

Estimating parameters of the nonlinear cloud and rain equation from a large-eddy simulation

Lunderman, Spencer; Morzfeld, Matthias; Glassmeier, Franziska; Feingold, Graham

DOI

[10.1016/j.physd.2020.132500](https://doi.org/10.1016/j.physd.2020.132500)

Publication date

2020

Document Version

Accepted author manuscript

Published in

Physica D: Nonlinear Phenomena

Citation (APA)

Lunderman, S., Morzfeld, M., Glassmeier, F., & Feingold, G. (2020). Estimating parameters of the nonlinear cloud and rain equation from a large-eddy simulation. *Physica D: Nonlinear Phenomena*, 410, Article 132500. <https://doi.org/10.1016/j.physd.2020.132500>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Estimating parameters of the nonlinear cloud and rain equation from a large-eddy simulation

Spencer Lunderman^a, Matthias Morzfeld^b, Franziska Glassmeier^{c,d,e}, and Graham Feingold^d

^aDepartment of Mathematics, University of Arizona

^bInstitute of Geophysics and Planetary Physics, Scripps Institution of Oceanography,
University of California, San Diego

^cDepartment of Environmental Sciences, Wageningen University

^dChemical Sciences Laboratory/Earth System Research Laboratories, NOAA

^eCurrently at: Faculty of Civil Engineering and Geosciences Department of Geoscience and
Remote Sensing, Delft University of Technology.

April 10, 2020

Abstract

Predator-prey dynamics have been suggested as simplified models of stratocumulus clouds, with rain acting as a predator of the clouds. We describe a mathematical and computational framework for estimating the parameters of a simplified model from a large eddy simulation (LES). In our method, we extract cycles of cloud growth and decay from the LES and then search for parameters of the simplified model that lead to similar cycles. We implement our method via Markov chain Monte Carlo. Required error models are constructed based on variations of the LES cloud cycles. This computational framework allows us to test the robustness of our overall approach and various assumptions, which is essential for the simplified model to be useful. Our main conclusion is that it is indeed possible to calibrate a predator-prey model so that it becomes a reliable, robust, but simplified representation of selected aspects of a LES. In the future, such models may then be used as a quantitative tool for investigating important questions in cloud microphysics.

Keywords

Predator-prey dynamics; Large-eddy simulation; Stratocumulus clouds; Bayesian inversion;
Markov chain Monte Carlo;

1 Introduction

Stratocumulus cloud decks can reach 1000s of km in scale and cover vast stretches of the subtropical oceans. These decks consist of a space-filling arrangement of convective cells, with clouds marking updraft regions. Depending on the environmental conditions like sea-surface temperature or atmospheric aerosol, stratocumulus occur in two configurations (Agee, 1984; Wood and Hartmann, 2006; Glassmeier and Feingold, 2017)

- (i) wide updraft areas coinciding with cloud cells (“closed-cells”) whose *cloud-free boundaries* form a honeycomb-like pattern;
- (ii) narrow updrafts and cloudy rings that outline a honeycomb-like pattern (“open cells”).

Due to its lower cloud fraction, the open-cell configuration is significantly less reflective than the closed-cell configuration. Since roughly one fifth of the Earth’s surface is covered by stratocumulus cloud decks (Wood, 2012), the radiative effects of stratocumulus have a large impact on the Earth’s energy budget. In fact, stratocumulus remain one of the main sources of uncertainty in quantifying climate change (Boucher et al., 2013; Myhre et al., 2013; Schneider et al., 2017).

Stratocumulus, and in particular transitions from closed-cell to open-cell configurations, have been studied numerically with a hierarchy of mathematical and computational models. Large eddy simulations (LES) resolve the governing equations of moist hydrodynamics down to the cloud scale and can faithfully represent the formation of stratocumulus and how they transition between the open- and the closed-cell configurations, see, e.g., Feingold et al. (2015). In addition to these detailed but computationally expensive models, drastically simplified, low-dimensional models have been proposed to capture the spatial configuration of stratocumulus. For example, dynamic cellular networks can be used to describe the patterns that are formed within stratocumulus cloud systems (Glassmeier and Feingold, 2017). Predator-prey models, where the rain acts as the predator of clouds, have been proposed as phenomenological models for stratocumulus (Koren and Feingold, 2011; Feingold and Koren, 2013; Koren et al., 2017).

The predator-prey models can reproduce two configurations that are relevant to stratocumulus clouds: oscillatory (limit cycle) and stationary solutions for cloud depth. The limit cycles model a scenario in which strong rain dissipates the cloud that created it, followed by renewed cloud build-up that proceeds until the cloud is again thick enough to produce strong rain, which restarts the whole process. A stationary cloud depth represents a situation in which the rain consumes the cloud at the same rate as the cloud replenishes.

We focus on one of the predator-prey models, the nonlinear cloud and rain equation of Koren et al. (2017), which we call KTF17 for short. Our primary goal is to build a mathematical and computational framework to convert KTF17 into a quantitative tool. We argue that this can be done by adopting a Bayesian approach, in which a posterior distribution over the parameters of KTF17 is defined based on cloud depth time series of stratocumulus. A natural data source for these time series would be observations of stratocumulus in the Earth’s atmosphere, e.g., derived from the Geostationary Operational Environmental Satellite-R Series (GOES-R). KTF17 does not account for horizontal advection which is usually present in satellite derived observations. Using observational data would thus require tracking stratocumulus patches within a larger cloud system over time to “remove” advection, see, e.g., Koren and Feingold (2013). To avoid these technicalities, we use LES output, generated in the absence of advection, as “data” in place of observations. The resulting KTF17 model, with stochastic parameters distributed according to a posterior distribution, is thus a quantitative, but simplified representation of selected aspects of cloud systems that are realistically represented by LES. Our approach thus connects the extreme ends of the hierarchy of cloud models and may be used to obtain new insights into complex cloud and rain interactions. Given the example of the predator-prey-based parameterization of Nuber and Graf (2005) to represent convection, simple predator-prey models, “calibrated” to a LES via a parameter estimation, may eventually even prove useful for representing some aspects of cloud systems in climate models. We focus, however, on establishing a suitable mathematical and computational framework for the task of “calibrating” a predator-prey model with LES data.

More specifically, we describe how parameters of KTF17 can be estimated from a LES by a “Bayesian inversion”. The inversion is based on two distributions: a prior distribution, that represents knowledge about the model parameters, *without* taking the data into account, and a likelihood, that describes the probability of the data, given a set of parameters, see, e.g., Reich and Cotter (2015); Asch et al. (2017); Tarantola (2005). Jointly, the prior and likelihood define a posterior distribution over the parameters that represents our knowledge of the parameters and

their uncertainties in view of the data, our prior knowledge and assumed errors.

Typically, a likelihood is based on a point-wise mismatch of model outputs and data. In our context, the “data” are a time series of cloud depth of the LES, i.e., a 2D field that evolves over time (note that we refer to simulation outputs as data because we treat them as such). KTF17, however, does not have an associated spatial scale. Thus, it is not straightforward to compare KTF17 to LES output. We address this issue by using “feature-based” likelihoods (Maclean et al., 2017; Morzfeld et al., 2018). The basic idea is that compressing the data into suitable features can bridge gaps between drastically simplified models and complex processes. The feature we consider is a stochastic representation of cycles of growth and decay in cloud depth, derived from the LES, that can be compared directly (point-wise) to limit cycles of KTF17. Required error models of the features are constructed based on variations of the cloud cycles extracted from the LES. We solve the resulting, feature-based inverse problem numerically by a Markov chain Monte Carlo (MCMC) method. This means that we generate a (large) set of physically relevant “samples” (model parameters) that lead to KTF17 limit cycles that are comparable to the cloud cycles observed in the LES, to within the assumed errors. In particular, we observe an overall good fit in terms of the cycle’s periods, amplitudes and average growth and decay times. The Bayesian approach and MCMC implementation further provide information about posterior errors and uncertainties, which in turn depend on expected model errors. This allows us to assess, in hindsight, the validity of our assumptions about errors and error models. We further carefully test the robustness of our overall approach by numerical sensitivity studies. These tests of robustness and of the validity of error models are essential to being able to use KTF17 to make precise and definite statements. Finally, we illustrate how to use our technique to investigate cloud microphysics questions. Specifically, we compute sensitivity of model parameters to temporal changes in the morphology of the cloud system. We must emphasize that our results and conclusions with respect to cloud microphysics are limited, in part because our study is limited to one particular LES.

