

iGen 0.1: the automated generation of a parameterisation of entrainment in marine stratocumulus

D. F. Tang¹ and S. Dobbie²

¹School of Mathematics, University of Manchester, Manchester, UK

²School of Earth and Environment, University of Leeds, Leeds, UK

Received: 1 April 2011 – Published in Geosci. Model Dev. Discuss.: 5 May 2011

Revised: 4 September 2011 – Accepted: 7 September 2011 – Published: 16 September 2011

Abstract. In a previous paper we described a new technique for automatically generating parameterisations using a program called iGen. iGen generates parameterisations by analysing the source code of a high resolution model that resolves the physics to be parameterised. In order to demonstrate that this technique scales up to deal with models of realistic complexity we have used iGen to generate a parameterisation of entrainment in marine stratocumulus. We describe how iGen was used to analyse the source code of an eddy resolving model (ERM) and generate a parameterisation of entrainment velocity in marine stratocumulus in terms of the large-scale state of the boundary layer. The parameterisation was tested against results from the DYCOMS-II intercomparison of ERM models and iGen's parameterisation of mean entrainment velocity was found to be $5.27 \times 10^{-3} \pm 0.62 \times 10^{-3} \text{ m s}^{-1}$ compared to $5.2 \times 10^{-3} \pm 0.8 \times 10^{-3} \text{ m s}^{-1}$ for the DYCOMS-II ensemble of large eddy simulation (LES) models.

1 Introduction

In Tang and Dobbie (2011) we described a technique for automatically generating parameterisations of physical processes by using a newly developed computer program called iGen. In this companion paper we apply this technique to the problem of parameterising cloud-top entrainment in a stratocumulus topped boundary layer (STBL).

The large-scale structure and dynamics of the STBL was described in the landmark paper of Lilly (1968), more recent developments are described in Stevens (2002). Typically, there is a well mixed boundary layer from the surface up to the cloud top within which, due to strong turbulent mixing,

the total water and liquid water potential temperature is close to homogeneous. The boundary layer is capped at cloud top by a strong, well defined inversion leading up into a much warmer, dryer, stable free atmosphere. The turbulence in the boundary layer is driven partly by surface fluxes of heat and moisture but predominantly by strong radiative cooling at cloud top and, to a lesser extent, by radiative warming at cloud base from the warmer, underlying surface. This turbulence causes some of the stable, free-atmosphere air to be mixed, or “entrained” into the turbulent boundary layer. Given the rate of this entrainment, the large-scale dynamics of the system is easily calculated from budgets of mass, energy and moisture. However, no analytic derivation of this entrainment rate has been found. Lilly (1968) derives upper and lower bounds and Stevens (2002) gives details of various parameterisations. However, the simulation of marine boundary layer cloud remains a large source of uncertainty and error in existing climate models. Bony and Dufresne (2005) have shown that disagreement between climate models about the behaviour of marine stratocumulus is a major source of uncertainty in the estimation of climate sensitivity. They have also shown that it is in the simulation of the radiative forcing due to marine stratocumulus that climate models are most in error compared to present day observations. A more recent study (Dufresne and Bony, 2008) shows that this situation has not improved in more recent years.

In order to generate a parameterisation of entrainment we followed the method described in Tang and Dobbie (2011) (hereafter TD) which, in this case, consists of the following steps:

1. take a high-resolution, ERM model that is capable of simulating the STBL;
2. decide on a set of variables that define the large-scale state of the STBL and define the required output of the parameterisation: in this case mean and standard deviation of entrainment velocity;



Correspondence to: D. F. Tang
(daniel.tang@manchester.ac.uk)

3. “wrap” the ERM model by adding extra code so that its input is the large-scale state of the STBL and its output is the mean and standard deviation of the entrainment velocity;
4. feed the source code of the wrapped ERM into iGen. iGen then analyses the source code, applies appropriate approximations and automatically generates the source code of a parameterisation.

2 An ERM to simulate stratocumulus

A 2 dimensional ERM was written in C++ in order to simulate entrainment in stratocumulus under nocturnal, non-precipitating conditions. A new ERM was written, rather than using an existing model, since iGen can at present only analyse C++ programs, while most existing models are written in Fortran. Writing a new model also gave us much more freedom to see how iGen performed with different schemes and algorithms. We chose to make the model 2 dimensional so that the simulations and analysis could be performed in reasonable time on a modest desktop computer, thus avoiding the effort and cost associated with using a supercomputer. Making the model 3 dimensional would increase execution time by 2 to 3 orders of magnitude; a similar figure to the increase in processing power available to a job on a supercomputer compared to a desktop. Thus, as a rule of thumb, what can be done on a desktop computer in 2 dimensions can be done on a supercomputer in 3 dimensions in around the same time. For this reason we consider using iGen to analyse a 2 dimensional model on a desktop computer to be a valid test of its capability.

2.1 Numerical implementation

The model was based on that of Klemp and Wilhelmson (1978) with modifications detailed in Skamrock and Klemp (1994). A number of changes were made to the Klemp and Wilhelmson model to better suit our needs. It was found that the vertical advection scheme caused “ringing” effects around the steep gradients at the inversion, leading to unrealistic cooling below cloudtop and heating above. To deal with this, a flux limiting advection scheme was used instead. This calculated advection as a mix between a fourth order, centred finite difference scheme and an upstream scheme. The flux limiting function used was

$$\phi(r) = \begin{cases} 0 & \text{if } r < 0 \\ 2r & \text{if } 0 \leq r \leq \frac{1}{2} \\ 1 & \text{otherwise.} \end{cases} \quad (1)$$

where r is the upwind gradient divided by the downwind gradient. Figure 1 compares the original and modified advection schemes for the 1 dimensional advection of a top-hat function in a uniform velocity field. The original scheme clearly

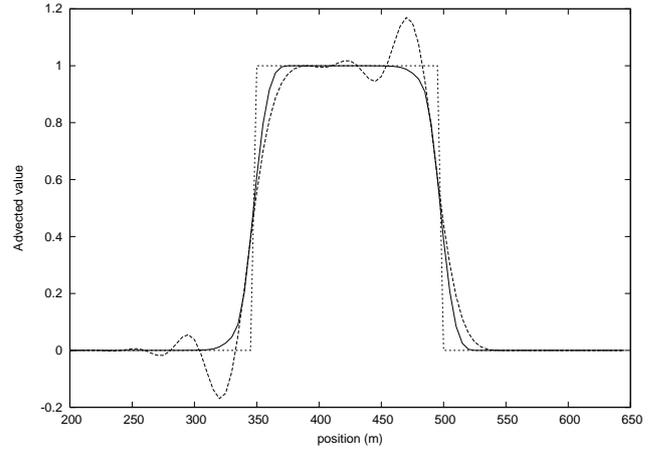


Fig. 1. Advection of a top hat function through 200 m using flux-limiting advection (solid line) and the Klemp-Wilhelmson scheme (long dashes). The short-dashed line shows the exact solution. Grid spacing was 5 m and velocity was 1 m s^{-1} .

shows the ringing effect at the sharp edges, which is removed in the flux-limited scheme.

