta B)^q
\end{aligned}
$$

Equation 1. System of differential equations describing the interactions between surface water (H), soil water (W) and vegetation biomass (B).

l_1 = evaporation rate	d_1 = surface water diffusion rate
l_2 = seepage rate	d_2 = soil water diffusion rate
l_3 = mortality rate	d_3 = vegetation diffusion rate
r_1 = water uptake rate	j = water-use efficiency
r_2 = water infiltration rate	η = self-facilitated water uptake
p = rainfall	q = nonlinearity of self-facilitation

Training

- We initialize a model with random parameters
- The model runs a 10-year simulation, starting with a 2013 satellite image of vegetation density as initial condition
- The loss between the model prediction and a satellite image of vegetation density is evaluated every year
- The cumulative loss is backpropagated to update the parameters of the model

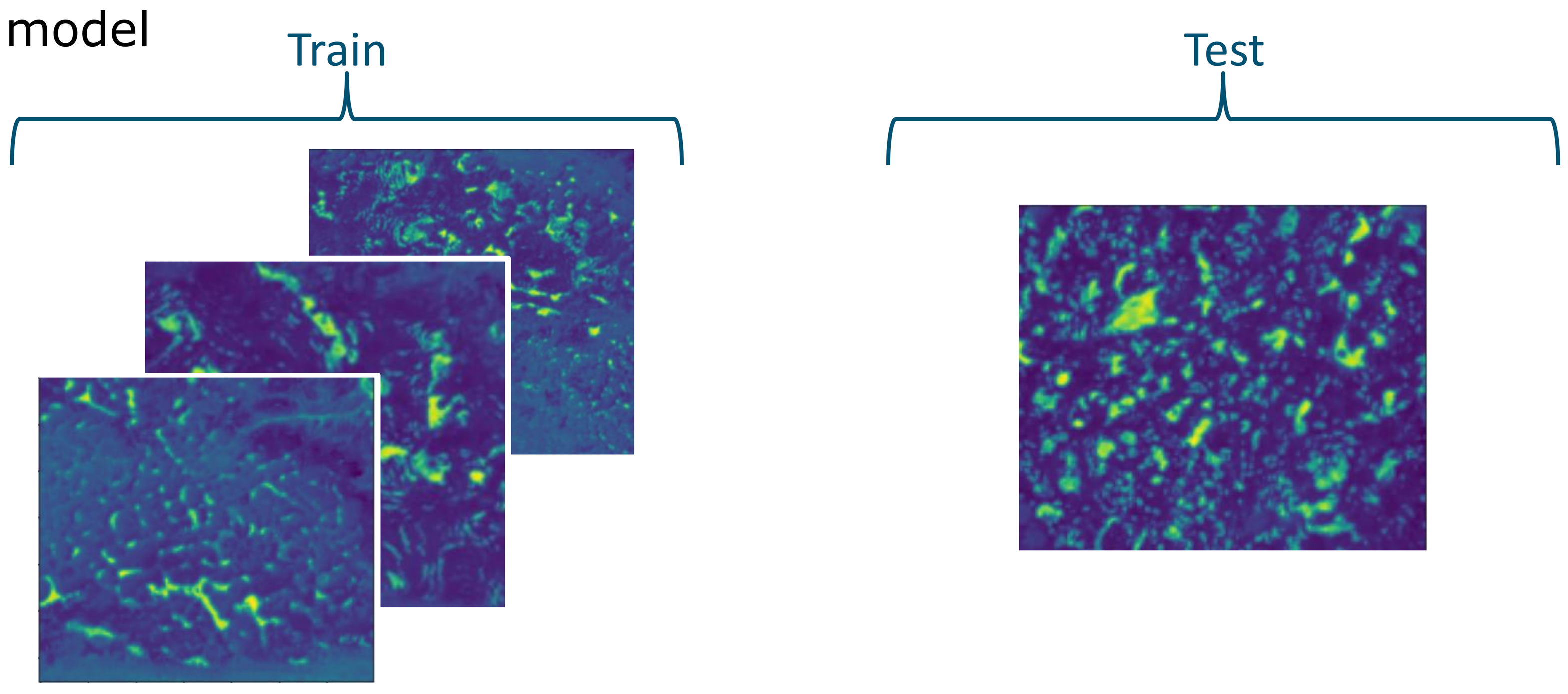

Figure 2. The training is carried out in a 4-fold cross validation of four sites located close to each other.