2 Background: the nonlinear cloud and rain equation, the LES and feature-based Bayesian inversion

We use a Bayesian approach to combine information from a LES with a simplified predator-prey model of stratocumulus clouds. In this section, we describe the predator-prey model and the LES. We then provide background and notation for Bayesian inversion and feature-based Bayesian inversion.

2.1 The nonlinear cloud and rain equation (KTF17)

The coevolution of cloud and rain can be captured, qualitatively, by predator-prey type dynamics and, more specifically, by differential equations with a delayed sink term (Koren and Feingold, 2011; Feingold and Koren, 2013; Koren et al., 2017). The delay stems from the fact that the predator (rain) is produced by the cloud (prey) with a delay that is associated with the time required for cloud droplets to coalesce to form larger raindrops. This delay time is a function of the amount of cloud water and the cloud drop concentration and is typically on the order of 15 minutes.

The predator-prey models are capable of reproducing two different dynamical regimes that are relevant to stratocumulus clouds. When the predator-prey models exhibit a constant cloud depth, the rain consumes the cloud at the same rate as cloud replenishment. When the predator-prey models exhibit oscillations (limit cycles), strong rain nearly depletes the cloud and then dissipates until the cloud is thick enough to again produce rain.

We consider the “nonlinear cloud and rain equation” (Koren et al., 2017), subsequently called KTF17:

$$\frac{dH}{dt} = \frac{H_0 - H}{\tau} - \frac{\alpha}{\sqrt{N}} H^2 (t - T). \quad (1)$$

Here, H (in m) is cloud depth, H_0 (in m) is the cloud depth carrying capacity, τ (in days) is the characteristic time to reach carrying capacity, T (in days) is the delay associated with the time it takes to generate rain and N (in cm^{-3}) is the droplet concentration; the scaling factor α (in $\text{day}^{-1}\text{m}^{-2.5}$) links the cloud depth, droplet concentration and rain rate (see Koren et al. (2017) for further detail).

In summary, the parameters of the KTF17 model are the delay, the carrying capacity, the characteristic time and the scaling factor. For a given set of parameters and initial conditions, we solve (1) numerically by a 4th order Runge-Kutta method with time step $\Delta t = 0.1$ min. The numerical integration requires that we prescribe the cloud depth $H(t)$ during “negative times” on the interval $t \in [-T, 0]$ and we assume that $H(t)$ is constant during this interval. The result of a numerical solution of KTF17 is a time series of cloud depth.

We note that KTF17 assumes that droplet concentration be fixed. This is justified when there is an approximate balance between replenishment of aerosol particles, which form the nuclei for new droplets, and consumption of droplets/particles via coalescence and their removal by rain. Below, we use values between $N = 16 \text{ cm}^{-3}$ and $N = 45 \text{ cm}^{-3}$, which are typical of the drop concentrations in clean marine environments associated with open cellular convection and which are also in line with the values of N in the LES we consider (see Section 2.2). Nonetheless, droplet concentration may not be constant in a stratocumulus cloud system or in an LES ((Yamaguchi et al., 2017)). Thus, the fixed droplet concentration may limit the usefulness of the KTF17 model in certain conditions. We discuss these issues in more detail below.

2.2 Description of the LES

A LES is a detailed model of a cloud system in space (3D) and time. It solves the anelastic Navier–Stokes equations on an Eulerian spatial grid, resolving convection and clouds, and in the current work, also simulates microphysical processes such as the formation of droplets on suspended particles (condensation nuclei), their growth by coalescence, and their removal by rain. We use the LES output to estimate the parameters of KTF17, which produces a times series of cloud depth (H in (1)). During Bayesian inversion, we will connect KTF17 to the LES by extracting time series of cloud depth from the LES (see Section 3.2).

We use the LES described in Feingold et al. (2015), with modifications. The atmospheric conditions derive from a well studied drizzling stratocumulus case, but unlike Feingold et al. (2015), the initial concentration of particles on which drops can form is about 100 cm^{-3} but decreases naturally due to droplet coalescence and rain removal processes.

The spatial domain of the LES is 40 km by 40 km wide and 1.6 km high with a grid spacing of 200 m in the horizontal and 10 m in the vertical. The simulation covers a total of 12 hrs with a time step of one second. Simulation output is available every one minute. We disregard the first 4.5 hours of the LES during which the system rapidly transitions from a closed-cell to an open-cell state. We thus only consider 7.5 hours, or 450 minutes, of simulation of an open-cell system for the Bayesian inversion.

The KTF17 model describes cloud depth as a function of time, but not any other quantities of the LES. For this reason, we consider cloud depth of the LES and disregard most other simulation outputs with the exception of droplet concentration N and column liquid water path, (see




Figure 1: Snapshots, taken every 30 mins, of the 2D cloud depth field of the LES. Examples of cloud cycles, i.e., an increase in cloud depth, followed by a decrease in cloud depth, are highlighted by blue and red squares. After Feingold et al. (2015).

Section 3.2). Figure 1 shows snapshots of the cloud depth field over the span of 7 hrs, sampled every 30 minutes, which is approximately the decorrelation timescale of the cloud field.

Figure 1 illustrates that, during the first 3.5 hrs, the system is characterized by a relatively dense collection of clouds with high average cloud fraction; a gradual transition to a lower cloud fraction then occurs as the cloud system self-organizes into a sparse collection of cloudy rings that outline a honeycomb-like pattern of cloud-free cells. We will refer to the first 3.5 hrs of simulation as the “dense phase” and to the remaining 4 hrs of simulation as the “sparse phase” (see Section 4.2). The droplet concentration falls from about 45 cm^{-3} down to 16 cm^{-3} during the course of the 7.5 hrs of simulation, as illustrated in Figure 2. We compute the droplet concentration over cloudy parts of the domain by averaging N vertically and horizontally over the entire domain and scale this average by the average cloud fraction. The consequences for parameter estimation with KTF17, which assumes a constant N , will be discussed in detail below.

2.3 Bayesian inversion

Bayesian inversion means inferring information about model parameters from data. This is done as follows. We denote the model parameters by the vector θ and we write the model as $\mathcal{M}(\theta)$. The function $\mathcal{M}$ could, for example, involve solving the KTF17 model numerically to produce time series of cloud depth (see below). A priori, one may know a few things about the parameters. For example, one may know that certain parameters must be positive to be physically relevant. In the Bayesian framework, such “prior knowledge” is expressed as a prior distribution $p_0(\theta)$. Priors are often uniform distributions. For example, if bounds on the parameters are known, then the prior can be chosen uniformly within the bounds.