Other changes are as follows:

- A more accurate version of Teten’s formula was used (Emanuel, 1994).
- The prognostic variable and prognostic equation for temperature was reformulated in terms of liquid water potential temperature.
- The prognostic variable and prognostic equation for moisture was reformulated in terms of total water mixing ratio, cloud being diagnosed when this exceeds saturation.
- In order to simulate longwave radiative heating/cooling, the radiation scheme described in Larson et al. (2007) was added.
- The prognostic variable and equation for rain was removed.
- Surface fluxes of heat and moisture were added.
- A homogeneous divergence was added in order to simulate large-scale subsidence.

The left and right boundaries were periodic in all variables. The upper and lower boundaries each lay on the vertical velocity points of the staggered grid. At the ground, horizontal and vertical velocity were constrained to zero. Other variables had the condition that $\frac{\partial}{\partial z}$ goes to zero at the ground. At the top of domain boundary, vertical velocity $v = -Dh$ where D is the large-scale divergence and h is the domain height; horizontal velocity u is the value of geostrophic wind; pressure perturbation from equilibrium is zero, liquid water

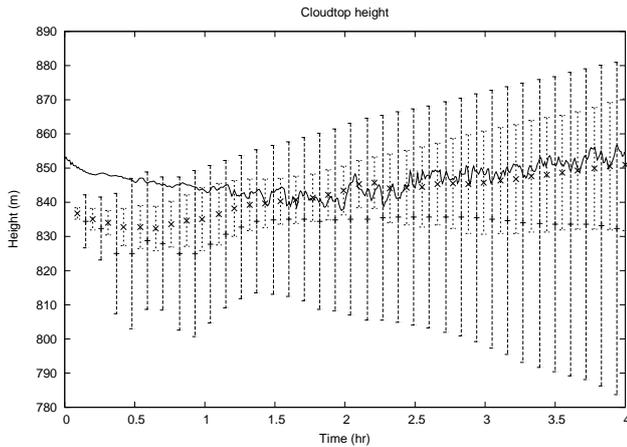


Fig. 2. Cloudtop height of DYCOMS-II simulation: the solid line shows results from our 2-D ERM. The inner error bars show the first and third quartiles of the ensemble of LEMs in the Stevens et al. (2005) intercomparison, the outer error bars show the maximum and minimum values of the ensemble. The mid-points of the error bars are marked by crosses and plus signs, respectively.

and temperature go to the large-scale free atmosphere values and turbulent kinetic energy has the boundary condition $\frac{\partial K_m}{\partial z} = 0$ in order to ensure that there is no sub-grid turbulent flux of turbulence across the boundary. For all experiments, the gridbox size at the inversion was 5 m vertically and 11 m horizontally. The surface fluxes of latent and sensible heat were calculated using a simple bulk aerodynamic formulation described in Krishnamurti and Bounoua (1995). Fluxes were added to the lowest gridbox of each column according to

$$\frac{\partial \theta}{\partial t} \Big|_{\text{surf}} = \frac{1}{\Delta z} \|u_{10}\| C_h (T_{\text{sst}} - T) \quad (2)$$

and

$$\frac{\partial q_t}{\partial t} \Big|_{\text{surf}} = \frac{1}{\Delta z} \|u_{10}\| C_q (q_{\text{sat}} - q_t) \quad (3)$$

where T and q_t are the temperature and total water of the lowest gridbox, respectively, Δz is the height of the lowest thermodynamic grid-point and

$$u_{10} = u_x \frac{\log\left(\frac{10.0}{z_0}\right)}{\log\left(\frac{\Delta z}{2z_0}\right)} \quad (4)$$

where u_x is the horizontal velocity at the lowest gridpoint, Δz is the height of the lowest gridpoint and z_0 is the roughness length, which was taken to have a constant value of 5×10^{-4} m based on figures in Stull (1988). The exchange coefficients were set constant at $C_h = 1.4 \times 10^{-3}$ and $C_q = 1.6 \times 10^{-3}$ based on figures in Krishnamurti and Bounoua (1995).

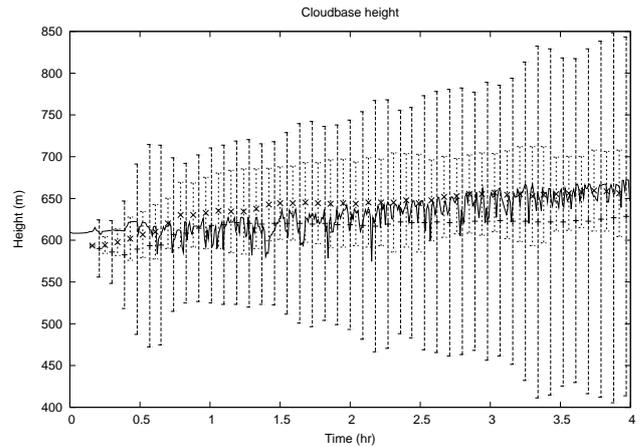


Fig. 3. Cloudbase height of the DYCOMS-II simulation: the solid line shows the results from the 2-D ERM. The inner error bars show the first and third quartiles of the ensemble of LEMs in the Stevens et al. (2005) intercomparison, the outer error bars shows the maximum and minimum values of the ensemble.

2.2 Testing the ERM

The ERM was compared against observations and other cloud resolving models by performing a simulation based on the first research flight of the DYCOMS-II study. This case was chosen as it has been used in an intercomparison study of LES models (Stevens et al., 2005) for which a detailed specification of an idealised simulation was given, and results were collected from an ensemble of models from ten different modelling centres. This allowed our model to be tested against a wide selection of commonly used models as well as against observations.

Our model showed a longer spin-up period than the models in the intercomparison (Fig. 2) and this was attributed to the 2 dimensional turbulence of the model, compared to the 3 dimensional turbulence of the models in the intercomparison since the cascade of turbulent kinetic energy and vorticity is known to be different in 2 and 3 dimensions (Kraichnan, 1967), this is discussed at greater length in Sect. 4.3. During this spin-up period, the low turbulent kinetic energy led to low entrainment and so the prescribed large-scale subsidence caused the cloudtop to descend. In order to account for this descent during the spin-up period, the initial cloudtop height was raised by 10 m, this had the effect of bringing the cloudtop height in-line with the other models at 2-h into the simulation when the spin-up period was over.

From 2-h into the simulation to the end of the simulation the model was in good agreement with both observation and the models of the intercomparison. Cloudtop height, and therefore entrainment, was very close to the ensemble average. Cloudbase height was also very close to the ensemble average (see Fig. 3).

