CARBON ATMOSPHERIC TRACER RESEARCH TO IMPROVE NUMERICAL SCHEMES AND EVALUATION

CATRINE

Carbon Atmospheric Tracer Research to Improve Numerics and Evaluation

D3.3: Evaluation and Validation of Resolved and Parameterised Plume Rise in Large-Eddy Simulations of Buoyant Emissions.

Due date of deliverable	30-06-2025
Submission date	18-06-2025
File Name	CATRINE-D3-3-V1.0
Work Package /Task	WP3, T3.3: plume rise and chemistry
Organisation Responsible of Deliverable	Wageningen University & Research (WUR)
Author name(s)	Bart van Stratum, Maarten Krol
Revision number	1
Status	Issued
Dissemination Level / location	Public

Funded by the European Union

The CATRINE project (grant agreement No 101135000) is funded by the European Union.
Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Commission. Neither the European Union nor the granting authority can be held responsible for them.

1 Executive Summary

Pollutant and greenhouse gas emissions from stacks – including those at power plants, steel mills, and other combustion facilities – are often accompanied by an emission of heat, resulting in buoyancy-driven plume rise. This modifies the vertical distribution of the emitted pollutants, influencing both ground-level concentrations and long-range transport.

There are different methods to account for plume rise in atmospheric models. High-resolution turbulence resolving models like Large-Eddy Simulation (LES) can explicitly resolve plume rise by co-emitting heat. However, the accuracy likely depends on the combination of the emission strength and model resolution. Lower-resolution models that do not resolve individual plumes must rely on plume-rise parameterisations that provide a measure of the effective emission height.

This report investigated two aspects related to plume rise: the resolution sensitivity of resolved plume rise in LES, and the accuracy of two commonly used (Briggs) plume rise parameterisations. This led to the following two conclusions and recommendations:

1. Although resolved plume rise in LES is resolution sensitive, resolving plume rise directly in LES simulations is still more accurate and preferable to using a plume rise parameterisation.
2. The plume rise parameterisation for straight and bent plumes from Briggs (1984) is – despite its age and relative simplicity – still a solid choice for models that cannot resolve plume rise.

Contents

2 Introduction

2.1 Background

The emission of greenhouse gases and pollutants from stacks – including those at power plants, steel mills, and other combustion facilities – are often accompanied by an emission of heat and moisture, resulting in buoyancy driven plume rise. As a result, the vertical distribution of the pollutants is modified, with the downwind peak concentration not occurring at the stack height z_s, but at some effective emission height $z_s + \Delta h$, where Δh is the inertia and buoyancy-driven plume rise. Depending on the initial inertia, the plume's heat content, and environmental conditions such as the atmospheric stability and wind speed, Δh can be hundreds of meters or more (Briggs, 1982; Weil, 1988). This redistribution affects the pollutant concentration height, influencing both ground-level concentrations and long-range transport (e.g Essa et al., 2006; Guan et al., 2008; Wells, 1917).

Most atmospheric models do not explicitly resolve individual plumes and therefore require a parameterization to calculate Δh. In contrast, Large-Eddy Simulation (LES) can – at least in theory – resolve plume rise explicitly, by including a source of heat together with the emitted pollutant (e.g. Nieuwstadt & de Valk, 1987). However, the skill of LES to accurately capture plume rise depends on several factors, like the source's heat output and spatial extent of the emission, relative to the spatial resolution of the LES model. Especially for small and buoyantly weak sources, simulated at a coarse LES resolution, it is questionable whether the resolved plume rise is accurate, as the released heat is artificially diluted over at least the volume of a single LES grid point. When using LES models to study buoyant plumes or generate synthetic data for improving or validating plume rise parameterisations for large-scale models, it is crucial to understand these potential limitations of resolving plume rise in LES.

For the parameterisation of Δh in large-scale models, there is a wide variety of parameterisations, often based on the seminal work on plume rise by Gary Briggs. This family of models contains both relatively simple ones which require few input parameters, up to entraining plume models that require information on the vertical thermodynamic state of the atmosphere (Briggs, 1984). In one LES study by Karagodin-Doyennel et al. (2024), such a parameterisation was preferred over resolved plume rise, although it was not clearly motivated why. This raises the question of how well plume rise models perform compared to resolved plume rise in LES, and whether parameterised plume rise can serve as a viable alternative.

2.2 Scope of this deliverable

2.2.1 Objectives of this deliverable

The aim of this report is twofold. First, we will systematically study plume rise in LES across a broad parameter space. This parameter space includes different stability regimes (stable, neutral, unstable), wind speeds, heat emission strengths, and model resolutions. Using the highest-resolution LES as a reference, this allows us to study the impact of resolution on plume rise. Second, we will use the library of LES simulations to validate two commonly used plume rise parameterisations with different complexities (Briggs, 1971, 1984)

2.2.2 Work performed in this deliverable

- The development and implementation of a new emission module in MicroHH, which is more efficient for many points sources (Section 3.1.1).

- The implementation of resolved plume rise in MicroHH by adding the option to specify the emission temperature for each point source (Section 3.1.2).
- The simulation of 252 idealised LES cases to test the sensitivity of plume rise to model resolution, for a range of atmospheric stabilities and wind speeds.
- The validation of two frequently used plume rise parameterisations (Briggs, 1971, 1984) with the library of LES experiments.

2.2.3 Deviations and counter measures

2.2.3.1 GPU chemistry

Part of this work package and task was to port the KPP-generated non-linear chemistry solver in MicroHH (Krol et al., 2024) from the CPU to GPU. We decided to focus our efforts primarily on plume rise instead of porting the chemistry to the GPU, for several reasons.

After analysing the KPP-generated code, we concluded that porting it entirely to the GPU was not feasible within the available time. KPP solvers are kept flexible by their sequential nature, where the KPP routine is called independently for each grid point, performing a series of *N* complex calculations that results in the updated concentrations for that single grid point. This design makes it easy to implement the KPP-solver in different models, regardless of their data structures or griFd layout. This sequential nature can be visualised as a matrix of grid points and calculations (*Figure **1***), through which the KPP-code steps row-by-row. For efficient GPU code this order needs to be transposed – the solver would have to be reorganised into *N* individual sequential functions (GPU-kernels), each applied to a sufficiently large number of grid points in parallel, thus stepping column-to-column through the matrix in Figure 1. Achieving this requires substantial changes to the existing CPU-KPP code.

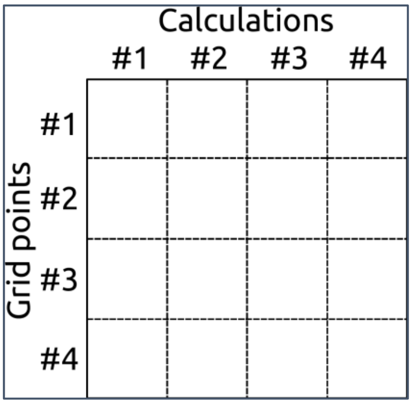

Figure 1. Illustration sequential vs. parallel code. KPP steps through this matrix row-by-row, efficient GPU code needs to traverse the matrix column-by-column.

Another challenge is how to approach the development of a GPU-based KPP solver. The beauty of KPP lies in its flexibility to generate different solvers, and porting a single instance to the GPU would remove most of that flexibility. The alternative is to modify KPP itself such that it can output GPU code (e.g. CUDA or OpenACC), but this is likely an enormous amount of work. An intermediate approach is to manually port the KPP-generated code to the GPU, while aiming to maintain flexibility for using different KPP schemes. This, however, requires a careful examination of the shared and overlapping components in different KPP-generated solvers.

In short, porting the KPP-chemistry to the GPU is a non-trivial task that requires more thought, planning, and time.

2.2.3.2 Non-linear chemistry and comparison with satellites

A direct result of the unfeasibility of porting the chemistry solver to the GPU is that we could not run our large ensemble of LESs with non-linear chemistry. Running this ensemble on the CPU with non-linear chemistry would have been prohibitively expensive. As a result, we excluded non-linear chemistry from all experiments and could not study the impact of plume rise dynamics on chemical transformations.

In addition, both the lack of non-linear chemistry and our choice to focus on idealized experiments, made a comparison of plume characteristics with satellites (TROPOMI, GEMS) impossible.

Currently, we are conducting MicroHH simulations for Rotterdam according to the intercomparison protocol (D3.1), and the simulation for Paris will follow. In these simulations, we will account for non-linear chemistry, allowing the comparison of the simulations with TROPOMI.

3 Methods

3.1 LES model description

We used the GPU version of MicroHH (van Heerwaarden et al., 2017) for all LES experiments. With the performance and efficiency of modern (NVIDIA H100) GPUs, this allowed us to run the large ensemble of 252 LES cases described in the next section, with limited computational costs.