For a given θ , the numerical model can be simulated and its output can be compared to data, $\mathbf{y}$. Model and data are thus connected by

$$\mathbf{y} = \mathcal{M}(\theta) + \eta, \quad (2)$$




Figure 2: Droplet concentration N , scaled by average cloud cover, as a function of time. The dashed vertical line indicates the separation of the dense and sparse phases of the simulation

where $\boldsymbol{\eta}$ represents discrepancies between the model and data, and is typically assumed to be Gaussian distributed with mean zero and covariance matrix $\mathbf{R}$. Equation (2) then defines the “likelihood”

$$p_l(\mathbf{y}|\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}\|\mathbf{R}^{-1/2}(\mathbf{y} - \mathcal{M}(\boldsymbol{\theta}))\|^2\right), \quad (3)$$

where $\mathbf{R}^{1/2}$ is a matrix square root and where the vertical bars denote the Euclidean norm. The prior and likelihood jointly define the posterior distribution

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto p_0(\boldsymbol{\theta}) p_l(\mathbf{y}|\boldsymbol{\theta}), \quad (4)$$

which describes our knowledge of the parameters and their uncertainties in view of the data. This means, in particular, that a numerical model with parameters distributed according to the posterior distribution, is “calibrated” to the data in the sense that simulations lead to model outputs compatible with the data up to the assumed errors.

2.4 Feature-based Bayesian inversion

In many Bayesian inverse problems, the model $\mathcal{M}$ is an accurate and detailed representation of the physical process that generates that data. For example, atmospheric models used for “data assimilation” and global numerical weather prediction, generate the full 3D atmospheric state. In this case, Equation (2) directly connects model outputs to measurements of the atmospheric state (data). This means that the likelihood (3) is a measure of the “point-wise” model-data mismatch, e.g., describing the differences between the observations of the atmospheric states and the predictions of the atmospheric model. Below, we will use Bayesian inversion to connect the outputs of a LES with a very simple, phenomenological predator-prey model of stratocumulus clouds without an associated spatial scale. The more common, point-wise definition of a likelihood is thus not useful for our purposes and we use a “feature-based” approach. The idea is that while a simplified model may not be able to reproduce the data in their entirety, it may be able to reproduce selected aspects of the data, see Morzfeld et al. (2018). The selected aspects that *are* reproducible by the model are called “features”. A feature-based inverse problem thus requires that we define features that are comparable in the more usual “point-wise” sense.

Specifically, we define $\mathcal{F}_{\mathcal{M}}(\boldsymbol{\theta})$ as a function that extracts the feature from the model and denote by $\mathbf{f}_y$ the feature extracted from the data. Assuming that discrepancies between the model feature and the data feature can be accounted for by a random variable, we write

$$\mathbf{f}_y = \mathcal{F}_{\mathcal{M}}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}. \quad (5)$$

If $\boldsymbol{\varepsilon}$ is Gaussian distributed with mean zero and covariance $\mathbf{R}$ (with slight abuse of notation because $\mathbf{R}$ was used above for another covariance matrix), the feature-based likelihood is

$$p_{l,f}(\mathbf{f}_y|\boldsymbol{\theta}) \propto \exp\left(-\frac{1}{2}\|\mathbf{R}^{-1/2}(\mathbf{f}_y - \mathcal{F}_{\mathcal{M}}(\boldsymbol{\theta}))\|^2\right). \quad (6)$$

We emphasize that the feature-based likelihood is defined by the Euclidian norm of the differences of the model feature and the data feature. The feature-based likelihood measures the point-wise mismatch of the features of model and data in the same way as the “usual” likelihood measures the point-wise mismatch between the model outputs and the data (see above). Assuming that a prior $p_0(\boldsymbol{\theta})$ for the model parameters is given, the feature-based posterior distribution is

$$p_f(\boldsymbol{\theta}|\mathbf{f}_y) \propto p_0(\boldsymbol{\theta}) \exp\left(-\frac{1}{2}\|\mathbf{R}^{-1/2}(\mathbf{f}_y - \mathcal{F}_{\mathcal{M}}(\boldsymbol{\theta}))\|^2\right). \quad (7)$$

In summary, a model with parameters distributed according to the feature-based posterior, produces features that are compatible with the features extracted from the data, up to the assumed errors. We emphasize that the prior has a direct influence on the shape of the posterior distribution, which is just the product of prior and likelihood. Since parameter estimates are based on the posterior distribution, a different choice of prior will ultimately result in different parameter estimates.

2.5 Markov chain Monte Carlo for the numerical solution of Bayesian inverse problems

Monte Carlo methods can be used to numerically implement the (feature-based) Bayesian inversion. The idea is to draw samples from the posterior distribution in such a way that averages over the samples converge to expected values with respect to the posterior distribution when the number of samples, N_e goes to infinity, see, e.g., Chorin and Hald (2013). In this sense, the samples, generated by the Monte Carlo method, approximate the posterior distribution and can be used for inferences, e.g., for computing the posterior mean and covariance matrix.

We use Markov chain Monte Carlo (MCMC) to draw posterior samples. A MCMC sampler operates as follows. A sample is proposed by drawing from a proposal distribution and the proposed sample is accepted with a probability that ensures that the stationary distribution of the Markov chain is the targeted posterior distribution, see, e.g., Gilks et al. (1996). The various MCMC samplers in the literature use different proposal mechanisms to speed up convergence, often by exploring specific characteristics of the sampling problem. If one does not know of a particular problem structure to exploit, one should use “general purpose” ensemble samplers, e.g., the affine invariant MCMC ensemble sampler of Goodman and Weare (2010) or the t-walk of Christen et al. (2010). These samplers are known to be effective for low-dimensional, nonlinear/non-Gaussian problems and efficient implementations are also available.

To assess the accuracy of the MCMC solution one computes the integrated auto-correlation time (IACT), see, e.g., Sokal (1996); Wolff (2004). The idea is that, while MCMC samples are generally not independent, one can estimate an effective number of independent samples by

$$N_{s,\text{eff}} = \frac{N_s}{\text{IACT}}, \quad (8)$$

where N_s is the number of samples from the MCMC sampler. The reasoning is that if one has, e.g., 10^6 samples, and one has computed IACTs of a few hundred, then one should expect an accuracy that is comparable to that computed with thousands of independent samples.

3 Feature-based Bayesian inversion of the LES

The KTF17 model parameters are the carrying capacity H_0 , the delay T , the characteristic time τ and the scaling factor α . We combine these four parameters in the parameter vector $\boldsymbol{\theta} = [H_0, \tau, T, \alpha]^T$. Our goal is to compute the model parameters $\boldsymbol{\theta}$ by a feature-based inversion of the LES output. As described in Section 2.4, a feature-based inversion requires that we define a prior distribution and a feature-based likelihood. We now describe in detail how these distributions are constructed. The feature-based posterior follows from these two distributions and is used for inferences, numerically implemented by MCMC.

3.1 Prior distribution

The prior distribution describes our a priori knowledge of the KTF17 model parameters. We define this prior to be a uniform distribution over the set of parameters that are (i) physically relevant (positive); and (ii) lead to physically relevant limit cycles in KTF17. All parameters that satisfy these conditions, receive the same nonzero prior probability while all other parameters receive zero prior probability. Thus, a parameter vector $\boldsymbol{\theta} = [H_0, \tau, T, \alpha]^T$ must satisfy the following four conditions in order to receive non-zero prior probability.

1. All four model parameters must be positive.
2. The characteristic time to reach carrying capacity is longer than the delay time.
3. The parameter vector must produce solutions that are limit cycles.
4. Cloud depth must be positive.