Since our primary aim in this paper is to demonstrate that iGen can analyse a model of realistic complexity, it is not strictly necessary to demonstrate that our ERM resembles reality, just that its complexity is comparable to one that does. However, showing that our ERM is realistic is informative for two reasons. Firstly, it demonstrates that the ERM has a complexity comparable to that of a realistic model, since it is a realistic model itself. Secondly, it allows a dual interpretation of the parameterisation that iGen generates: on the one hand it is a demonstration of iGen's ability to analyse the ERM, while on the other hand it is a realistic parameterisation of the physical process of entrainment.

3 Defining the inputs and outputs of the parameterisation

Given a model that resolves entrainment, the next step in our method is to define the inputs and outputs we require of the parameterisation and decide on the ranges of the inputs over which the parameterisation should be valid. Since the parameterisation is intended to be a closure of the large-scale dynamics, the input should specify the large-scale state of the boundary layer as defined in Lilly (1968). We chose the input variables:

$q_{l,ct}$ Specific liquid water content at cloud top

Δq_t Jump in specific total water at cloud top

ΔB Jump in buoyancy at cloud top

F_0 Down-welling longwave radiation just above cloud top

F_1 Up-welling longwave radiation just below cloud base

The output we require of a parameterisation is the total entrainment in a 3 dimensional GCM gridbox over a single GCM timestep. Since the entrainment is caused by the action of a finite number of sub-grid, turbulent events, a high-resolution evolution of a gridbox cannot be predicted given only the large-scale state. So there is a certain amount of uncertainty associated with the mean entrainment. That this uncertainty is significant is shown in the experiment described in Sect. 4.1. In order to deal with this uncertainty we choose to define the required output to be the mean entrainment velocity \bar{E} and the standard deviation of the entrainment velocity σ_E averaged over a column of cross sectional area the size of a GCM gridbox and over the duration of a GCM timestep. Given these values, a deterministic parameterisation would just return \bar{E} while a stochastic parameterisation can also make use of the standard deviation. Numerical experiments on the ERM showed that the distribution of entrainment averaged over a typical GCM timestep and gridbox was Gaussian, so it is not necessary to calculate higher moments.

Using the wrapping technique described in TD, the outputs can be expressed in terms of a high-resolution simulation. Let $g(\Psi, s)$ be a function that turns a large-scale state Ψ into a high-resolution state ψ such that

$$P(g(\Psi, s) = \psi) = P(\psi | \Psi) \quad (5)$$

where $P(x)$ denotes the probability of event x and s is a random variable with uniform probability in the range $[0 : 1]$. That is, if we choose an s at random, the probability that $g(\Psi, s) = \psi$ should be the conditional probability of finding the system in a high-resolution state ψ given that it is in large-scale state Ψ . Since there exist many high-resolution states that conform to a given large-scale state, s can be thought of as defining the choice of high-resolution state from among the possibilities (if we think of g as a computer program, s can be thought of as a seed to a pseudo-random number generator). Now let $f(\psi_0, t)$ be the 3 dimensional high-resolution state of the system integrated over time t from a start state ψ_0 and let $E_{\Delta A}(\psi)$ be the instantaneous entrainment velocity of the high-resolution state ψ averaged over a column of area ΔA . The required outputs can now be expressed as

$$\bar{E}(\Psi) = \int_0^1 F(\Psi, s) ds \quad (6)$$

and

$$\sigma_E(\Psi) = \sqrt{\int_0^1 F(\Psi, s)^2 ds - \bar{E}(\Psi)^2} \quad (7)$$

where

$$F(\Psi, s) = \frac{\int_0^{\Delta T} E_{\Delta A}(f(g(\Psi, s), t)) dt}{\Delta T} \quad (8)$$

One complication in this definition is that F is dependent on ΔA and ΔT , the horizontal cross-sectional area of a GCM gridbox and the duration of a GCM timestep. After integrating over s this dependency will disappear for \bar{E} but not for σ_E . Since we would like a parameterisation that can be used in GCMs of any gridbox size and timestep, we now derive a scaling rule for σ_E .

We suppose that the entrainment consists of a number of random, independent entrainment events. As the GCM timestep and gridbox size changes, so does the number of entrainment events we expect to be averaging over in one gridbox and timestep. The standard deviation of N independent events scales as $N^{-\frac{1}{2}}$, so if σ_E is the standard deviation of N events and σ_G is the standard deviation of M events from the same process, the relationship between the standard deviations would be

$$\sigma_G = \sigma_E \sqrt{\frac{N}{M}} \quad (9)$$

Since the number of entrainment events in a column over a GCM timestep scales as $\Delta A \Delta T$, when averaging over a

column of cross sectional area ΔA_G and timestep ΔT_G the standard deviation will be

$$\sigma_G = \sigma_E \sqrt{\frac{\Delta A \Delta T}{\Delta A_G \Delta T_G}} \quad (10)$$

So, given σ_E for some column of standard ΔA and ΔT , this scaling rule can be used to find the standard deviation for any gridbox size and timestep.

The ranges of the large-scale variables over which the parameterisation should be valid were calculated from the results of a number of field campaigns and idealised cases of nocturnal marine stratocumulus as shown in Table 1. Based on these results, the ranges used for iGen's analysis of the wrapped ERM were:

- $1 \times 10^{-4} \leq q_{l,ct} \leq 1 \times 10^{-3} \text{ Kg Kg}^{-1}$
- $-8.0 \times 10^{-3} \leq \Delta q_t \leq -2.0 \times 10^{-3} \text{ Kg Kg}^{-1}$
- $0.065 \leq \Delta B \leq 0.5 \text{ m s}^2$
- $20 \leq F_0 \leq 110 \text{ W m}^2$
- $7 \leq F_1 \leq 33 \text{ W m}^2$.

When generating the parameterisation, iGen uses these ranges to define the domain over which the approximation errors should be kept small. Outside of this range there is no guarantee how accurate the resulting parameterisation will be.

4 Wrapping the ERM

The next stage in our method is to “wrap” the ERM so that its inputs and outputs are those of the parameterisation defined in the previous stage. That is, to devise an effective procedure to calculate Eqs. (6), (7) and (8).