To limit the complexity of the experiments, we used dry thermodynamics in MicroHH, which uses the potential temperature as the only prognostic variable influencing buoyancy. In addition, we excluded radiation, microphysics, and the interactive land-surface scheme, and drive all cases by a spatially homogeneous surface heat flux. As indicated in Section 2.2.3, the unfeasibility of porting the KPP-based non-linear chemistry solver to the GPU forced us to exclude non-linear chemistry from all experiments.

3.1.1 Point source emissions

Earlier versions of MicroHH did not include higher-order flux-limited (monotonic) advection schemes. For that reason, point sources were distributed over three-dimensional Gaussian volumes, to prevent Gibbs oscillations from producing spurious negative concentrations near sharp gradients. The use of these Gaussian volumes leads to an artificial spatial dilution of point sources, which may affect the peak concentrations, chemical transformations, and/or plume rise and shape. With the availability of a monotonic advection scheme in MicroHH (Koren, 1993), this approach is no longer needed, and we can simply emit point sources into a single LES grid point.

A second drawback of using Gaussian volumes is that they require a three-dimensional loop over each volume (containing computationally expensive mathematical functions) at every time step. While this is manageable for a small number of sources, it becomes costly when dealing with large numbers of point sources, such as the simulations in T3.1 over Rotterdam, which includes hundreds of point sources.

To address these issues, we implemented a new emission module in MicroHH that reads three-dimensional emission fields (with a limited vertical extent). This approach offers full flexibility to prescribe time and/or space varying emissions in any shape or size required by the user. Our newly developed Python package *microhhpy*[1] includes a help class for defining emission fields, adding emissions, and writing them in the input format required by MicroHH.

3.1.2 Resolved plume rise

The new emission module in MicroHH contains an option to resolve plume rise by providing the model with a stack emission volume flux V_s (m^3 s^{-1}) and temperature T_s (K). The volume of air entering the model is assumed to mix instantaneously within a grid cell, resulting in a local potential temperature tendency:

$$\frac{\partial \theta}{\partial t} = \frac{V_s}{\Delta x \Delta y \Delta z}(T_s - T_m)\left(\frac{p_0}{p_m}\right)^{\mathrm{R_d}\,/\,\mathrm{c_p}} \quad (1)$$

where Δx, Δy, Δz are the grid dimensions (m), T_m and p_m are the absolute temperature and pressure of the LES grid point, respectively, $p_0 = 10^5$ Pa is a reference pressure, $R_d = 287.04$ J kg^{-1} K^{-1} is the gas constant for dry air, and $cp = 1005$ J kg^{-1} K^{-1} is the isobaric specific heat of air.

Plume rise can be further enhanced if the plume has a significant stack exit velocity. However, buoyancy is typically the primary driver of plume rise (Weil, 1988) and including momentum generally has little impact on plume rise (e.g. Gordon et al., 2018). Therefore, plume rise due to momentum is currently not accounted for in MicroHH.

3.2 Plume rise parameterisations.

Plume rise parameterisations can broadly be divided in two categories. First, there are models which given the emission temperature, environmental temperature, and some metric of the atmospheric stability, provide a single value for Δh. These models are most useful when detailed information about the atmospheric state is not available, or when a rough estimate of Δh is sufficient. We will refer to those as *empirical plume rise models*. Second, there are models that incorporate more details of the atmospheric state, like a variation of stability with height. Such models often use an entraining plume model to calculate the change in a parcel's buoyancy with height, until the buoyancy excess over the environment is zero. We will refer to this approach as *dynamic plume rise models.*

In this report we validate both approaches. For the empirical model, we use the Dutch *Operational Priority Substances* (OPS) model, which is based on Briggs (1971) . For the dynamic approach, we will use the model for straight and bent plumes from Briggs (1984).

3.2.1 Empirical model

The OPS model (Briggs, 1971; Sauter et al., 2023) defines the stack buoyancy flux (m^4 s^{-3}) as:

$$F_{b1} = \frac{g}{\pi} V_s \left(1 - \frac{T_e}{T_s}\right) \quad (2)$$

Where $g = 9.81$ m s^{-2} is the acceleration due to gravity, and T_e is the ambient temperature (K). For convective and neutral conditions, plume rise is defined by:

$$\Delta h = 38.8 \frac{F_{b1}^{3/5}}{U_s} \text{ for } \mathrm{F_{b1}} \geq 55 \quad (3)$$

$$\Delta h = 21.1 \frac{F_{b1}^{3/4}}{U_s} \text{ for } \mathrm{F_{b1}} < 55 \quad (4)$$

[1] https://pypi.org/project/microhhpy/

where U_s (m s^{-1}) is the absolute wind speed at stack height. In more recent versions of OPS, U is evaluated at z_e + ½Δh, which thus requires an iterative approach. In this report, we substitute U_s with the prescribed background wind. For stable conditions the final plume rise is given by:

$$\Delta h = 2.6\left(\frac{F_{b1}}{s\,U_s}\right)^{1/3} (4) \tag{5}$$

with s is the stability parameter:

$$s = \frac{g}{T_e}\frac{\partial \theta}{\partial z} \tag{6}$$

where ∂θ/∂z is the potential temperature gradient at stack height. In the OPS model, ∂θ/∂z is fixed at 0.006 K m^{-1}.

3.2.2 Dynamic model

In Briggs's plume rise model for straight and bent plumes (Briggs, 1984, their section 8-4.4), the buoyancy at stack height is defined as:

$$F_{b2} = \frac{g}{\pi} V_s \left(\frac{T_s - T_e}{T_e}\right) \tag{7}$$

This buoyancy excess is decreased level-by-level until the parcel is in equilibrium with the environment. For straight vertical plumes, the buoyancy decrease from a level k to $k+1$ is described by:

$$F_{k+1} = F_k \;-\; 0.015\; s_k\, F_{b2}^{1/3} \left({z'}_{k+1}^{8/3} - {z'}_{k}^{8/3}\right) \tag{8}$$

where s_k is the stability parameter (Eq. 6), and $z' = z - z_s$. Note that several papers (Akingunola et al., 2018; Karagodin-Doyennel et al., 2024) substitute F_{b2} with F_{k-1}. It is unclear where the difference with Briggs' original work stems from. Since this modification leaves F_{k-1} undefined at stack height, we follow the original formulation.

For bent plumes, the buoyancy decrease follows:

$$F_{k+1} = F_k \;-\; 0.053\; s_k\, U_k \left({z'}_{k+1}^{3} - {z'}_{k}^{3}\right) \tag{9}$$

Where U_k is the absolute horizontal wind speed at the k-th level. For mixed conditions where straight and bent-over plumes occur, Briggs recommends choosing the formulation which results in the largest buoyancy decrease for a given level.

In our analysis, Eqs. 8 and 9 are used to decrease F_k until it crosses zero, and the final plume rise height is determined by linearly interpolating to the height where F equals zero.

3.3 Case setup

We used an idealised setup to systematically study a wide range of conditions, including different atmospheric conditions, and different emission and heat strengths. Our emission parameter space is loosely based on the Dutch Electronic Environmental Annual Report (e-MJV), obtained through the Dutch Open Government Act ("WOO").

For each heat emission $Q = \{0, 0.5, 1, 5, 10, 30\}$ MW, we simulated three different atmospheric stabilities (convective, neutral, stable) with wind speeds in the range $U = \{1, 5, 10\}$ m/s, at horizontal resolutions $\Delta_{xy} = \{10, 25, 50, 100, 200\}$ m, with $\Delta_z = \min(\Delta_{xy}, 20)$ m. The upper limit of Q = 30 MW represents typical heat emission rates for power plants across Europe. In the absence of observations, the highest resolution experiment is used as a reference. For the stable case, we omitted the 1 m/s experiment as it results in strongly stable conditions where the highest resolution setup does not resolve turbulence. Each emission is accompanied by a fixed emission of CO_2, set to $5{\cdot}10^9$ kg year^{-1}.

The initial vertical potential temperature profiles are described by an idealised profile, with a surface value (θ_s), a boundary layer depth (z_i), fixed lapse-rates in the boundary layer ($\gamma_{\theta;abl}$) and free-troposphere ($\gamma_{\theta;ft}$), with a temperature jump ($\Delta\theta$) in between. The inversion layer is smoothed over a certain depth (Δz_i) using the error function. The neutral and stable cases share the same initial well-mixed profile, but the neutral case remains neutral by not applying a surface heat flux, while the convective case is heated by a 0.2 K m s-1 heat flux. To not further stabilise the stable case and prevent issues with using flux boundary conditions with negative heat fluxes (Basu et al., 2008), no surface flux is applied to the stable case. The parameters of the vertical profiles are described in Table 1, and the profiles are shown in Figure 2.