For condition 3, we rely on the linear stability analysis in Koren et al. (2017). The parameters that lead to limit cycles in KTF17 are characterized by the real part of a dimensionless parameter β being positive. Here,

$$\beta = \frac{\tau}{T} W \left(-2 \left[\sqrt{\frac{1}{\mu} + \frac{1}{4}} - \frac{1}{2} \right] \frac{T}{\tau} \exp \left[\frac{T}{\tau} \right] \right) - 1, \quad (9)$$

where $\mu = \sqrt{N}/(\alpha\tau H_0)$ and $W(\cdot)$ is the Lambert-W function. In other words, limit cycles occur only if $\text{Re}(\beta) > 0$. Condition 4, i.e., checking for negative cloud depth, requires a simulation. For a given $\boldsymbol{\theta}$, we solve KTF17 numerically and if we detect negative cloud depth, the parameter vector receives zero prior probability. To streamline computations, we check for negative cloud depth after checking conditions 1-3.

The prior is illustrated in terms of a “triangle plot” in the left panels of Figure 3. A triangle plot contains histograms of all one and two-dimensional marginals of a given distribution, arranged in the form of a triangle; each marginal is normalized so that the integral (area under the graph) is equal to one. A triangle plot is, thus, a qualitative tool that illustrates regions in parameter space that receive a large probability. Recall that the prior contains the information we have about model parameters before the data are taken into account. Per our construction of the prior, this means that a triangle plot of the prior illustrates regions in parameter space that lead to physically relevant limit cycles of cloud depth.




Figure 3: Left: Triangle plot of 10^5 samples of the prior distribution. Blue indicates a low probability while red indicates a high probability. Right: 10^3 limit cycles of KTF17 corresponding to 10^3 parameter vectors drawn at random from the prior. The cycles are aligned to reach their peak depths at the same time. Five examples of cloud cycles are highlighted in purple.

Generating the triangle plot requires that we draw samples from the prior which we do via “importance sampling” with a proposal distribution that is uniformly distributed over the four-dimensional hyper-cube defined by the lower and upper parameter bounds listed in Table 1 (for

	H_0 , m	τ , min	T , min	α , days $^{-1}$ m $^{-2.5}$
Lower bound	0	0	0	100
Upper bound	4000	288	288	2000
Prior mean	1650	137	43	836
Prior std. dev.	1067	61	27	495

Table 1: Mean and standard deviations computed from 10^5 samples of the prior.

more details about importance sampling, see, e.g., Owen (2013); Chorin and Hald (2013)). The samples that constitute the triangle plot can also be used to compute prior means and standard deviations, listed in Table 1. We note that the standard deviations are between 40%-60% of the corresponding mean values, which indicates that the prior is “broad”, i.e., large parts of the parameter space receive non-zero prior probability.

The “broadness” of the prior is further illustrated in the right panel of Figure 3, which shows one period of 10^3 limit cycles of KTF17 corresponding to 10^3 prior samples of the prior (see Section 3.2 for details of how we compute these limit cycles). The limit cycles are arranged so that their maxima occur at the same instant. We observe a large variance in the period and amplitude of the cloud cycles. This means in particular that, a priori, we do not know the typical period or amplitude of a cloud depth cycle. The goal of a Bayesian inversion is to refine the prior distribution to a posterior distribution, which reduces variations in the cloud cycles via reducing variance in the parameters; the reduction of variance of the parameters is achieved by taking the LES into account.




Figure 4: 2D cloud depth field at $t = 4$ hrs at full resolution (256×256 , left) and spatially averaged cloud depth (32×32 , right). Time series of cloud depth for the locations encircled in red and orange are shown in Figure 5.

3.2 Feature-based likelihood

A feature-based likelihood requires that we define features of the model that can be compared to features extracted from the data. We now describe how we construct these features and an associated Gaussian error model.

3.2.1 Data feature

The data-feature is derived from the time varying 2D cloud depth field of the LES which defines 256×256 time series of cloud depth at each grid point (with no advection present). These time series, however, are noisy. To reduce the effects of this noise, we spatially average the 2D cloud depth field over small, square “tiles” that contain a few grid points. We average cloud depth only over regions where cloud exists, which we define by a positive integral of the liquid water content over the depth of the cloud (liquid water path), taken from the LES. We considered several tile sizes for the averaging and settled on tiles containing 8×8 grid points (see also Section 4.1). With a horizontal grid spacing of 200 m, this results in a “filter” length of about 1,600 m, which is large enough to smooth out noise, but retains the main aspects of the cellular structure. The full resolution (256×256) and the spatially averaged (32×32) cloud depth fields are illustrated in Figure 4.

The spatial averaging yields 1024 time series of cloud depth, H , over 7.5 hrs. We extract cycles of growth and decay from these time series as follows. We first apply a temporal smoothing by applying a Gaussian filter with a standard deviation of 10 minutes. We then compute local extrema of the filtered time series via finite differencing. Two consecutive local minima define one cycle and each cycle (without temporal smoothing) is stored. With this procedure, we extract 297 cycles from the LES.

The procedure of the feature extraction is illustrated in Figure 5. Panels (a) and (b) show H , after temporal and spatial smoothing, at the locations encircled in red and orange in Figure 4. Also shown are the extracted cycles (without temporal smoothing). The 297 cycles we compute are shown in light blue in panel (c); the four cycles, shown in panels (a) and (b), are also shown (in thicker purple, pink, brown and yellow lines). We align all cycles so that they reach their peaks at the same instant and pad shorter cycles with zeros, so that all cycles have a duration of 270 minutes (see also Section 4.1). The feature $\mathbf{f}_y$ is the average of the 297 cycles, shown as a thick dark blue line in panel (c).




Figure 5: (a) Cloud depth time series after temporal smoothing (red) for the location encircled in red in the right panel of Figure 4. Shown in purple and pink are the two cycles extracted from this time series (without temporal smoothing). (b) Cloud depth time series after temporal smoothing (orange) for the location encircled in orange in the right panel of Figure 4. Shown in brown and yellow are the two cycles extracted from this time series (without temporal smoothing). (c) 297 cycles, extracted from the LES (without temporal smoothing), are shown in light blue. The cloud cycles from panels (a) and (b) are shown in a thicker purple, pink, brown and yellow lines. The dark blue line is the average of the 297 cycles.

3.2.2 Model feature

The model-feature is defined as one limit cycle of KTF17. The limit cycle and the time needed to reach it depend on the value of the model parameters and the initial condition. During the feature-based Bayesian inversion, implemented by an MCMC sampler (see below), the initial conditions are fixed, but we need to find limit cycles corresponding to different parameter values (all with non-zero prior probability).

To robustly compute limit cycles we use the following iterative scheme. We first solve KTF17 numerically for one day (the initial condition is $H(t) = 0.1$ m for $t \leq 0$) and approximate the time derivative of $H(t)$ by finite differences to find the extrema of the cloud depth time series. The time instances of two consecutive local minima define one cycle of growth and decay (note that the data feature is defined in the same way). To check if a limit cycle is reached, we compare the root mean square error (RMSE) between the last two cycles and, if RMSE is less than 1 m, we stop the numerical solution and conclude that the system has reached its limit cycle. Otherwise, we continue the numerical solution of KTF17 for an additional day and, again, find local extrema to define cloud cycles and compute RMSE of the last two cycles. We repeat this process until two consecutive cycles are characterized by an RMSE of less than 1 m. The model-feature is then defined to be the last cycle of the cloud depth time series.

We align the peaks of the model- and data features and modify the model-feature to have the same duration (270 mins) as the data feature. Specifically, if the model feature has a shorter duration than the data feature, we pad the model feature with zeros (symmetrically before and after its peak). If the model feature is longer than the data feature, we truncate it (symmetrically before and after its peak).

Finally, we note that we are not aware of a proof that KTF17 has only *one* limit cycle for a given set of parameters with non-zero prior probability. Extensive numerical experiments, however, suggest that this is indeed the case. In particular, we performed a large number of simulations for several parameter vectors, drawn from the prior, starting at different initial conditions conditions $0 < H(0) < 500$ m (with $H(t) = H(0)$ for $t < 0$) and, for each parameter vector, found only one limit cycle, independent of the initial conditions.

