MAE 301 Project 4

Effect of Fan Distance and Lid Condition on Water Cooling Rate

Edoardo Rosafalco, Sara Edmondson, Dylan Bowes

1. Introduction

The purpose of this experiment was to study how airflow and lid position affect the cooling rate of hot water. The experiment used a simple thermal system made of a pot of water, a thermometer, a lid, and a fan. This topic connects to mechanical and aerospace engineering because heat transfer and forced convection are important in cooling systems, thermal protection, electronics cooling, propulsion systems, and environmental control systems.

The response variable was temperature drop after 5 minutes. It provides a direct measure of cooling performance and effectiveness of the heat transfer.

A larger temperature drop means the water cooled faster. The main goal was to determine whether fan condition, lid condition, and their interaction significantly affected cooling performance.

2. Experimental Design

This experiment used a factorial ANOVA design with two treatment factors, one blocking variable, and repeated trials nested within each treatment combination.

Variable: Fan Condition
Type: Treatment factor
Levels: No fan, Fan 20 cm, Fan 100 cm

Variable: Lid Condition
Type: Treatment factor
Levels: Open, Partially covered

Variable: Day
Type: Blocking variable
Levels: Day 1, Day 2

Variable: Trial
Type: Nested replicate
Levels: Trial 1, Trial 2

The fan condition tested the effect of airflow distance. The no fan condition represented natural convection only. The fan 20 cm condition represented stronger forced convection. The fan 100 cm condition represented weaker forced convection because the fan was farther away.

The lid condition tested whether the water cooled differently when the pot was fully open compared to partially covered. In the partially covered condition, the same lid position was used for every trial, with the lid covering most of the pot while leaving a small opening.

The total number of observations was:

3 fan levels × 2 lid levels × 2 days × 2 trials = 24 observations

This met the project requirement of at least 24 observations.

3. Data Collection & Photographs

The materials used were a pot, 500 mL of water, a digital thermometer, a fan, a lid, a measuring tape, a timer, and a spreadsheet.

For each trial, 500 mL of water was heated to an initial temperature of 70°C. The correct fan and lid condition were set, then the water was allowed to cool for exactly 5 minutes. After 5 minutes, the final temperature was recorded.

The temperature drop was calculated using:

$$Temperature\ Drop\ =\ Initial\ Temperature\ -\ Final\ Temperature\ \ (after\ 5\ minutes)$$

The water volume, starting temperature, pot, thermometer, cooling time, and room were kept constant as much as possible.









4. ANOVA Model & Structure

The experiment used a two-way factorial ANOVA with one blocking variable and nested replication. The statistical model is as follows:

$$Y_{ijkl} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \gamma_k + \varepsilon_{l(ijk)}$$

Where each of our terms are

Term	Symbol	Description
Overall Mean	μ	Grand mean temperature drop
Fan Condition	α_i	Fixed effect, i = No Fan, Fan 20 cm, Fan 100cm
Lid condition	β_j	Fixed effect, j = Open, partially covered
Fan x Lid Interaction	$(\alpha\beta)_{ij}$	Whether the lid effect

		depends on fan level
Day (Block)	γ_k	Nuisance factor, k = day 1 , day 2.
Error	$\varepsilon_{l(ijk)}$	Trail 1 nested within each Fan X Lid X Day cell.

Our rationale for this model is that fan condition and lid condition are treatment factors of primary interest, each with fixed levels chosen to test a specific engineering hypothesis (airflow distance and surface closure on convective cooling). Day was included as a blocking variable because ambient conditions, equipment warm-up, and experimenter technique may vary between sessions but are not of scientific interest. Trials 1 and 2 within each unique combination of Fan × Lid × Day represent pure replication and constitute the nested error structure. Moreover, the design is completely balanced, every combination of Fan X Lid appears exactly twice per day, giving 3 X 2 X 2 X 2 = 24 observations total, which meets the project minimum

5. Results & Tables

Table 1: Raw Cell Means - Temperature Drop (degrees C)