Wind is initialised as purely zonal flow, constant with height. Surface friction then reduces wind near the surface until a realistic wind profile is formed. In the results section, when referring to a wind speed, we refer to the initial or free-stream velocity.

Each simulation runs for a total of six hours, with the first four hours discarded as spin up.

Table 1. Input parameters idealised initial profiles

Quantity	Stable	Neutral/convective	Units
θ_s	290	290	K
z_i	200	1250	m
$\Delta\theta$	4	4	K
$\gamma_{\theta;abl}$	0.006	0	K m-1
$\gamma_{\theta;ft}$	0.003	0.006	K m-1
Δz_i	100	200	m

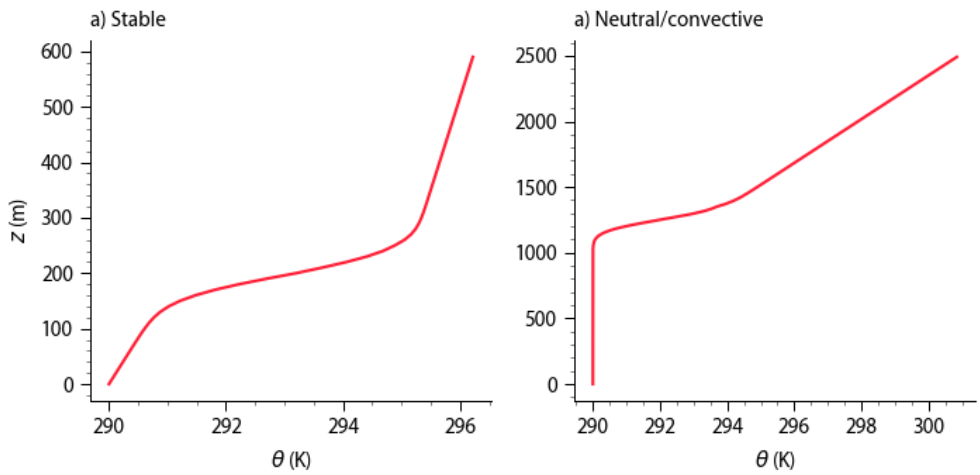

Figure 2. Idealised input profiles LES for stable and neutral/convective conditions

3.4 Post-processing LES experiments

We determine the plume rise in LES as the maximum height of a distinct peak in concentration. This method is illustrated in Figure 3. Just after release, the concentration near the stack shows a distinct peak close to stack height. Further downwind, the plume rises and disperses until turbulence and diffusion cause the peak to dissipate. We define plume rise as the highest downwind height at which a concentration peak is present, with the condition that the peak concentration must be at least 5% greater than the boundary-layer averaged concentration. This threshold prevents the algorithm from selecting a random height as the peak in a well-mixed profile. As these plumes often exhibit chaotic meandering behaviour, we used concentration profiles which are averaged in the spanwise direction and over the last two hours of the experiment.

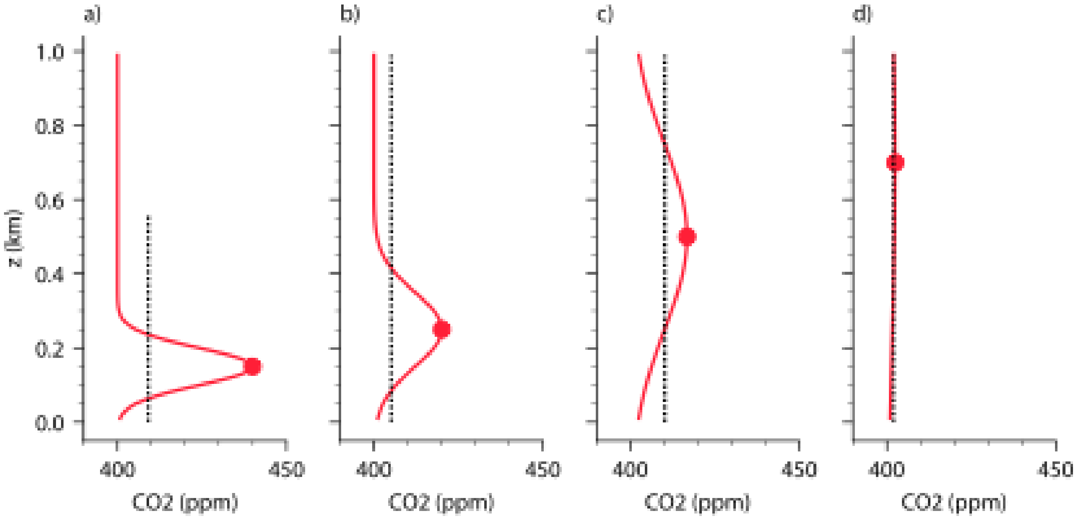

Figure 3. Idealised example of plume height determination

4 Results

We start this section by giving a general overview of the LES experiments in Section 4.1, followed by a more detailed evaluation of the vertical distribution of the plumes in Section 4.2. Finally, we evaluate the plume rise parameterisations with the LES data in Section 4.3.

4.1 General description plume rise

To get an initial understanding of the different LES experiments and their sensitivity to resolution, we begin by showing vertical cross-sections of the CO_2 plumes. These cross-sections have been averaged in the spanwise direction and over the last two hours of the LES experiments. Without the averaging, the turbulent meandering nature of the plume makes comparing experiments and resolutions difficult. The averaging process, however, makes the plumes appear artificially smooth. We limit the comparison to a select number of experiments that illustrate the behaviour of the LESs.

4.1.1 Stable conditions

Figure 4 shows the experiments for the case with the weakest heat emission of 0.5 MW. Each row presents a different wind speed, each column a different resolution. The dashed line shows the maximum plume rise height determined with the algorithm described in Section 3.4. For both wind speeds, the plume rise captured by the 10 m resolution experiment decreases with decreasing resolution. This is likely the result of the spatial dilution of buoyancy; in the highest resolution case the heat is distributed over a 1000 m^3 volume, whereas the 50 × 50 × 20 m^3 resolution case has a grid point volume of 50’000 m^3, resulting in a factor 50 dilution.

Next, Figure 5 shows the same for the experiments with a 5 MW emission. At this heat output, the lowest wind speed experiments show little sensitivity to resolution, with the low-resolution LESs plume rise being only marginally lower than the reference run. However, at the highest wind speed, there is still a clear dependence of plume rise on resolution, although at least all cases now capture some plume rise.

Finally, Figure 6 shows the Q=30 MW experiments. At such a heat output, most sensitivity of plume rise height on resolution has vanished. However, plume rise is still sensitive to resolution, as the lower resolution runs clearly show a slower rising plume.

All experiment from Q=0.5 MW to Q=30 MW show an additional issue when using low-resolution LES under stable conditions: as the resolution is coarsened and less turbulence is resolved, the models become more diffusive (van Stratum & Stevens, 2015), which is most evident in Figure 4e. These issues are further discussed in Section 5.

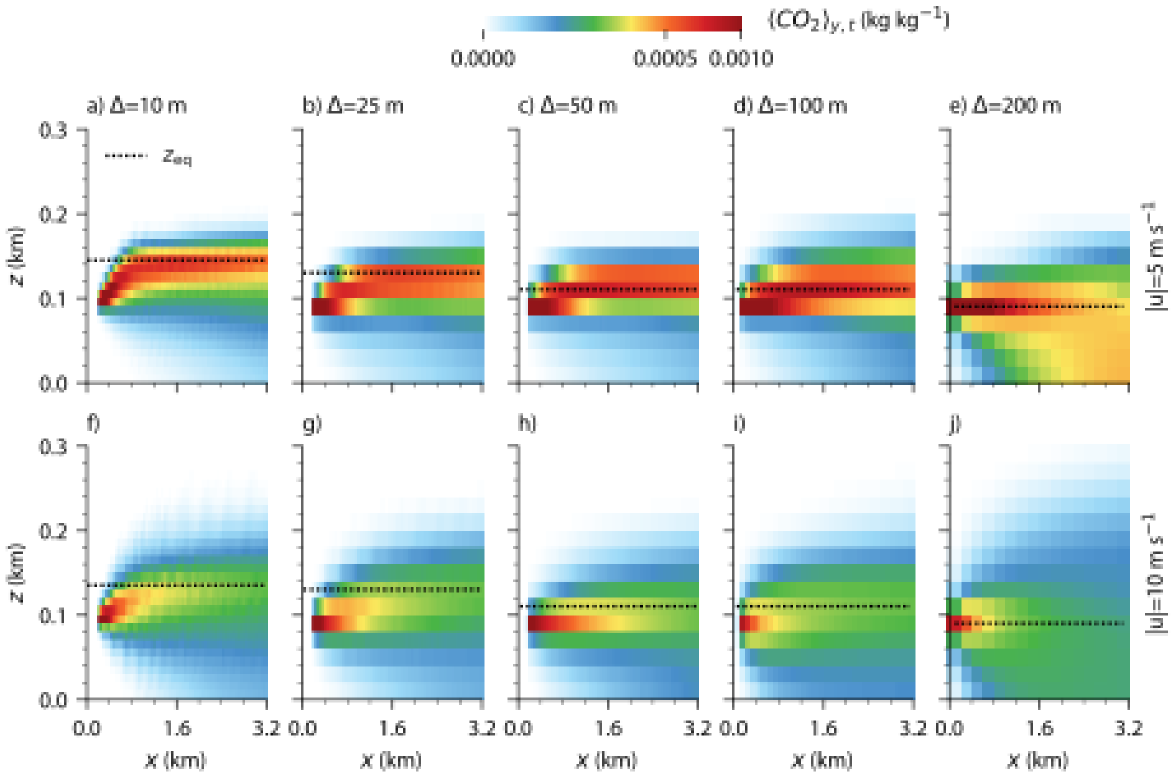

Figure 4. Spanwise and time averaged cross-section of CO2 for the stable case with Q=0.5 MW. Each row shows a different wind speed, each column a model resolution. The dashed black line indicates the plume equilibrium height.