3.2.3 Gaussian error model

To finish the construction of the feature-based likelihood we need to define the errors $\boldsymbol{\varepsilon}$ in Equation (5). As is customary, we assume a Gaussian distribution with a mean of zero. The covariance matrix that defines the error model is computed based on the variations of the 297 cloud cycles extracted from the LES. Specifically, we define the covariance $\mathbf{P}$ as the sample covariance of the 297 cycles and then choose the covariance $\mathbf{R}$ of $\boldsymbol{\varepsilon}$ in (5) as

$$\mathbf{R} = \mathbf{P} + \sigma^2 \mathbf{I}, \quad (10)$$

where $\mathbf{I}$ is the identity matrix and $\sigma = 100$ m. Note that $\mathbf{P}$, $\mathbf{R}$ and $\mathbf{I}$ are matrices of size 270×270 , because each (padded) cloud cycle has a duration of 270 mins and the time step is one minute. We use an additive “inflation” of the covariance $\mathbf{P}$ because the padding leads to small variances at the beginning and end of the 270 min time interval. We will assess, in hindsight, our assumptions about errors in the features as well as how the padding with zeros affects the results in Section 4.

Each element of Equation (5) is now defined, which implies the feature-based likelihood by (6). Together with the prior, the feature-based likelihood defines the feature-based posterior, which can be written as

$$p_f(\boldsymbol{\theta}|\mathbf{f}_y) \propto \begin{cases} 0 & \text{if } p_0(\boldsymbol{\theta}) = 0, \\ \exp\left(-\frac{1}{2}\|\mathbf{R}^{-1/2}(\mathbf{f}_y - \mathcal{F}_{\mathcal{M}}(\boldsymbol{\theta}))\|^2\right) & \text{otherwise,} \end{cases} \quad (11)$$

where $\mathbf{R}$, $\mathbf{f}_y$ and $\mathcal{F}_{\mathcal{M}}(\boldsymbol{\theta})$ are as above.

3.3 Numerical solution by MCMC

We use the python implementation of the t-walk (see <https://www.cimat.mx/~jac/twalk/>) and the python implementation “emcee” of the affine invariant ensemble sampler (Foreman-Mackey et al., 2013). Below we only show results obtained by emcee, but results obtained by the t-walk are qualitatively and quantitatively similar. The emcee sampler requires an ensemble of N_e “walkers”, where

$$N_e \geq 2 \times (\text{number of model parameters}) = 8. \quad (12)$$

We chose an ensemble size of $N_e = 20$, because larger ensemble sizes are preferable (Foreman-Mackey et al., 2013). The initial ensemble is generated as follows. We draw 10^3 samples from the prior distribution and, for each one, evaluate (7), which is proportional to the posterior probability. The 20 samples with the highest values, which also correspond to the samples with the highest posterior probabilities, are the initial ensemble used in emcee.

Our code can be found at <https://github.com/SpenceLunderman/LMGF20> and can generate 10^5 samples in about 10 hrs and 10^6 samples in about 4 days (on a single core). For the results shown below, we discard the first N_{discard} samples as “burn-in”, where $N_{\text{discard}} = 5 \cdot \max \text{IACT}$, and $\max \text{IACT}$ is the largest IACT of the four parameters. Based on $2 \cdot 10^6$ samples, we compute IACTs of a few hundred (see below), which indicates that the number of samples we generate is sufficiently large (accuracy comparable to thousands of independent samples).

4 Results and discussion

We perform the feature-based inversion, as described above, using a constant droplet concentration of $N = 25 \text{ cm}^{-3}$, which is the time-average of N during the 7.5 hrs of simulation considered. In this context, it is important to realize that the effect of a varying N over the range encountered in the LES has a minor effect. The reason is that Equation (1) implies that changes in N result in a scaling




Figure 6: Left: Triangle plot of the posterior distribution ($2 \cdot 10^6$ samples). Right: Shown in green are the limit cycles of KTF17 corresponding to 10^4 parameter vectors drawn at random from the posterior. The LES feature (average of 297 LES cloud cycles) is shown as a dark blue line. The light blue shaded region represents two sample standard deviations of the cloud cycles at each time instant (representing variations in the cloud cycles extracted from the LES).

of α with the square root of N , but all other parameters are independent of the value of N . In particular, if α_0 is estimated by assuming $N = N_0$, then setting $N \rightarrow N_1$ results in $\alpha_1 = \alpha \sqrt{N_1/N_0}$.

The results of the feature-based inversion, based on an MCMC chain with $2 \cdot 10^6$ samples, are illustrated in Figure 6. The left panel shows a triangle plot of the posterior samples, obtained via the MCMC, and the right panel shows 10^4 limit cycles of KTF17, corresponding to 10^4 parameter vectors drawn at random from the posterior. Also shown are the LES feature and the variations in the cloud cycles extracted from the LES. This figure should be compared to Figure 3, which shows the same information *before* the Bayesian inversion, i.e., based on the prior distribution. We note that the posterior distribution is more sharply peaked than the prior (note the different axes in the triangle plots of Figures 3 and 6), which indicates that the LES derived feature indeed constrains all four parameters of KTF17.

The sharpening of the prior to a feature-based posterior distribution can also be seen by computing the sample mean and sample standard deviations, listed in Table 2. We note a shift in the

	H_0 , m		τ , min		T , min		α , $\text{days}^{-1}\text{m}^{-2.5}$	
	Prior	Posterior	Prior	Posterior	Prior	Posterior	Prior	Posterior
Mean	1650	2063	137	120	43	33	836	548
Std.	1067	722	61	48	27	7	495	176
MAP	-	2062	-	131	-	36	-	450

Table 2: Mean and standard deviations of the prior and posterior distributions. The MAP of the posterior is also listed. Posterior quantities are computed from a MCMC chain with $2 \cdot 10^6$ samples; prior quantities are computed from 10^5 samples of the prior.

sample mean and a reduction in sample standard deviations from the prior to posterior distribution. Table 2 further lists the maximum a posteriori (MAP) estimates, i.e., the sample with the largest

posterior probability ¹. We note that the MAP and mean are not equal, which indicates that the posterior distribution is not nearly Gaussian. In this context, it is also important to realize that the posterior mean is *not* a posterior sample, i.e., its posterior probability can be zero (because it may not satisfy all four prior constraints). For this reason, the MAP may be a more useful estimate of the KTF17 parameters than the posterior mean.

The left panel of Figure 6 illustrates that cycles of KTF17, obtained by numerical solution of KTF17 with parameters sampled from the posterior, are well within the variations of the cloud cycles extracted from the LES. This indicates that our error model and the error covariance matrix $\mathbf{R}$ are reasonable. Here, we tuned, to some extent, the additive inflation defined by σ in (10). Recall that error models are notoriously difficult to come by because error models represent “what we do not know” about the system. Our approach here is to introduce a tunable covariance inflation factor, σ , that is selected so that the posterior uncertainties, as illustrated by the trajectory ensemble in the right panel of Figure 6, are reasonable, and within the expected uncertainties, derived directly from the LES.

