

BEE 4750/5750

Homework 5: Gaussian Plumes and Model Validation

2026-10-14

Due Date

Thursday, 10/22/26, 9:00pm

Overview

Instructions

- Problem 1 asks you to evaluate the plume equation and reason about stability class.
- Problem 2 asks you to analyze residuals to interpret goodness-of-fit.
- Problem 3 asks you to calibrate a working plume model against a set of monitor observations and decide whether it is the right model for the setting.

Load Environment

```
import Pkg
Pkg.activate(@_DIR_)
Pkg.instantiate()
```

```
using CSV
using DataFrames
using Statistics
using Optim
using Plots
```

Problems (Total: 50 Points)

Problem 1 (10)

The ground-level concentration downwind of an elevated point source, on the plume centreline and accounting for reflection at the ground, is

$$C(x, y) = \frac{Q}{\pi u \sigma_y(x) \sigma_z(x)} \exp\left(\frac{-y^2}{2\sigma_y(x)^2}\right) \exp\left(\frac{-H^2}{2\sigma_z(x)^2}\right).$$

A stack emits $Q = 100$ g/s at an effective height of $H = 41$ m into a wind of $u = 2.5$ m/s.

Problem 1.1 (5)

At 2 km downwind under the stability class in question, $\sigma_y = 62$ m and $\sigma_z = 23$ m.

1. Compute the ground-level concentration on the centreline at 2 km, in $\mu\text{g}/\text{m}^3$.

2. Compute it again at a receptor 60 m off the centreline at the same distance. What fraction of the centreline value is it?

Problem 1.2 (5)

Stability classes run from A (very unstable, strong daytime heating) through F (stable, clear calm night). In lecture, this same stack produced a higher ground-level peak under class B (a sunny afternoon) than under class F (a clear night), much closer to the stack. The dispersion coefficients at two distances downwind are:

Distance	Class B: σ_y, σ_z (m)	Class F: σ_y, σ_z (m)
1 km	156, 110	34, 14
8 km	1001, 1064	218, 42

A school sits 1 km directly downwind of the stack, and a town 8 km downwind. The 1-hour SO_2 standard is $196 \mu\text{g}/\text{m}^3$.

1. Compute the ground-level centreline concentration at the school and at the town under each class. Which class is the worst case at each receptor?
2. The facility's consultant proposes a permit analysis that models only class B, since it produces the highest peak. What would that analysis miss? Which conditions should a permit analysis model instead? Use your table to say why.

Problem 2 (5)

A dissolved-oxygen model is compared with measurements at six stations along a river, numbered in order downstream. Three versions of the model leave these residuals (model minus observation, in mg/L):

Station	1	2	3	4	5	6
Model A	+0.7	-0.5	-0.1	+0.5	-0.7	+0.1
Model B	+0.5	+0.5	+0.5	+0.5	+0.5	+0.5
Model C	-0.7	-0.5	-0.1	+0.1	+0.5	+0.7

Problem 2.1 (1)

Compute the RMSE of each model.

Problem 2.2 (3)

For each model, say in a sentence what its residuals show that the RMSE does not, and what kind of error that suggests.

Problem 2.3 (1)

A parameter that shifts every prediction up or down by the same amount (for example, a sensor's calibration offset) is changed until the fit is as good as it can be. Which model or models could that fix?

Problem 3 (35)

A smelter sits on the floor of a narrow valley that runs east–west. One clear night there was almost no wind, and a **temperature inversion** formed over the valley floor.

Air normally gets colder with height, so air warmed near the ground rises and carries pollution up and away. On a clear, calm night the ground cools quickly and chills the air just above it, which leaves a pool of cold air sitting under warmer air. This is called an **inversion**. The cold air is heavier than the air above it and cannot rise, so the warmer air acts as a lid. That night the smelter's SO_2 stayed in the pool of cold

air on the valley floor. With almost no wind to carry it out of the valley, its concentration built up through the night.

The facility's anemometer averaged 1 m/s from the west over the night, but the wind was light and swung through every direction.

Twelve monitors on the valley floor recorded the night's average SO₂ concentration. Their readings are in `data/receptors.csv`, with positions measured from the stack: `downwind_km` is the distance east of the stack (negative values are west of it), and `crosswind_m` is the distance north.

The facility's consultant modelled the night with the Gaussian plume model below, using stability class F and the night's average wind, 1 m/s from the west. The stack's effective height, 41 m including plume rise, is known. Your job is to decide whether that model is fit to be used.