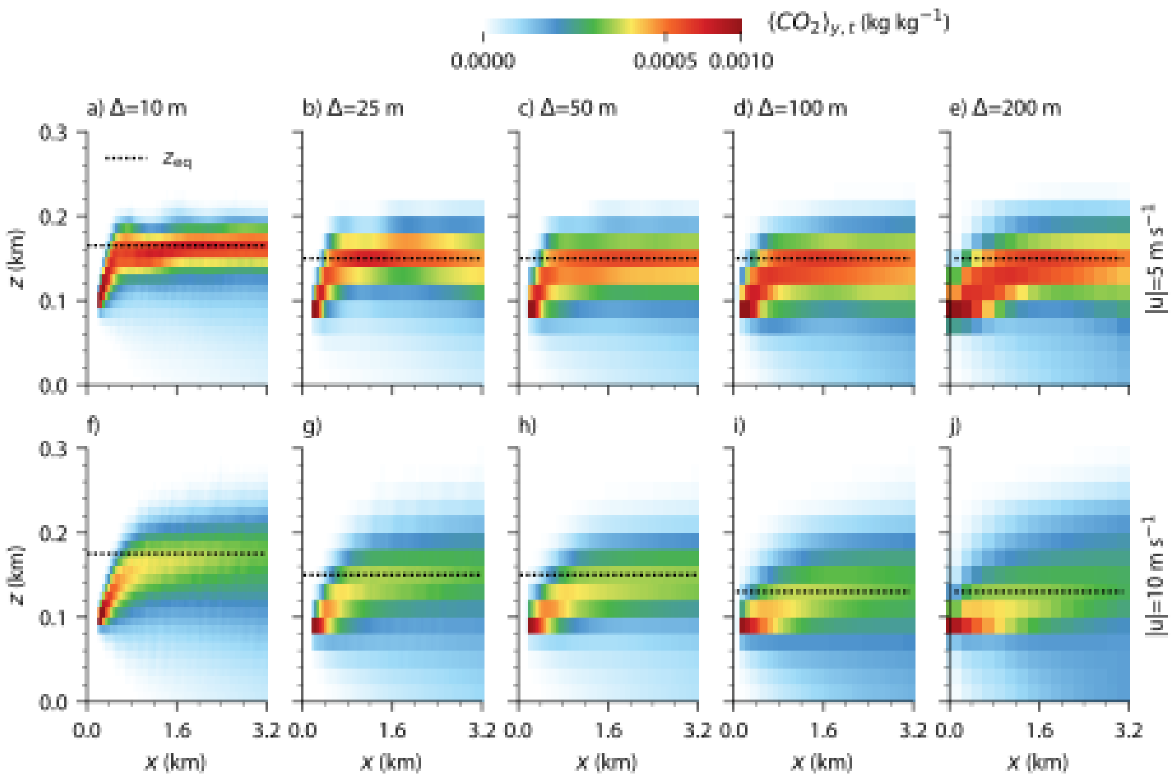

Figure 5. Like previous figure, for Q=5 MW

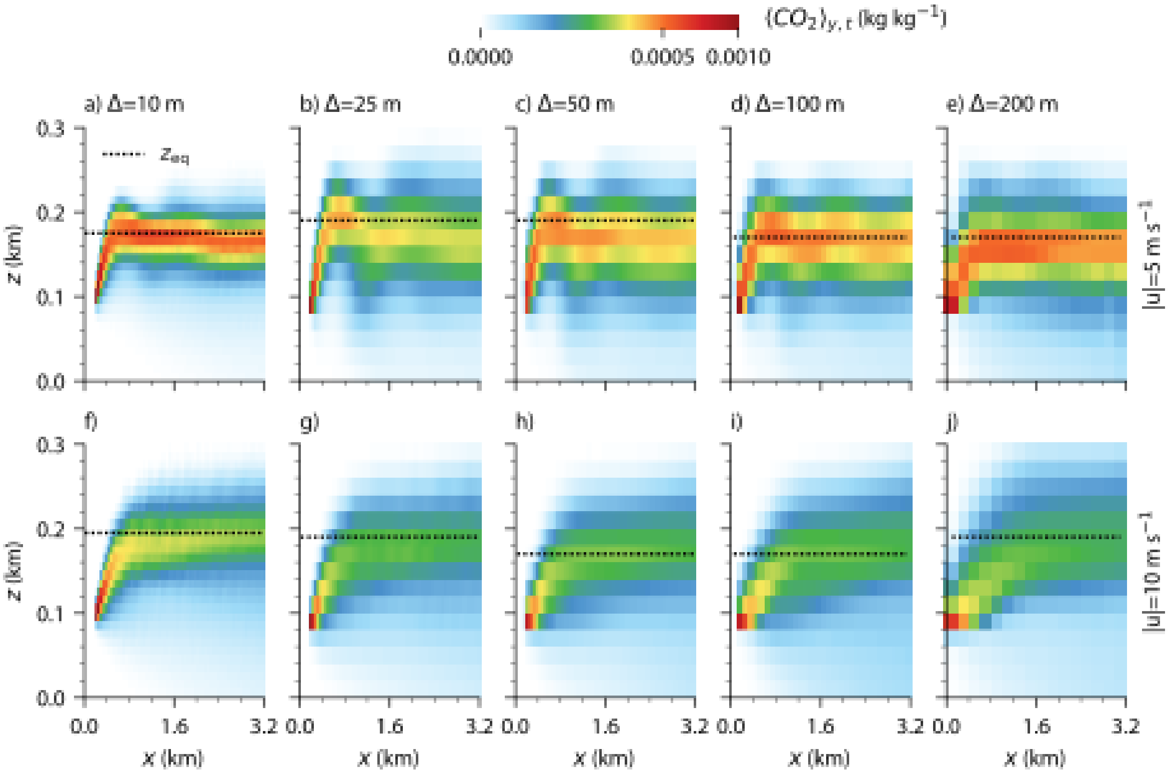

Figure 6. Like previous figure, for Q=30 MW

4.1.2 Neutral conditions

Under neutral conditions, plume rise behaves completely differently, as the rising air is not directly inhibited by an inversion layer. Figure 7 shows the LES experiment for the weakest (Q=0.5 MW) emission. At the lowest wind speed, the reference experiment shows that the plume can penetrate all the way to the capping inversion, which begins at ~1 km height. The lower resolution experiments show significantly less plume rise, with the lowest resolution simulation resolving almost none. At higher wind speeds, the heat output is insufficient to cause any significant plume rise in any of the cases.

For Q>1 MW, all experiments show little sensitivity to resolution. Figure 8 shows the experiment with Q=10 MW. Although plume rise itself is clearly dependent on wind speed, all experiments at all resolutions show approximately the same behaviour. Plume rise in the 1 to 5 m s^{-1} experiments is inhibited by the inversion layer, while the plumes in the 10 m s^{-1} case level off at approximately half the boundary layer depth. The diagnosed plume rise differs between experiments, but without a clear trend from the highest to lowest resolution.