We can use the results of the feature-based inversion to investigate if the cycles of KTF17 have similar properties as the cycles extracted from the LES. Specifically, we can consider the period, amplitude, and growth and decay times of the KTF17 and LES derived cycles. Here, the period is the duration of the cloud cycle (without zero padding); the amplitude is the difference between the maximum and minimum cloud depth reached during a cycle ². The cycle growth time describes how long it takes a cloud to build up to its maximum cloud depth, and the decay time describes how long it takes to decay from maximum cloud depth to its minimum (equivalently, the decay time is equal to the period minus the growth time). These four properties are computed for each cloud cycle extracted from the LES and for 10^4 KTF17 limit cycles, defined by parameters that are drawn from the posterior distribution. The means and standard deviations of the four cycle properties are listed in Table 3. We note that the mean of each cycle property, computed from KTF17, is

	Period, min		Amplitude, m		Growth time, min		Decay time, min	
	LES	KTF17	LES	KTF17	LES	KTF17	LES	KTF17
Mean	137	119	619	591	69	66	69	55
Std. Dev.	23	26	142	102	23	15	19	12

Table 3: Mean and standard deviations of cloud cycle properties of the LES and KTF17. LES results are computed from 297 cycles and KTF17 results are computed from 10^4 simulations with parameters drawn from the posterior distribution.

within one standard deviation of the mean of the corresponding property computed from the LES. Moreover, the standard deviations of the LES and KTF17 cycle properties are also comparable, which suggests an overall good “fit” of KTF17 to the LES in terms of these cycle properties.

To report on the statistical accuracy of the MCMC solution, we list the IACTs, estimated from the $2 \cdot 10^6$ samples, of all four parameters in Table 4. The IACTs are less than 10^3 , which indicates that the number of samples is sufficient to accurately compute posterior means, standard deviations

¹It is important to remember that marginal distributions, shown in the form of histograms in the triangle plots, are not “projections” of the multivariate probability distribution. For this reason, the maxima of the posterior marginals (histograms) do not correspond to the mode of the multivariate posterior distribution (MAP).

²We emphasize that the blue line, shown in Figure 6, is the average of the LES cycles, but taking into account the zero padding, and stitching the cloud cycles together at their maximum value. This means that the maximum value of the blue line in Figure 6 equals the average maximum cloud depth over all cycles, which is different from the average amplitude in Table 3. The same reasoning explains why the average amplitude of KTF17, reported in Table 3, is different from what one might expect by visually taking the average of the green lines in Figure 6.

H_0 , m	τ , min	T , min	α , days $^{-1}$ m $^{-2.5}$
620	326	539	665

Table 4: Integrated autocorrelation times (computed from the $2 \cdot 10^6$ samples).

and the MAP, with an effective sample size in the thousands.

4.1 Robustness of the LES feature

The computational framework we describe, and in particular the construction of the LES feature, relies on several assumptions and modeling choices. The Bayesian approach and MCMC implementation allow us to investigate, numerically, the validity of our assumptions and choices. We already described the effects of the error model and our choice of additive covariance inflation (see Figure 6). We now investigate the robustness of the LES feature to two other modeling choices: the spatial smoothing and the zero-padding of the cloud cycles (see Section 3.2.1).

4.1.1 Robustness to spatial smoothing

While it is difficult to determine the precise amount of spatial smoothing, it is clear that (i) smoothing is necessary, or else the cloud depth time series are too noisy; and (ii) that there is a maximum amount of smoothing that should not be exceeded, or else the effects of cloud entities are averaged out. We investigate this issue by performing the feature-based Bayesian inversion for three spatial averages over “tiles” consisting of 4×4 , 8×8 and 16×16 grid points respectively. With each spatial averaging, we compute the data-feature and perform the Bayesian inversion via MCMC, generating $2 \cdot 10^6$ samples in each configuration. In all three cases, the prior distribution is the same as above, because the prior is independent of the definition of features, or, equivalently, the likelihood. We also keep all other aspects (covariance inflation, temporal smoothing etc.), that define the data-feature, as above.

Table 5 lists the posterior mean, standard deviation and MAP estimates for three spatial smoothings, computed from three MCMC runs with $2 \cdot 10^6$ samples. We note that the parameter estimates

Tile width	H_0 , m			τ , min			T , min			α , days $^{-1}$ m $^{-2.5}$		
	4	8	16	4	8	16	4	8	16	4	8	16
Mean	2344	2063	1776	114	120	141	29	33	41	560	548	458
Std.	989	722	544	48	48	49	7	7	6	183	176	83
MAP	2147	2062	2215	119	131	188	34	36	42	453	450	414

Table 5: Posterior means, standard deviations, and MAP estimates for the four parameters of KTF17 and for the three configurations which differ in their spatial smoothing of the LES cloud depth field. Posterior means, the MAP, and standard deviations are computed from the MCMC chain with $2 \cdot 10^6$ samples.

for the three configurations are within a standard deviation of each other, independently of which standard deviation one chooses to use. The only exception is the parameter T , where the estimates for the 16×16 case are within two standard deviations of the 4×4 or 8×8 scenario. A smoothing over 16×16 grid points may, therefore, be labeled as excessive.

Nonetheless, averaging over tiles of size 4×4 or 8×8 gives nearly identical results, which indicates some robustness of our approach with respect to spatial smoothing. We emphasize, however, that

a significantly larger amount of smoothing (tiles consisting of more than 16×16 grid points) does not lead to reasonable parameter estimates because the effects of cloud entities are averaged out.

4.1.2 Robustness to padding of LES cycles

In the construction of the LES feature, cycles are aligned at their peaks of cloud depth. The cycles are then “padded with zeros” so that all cycles have the same duration (270 min). The LES feature is simply the average of the padded cycles. Below, we call this construction “Version (a)”. We now investigate the consistency of the parameter estimation results when we choose another method to derive the LES feature that does not make use of zero padding.

In “Version (b)”, we again align all cloud cycles at their peaks, but rather than padding with zeros, we average only those cycles that “exist”, i.e., which have non-zero cloud depth at a given time instant. We further exclude all instances where less than 10 cycles exist. The error covariance of ϵ for Version (b) is computed in the same way as in Version (a). Figure 7 illustrates the LES feature of versions (a) and (b). Note that the duration of the average in Version (b) is shorter than in Version (a) because we only consider instances when at least 10 cycles are non-zero and longer cycles occur less frequently. We further note that near the peak, these two versions are equal because, at peak times, zero padding in Version (a) has no effect and no cycles are excluded in Version (b) because more than 10 cycles exist.

We perform a feature-based Bayesian inversion for LES features constructed using versions (a) and (b) and, as before, generate $2 \cdot 10^6$ samples by MCMC. Results are shown in Table 6.

	H_0 , m		τ , min		T , min		α , days $^{-1}$ m $^{-2.5}$	
	Version (a)	Version (b)	V. (a)	V. (b)	V. (a)	V. (b)	V. (a)	V. (b)
Mean	2063	2008	120	126	33	35	548	526
Std.	722	684	48	50	7	9	176	180
MAP	2062	2321	131	164	36	38	450	430

Table 6: Posterior means, standard deviations and MAP estimates for the four parameters of KTF17 and for the two configurations which differ in their calculation of the LES feature. Version (a) and Version (b) correspond to the two LES features shown in Figure 7. Posterior means, the MAP and standard deviations are computed from the MCMC chain with $2 \cdot 10^6$ samples.

We note that the parameter estimates resulting from versions (a) and (b) are not significantly different. The reason is that most KTF17 cycles occur between 50 and 200 minutes (see Figure 6), i.e., when the two LES features of versions (a) and (b) are similar. The similarities between these two posterior distributions can also be seen in the marginal distributions in Figure 8. The left panel shows a triangle plot of the posterior distribution of Version (a) and the right panel shows a triangle plot of the posterior distribution of Version (b). In both panels, the plots on the diagonals show the one-dimensional marginal distributions of both posteriors in black (Version (a)) and blue (Version (b)). In summary, the similarity in the parameter estimates and posterior distributions of Versions (a) and (b) suggests that estimation framework we describe is robust to small changes in the details of how one calculates the data feature.