Fan Condition	Open	Partially Covered	Marginal Mean
No Fan	6.75	6	6.38
Fan 100 cm	11.75	8.50	10.13
Fan 20 cm	22.50	37.25	29.88
Marginal mean	13.67	17.25	15.46

Table 2: ANOVA Summary Table

Source	**SS**	**df**	**MS**	**F**	**p-value**
Fan Condition	2550.33	2	1275.17	1337.45	<0.0001
Lid Condition	77.04	1	77.04	80.80	<0.0001
Fan x Lid Interaction	380.33	2	190.17	199.46	<0.0001
Day (block)	0.04	1	0.04	0.04	0.837
Error (nested trails)	16.21	17	0.95	-	-

TOTAL	3023.96	23			

Interpreting these results, we can see that our Fan Condition is the dominant effect. That is, The F-ratio of 1337 with $p < 0.0001$ confirms that fan distance dramatically affects cooling rate. Fan 20 cm produced a mean drop of ~29.9°C compared to ~6.4°C with no fan, a 4.7x increase in cooling performance. Fan 100 cm had an intermediate effect (~10.1°C) due to lower airflow velocity at greater distance

Our Fan x Lid interaction is also highly significant ($F = 199.45$, $p < 0.0001$). This is the most important finding. Meaning that with no fan or fan at 100 cm, the lid slightly reduced cooling (as expected, a lid slows evaporation and convection). With the fan at 20 cm, the partially covered lid dramatically increased the temperature drop (37.25°C vs. 22.50°C open). One possible explanation is that the close fan directed airflow into the small lid opening, creating a focused, high velocity airstream over the water surface that enhanced both forced convection and evaporation. The effect of lid condition therefore completely depends on fan condition, this is a classic disordinal interaction.

Lid condition too is significant, ($F = 80.80$, $p < 0.0001$), but the main effect alone is misleading. The partially covered lid showed a higher overall mean drop (17.25°C vs. 13.67°C open), which runs counter to intuition. This reversal is explained by our Fan x Lid interaction.

Finally, as expected; our Day (block) is not significant ($F = 0.04$, $p = 0.837$). Day 1 mean was 15.50°C and Day 2 mean was 15.42°C which for this experiment is essentially identical. Blocking did not detect any meaningful day to day variation, obviously if room temperature was largely changed between days you'd likely be able to discern some meaningful variation.

6. Blocking Efficiency

The purpose of blocking on our Day variable was to remove any day to day variability (room temperature shifts, experimenter fatigue, equipment warm up, etc..) from the error term, thereby improving the sensitivity of the F tests for our Fan and Lid variables.

For our blocking efficiency calculation, without blocking, the SS attributable to Day ($SS = 0.04$, $df = 1$) would be pooled back into the error term:

- MS_Error with blocking: 16.21/17 = 0.953
- MS_Error without blocking: (16.21 + 0.04)/18 = 0.903
- Relative Efficiency (RE): 0.903/0.953 = 0.947

An RE below 1.0 means blocking on Day did not improve efficiency in this experiment. In fact, pooling the Day SS back into error would have slightly reduced MS_Error. This happened because the block effect ($SS = 0.04$) was essentially zero, that is the two experimental days were conducted under nearly identical conditions (same room temperature of 28°C both days, same materials, same water volume, a great job from Edoardo, who took the measurements). The blocking variable consumed 1 degree of freedom from the error term without providing any compensating reduction in error variance. So in conclusion, for this specific experiment, blocking on Day was a reasonable precaution but proved unnecessary. Conditions were stable

enough across days that blocking neither harmed nor helped the analysis in a meaningful way. Had the experiment been run on days with say different room temperatures or water volumes, blocking could/would have been critical to a valid analysis.