The function $g(\Psi, s)$ in Eq. (8) can be calculated using the ERM by deterministically starting in a canonical, high-resolution state that conforms to Ψ , then returning the state after a spin-up of duration $ks + t_m$, where t_m is some minimum spin-up and k is large enough to ensure that Eq. (5) is satisfied. The canonical state should be some stable, low-energy state from which the spin-up can quickly relax to the system's attractor. During the spin-up the entrainment will give rise to a flux of mass, heat and moisture into the boundary layer, and this will tend to make the large-scale state drift away from the prescribed state. To counteract this, an external forcing should be added which maintains the large-scale state. If we let $f_c(\psi_0, t)$ be the state of the system integrated over time t from a start state ψ_0 while holding the large-scale state constant, and let $\gamma(\Psi)$ be a function that returns the canonical high-resolution state given Ψ then

$$g(\Psi, s) = f_c(\gamma(\Psi), ks + t_m) \quad (11)$$

Once a start-state has been calculated using g , F can be calculated by performing a high-resolution integration over the duration of a GCM timestep and calculating the average entrainment. If we note that the large-scale state changes slowly compared to the duration of a GCM timestep then we can maintain the flux applied during the spin-up period in g and expect that the result will not be significantly affected. By doing this, we can calculate both the spinup and integration parts with f_c , giving

$$F(\Psi, s) \approx \frac{1}{\Delta T} \int_0^{\Delta T} E_{\Delta A}(f_c(\gamma(\Psi), t_m + ks), t) dt \quad (12)$$

$$= \frac{1}{\Delta T} \int_0^{\Delta T} E_{\Delta A}(f_c(\gamma(\Psi), t_m + ks + t)) dt$$

So, $F(\Psi, s)$ can be approximated by starting in the canonical state of Ψ , spinning up for a duration of $t_m + ks$, then simulating for ΔT and measuring the average entrainment, keeping the large-scale state constant all the time. Substituting this into Eq. (6) gives, on the assumption that ΔT is much less than k

$$\bar{E}(\Psi) \approx \int_0^1 \frac{1}{\Delta T} \int_0^{\Delta T} E_{\Delta A}(f_c(\gamma(\Psi), t_m + ks + t)) dt ds \quad (13)$$

$$\approx \frac{1}{k} \int_0^k E_{\Delta A}(f_c(\gamma(\Psi), t_m + t)) dt$$

which means that the integral over s can be replaced by a time average from t_m to $t_m + k$. Making the same substitution for the standard deviation gives

$$\sigma_E(\Psi)^2 \approx \int_0^1 \left(\frac{1}{\Delta T} \int_0^{\Delta T} E_{\Delta A}(f_c(\gamma(\Psi), t_m + ks + t)) dt \right)^2 ds - \bar{E}(\Psi)^2 \quad (14)$$

$$= \frac{1}{k} \int_0^k \left(\frac{1}{\Delta T} \int_0^{\Delta T} E_{\Delta A}(f_c(\gamma(\Psi), t_m + t + t')) dt' \right)^2 dt - \bar{E}(\Psi)^2$$

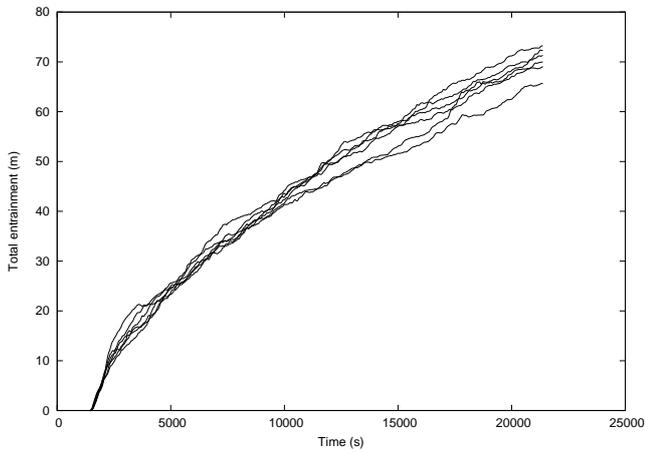
$$\approx \left[\frac{k}{\Delta T} \right]^{-1} \sum_{n=0}^{\lfloor \frac{k}{\Delta T} \rfloor} \left(\frac{1}{\Delta T} \int_0^{\Delta T} E_{\Delta A}(f_c(\gamma(\Psi), t_m + n\Delta T + t')) dt' \right)^2 - \bar{E}(\Psi)^2 \quad (15)$$

so that the sum over s can be replaced by the square of contiguous averages of length ΔT between times t_m and $t_m + k$.

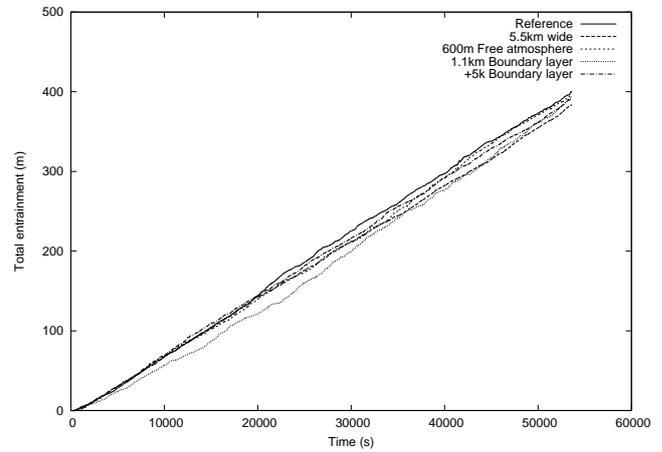
To implement this, we chose t_m to be 2 h, based on the spinup periods seen when testing the ERM (Fig. 2), and when performing sensitivity tests (Figs. 4 and 5). k was chosen to be 4 h, this was chosen as a trade off between reducing the standard error in \bar{E} and reducing the computational cost of the analysis. When used in a parameterisation for a GCM timestep of 30 min (a typical value), setting k to 4 h means that the standard error in \bar{E} is $\frac{1}{\sqrt{8}}$ times the standard deviation σ_G . That is, the standard error due to our choosing $k = 4$ is around a third of the inherent uncertainty due to the GCM's finite resolution. ΔT was chosen to be 216 s, this too was a trade off. On the one hand we would like to collect a large number of samples over the 4 h simulation, while on the other, the duration should be long enough to ensure that contiguous samples can be modelled as independent, random

Table 1. The ranges of large-scale boundary layer values found in various published sources.

Reference	$\theta_{l,bl}$ (K)	$\Delta\theta_l$ (K)	$q_{t,bl}$ (kg kg^{-1})	Δq_t (kg kg^{-1})	H (m)	F_0 (W m^{-2})	$q_{l,ct}$ ($\times 10^{-4} \text{ kg kg}^{-1}$)
Stevens et al. (2005)	289.0	8.5	9.0	-7.5	840	70	4.75
Moeng et al. (1996)	288.0	5.5	8.0	-4.6	662.5 : 690	15 : 30	4.4 : 6.6
Albrecht et al. (1988)	289	5 : 11	7.0	-6.0	500 : 1100	40	—
Bretherton and Pincus (1995)	292.0	2 : 4	10.0	-5 : -8	435 : 1358	—	2.2 : 5.2
Bretherton et al. (1995)	290 : 294.5					—	
Bretherton and Pincus (1995)	291 : 292	11	9.0	-3 : -8	1800	—	0 : 10.0
Bretherton et al. (1995)	291 : 293.5						
Klein and Hartmann (1993)						100	

**Fig. 4.** Total entrainment against simulated time for six executions of the wrapped ERM differing only in a 0.0025 K perturbation to the high-resolution start state.

events. The ERM was used to calculate f_c on a domain of size 770 m high by 1166 m wide, with the inversion held at 600 m. Experiments showed that entrainment was not affected by increases in the size of the domain (see Sect. 4.2) or height of the inversion. We defined γ , (the canonical, high-resolution state of the atmosphere for a given large-scale state) to consist of a homogeneous boundary layer and homogeneous free atmosphere separated by a linear transition of 25 m height. Initial velocities were zero everywhere and there was no sub-grid turbulent kinetic energy. Pressure was initialised to the hydrostatic value with sea level pressure set to $1 \times 10^5 \text{ N m}^{-2}$. In order to break symmetry, a random perturbation of $\pm 0.0025 \text{ K}$ was added to each gridbox below 100 m and within 100 m below the inversion. Geostrophic winds were not included for the same reason as given in Moeng et al. (1996): if we are to include geostrophic winds, this raises the question of the orientation of the 2-D domain in relation to the wind direction. Since roll motions tend to be aligned closely to the wind direction, the natural choice would be perpendicular to the wind direction, meaning no geostrophic wind across the domain.