4.2 Studying changes in cloud system morphology

The Bayesian inversion and the KTF17 model will prove useful if one can map meteorological conditions to changes in the parameters of KTF17. We illustrate how to do this with a simple example in which we start to investigate the effects of large-scale changes within the cloud field on




Figure 7: Cloud cycles and data feature of Version (a) (left) and Version (b) (right). Light blue: LES cycles. Thick blue: LES feature.




Figure 8: Left: triangle plot of the posterior distribution of Version (a). Right: triangle plot of the posterior distribution of Version (b). The diagonal plots of each panel show the one-dimensional marginals of both distributions (in black for Version (a), in blue for Version (b)). All plots are based on $2 \cdot 10^6$ MCMC samples.

the parameters of KTF17. We base this investigation on only one LES, which represents one cloud system and, for that reason, our results and conclusion are limited.

We note that the cloud system undergoes a change in its morphology from a relatively dense cloud configuration with a higher average cloud fraction to a sparse coverage with a lower average cloud fraction (see Figure 1). During this transition, the droplet concentration also decreases (see Figure 2). The transition occurs roughly at the 3.5 hour mark and aligns with a change in the thickness of the boundary layer, whose thickness increases until about 3.5 hours, and then decreases.

To investigate the effects of the morphological change in the macro-structure of the cloud system on the parameters of KTF17, we perform two feature-based inversions as follows. We separate the cloud cycles, extracted from the LES, into two groups: cycles occurring before and after the transition from the dense to the sparse cloud cover, i.e., before or after 3.5 hr. For example, the cloud cycles shown in purple and brown in panels (a) and (b) of Figure 5 occur before the transition (dense phase), but the cycles shown in pink and yellow occur after the transition (sparse phase). In this way, we obtain 166 cycles during the dense phase and 131 during the sparse phase, shown along with their averages (using zero-padding) in Figure 9.




Figure 9: Cloud cycles and data feature for the dense (a) and sparse phases (b). Light blue: LES cycles. Thick blue: LES feature.

We compute data-features separately for the dense and sparse phases using the techniques described above (using the default spatial smoothing over tiles consisting of 8×8 grid points and Version (a), i.e., zero padding of the cycles). In this way, we define feature-based likelihoods for the dense and sparse phases. We use the same prior for the dense and sparse phases to define two posterior distributions. We assign the time average of the droplet concentration in Figure 2, computed separately over the dense and sparse phases, as the values used for N in the Bayesian inversion. Specifically, we chose $N = 31 \text{ cm}^{-3}$ for the dense phase and $N = 20 \text{ cm}^{-3}$ for the sparse phase. As before, we use the MCMC sampler to draw $2 \cdot 10^6$ samples from the posterior distributions associated with the dense and sparse phases.

Table 7 lists parameter estimation results for the two dense and sparse phases. We note that

	H_0 , m		τ , min		T , min		α , $\text{days}^{-1}\text{m}^{-2.5}$	
	Dense	Sparse	Dense	Sparse	Dense	Sparse	Dense	Sparse
Mean.	2028	1886	122	110	36	32	525	535
Std.	615	616	45	44	6	7	117	166
MAP	2112	2408	130	165	36	36	483	405

Table 7: Maximum a posterior (MAP) estimate, posterior mean and posterior standard deviation for the dense and sparse phases, computed from MCMC chains with $2 \cdot 10^6$ samples.

the parameter estimates (posterior mean and MAP) are within one standard deviation of each other. Furthermore, the parameter estimates listed in Table 7 are comparable with the parameters in Table 5, which are estimated based on *all* cloud cycles extracted from the LES (i.e., cycles in dense and sparse configurations). Similarities and differences in the parameter estimates can also be illustrated by triangle plots of the two posterior distributions, shown in Figure 10. The left panel shows the posterior distribution associated with the dense phase; the right panel shows the posterior distribution associated with the sparse phase. It is apparent that the posterior distributions are quite similar, but it is also apparent that there are differences, especially in the delay T and the scaling factor α .

It is difficult to determine whether or not the differences in the parameter estimates are significant. Taking into account the standard deviations as an indicator of uncertainty, one may be tempted to conclude that the differences are not significant. One can study this further by comparing the differences in parameter estimates induced by the dense and sparse phases, with differences induced by variations in the smoothing or paddings. Figure 11 illustrates this point and shows 1D




Figure 10: Left: triangle plot of the posterior distribution associated with the dense phase. Right: triangle plot of the posterior distribution associated with the sparse phase. The diagonal plots of each panel show the one-dimensional marginals of both distributions (in black for the dense, in blue for the sparse phase). All plots are based on $2 \cdot 10^6$ MCMC samples.




Figure 11: Marginal posterior distributions over the four parameters. Left to right: H_0 , τ , T and α . Green: 8×8 spatial smoothing, zero padding, all cycles. Blue: 4×4 spatial smoothing, zero padding, all cycles. Red: 8×8 spatial smoothing, no padding, all cycles. Orange: 8×8 spatial smoothing, zero padding, dense cycles. Purple: 8×8 spatial smoothing, zero padding, sparse cycles.

posterior marginals over the four parameters for five of the cases considered. Three of the cases (green, blue and red in Figure 11) indicate uncertainty in parameter estimates induced by variations in the numerical setup. Variation in the posterior distributions indicates the variability one should expect due to different choices in the numerics. We then overlay the posterior distributions defined by only the dense or only the sparse phase cycles (orange and purple in Figure 11). The variation of these distributions indicates variability in the parameter estimates caused by changes in the large scale structure of the cloud system. We note, as before, the largest differences in the delay T and the scaling factor α .

One can further investigate how differences in the parameter estimates propagate to characteristics of the cloud cycles, such as their period, amplitude, and growth and decay times (see above for definitions). We compute the period, amplitude, and growth and decay times based on the LES for three cases (i) using all cycles; (ii) using only cycles of the dense phase; (iii) using only cycles of the sparse phase. We then repeat this procedure for the KTF17 model with parameters drawn from the posterior distributions corresponding to the above three cases. Figure 12 illustrates this point. Here we plot the average and standard deviation of the cloud cycle properties for the LES (left) and KTF17 (right) for the three cases; all quantities are scaled by the associated mean value




Figure 12: Mean and standard deviation of the period, amplitude, growth time and decay time of cloud cycles. Left: LES. Right: KTF17. Blue diamond – all cycles. Orange circle – dense cycles. Green square – sparse cycles. All quantities are scaled by the average values computed from all cycles of the LES.

of the LES case (i). We note that the properties of the LES do not change dramatically when moving from dense to sparse phases. Moreover, the cycle properties of KTF17 are comparable with those of the LES, but we observe a larger spread in the amplitude. Since the cycle properties do not change much during the transition from the dense to the sparse phase, one might expect that parameters of the KTF17 model should also be largely unaffected by this transition. Taking also the similarities in the parameter estimates and posterior distributions into account, one might conclude that the cycles of cloud patches within a cloud system may not necessarily be affected by changes in the macro-structure of the cloud system. This conclusion, however, is based on a single LES, which represents a case study with specific large-scale and thermodynamic boundary conditions. It is conceivable that KTF17 parameters will be sensitive to these boundary conditions.

5 Summary and conclusions

Stratocumulus clouds are an important part of the Earth system and have a large effect on Earth’s overall radiative balance and climate. For these reasons, stratocumulus cloud systems are studied computationally by a hierarchy of models ranging from simplified, phenomenological models to cloud resolving simulations of the atmosphere. We described a conceptual and computational strategy for turning a simplified, phenomenological model into a quantitative tool. Specifically, we use the nonlinear rain equation (KTF17) and estimate its parameters from the outputs of a large eddy simulation (LES).