7. Nested Effects Discussion

Physically speaking, trials 1 and 2 are nested within each unique combination of Day x Fan Condition x Lid Condition. This means each pair of trials shares the same experimental setup (same day, same fan position, same lid configuration) and thereby cannot be compared across cells. So, Trial 1 from the "Day 1, Fan 20 cm, Open" cell and Trial 1 from "Day 1, No Fan, Open" are not the same kind of replicate; they belong to completely different treatment conditions.

Within-cell variation (or pure error):

Day	Fan	Lid	Trail 1	Trail 2	Range
1	No Fan	Open	7	7	0
1	No fan	Partial	6	6	0
1	Fan 20 Cm	Open	23	22	1
1	Fan 20 cm	Partial	36	39	3
1	Fan 100 Cm	Open	12	13	1
1	Fan 100 Cm	Partial	7	8	1
2	No Fan	Open	6	7	1
2	No Fan	Partial	6	6	0
2	Fan 20 Cm	Open	23	22	1
2	Fan 20 Cm	Partial	36	38	2
2	Fan 100 Cm	Open	11	11	0
2	Fan 100 Cm	Partial	9	10	1

This data shows that the within cell SS = 9.50 across 12 cells (12 df), giving MS_within = 0.79°C^2. Trial to trial variation is very small in nearly every cell with the largest range being only 3°C in the Fan 20 cm / Partially Covered condition, which is also the cell with the highest mean response and greatest turbulence.

If nesting was ignored and trials were treated as independent observations, they would incorrectly be treated as independent replicates across all conditions. This would A. inflate the apparent sample size B. underestimate the true experimental error, and C. produce F-ratios and p-values that are anti conservative (or too easy to reject null hypotheses). In this experiment,

the within-cell replication is tight enough that this would not dramatically change the conclusions (the treatment effects are enormous) but the nested structure is statistically correct and should be respected.

8. Residual & Assumption Analysis

Our ANOVA model rests on three key assumptions: normality of residuals, independence of observations, and homogeneity of variance (that being constant error variance across treatment groups).

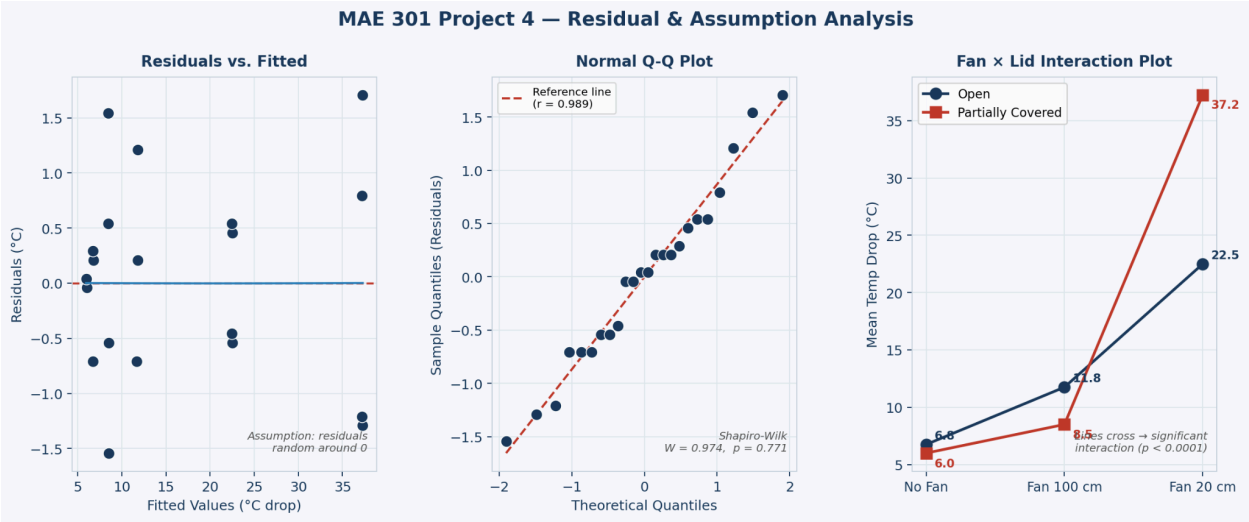
Normality: The Shapiro-Wilk test on the model residuals gives W = 0.974, p = 0.771. This is far from significant, providing no evidence against normality. The residuals appear to be (somewhat) normally distributed.