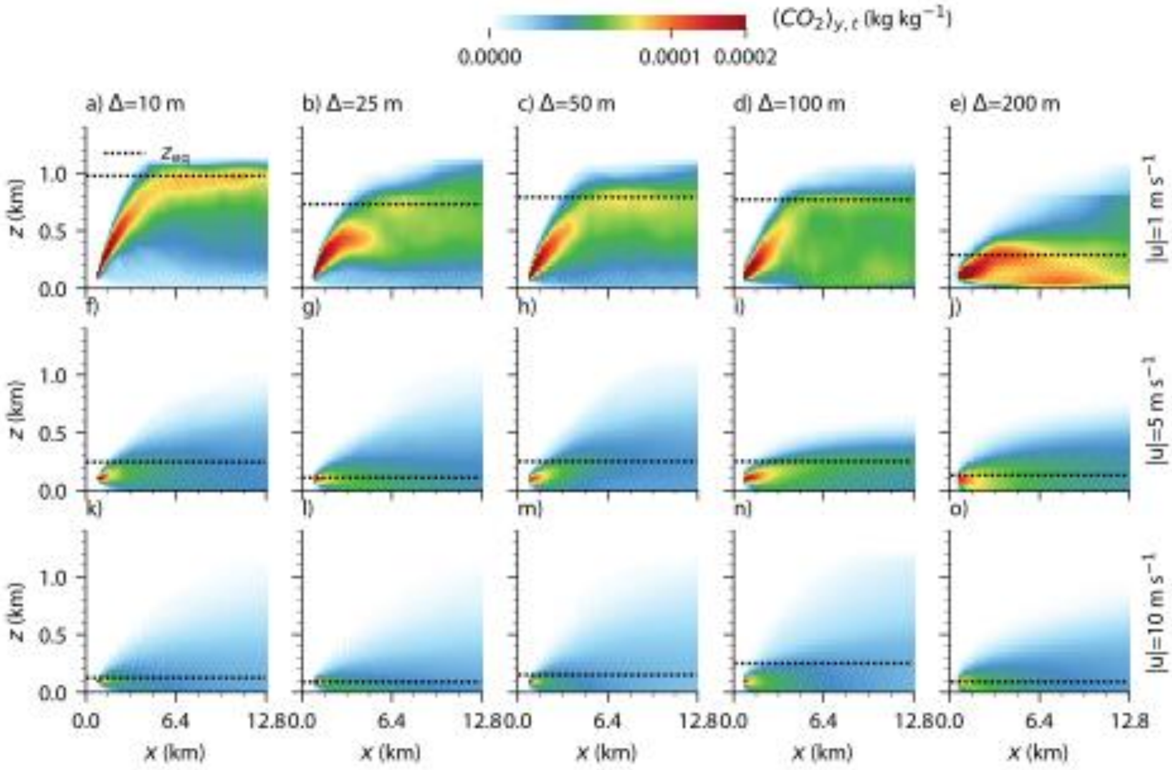

Figure 7. Like the previous figures, for the neutral case with Q=0.5 MW

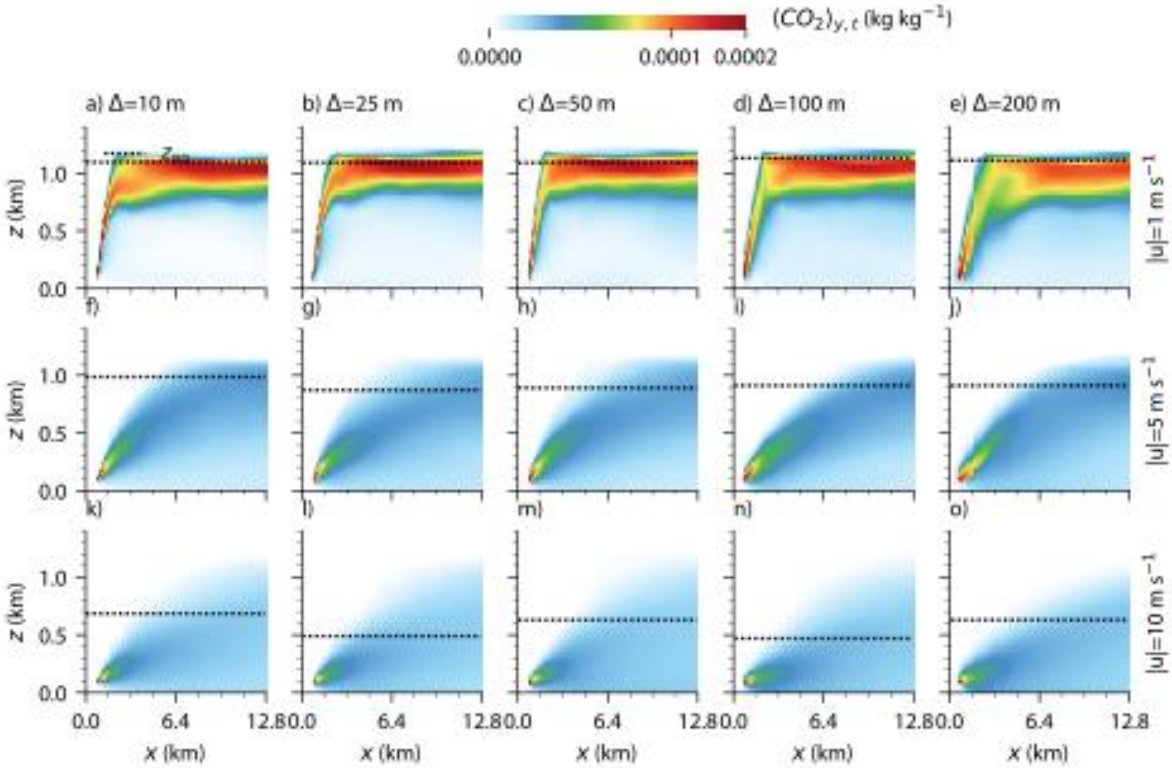

Figure 8. Like previous figure, for Q=10 MW

4.1.3 Convective conditions

Plume rise is most difficult to quantify in the convective case, as natural convection quickly mixes pollutants into a well-mixed state. To further complicate the diagnosis of plume rise, even the control experiment with Q=0 MW (not shown) shows a concentration peak in the inversion layer. This is likely caused by pollutants getting trapped as natural convection penetrates the inversion. Furthermore, the combination of natural and anthropogenic convection causes the statistics to converge slower in time.

Figure 9 shows the result for the Q=0.5 MW experiments. Visually, all experiments show the same behaviour where an initial peak in concentration is rapidly mixed away by convection. The diagnosed plume rise is also similar in all cases, although there is some resolution independent variability in the lowest wind speed case. Similar results are observed at higher heat outputs, as shown in Section 4.2. Overall, the mixing by natural convection reduces the significance of plume rise in convective conditions.

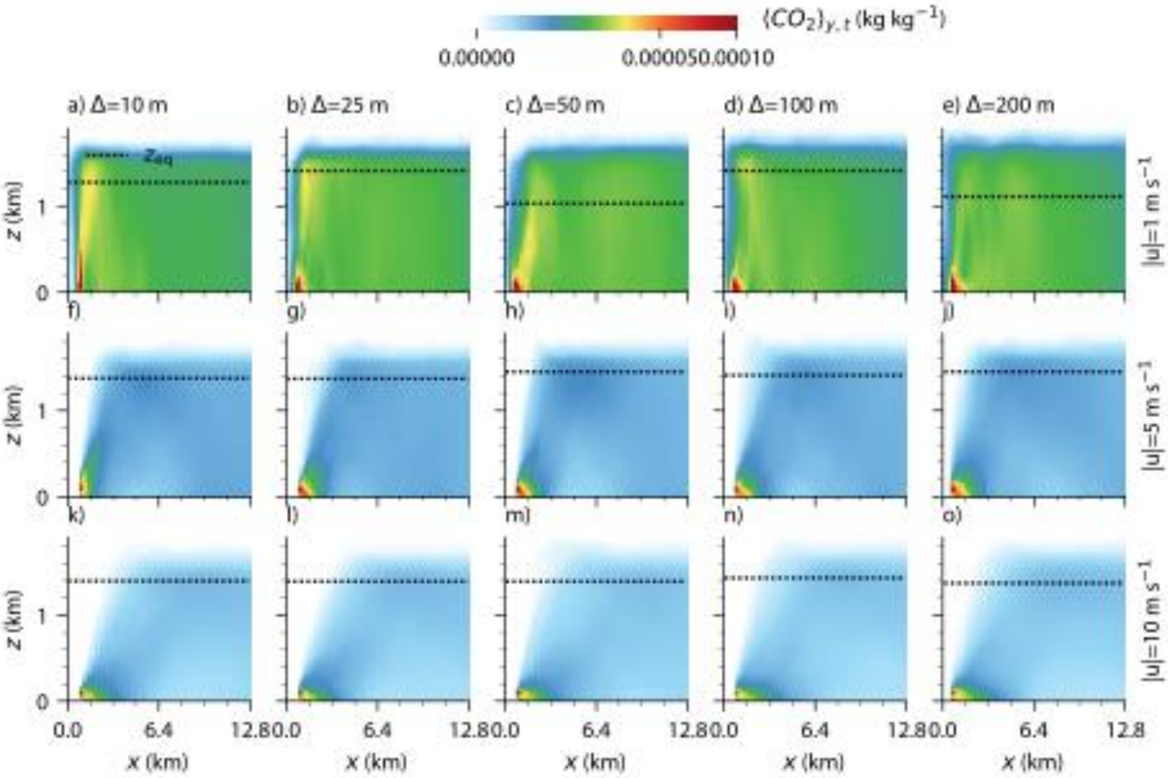

Figure 9. Like previous figures, for the convective case with Q=0.5 MW

4.2 Vertical distribution pollutants

To further study plume rise and differences in the vertical pollutant distribution as a function of resolution, we continue by examining vertical profiles. For each case we use the spanwise and time averaged profile at the downwind location where the highest (in terms of height, not concentration) peak was observed. Given the large differences in absolute concentrations between different cases, we normalise all concentrations with the peak concentration of the reference case.