The main technical difficulty for such a parameter estimation is that the phenomenological model and the LES operate in vastly different regimes in terms of what the two models are actually capable of. In particular, the LES has temporal and spatial scales, whereas KTF17 has no associated spatial scale. We overcame these difficulties by realizing that the KTF17 model produces cycles of cloud growth and decay that are comparable to cycles within the LES. We use cycles of growth and decay to define “features” and base the parameter estimation on these features. This includes deriving error models for the features which in turn allows us to formulate the parameter estimation problem within the Bayesian framework. The resulting Bayesian inverse problem is solved numerically by a

Markov chain Monte Carlo method, which allows us to assess posterior uncertainties.

We carefully studied the validity of our assumptions and modeling choices. The error model, which is notoriously difficult to construct because it represents “what we do not know”, was set up to have one tunable parameter (defining an additive covariance inflation). This parameter is tuned so that posterior uncertainties are reasonable and match the variability in the cycles derived from the LES. In addition, we investigated the robustness of our approach to the details of the construction of the features (spatial smoothing of the LES cloud depth fields and zero-padding of resulting cloud cycles). Stringent tests of this type are necessary to show that the estimated parameters are precise enough for drawing conclusions. Our numerical experiments indeed suggest that the KTF17 model, with parameters distributed according to the feature-based posterior distribution, is robustly capable of representing cloud cycle properties of a LES.

As an illustration of how one may use a simplified model as a quantitative tool, we investigated the sensitivity of the KTF17 parameters to morphological changes within the cloud system simulated by the LES. The system evolves from relatively dense cloud configuration to a sparse coverage (see Figure 1). The KTF17 parameters do not change significantly during the morphological transition of the system, which suggests that cycles of cloud growth and decay of cloud patches may be independent of the large-scale behavior of the system. This result, however, is conditional on the one LES we considered and it is likely that the KTF17 parameters are indeed sensitive to changes in other meteorological conditions, e.g., in the boundary conditions. Future work will explore this idea with a range of LES in different meteorological conditions.

Acknowledgments

We thank Prof. Joceline Lega of the University of Arizona for interesting and inspiring discussion. GF acknowledges a long collaboration with Ilan Koren on the predator prey model as applied to cloud systems.

SL and MM gratefully acknowledge support by the National Science Foundation under grant DMS-1619630. FG acknowledges support by a National Research Council Research Associateship award at the National Oceanic and Atmospheric Administration (NOAA), by the Innovational Research Incentives Scheme Veni of the Netherlands Organisation for Scientific Research and by The Branco Weiss Fellowship – Society in Science, administered by the ETH Zurich.

References

- Agee, E. M. (1984). Observations from space and thermal convection: A historical perspective. *Bulletin of the American Meteorological Society*, 65(9):938–949.
- Asch, M., Bocquet, M., and Nodet, M. (2017). *Data Assimilation: Methods, Algorithms, and Applications*. SIAM, Philadelphia.
- Boucher, O., Randall, D., Artaxo, P., Bretherton, C., Feingold, G., Forster, P., Kerminen, V.-M., Kondo, Y., Liao, H., Lohmann, U., Rasch, P., Satheesh, S., Sherwood, S., Stevens, B., and Zhang, X. (2013). Clouds and aerosols. In Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P., editors, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to IPCC AR5: Clouds and Aerosols*. Cambridge.

- Chorin, A. and Hald, O. (2013). *Stochastic Tools in Mathematics and Science*. Springer, New York, third edition.
- Christen, J. A., Fox, C., et al. (2010). A general purpose sampling algorithm for continuous distributions (the t-walk). *Bayesian Analysis*, 5(2):263–281.
- Feingold, G. and Koren, I. (2013). A model of coupled oscillators applied to the aerosol cloud precipitation system. *Nonlinear Processes in Geophysics*, 20(6):1011–1021.
- Feingold, G., Koren, I., Yamaguchi, T., and Kazil, J. (2015). On the reversibility of transitions between closed and open cellular convection. *Atmospheric Chemistry and Physics*, 15(13):7351–7367.
- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J. (2013). Emcee: the MCMC hammer. *Publications of the Astronomical Society of the Pacific*, 125(925):306.
- Gilks, W., Richardson, S., and Spiegelhalter, D. (1996). Introducing Markov chain Monte Carlo. In Gilks, W., Richardson, S., and Spiegelhalter, D., editors, *Markov Chain Monte Carlo in Practice*, chapter 1, pages 1–20. Springer-Science+Business Media.
- Glassmeier, F. and Feingold, G. (2017). Network approach to patterns in stratocumulus clouds. *Proceedings of the National Academy of Sciences*, 114(40):10578–10583.
- Goodman, J. and Weare, J. (2010). Ensemble samplers with affine invariance. *Communications in applied mathematics and computational science*, 5(1):65–80.
- Koren, I. and Feingold, G. (2011). Aerosol–cloud–precipitation system as a predator-prey problem. *Proceedings of the National Academy of Sciences*, 108(30):12227–12232.
- Koren, I. and Feingold, G. (2013). Adaptive behavior of marine cellular clouds. *Scientific Reports*, 3:2507.
- Koren, I., Tziperman, E., and Feingold, G. (2017). Exploring the nonlinear cloud and rain equation. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(1):013107.
- Macleán, J., Santitissadeekorn, N., and Jones, C. K. (2017). A coherent structure approach for parameter estimation in Lagrangian data assimilation. *Physica D: Nonlinear Phenomena*, 360:36–45.
- Morzfeld, M., Adams, J., Lunderman, S., and Orozco, R. (2018). Feature-based data assimilation in geophysics. *Nonlinear Processes in Geophysics*, 25(2):355–374.
- Myhre, G., Shindell, D., Bréon, F.-M., Collins, W., Fuglestad, J., Huang, J., Koch, D., Lamarque, J.-F., Lee, D., Mendoza, B., Nakajima, T., Robock, A., Stephens, G., Takemura, T., and Zhang, H. (2013). Anthropogenic and natural radiative forcing. In Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S., Boschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P., editors, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to IPCC AR5: Anthropogenic and Natural Radiative Forcing*. Cambridge.
- Nober, F. J. and Graf, H. F. (2005). A new convective cloud field model based on principles of self-organization. *Atmospheric Chemistry and Physics*, 5:2749–2759.
- Owen, A. (2013). *Monte Carlo Theory, Methods and Examples*. <https://statweb.stanford.edu/~owen/mc/>.

- 668 Reich, S. and Cotter, C. (2015). *Probabilistic Forecasting and Bayesian Data Assimilation*. Cam-
669 bridge University Press, Cambridge.
- 670 Schneider, T., Teixeira, J., Bretherton, C. S., Brient, F., Pressel, K. G., Schär, C., and Siebesma,
671 A. P. (2017). Climate goals and computing the future of clouds. *Nature Climate Change*, 7(1):3.
- 672 Sokal, A. (1996). Monte Carlo methods in statistical mechanics: Foundations and new algorithms.
673 In C., D.-M., P., C., and A., F., editors, *Functional Integration*, volume 361. NATO ASI Series
674 (Series B: Physics).
- 675 Tarantola, A. (2005). *Inverse Problem Theory and Model Parameter Estimation*. SIAM, Philadel-
676 phia.
- 677 Wolff, U. (2004). Monte Carlo errors with less errors. *Computer Physics Communications*,
678 156(2):143–153.
- 679 Wood, R. (2012). Stratocumulus clouds. *Monthly Weather Review*, 140(8):2373–2423.
- 680 Wood, R. and Hartmann, D. L. (2006). Spatial variability of liquid water path in marine low cloud:
681 The importance of mesoscale cellular convection. *Journal of Climate*, 19(9):1748–1764.
- 682 Yamaguchi, T., Feingold, G., and Kazil, J. (2017). Stratocumulus to cumulus transition by drizzle.
683 *Journal of advances in modeling earth systems*, 9:2333–2349.