**Fig. 5.** Total entrainment against simulated time for executions of the wrapped ERM with different geometries and different boundary layer temperature.

Given the 5 large-scale variables, it was found that the dependency of entrainment on boundary layer temperature was very weak over the range of values we expect to experience (see Sect. 4.2). In light of this insensitivity it was decided to set average boundary layer liquid water potential temperature to 290 K, the centre of its range. Sea surface temperature was held fixed at 291 K.

In order to calculate f_c , the large-scale state must be held constant in order to ensure that it does not drift away from the values supplied as the input of the parameterisation. The effect of allowing the large-scale state to drift can be seen in Fig. 4. This shows the total entrainment plotted against time, the gradient of this line at any point gives the instantaneous entrainment rate. As the simulation proceeds, the large scale state drifts and the entrainment rate slows. Compare this with Fig. 5 where the large scale state is held constant. Here, the total entrainment has the form of a line of constant gradient, with random noise added. The gradient of a least squares fit of this line would give a reasonable value for the average entrainment rate over the entire simulation.

In order to hold the large scale state constant, a set of fluxes were calculated every 12 simulated seconds and added at each timestep. In order to keep boundary layer temperature and moisture constant a total water flux and temperature flux was added. Total water flux was added to the sub-cloud portion of the boundary layer. This included a flux that tended to homogenise the field and was calculated at each gridpoint as

$$\frac{\partial q_t}{\partial t} = m_q + \frac{q_{tbl} - q_t}{8\Delta t} \quad (16)$$

where m_q is the gradient of the least squares linear fit of the total flux from the beginning of the simulation, q_{tbl} is the large-scale total water in the boundary layer and q_t is the field of actual total water. The homogenisation is not physical but is justified on the grounds that we want to find a formula for entrainment in order to close the large-scale dynamics of the boundary layer. However, the large-scale dynamics is only valid under the assumption of a homogeneous boundary layer so we are merely enforcing the assumption made by the large-scale dynamical view.

The flux of liquid water potential temperature was calculated so as to add a constant buoyancy to the whole boundary layer from the ground up to the isoline of temperature half way between the large-scale boundary layer and free atmosphere values. In this way, the dynamics of the boundary layer is not affected by the flux. The calculation was performed by first calculating a homogeneous buoyancy flux

$$\frac{\partial B}{\partial t} = m_b + \frac{\theta_{tbl} - \bar{\theta}_l}{30\theta_{tbl}\Delta t} \quad (17)$$

where $\bar{\theta}_l$ is the average liquid water temperature between 200 m and 100 m below cloud top and m_b is the gradient of the least squares linear fit of the total flux of buoyancy since the beginning of the simulation.

The flux of liquid water potential temperature necessary to achieve a given change in buoyancy ΔB over a single timestep, given a change in total water Δq_t , was calculated and added at the end of each timestep. The change in liquid water potential temperature $\Delta\theta_l$ at each gridbox was calculated using the following procedure: in the absence of liquid water

$$\Delta\theta_{l,dry} = \theta_{tbl} (\Delta B - 0.61\Delta q_t) \Delta t \quad (18)$$

in the presence of liquid water

$$\Delta\theta_{l,wet} = \theta_{tbl} \left(\Delta B - \Delta q_t \left(\frac{\gamma}{\theta_{tbl}} - 1 \right) \left(1 - \frac{\partial q_{sat}}{\partial q_t} \right) \right) \quad (19)$$

and $\frac{\partial q_{sat}}{\partial q_t}$ is the rate of change of saturation with q_t at constant θ_l . In the case that the flux causes a transition between clear sky and cloud, it is necessary to calculate the fraction of buoyancy and q_t change that occurs in cloud and the fraction in clear sky and to add these contributions separately. When

going from clear sky to cloudy, the fraction in clear sky is given by

$$m = \frac{q_{sat} - q_t}{\Delta q_t - \frac{\partial q_{sat}}{\partial \theta_l} \Delta\theta_{l,dry}} \quad (20)$$

where $\frac{\partial q_{sat}}{\partial \theta_l}$ is the rate of change of saturation with θ_l at constant q_t . When going from cloudy to clear, the fraction in cloudy sky is

$$m = \frac{q_t - q_{sat}}{\left(1.0 - \frac{\partial q_{sat}}{\partial q_t} \right) \Delta q_t - \frac{\partial q_{sat}}{\partial \theta_l} \Delta\theta_{l,wet}} \quad (21)$$

In addition to maintaining the temperature and moisture of the boundary layer, the height of the inversion was kept constant by adding a homogeneous, large-scale divergence. This ensured that the cloud top did not come close to the upper boundary of the domain. The divergence was calculated according to:

$$\nabla \cdot v = \frac{m_d + \frac{h-H}{5\Delta t}}{H} \quad (22)$$

where m_d is the gradient of the least squares linear fit to the total entrainment over the duration of the simulation so far, h is the measured height of the boundary layer, Δt is the time between updates (12 s) and H is the required height. The height of the inversion was defined to be the average height of the isoline of total water content half way between the large-scale boundary layer and free atmosphere values.

Adjusting the large scale divergence in this way will lead to a large scale downward velocity of air through the inversion at each timestep. In the wrapped model, this velocity is integrated over time as the simulation proceeds and the running total is recorded every 12 s. So, by the end of the simulation there is a record of total entrainment, in meters, from the beginning to the end of the simulation. The mean and standard deviation of the entrainment velocity was calculated from the final 4 hours of this record (after the 2 h spin-up). The mean was calculated from the gradient of a least-squares fit of the total entrainment record and the standard deviation was calculated from the average entrainment velocities of contiguous 216 s intervals.