Figure 10 shows the vertical profiles for the stable case. Somewhat worryingly, there seems to be little convergence of the profiles with increasing resolution. For both wind speeds, the reference case typically has higher plume rise and a skewed Gaussian shape, where the lower resolution cases have an approximately non-skewed shape. In lower-resolution cases, the plume is also broader, likely caused by increased vertical diffusion at (too) low-resolution LES of stable conditions.

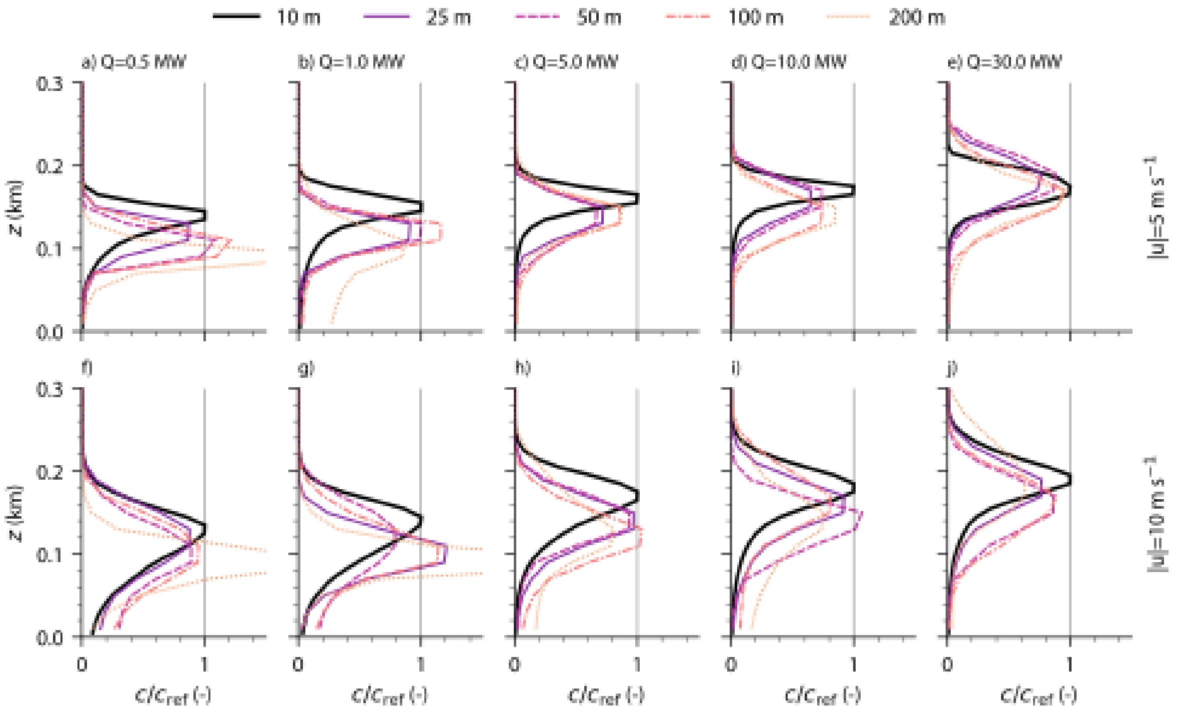

Figure 10. Vertical profiles at location of highest (z) peak concentration for the stable case. The concentrations are normalised with the concentration of the reference simulation.

Figure 11 shows the same profiles for the neutral case. Under these neutral conditions turbulence is generated both by wind shear and the plume's buoyancy. Therefore, turbulence increases (not shown) both when going from left-to-right (increase buoyancy) and/or top-to-bottom (increase wind shear) through the figure. At the two lowest heat emissions, the amount of resolved turbulence depends strongly on the resolution, and as a result there are large differences between the different experiments. In addition, the algorithm used to determine plume rise height seems to be imperfect, as (for the strongest wind speed), most cases have their peak concentration at the emission height, while two of the lower-resolution cases have their peak concentration at a greater height, resulting in completely different profile shapes. At the three highest heat emissions, the results converge better with resolution. In most cases the shape of all profiles is similar, and typically plume rise is captured well.

Figure 12 shows the vertical profiles for the convective case. As mentioned earlier, natural convection creates more turbulent and chaotic simulations where the statistics converge slower in time, and as a result, many of the profiles are far from being smooth. The lowest wind speed cases are on average well mixed since convection has more time to mix the pollutants before leaving the domain. The two highest wind speeds cases do show a clear peak in concentration near the inversion layer, which is captures well by all experiments and resolutions.

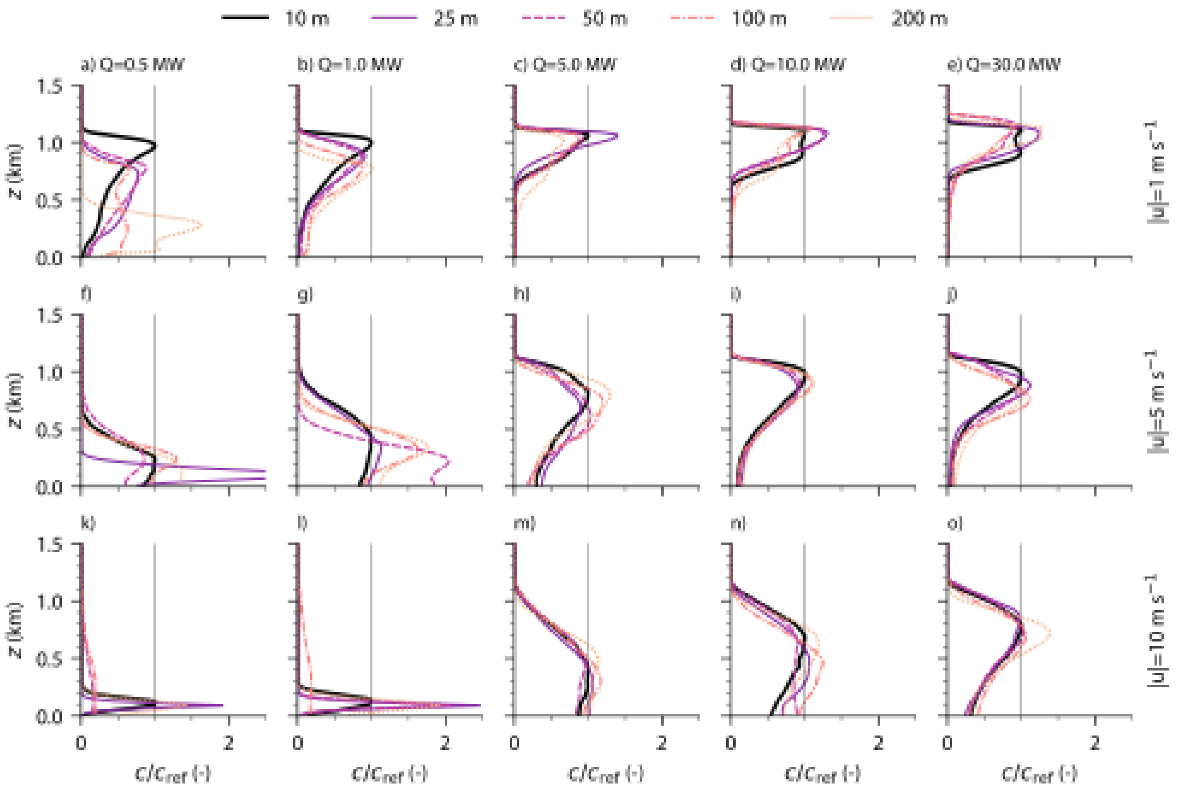

Figure 11. Like the previous figure, for the neutral case.