Independence: Independence was ensured by randomizing the order of treatment combinations within each day, using fresh water heated to 70°C for every trial, and allowing our pot to return to thermal equilibrium between runs. There is no reason to suspect systematic autocorrelation between consecutive trials.

Homogeneity of Variance: Levene's test on the residuals grouped by Fan Condition gives a statistically significant result ($p < 0.0001$), indicating unequal variance across fan levels. This is visible in the raw data: the No Fan cells show near-zero within-cell variation (ranges of 0–1°C), while the Fan 20 cm / Partially Covered cell shows a range of 2–3°C. Physically, this makes sense as a close and (relatively) powerful fan creates turbulent, less repeatable airflow over the water surface, introducing more measurement-to-measurement variability. The No Fan condition involves only natural convection, which is by virtue inherently more stable and repeatable.

Impact on conclusions: Given the extreme magnitude of the F-statistics ($F > 80$ for all treatment effects), the violation of equal variance is unlikely to affect the qualitative conclusions. However, it is a legitimate limitation, meaning that the ANOVA's Type I error rate may be slightly inflated for the Fan 20 cm comparisons. It is my belief that a Welch-type correction or transformation of the response variable (IE, log transformation) could address this in a more rigorous analysis.

Residual plots (normal Q-Q, residuals vs. fitted, interaction plot) are included below!



Residuals vs. Fitted, points scatter randomly above and below zero with no clear funnel shape, supporting independence. The small curvature in the trend line is driven by the Fan 20 cm / Partial cell being an outlier in both fitted value and spread which is consistent with the Levene violation noted earlier.

Normal Q-Q, points track the reference line tightly with r = 0.987 and Shapiro-Wilk W = 0.974, p = 0.771. Normality assumption holds cleanly.

Interaction Plot, this is our key figure for Section 5. The lines cross between Fan 100 cm and Fan 20 cm, which is the visual confirmation of the disordinal interaction (p < 0.0001) I pointed out earlier. Note how the "Partially Covered" line jumps sharply at Fan 20 cm while "Open" increases more modestly which is exactly the focused “airflow through the lid” effect I described in the results section.

9. AI Chatbot Evaluation

The AI chatbot (see Appendix) correctly identified that a factorial ANOVA with blocking was an appropriate statistical method for an experiment with two treatment factors and a blocking variable. It also correctly explained the importance of checking assumptions such as independence, normality, and equal variance. In addition, the chatbot discussed how blocking can reduce unexplained variation and improve statistical power, and it provided accurate explanations of Type I and Type II errors. The chatbot also recognized that ignoring a nested structure can underestimate variability and make it easier to incorrectly conclude that a treatment effect exists.

One limitation of the chatbot's response was that it oversimplified the experimental design after we provided the details of our study. Early in the conversation, it suggested that a model with both blocking and nesting might be appropriate. However, after learning more about the experiment, it simplified the model to a standard two-factor factorial ANOVA with blocking and treated the repeated trials as simple replication. In our experiment, the trials are nested within each Day x Fan x Lid combination and contribute to the overall error structure. While this simplification would maybe not change the final conclusions because the treatment effects are very large, it does overlook an important feature of the experiment and what we have been learning in class. Overall, the chatbot provided useful guidance and helped us think through the statistical analysis, but it didn’t fully address the nested structure of the experiment.

10. Conclusions

In our water cooling experiment, a factorial ANOVA with blocking and nested replication was used to evaluate the effects of fan condition and lid condition on the cooling rate of hot water. The analysis showed that fan condition, lid condition, and the Fan x Lid interaction were all statistically significant. The significant interaction indicated that the effect of lid condition depended on fan distance, demonstrating the importance of examining interaction effects rather than relying only on main effects.