4.1 Sensitivity of the wrapped ERM to initial conditions

Since knowledge of the large-scale state of the system gives us incomplete knowledge of the high-resolution state, there is an inherent uncertainty in the mean entrainment when averaged over a finite area and time. In order to show that this uncertainty is significant, a numerical experiment was performed on the wrapped ERM to test the sensitivity of entrainment rate to an initial random perturbation of ± 0.0025 K to each gridbox in the lowest 100 m of the boundary layer and within 100 m below the inversion. The large-scale state was chosen to be around the centre of the expected ranges of each value:

Net radiation flux above inversion

$$\text{Boundary layer } q_t = 8 \times 10^{-3} \text{ Kg Kg}^{-1}$$

$$\Delta\theta_1 \text{ at inversion} = 8.5 \text{ K}$$

$$\Delta q_t \text{ at inversion} = -6 \times 10^{-3} \text{ Kg Kg}^{-1}$$

$$\text{Net radiation flux above inversion} = -55 \text{ W m}^{-2}$$

$$\text{Net radiation flux at cloud base} = 22 \text{ W m}^{-2}$$

The domain size was 1166 m horizontally and 770 m vertically. The inversion height was 600 m above the bottom of the domain.

Six simulations were made with different random perturbations. The fluxes of heat and moisture which keep the boundary layer at a constant, large-scale state were turned off in order to discount them as amplifiers of sensitivity.

The total entrainment of each simulation was output at 12 s intervals to show the time evolution of entrainment as the simulation progressed. The results are shown in Fig. 4. After 6 h there was a 10 % spread in total entrainment between simulations, showing that there is significant sensitive dependence on initial conditions under these conditions.

Simulations were also made with the large-scale divergence feedback turned off in order to discount this as a possible amplifier of sensitivity. Results still showed sensitivity to initial conditions. Different domain geometries did not show any reduction in sensitivity after scaling according to Eq. (23). Sensitivity was reduced to around 5 % under the same large-scale state when the large-scale boundary layer state was held constant by turning on the fluxes of heat and moisture.

4.2 Sensitivity of the wrapped ERM to domain geometry

Numerical experiments were performed to find the sensitivity of the wrapped ERM output to the domain geometry of the ERM. The reference geometry was 770 m vertical by 1166 m horizontal, with the inversion at 600 m. The following perturbations to the reference geometry were tested:

- domain width of 5500 m,
- domain height of 1200 m,
- inversion at 1100 m with domain height of 1270 m.

In all cases, the values of the large-scale inputs to the model were chosen to be as follows:

$$q_{l,ct} = 5.5 \times 10^{-4} \text{ Kg Kg}^{-1}$$

$$\Delta q_t = -6.0 \times 10^{-3} \text{ Kg Kg}^{-1}$$

$$\Delta B = 0.215 \text{ m s}^{-2}$$

$$F_0 = 55 \text{ W m}^{-2}$$

Table 2. The least squares fit of the rate of entrainment for different domain geometries.

Simulation	Entrainment (m s^{-1})
Reference	7.45×10^{-3}
Wide	7.28×10^{-3}
Free atmosphere	7.59×10^{-3}
1100 m boundary layer	7.31×10^{-3}
$\theta_{l,bl} = 295 \text{ K}$	7.50×10^{-3}

$$F_1 = 22 \text{ W m}^{-2}$$

$$\theta_{l,bl} = 290 \text{ K}$$

$$\Delta T_{\text{sst}} = 1 \text{ K}.$$

In addition, sensitivity to boundary layer liquid water potential temperature (with all other variables fixed) was tested by performing a simulation at 295 K, the upper limit of the expected range.

The simulations lasted 15 simulated hours and the initial spin-up period was 9 h. The total entrainment of the simulations were recorded at 12 s intervals throughout the duration of the simulations, these are shown in Fig. 5 as a function of time. The mean entrainment velocity was calculated from the gradient of a least-squares linear fit to the total entrainment values after discounting the values from the spin-up period. The results, displayed in Table 2, show that the reference geometry, although small, gives values for entrainment that agree well with those of different geometries, when the the intrinsic standard deviation of entrainment is taken into account.

4.3 Accounting for 2-D simulation

Since our ERM is 2 dimensional and the parameterisation is intended for a 3 dimensional GCM, some discussion of the implications of this is necessary. The results shown in section 2.2 show that the entrainment velocity of our 2 dimensional ERM is comparable to that of 3 dimensional models, despite the difference in the cascade of turbulent kinetic energy in the 2 dimensional case compared to that of 3 dimensions (Kraichnan, 1967). This is backed up by similar results in Moeng et al. (1996). Grabowski et. al. (1998) also performed comparisons of 2 and 3 dimensional models of the turbulent atmosphere and found them to give comparable results in many respects. Exactly why this is so in the case of entrainment is not entirely clear. One possible explanation is that, even at the 5m resolution used in the wrapped ERM, the entrainment predominantly occurs at the sub-grid scale by way of the sub-grid turbulence parameterisation. Since the turbulence parameterisation is largely unchanged between our 2 dimensional ERM and a 3 dimensional model we would, in this case, expect similar results.

The extension from 2 dimensional to 3 dimensional standard deviations of entrainment also needs to be addressed. Equation (10) cannot be used to scale the standard deviation measured in a 2 dimensional ERM to that of a 3 dimensional gridbox because the ERM has no horizontal cross sectional area, ΔA . To deal with this we suppose that the entrainment in the ERM is caused by the action of a number of entrainment events of characteristic length scale l . In this case, the number of entrainment events over the simulated domain of the ERM would scale as $\frac{x\Delta T}{l}$, where x is the width of the domain and ΔT is the duration of the ERM simulation. Similarly, the number of events in a GCM column of cross-sectional area ΔA_G will scale as $\frac{\Delta A_G \Delta T_G}{l^2}$, where ΔT_G is the GCM timestep. So, we would expect the standard deviation of the ERM, σ_E , to scale as

$$\sigma_G = \sigma_E \sqrt{\frac{lx\Delta T}{\Delta A_G \Delta T_G}} \quad (23)$$

where σ_G is the standard deviation over a GCM gridbox.

An order of magnitude figure for the characteristic length scale can be calculated by noting that since the entrainment predominantly occurs in only one direction, each entrainment event must predominantly mix free atmosphere air into the boundary layer. For this to be the case, the mean entrainment must remain larger than the standard deviation, so the scale of the process of entrainment must be no smaller than that where the standard deviation of each event equals its mean. Taking data from a 15 h simulation with the same large-scale state as the previous section and supposing a characteristic velocity of 1 m s^{-1} gives a characteristic length scale of 15 m.

5 Analysing the wrapped ERM using iGen

The wrapped ERM now has the inputs and outputs of a parameterisation of entrainment and, indeed, could be used as a very slow and computationally inefficient parameterisation. The next stage in the method is to use iGen to analyse the source code of the wrapped ERM in order to generate a more efficient parameterisation.