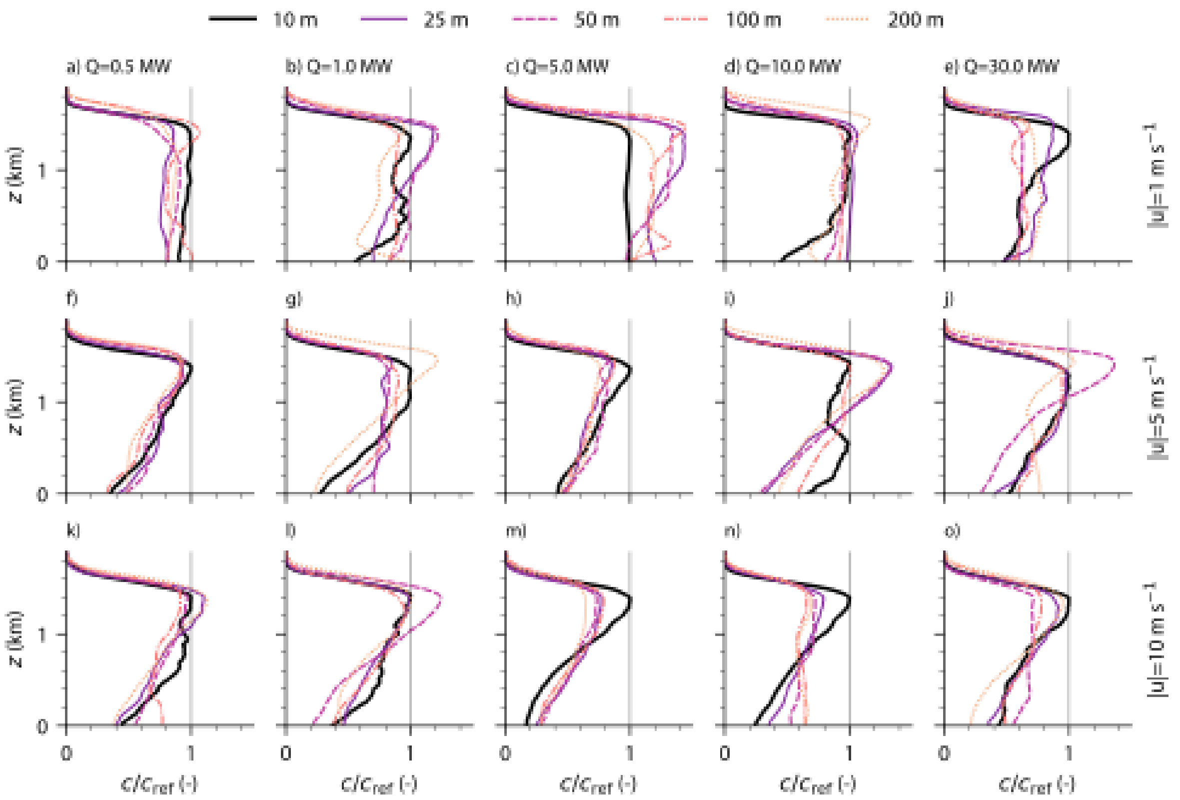

Figure 12. Like the previous figure, for the convective case.

4.3 Validation Briggs plume rise parameterisations

This section validates the two plume rise parameterisations described in Section 3.2 with the high-resolution reference LES experiments of each case. To reiterate, the model used by OPS is the plume rise model from Briggs (1971, hereafter called B71), which only requires knowledge of the change in wind with height, and makes general assumptions about the thermodynamic structure of the atmosphere for stable and neutral/convective cases. On the other hand, the plume rise model from Briggs, (1984, hereafter B84) is an entraining plume model that uses knowledge of the thermodynamic structure from an external source – in our case LES.

Figure 13 shows the comparison of the plume rise height Δh from both parameterisations with LES. Starting with the stable case, both B71 and B84 perform well for emissions up to 5 MW, although on average they slightly overestimate plume rise. For stronger emissions, the B71 model starts to significantly overestimate plume rise, whereas the B84 model – with its knowledge of the temperature stratification from LES – does a better job, although plume rise is still consistently overestimated.

For neutral conditions the comparison is more complex. At the two highest wind speeds, the B71 model strongly underestimates plume rise for $Q > 1$ MW. Depending on the wind speed, the plumes in LES rise to half the inversion height (~ 1km) or more, whereas B71 predicts a plume rise of only a few hundred meters. At the lowest wind speed, B71 only provides a reasonable estimate $Q = 30$ MW.

The B84 model performs much better. At the lowest wind speed the predicted plume rise is accurate for $Q \geq 5$ MW and underestimated by 100-200 m at weaker emissions. At the intermediate wind speed the same pattern emerges, although the underestimation is larger for weaker emissions, with differences of 400 m for $Q = 0.5$ MW. B84 performs the worst at the highest wind speed: for weak emissions plume rise is overestimated by ~500 m, a difference that decreases with the emissions strength, but never gets smaller than ~150 m.

Finally, for the convective case, B71 again clearly performs the worst. Only at the lowest wind speed and strongest emission the model gives a good estimate of plume rise. The B84 model performs much better, although it consistently underestimates plume rise with 200 – 400 m. This might be because B84 accounts only for buoyancy and neglects inertia – both at the initial stack release and inertia gained in the atmosphere through positive buoyancy – so plumes cannot overshoot the height where buoyancy becomes zero. However, as shown in Section 4.1.3, the rapid mixing of pollutants by natural convection reduces the significance of these biases in plume rise under convective conditions.

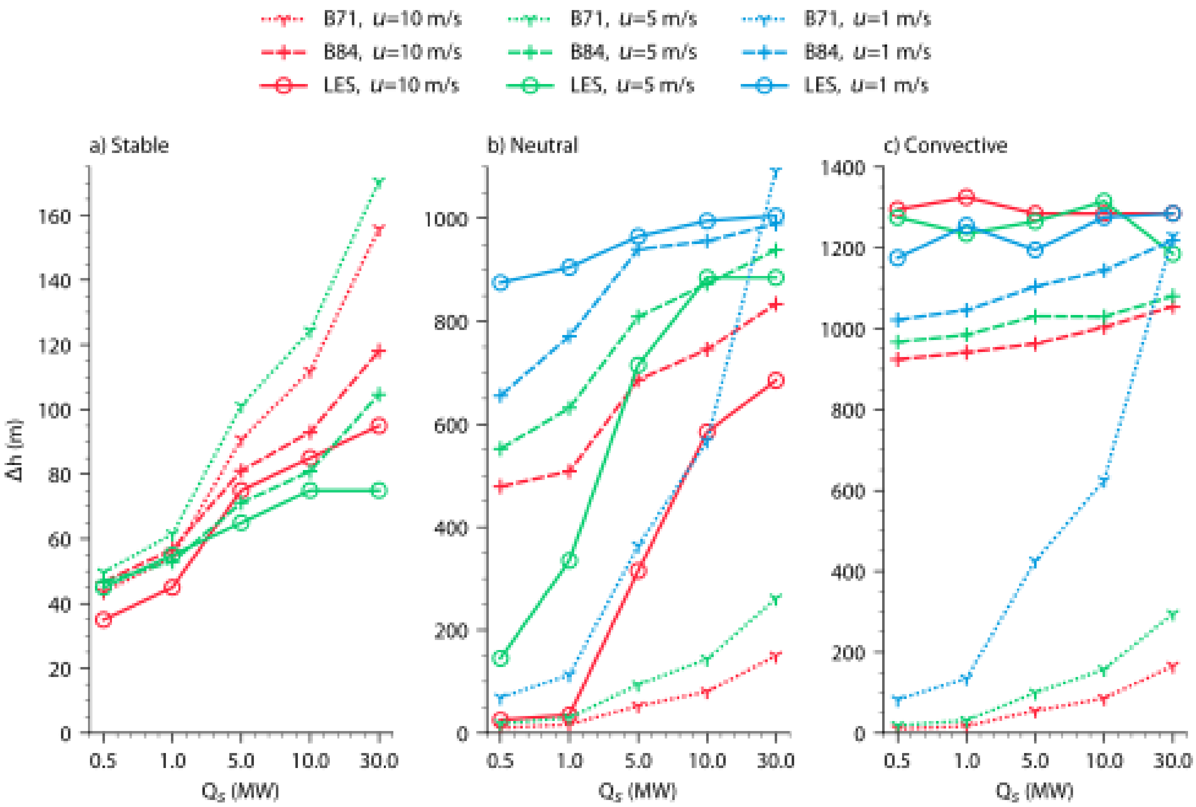

Figure 13. Validation plume rise parameterisations with LES

5 Discussion

This report examined the influence of resolution on resolved plume rise in LES, using idealised experiments. This idealisation introduces some simplifications compared to real-world conditions. First, our experiments only included emissions of heat, while stacks often emit significant amounts of moisture, which can influence the plume's thermodynamics and rise (e.g. Fathi et al., 2025; Schatzmann & Policastro, 1984). The influence of moisture can be further amplified if condensation occurs, where the latent heat release can further enhance plume rise. Second, we only studied the influence of buoyancy on plume rise. However, if the stack exit velocity is substantial, the initial inertia of the plume can add to the inertia generated by buoyancy, further enhancing plume rise, which is an effect we did not account for.