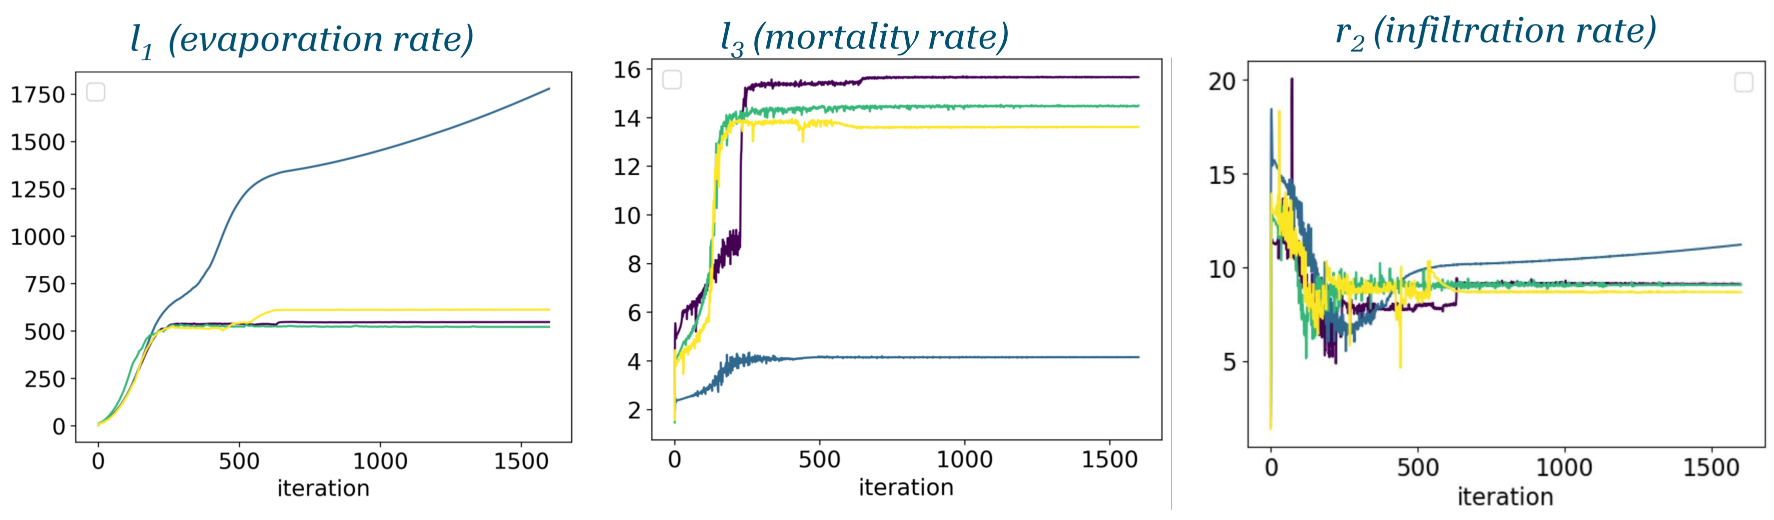

Figure 3. Training results of the 4-fold cross validation for three selected parameters. Each line represents a different combination of sites used as training data.

Testing

The performance of the model is evaluated by running a simulation on the test site and comparing the predicted change in vegetation density with the measured changed in vegetation density