The wrapped ERM was analysed by iGen on a desktop computer with 1.8 GHz Intel Core-Duo. For this analysis iGen was configured to use approximation by Lagrange interpolation when performing polynomial arithmetic. The analysis proceeded according to the mechanism described in TD: the inputs to the wrapped ERM (i.e. the 5 variables of the large-scale state) were interpreted as unknown, independent variables and each statement of the wrapped ERM's source code was interpreted as an operator on multivariate polynomials of these independent variables. In addition to the programming structures described in TD, the wrapped ERM contained a number of other structures, but these did not require any significant changes to the analysis technique. C++ objects were interpreted as objects whose members were

polynomials, in an obvious extension to way that variables were interpreted as polynomials. Calls to methods and sub-routines were interpreted as if expanded in-line, avoiding the need for any special treatment (recursive functions were not encountered in the wrapped ERM, but these could have been dealt with by, for example, recursively expanding in-line until the recursion terminates for all possible inputs within the specified range). Calls to the standard library functions `exp` and `pow` were required during evaluation of Teten's formula and calculation of the hydrostatic pressure respectively. Rather than analyse the C++ `math` library, these calls were overloaded to directly calculate the exponential and power of a polynomial by applying the respective function to each point in the polynomial's sparse grid representation. The wrapped ERM contained no instances of arrays being indexed by values that depended on the values of the inputs, so arrays could be represented simply as arrays of polynomials.

The majority of the wrapped ERM code involved only simple arithmetic operations. The wrapped ERM begins by calculating the canonical high-resolution state, this involves initialising the arrays that store the gridded prognostic variables. An analysis of this portion of the program would calculate polynomials for each member of these arrays. A timestep of the ERM consisted of finite difference calculations on the grids, and this was analysed as simple arithmetic operations on the polynomials at each grid-point. Calculation of pressure and velocity again consisted of a loop of finite difference calculations analysed as simple arithmetic operations on polynomials. Diagnosis of liquid water and the flux-limiting advection each involved an `if` statement, which was analysed using the Heaviside step function as described in TD.

The polynomials were represented as points on a sparse grid. This grid has the property that the points representing a polynomial of degree n form a subset of the points of a polynomial of degree $2n$. As a consequence of this, the results of a low degree analysis can be re-used to do a higher degree analysis allowing a simple form of adaptive grid refinement to be implemented (Gerstner and Griebel, 2003). So, iGen incrementally increased the degree of the analysis until the error between the parameterisation and the wrapped ERM was within the required limit.

The error-bounding mechanism used by iGen in this experiment was slightly different from that described in TD and made use of the information available from the analysis of the standard deviation σ_E . Conceptually, the output of the wrapped ERM can be viewed as a noisy signal $s + n$ where s is the value that the wrapped ERM would return in the limit that the length of the simulation, k , tends to infinity, and n is the noise due to the finite length of the simulation combined with the model's sensitivity to initial conditions. n can be modelled as a random perturbation with Gaussian distribution and standard deviation of $\sigma = \sigma_E \sqrt{\frac{\Delta T}{k}}$, where k is the duration of the wrapped ERM's simulation and $\Delta T = 216 \text{ s}$.

As the degree of the analysis is increased, if a point is reached where any further increase in the degree causes a perturbation that is not significantly larger than that expected from the noise alone, then the analysis should terminate as any higher degree of analysis would only resolve more noise.

The perturbation was calculated by converting the resulting entrainment polynomial to Chebyshev form and forming a “high-order polynomial” consisting of all the highest order terms (i.e. those for which all other terms have at least one variable of lower degree). The amplitude of the high-order polynomial was sampled at 10 000 randomly chosen points in the domain and the proportion of points that lay within 0.674 standard deviations of the mean was calculated. If this exceeded 48 % the analysis was terminated.

At this point, the error in the signal, s , due to the finite degree of analysis can be shown to be small by noting that the approximant of the signal converges as the degree of approximation is increased as long as there are no discontinuities and the total variation remains bounded¹ (Mastroianni and Szabados, 1995). This is ensured by the supposition that the problem is well conditioned. While there may be discontinuities in the *derivatives* of the signal caused by the onset of different regimes or different physical mechanisms of entrainment, we would not expect to see any discontinuities or unbounded variation in the signal itself.

6 Results

iGen was left running for 28 days and on return the analysis had terminated, returning a 10th total-degree polynomial for both mean entrainment and standard deviation. The resulting mean entrainment polynomial was confirmed to have converged by forming the “high-order polynomial” (see previous section) and evaluating it at 10 000 randomly chosen points in the input domain. The proportion of points at which the high-order polynomial was found to lie within 0.674 standard deviations was found to be 49.85 %.

The polynomials for mean and standard deviation that resulted from the analysis are given in the supplementary files. The polynomials are expanded in the Chebyshev basis, so x^n should be interpreted as the n -th Chebyshev polynomial (defined as $\cos(n\cos^{-1}(x))$). The variables of the polynomials have been offset and normalised according to $x' = \frac{2x - (x_h + x_l)}{x_h - x_l}$ where x lies in the range $[x_l : x_h]$ so that the ranges given in Sect. 3 are transformed to lie in the range $[-1 : 1]$. These can easily be converted into a program that evaluates the mean and standard deviation at any point in just over 2000 multiplications and additions by using Horner form evaluation. Approximations that require fewer operations could be created by Chebyshev truncation, by finding the minimax polynomial fit using Remez’ algorithm (Press

et al., 2007) or by finding the least squares fit by solving the appropriate set of linear equations (Press et al., 2007).

The polynomials were tested against the ensemble of cloud resolving models used in the DYCOMS-II intercomparison (Stevens et al., 2005). The ensemble-average large-scale state for the final hour of the simulations was used as input to the polynomial, and the entrainment over 1 h was predicted to be $5.27 \times 10^{-3} \pm 0.62 \times 10^{-3} \text{ m s}^{-1}$. This compares very well with the ensemble average of the LES’s entrainment rate which was $5.2 \times 10^{-3} \pm 0.8 \times 10^{-3} \text{ m s}^{-1}$.

7 Discussion

iGen’s analysis of the wrapped ERM took around 28 days to complete. This can be compared to the wrapped ERM’s execution time of around 30 min to show that, in this case, the analysis was around 1000 times slower than an execution. As mentioned earlier, this still means it would be possible to analyse a full, 3 dimensional model on a supercomputer. There remain many ways of increasing the speed of the analysis. It may be possible to make significant improvements in speed by incorporating a deeper analysis of loops, by having a more sophisticated adaptive refinement mechanism for polynomials and by performing automatic differentiation to apply greater approximations at points where these least affect the result.

Because iGen represents variables as polynomials during an analysis, the memory requirements of an analysis can be considerably larger than those of an execution. In this case the analysis required around 1000 times more memory than an execution. The memory requirements will increase with the order of the analysis, the number of input variables and the number of variables used in the wrapped model. This may become a problem when attempting to perform a very high order analysis of a very high resolution wrapped model. However, it would be possible to change the order in which iGen makes its calculations to allow intermediate results to be deleted once they are no longer needed. This would make the memory requirements independent of the order of the analysis and of the order of a few tens of times the requirements of an execution. This is left for future work.