The stability classes run from A (strong daytime heating) through D (neutral: overcast or windy) to F (a clear, calm night).

```
# Pasquill-Gifford coefficients for each stability class. Distances in km,  
# sigmas in metres.
```

```
stability_coefs = Dict(  
    "A" => (; a=213, c_near=440.8, d_near=1.941, f_near=9.27,  
              c_far=459.7, d_far=2.094, f_far=-9.6),  
    "B" => (; a=156, c_near=106.6, d_near=1.149, f_near=3.3,  
              c_far=108.2, d_far=1.098, f_far=2),  
    "C" => (; a=104, c_near=61, d_near=0.911, f_near=0,  
              c_far=61, d_far=0.911, f_far=0),  
    "D" => (; a=68, c_near=33.2, d_near=0.725, f_near=-1.7,  
              c_far=44.5, d_far=0.516, f_far=-13),  
    "E" => (; a=50.5, c_near=22.8, d_near=0.678, f_near=-1.3,  
              c_far=55.4, d_far=0.305, f_far=-34),  
    "F" => (; a=34, c_near=14.35, d_near=0.740, f_near=-0.35,  
              c_far=62.6, d_far=0.180, f_far=-48.6))  
class_names = ["A", "B", "C", "D", "E", "F"]
```

```
sigma_y(downwind_km, coef) = coef.a * downwind_km^0.894
```

```
function sigma_z(downwind_km, coef)  
    if downwind_km < 1  
        return coef.c_near * downwind_km^coef.d_near + coef.f_near  
    else  
        return coef.c_far * downwind_km^coef.d_far + coef.f_far  
    end  
end
```

```
# Ground-level concentration [µg/m³]. The wind blows along the downwind axis,  
# so the model puts nothing upwind of the stack.
```

```
function ground_concentration(downwind_km, crosswind_m, coef;  
                             emission=100, height=41, wind=1)  
    if downwind_km <= 0.05  
        return 0.0  
    end  
    spread_y = sigma_y(downwind_km, coef)
```

```

spread_z = sigma_z(downwind_km, coef)
if spread_z <= 0
    return 0.0
end
return (emission / (pi * wind * spread_y * spread_z)) *
    exp(-crosswind_m^2 / (2 * spread_y^2)) *
    exp(-height^2 / (2 * spread_z^2)) * 1e6
end

monitors = CSV.read("data/receptors.csv", DataFrame)

```

Problem 3.1 (6)

Run the consultant's model (class F, 1 m/s) at each monitor and compute the RMSE against the observations. Report it both in $\mu\text{g}/\text{m}^3$ and as a percentage of the mean observed concentration.

On that number alone, would you say the model fits?

Problem 3.2 (9)

Wind speed is the most uncertain input here, and the whole concentration field scales as $1/u$. It is the obvious thing to calibrate.

Keeping class F, fit the wind speed by minimising RMSE against the observations. Search between 0.3 and 20 m/s.

1. Report the fitted wind speed to the nearest 0.1 m/s and the RMSE after calibration.
2. Did calibration meaningfully improve the fit?
3. Explain the result in terms of what wind speed *can* and *cannot* change about the modelled field.

```

function model_rmse(coef, wind_speed)
    # TODO: run the model at every monitor with these stability coefficients
    # and this wind speed, and return the RMSE against the observed
    # concentrations.
    error("model_rmse has not been written yet -- replace this line")
end

```

💡 One-variable search in Optim

`Optim.optimize(f, lower, upper)` minimises a function `f` of one number between two bounds. Here `f` can be `u -> model_rmse(stability_coefs["F"], u)`. `Optim.minimizer` returns the best value found and `Optim.minimum` the function's value there.

Problem 3.3 (10)

Stability class is the other input the consultant chose. Calibrate it as well: for each of the six classes, fit the wind speed as in Problem 3.2.

1. Report the fitted wind speed and the RMSE for each class in a table. Which class fits best, and what is its RMSE as a percentage of the mean observed concentration?
2. For that best fit, list the modelled and observed concentrations at each monitor. At which monitors can no wind speed or stability class bring the model close to the observation, and why?

3. Was class F the wrong choice for this night? Answer from the conditions each class describes, not from the RMSE.

Problem 3.4 (10)

You are asked to certify the consultant's model as fit for deciding where to put a permanent compliance monitor in this valley.

1. Would you certify it? State your decision plainly.
2. Which assumptions of the plume model do not hold on a night like this one? Name two, and tie each to something in the setting or in your results.
3. The consultant replies that the model has now been calibrated, so it is "as good as it can be." Respond to that argument.

References

List any external references consulted, including classmates.