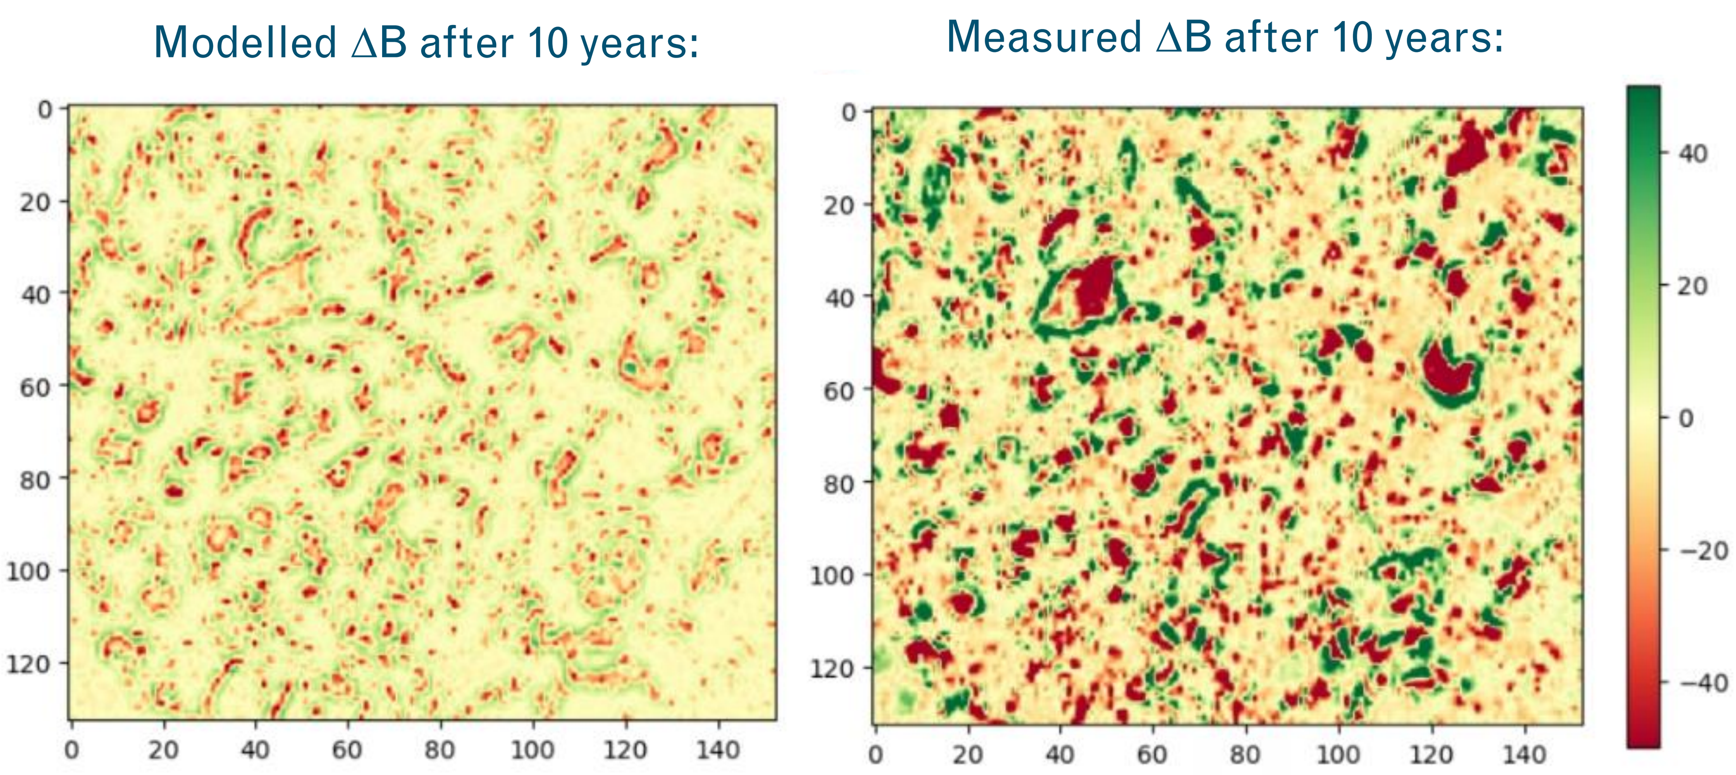

Figure 4. Modelled (left) and measured (right) change in vegetation density over a 10-year period. While the model struggles to predict the magnitude of both negative and positive changes, it does seem to reproduce structural features such as halos of regrowth around dying patches of vegetation

Wageningen University & Research
P.O. Box 123, 6700 AB Wageningen
Contact: jens.vanderzee@wur.nl