The error-bounding method used in this experiment required the assumption of bounded total variation and no discontinuities in the wrapped ERM’s output. This was plausible in this case but in order to make the analysis fully formal it would be desirable for iGen to use its analysis to prove these properties, and to identify discontinuities if they occur. iGen is currently being developed to do this.

8 Conclusions

iGen has analysed the source code of a wrapped, high-resolution eddy resolving model of entrainment in marine stratocumulus and from this has derived a parameterisation

¹Total variation is defined as the integral of the absolute value of the derivative.

of entrainment in terms of the large-scale state of the boundary layer.

The resulting parameterisation can be interpreted in two ways. Primarily it is a demonstration of iGen's ability to create parameterisations from models of realistic complexity. In addition, we have demonstrated that the ERM and the resulting parameterisation shows good agreement with an ensemble of LES models and could be incorporated into the boundary layer parameterisation scheme of a climate model. The biggest limitation of the parameterisation is that it is based on a 2 dimensional simulation. The similarity in results between our model and the 3 dimensional models in the DYCOMS-II case, however, would suggest that this may not adversely affect entrainment rates. This is in line with Moeng et al. (1996) who also found a similar insensitivity of entrainment rate to model dimensionality. This insensitivity may be a result of the finite resolution of the model, it is not clear whether the 5 m resolution of our ERM is enough to capture the processes involved in entrainment. It would be worthwhile repeating this experiment with a higher grid resolution and in 3 dimensions. It would also be worthwhile treating boundary layer temperature and sea surface temperature as input variables in order to formally show their functional role in entrainment.

iGen continues to be developed and improved but we have demonstrated that it can already generate a parameterisation of a physical process that hasn't previously been satisfactorily parameterised. We hope that in time iGen will become a widely used tool for generating parameterisations and models for numerical experiments.

Acknowledgements. We would like to thank Zen Internet Ltd for the loan of a computer and the Natural Environment Research Council for funding this research under award number NER/S/A/2006/14148. Thanks to J. Chrimes for supplying figures from the DYCOMS-II intercomparison study and for proof reading a manuscript of this paper. Thanks also to Wayne and Jane at North Tea Power for much needed coffee.

Edited by: A. Lauer

References

- Albrecht, B. A., Randall, D. A., and Nicholls, S.: Observations of marine stratocumulus clouds during FIRE, *B. Am. Meteorol. Soc.*, 69, 618–626, 1988.
- Bony, S. and Dufresne, J.: Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models, *Geophys. Res. Lett.*, 32, L20806, doi:10.1029/2005GL023851, 2005.
- Bretherton, C. S. and Pincus, R.: Cloudiness and marine boundary layer dynamics of the ASTEX Lagrangian experiments, Part I: Synoptic setting and vertical structure, *J. Atmos. Sci.*, 52, 2707–2723, 1995.
- Bretherton, C. S., Austin, P. and Siems, S. T.: Cloudiness and marine boundary layer dynamics in the ASTEX Lagrangian experiments. Part II: Cloudiness, drizzle, surface fluxes and entrainment, *J. Atmos. Sci.*, 52, 2724–2735, 1995.
- Dufresne, J. L. and Bony, S.: An assessment of the primary sources of spread of global warming estimates from coupled atmosphere-ocean models, *J. Climate*, 21, 5135–5144, 2008.
- Emanuel, K. A.: *Atmospheric Convection*, Oxford University Press, Oxford, 1994.
- Gerstner, T. and Griebel, M.: Dimension-adaptive tensor-product quadrature, *Computing*, 71, 65–87, 2003.
- Grabowski, W. W., Wu, X., Moncrieff, M. W., and Hall W. D.: Cloud-resolving modeling of cloud systems during phase III of GATE, Part II: Effects of resolution and the third spatial dimension, *J. Atmos. Sci.*, 55, 3264–3282, 1998.
- Klein, S. A. and Hartmann, D. L.: The seasonal cycle of low stratiform clouds, *J. Climate*, 6, 1587–1606, 1993.
- Klemp, J. B. and Wilhelmson, R. B.: The simulation of three-dimensional convective storm dynamics, *J. Atmos. Sci.*, 35, 1070–1096, 1978.
- Kraichnan, R. H.: Inertial ranges in two-dimensional turbulence, *Phys. Fluids*, 10, 1417–1423, 1967.
- Krishnamourti, T. N. and Bounouna, L.: *An Introduction to Numerical Weather Prediction Techniques*, CRC Press, Florida, 1995.
- Larson, V. E., Kotenberg, K. E., and Wood, N. B.: An analytic long-wave radiation formula for liquid layer clouds, *Mon. Weather Rev.*, 135, 689–699, 2007.
- Lilly, D. K.: Models of cloud-topped mixed layers under a strong inversion, *Q. J. Roy. Meteorol. Soc.*, 94, 292–309, 1968.
- Mastroianni, G. and Szabados, J.: Jackson order of approximation by Lagrange interpolation II, *Acta Math. Hungar.*, 69, 73–82, 1995.
- Moeng, C. H., Cotton, W. R., Bretherton, C., Chlond, A., Khairoutdinov, M., Krueger, S., Lewellen, W. S., Macvean, M. K., Pasquier, J. R. M., Rand, H. A., Siebesma, A. P., Stevens, B., and Sykes, R. I.: Simulation of a stratocumulus-topped planetary boundary layer: intercomparison among different numerical codes, *B. Am. Meteorol. Soc.*, 77, 261–278, 1996.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P.: *Numerical Recipes: the Art of Scientific Computing*, 3rd Edn., Cambridge University Press, Cambridge, UK, 2007.
- Skamrock, W. C. and Klemp, J. B.: Efficiency and accuracy of the Klemp-Wilhelmson time-splitting technique, *Mon. Weather Rev.*, 122, 2623–2630, 1994.
- Stevens, B.: Entrainment in stratocumulus-topped mixed layers, *Q. J. Roy. Meteorol. Soc.*, 128, 2663–2690, 2002.
- Stevens, B., Moeng, C., Ackerman, A. S., Bretherton, C. S., Chlond, A., DeRoode, S., Edwards, J., Golaz, J., Jiang, H., Khairoutdinov, M., Kirkpatrick, M. P., Lewellen, D. C., Lock, A., Muller, F., Stevens, D. E., Whelan, E., and Zhu, P.: Evaluation of large-eddy simulations via observations of nocturnal marine stratocumulus, *Mon. Weather Rev.*, 133, 1443–1462, 2005.
- Stull, R. B.: *An Introduction to Boundary Layer Meteorology*, Kluwer Academic Publishers, Netherlands, 1988.
- Tang, D. F. and Dobbie, S.: iGen 0.1: a program for the automated generation of models and parameterisations, *Geosci. Model Dev.*, 4, 785–795, doi:10.5194/gmd-4-785-2011, 2011.