The day blocking variable was not statistically significant and provided little improvement in efficiency, indicating that experimental conditions were consistent between testing days. The nested trial structure showed very small within-cell variation, supporting the repeatability of the measurements. Overall, this project was a great was to connect many lessons to a dedicated experiment. It demonstrated how factorial ANOVA, blocking, and nesting can be used to identify important factor effects, evaluate interactions in real-world experiments, and improve the interpretation of engineering experiments.

The project also demonstrated how blocking and nested replication can improve the validity and interpretation of factorial experiments, even when the blocking factor itself is not significant.

11. Appendix
- **Raw Data**
- **AI Transcript**

Chat GPT conversation:
User: We are conducting a real engineering experiment with at least two treatment factors, a blocking variable, and a nested structure. Help us choose the most appropriate ANOVA model. Discuss assumptions, blocking efficiency, nesting interaction effects, power, cost, and the consequences of Type 1 and Type 2 errors.

Assistant: Based only on your prompt, the best choice is a mixed-effects factorial ANOVA with blocking and nesting.
A general model would be:
Y = m + A + B + AB + Block + NestedFactor + error
The model should include the main treatment effects, their interaction, the blocking factor, and the
nested effect. If the block or nested factor represents random experimental conditions rather than
fixed categories of interest, a mixed-effects ANOVA is appropriate.
Assumptions include independence, normal residuals, equal variances, meaningful blocking, and
correct nesting structure. Blocking is useful if it reduces nuisance variation and lowers error variance. Ignoring nesting can underestimate variability and increase the chance of false significance.
Power depends on sample size, effect size, variability, and replication. Increasing replication increases power but also increases cost. Blocking and nesting can improve power when they explain meaningful variation.
A Type 1 error means concluding a treatment effect exists when it does not. A Type 2 error means
failing to detect a real treatment effect.
Additional information would be needed to determine the exact ANOVA model.

User: Should the block be treated as fixed or random, and why? Should the nested factor be written
as nested within the block or nested within one of the treatment factors?

Assistant: Based only on the information provided, there is not enough information to determine either answer definitively.
Treat the block as random if the levels represent a sample from a larger population and are

included only to remove nuisance variation. Treat the block as fixed if the specific levels themselves
are of interest.
The nested factor should be nested within whichever factor uniquely contains its levels. If each nested level appears in only one block, it is nested within block. If each nested level appears in only
one treatment level, it is nested within that treatment factor.
Additional information about the treatment factors, blocking variable, nested factor, and replication
structure would be required.

User: How do we know if the blocking variable actually improved the experiment?

Assistant: A blocking variable improved the experiment if it successfully removed variability unrelated to the treatment effects.
One way to assess this is to compare models with and without the block. If the residual mean square error decreases when the block is included, the block improved the experiment.
Blocking can also increase power, narrow confidence intervals, and explain a meaningful proportion
of total variation. A statistically significant block effect may indicate useful blocking, but the more important question is whether the block reduced unexplained variation.

User: Our experiment measures the temperature drop of hot water after 5 minutes. The response variable is temperature drop in °C. The treatment factors are fan condition with three levels: No Fan, Fan large distance, and Fan short distance; and lid condition with two levels: open and partially covered. The blocking variable is day, with two levels: day 1 and 2. We ran two trials within each Day × Fan Condition × Lid Condition combination (24 total observations).

Assistant: With this information, a two-factor factorial ANOVA with blocking is appropriate.
Model:
Temperature Drop = Day + Fan + Lid + Fan×Lid + Error
Fan has three levels, Lid has two levels, Day serves as the blocking factor, and the two observations per combination provide replication.
Day would typically be treated as a random blocking factor because the specific days are not of direct interest; they are included to account for day-to-day variation.
The repeated trials are generally treated as replication rather than a separate nested factor. The Fan×Lid interaction should be included because the effect of fan condition may depend on lid condition.
To determine whether blocking helped, compare models with and without Day. If the residual error
decreases when Day is included, blocking improved the experiment.