Furthermore, this report solely examined the influence of model resolution on plume rise. The model resolution is known to influence other plume properties, like the rate of dispersion, or non-linear chemical transformations inside the plume. This needs to be further examined in future work.

Finally, our LES simulations contain several experiments where the name No Eddy Simulation (NES) would be more appropriate, as the resolution is insufficient to resolve any turbulence. Especially for stable conditions, the resolution required to resolve turbulence is typically <10 m. However, for LES experiments which contain both convective daytime and stable nighttime conditions, resolution is often scarified in expense for the domain size requirements for daytime convection. For that reason, we opted to include these NES simulations in our analysis, as they are frequently employed in LESs of real-world weather.

6 Conclusions and recommendations

This report studied resolved plume rise in LES, and its sensitivity on resolution. Knowledge about this sensitivity is crucial when using LES to study buoyant plumes or generate synthetic data for improving or validating plume rise parameterisations for large-scale models.

Using different stabilities, wind speeds, and resolutions, we created a library of 252 LES simulations. This library was used to both study the dependence of plume rise on resolution, and to validate two commonly used plume rise parameterisations from Briggs.

From the LES results, we can conclude that resolved plume rise is sensitive to resolution. In LES, the heat source is always distributed over at least the volume of a single grid point. This results in a "dilution" of buoyancy, as for every factor two increase in grid spacing, the volume increases $2^3 = 8$ times. From a 10 m^3 to 100 x 100 x 20 m^3 resolution, this results in a factor 200 increase in release volume. Further difficulties are caused by using a low-resolution LES in neutral or stable conditions, where with decreasing resolution the amount of resolved turbulence decreases, influencing transport and mixing.

We validated two plume rise parameterisations from Briggs, the more simple one from (Briggs, 1971) as used by the OPS model, and the more complex entraining plume model from (Briggs, 1984). The latter clearly performs best, as it uses external stability information to describe stability with height, leading to an overall more robust model.

This leads us to the following recommendations:

1. <u>Resolved plume rise in LES remains preferred over parameterisations, despite its resolution sensitivity</u>. Not only is plume rise in coarse-resolution LES still more accurate than the parameterisations, it also does not require assumptions on the vertical distribution at the effective emission height (not discussed in this report). Furthermore, it naturally captures any bent plume structures that result from plume rise and wind shear.

2. <u>For large-scale models, Briggs' (1984) method remains a solid choice.</u> Despite the relative simplicity of Briggs' *"rise into irregular stability profiles"* model (their Section 8-4.4), the parameterisation gives a reasonable estimation of plume rise under a wide range of conditions, with a method that can easily and computationally cheaply be implemented in models where computational efficiency is of key importance.

7 Bibliography

Akingunola, A., Makar, P. A., Zhang, J., Darlington, A., Li, S.-M., Gordon, M., Moran, M. D., & Zheng, Q. (2018). A chemical transport model study of plume-rise and particle size distribution for the Athabasca oil sands. *Atmospheric Chemistry and Physics*, *18*(12), 8667–8688. https://doi.org/10.5194/acp-18-8667-2018

Basu, S., Holtslag, A. A. M., Van De Wiel, B. J. H., Moene, A. F., & Steeneveld, G.-J. (2008). An inconvenient “truth” about using sensible heat flux as a surface boundary condition in models under stably stratified regimes. *Acta Geophysica*, *56*(1), 88–99. https://doi.org/10.2478/s11600-007-0038-y

Briggs, G. A. (1971). *Some recent analyses of plume rise observations*. https://repository.library.noaa.gov/view/noaa/33598/noaa_33598_DS1.pdf

Briggs, G. A. (1982). Plume Rise Predictions. In D. A. Haugen (Ed.), *Lectures on Air Pollution and Environmental Impact Analyses* (pp. 59–111). American Meteorological Society. https://doi.org/10.1007/978-1-935704-23-2_3

Briggs, G. A. (1984). Plume Rise and Buoyance Effects. In *Atmospheric Science and Power Production*.

Essa, K. S. M., Mubarak, F., & Elsaid, S. E. M. (2006). Effect of the plume rise and wind speed on extreme value of air pollutant concentration. *Meteorology and Atmospheric Physics*, *93*(3), 247–253. https://doi.org/10.1007/s00703-005-0168-1

Fathi, S., Makar, P., Gong, W., Zhang, J., Hayden, K., & Gordon, M. (2025). The importance of moist thermodynamics on neutral buoyancy height for plumes from anthropogenic sources. *Atmospheric Chemistry and Physics*, *25*(4), 2385–2405. https://doi.org/10.5194/acp-25-2385-2025

Gordon, M., Makar, P. A., Staebler, R. M., Zhang, J., Akingunola, A., Gong, W., & Li, S.-M. (2018). A comparison of plume rise algorithms to stack plume measurements in the Athabasca oil sands. *Atmospheric Chemistry and Physics*, *18*(19), 14695–14714. https://doi.org/10.5194/acp-18-14695-2018

Guan, H., Chatfield, R. B., Freitas, S. R., Bergstrom, R. W., & Longo, K. M. (2008). Modeling the effect of plume-rise on the transport of carbon monoxide over Africa with NCAR CAM. *Atmospheric Chemistry and Physics*, *8*(22), 6801–6812. https://doi.org/10.5194/acp-8-6801-2008

Karagodin-Doyennel, A., Jansson, F., van Stratum, B., Denier van der Gon, H., Vilà-Guerau de Arellano, J., & Houweling, S. (2024). Carbon dioxide plume dispersion simulated at hectometer scale using DALES: Model formulation and observational evaluation. *EGUsphere*, 1–41. https://doi.org/10.5194/egusphere-2024-3721

Koren, B. (1993). *A robust upwind discretization method for advection, diffusion and source terms*. Centrum voor Wiskunde en Informatica Amsterdam.

Krol, M. C., van Stratum, B. J. H., Anglou, I., & Boersma, K. F. (2024). Evaluating NO_x stack plume emissions using a high-resolution atmospheric chemistry model and satellite-derived NO_2 columns. *Atmospheric Chemistry and Physics*, *24*(14), 8243–8262. https://doi.org/10.5194/acp-24-8243-2024

Nieuwstadt, F. T. M., & de Valk, J. P. J. M. M. (1987). A large eddy simulation of buoyant and non-buoyant plume dispersion in the atmospheric boundary layer. *Atmospheric Environment (1967)*, *21*(12), 2573–2587. https://doi.org/10.1016/0004-6981(87)90189-2

Sauter, F., Sterk, M., van der Swaluw, E., Wichink Kruit, R., de Vries, W., & van Pul, A. (2023). *The OPS-model: Description of OPS 5.1.1.0*. National Institute for Public Health and the Environment (RIVM).

Schatzmann, M., & Policastro, A. J. (1984). Plume Rise from Stacks with Scrubbers: A State-of-the-Art Review. *Bulletin of the American Meteorological Society*, *65*(3), 210–215. https://doi.org/10.1175/1520-0477(1984)065<0210:PRFSWS>2.0.CO;2

van Heerwaarden, C. C., Van Stratum, B. J., Heus, T., Gibbs, J. A., Fedorovich, E., & Mellado, J.-P. (2017). MicroHH 1.0: A computational fluid dynamics code for direct numerical simulation and large-eddy simulation of atmospheric boundary layer flows. *Geosci. Model Dev.*, *10*, 3145–3165. https://doi.org/10.5194/gmd-10-3145-2017

van Stratum, B. J. H., & Stevens, B. (2015). The influence of misrepresenting the nocturnal boundary layer on idealized daytime convection in large-eddy simulation. *Journal of Advances in Modeling Earth Systems*, *7*(2), 423–436. https://doi.org/10.1002/2014MS000370

Weil, J. C. (1988). Plume Rise. In A. Venkatram & J. C. Wyngaard (Eds.), *Lectures on Air Pollution Modeling* (pp. 119–166). American Meteorological Society. https://doi.org/10.1007/978-1-935704-16-4_4

Wells, A. E. (1917). Results of Recent Investigations of the Smelter Smoke Problem. *Journal of Industrial & Engineering Chemistry*, *9*(7), 640–646. https://doi.org/10.1021/ie50091a010

Document History

Version	Author(s)	Date	Changes
0.1 (Initial document created)	Bart van Stratum, Maarten Krol	04-06-2025	None
0.2 (After internal review)	Bart van Stratum, Maarten Krol	17-06-2025	Processed comments internal reviewers.
1.0 (Version after internal review and ready for submission to EC portal)	Bart van Stratum, Maarten Krol	18-06-25	Issued version

Internal Review History

Internal Reviewers	Date	Comments
Thomas Lauvaux	12-06-2025	
Giovanni Tumolo	12-06-2025	