

Lesson 48: Final Presentations

Data Science with Python – BSc Course

Data Science Program

BSc Course

45 Minutes

Previously in L47...

Bias enters at
every pipeline stage

Fairness metrics
often conflict

SHAP/LIME for
explainability

Today: Show the world what you built – with honesty, clarity, and pride.

Ethics is the last lesson before showtime – because it matters most

Learning Objectives

Today's Goals:

- Deliver a compelling 5-minute project presentation
- Demo your deployed application live
- Handle Q&A with confidence and honesty
- Reflect on 48 lessons of learning

Schedule:

- Quick presentation guidelines (5 min)
- Student presentations (5 min each + 2 min Q&A)
- Peer evaluation during each presentation
- Course wrap-up and next steps

This is your opportunity to showcase 48 lessons worth of learning

5 Minutes to Convince the Room Your Model Matters

The Challenge:

- You have 5 minutes. Not 6. Not 7. Five.
- In those 5 minutes, you must: explain the problem, show data, present results, demo the app
- If you cannot explain it in 5 minutes, you do not understand it well enough

The Audience Decides in 30 Seconds:


- Strong opening: “Every year, banks lose 2B to loan defaults. My model catches 89% of them.”
- Weak opening: “So, um, for my project I decided to look at credit data...”

The Grading Criteria (15 points):

- Clarity: Can a non-expert understand your project? (5 pts)
- Technical depth: Did you demonstrate ML understanding? (5 pts)
- Demo quality: Does the deployment work? (5 pts)

Practice tip: Time yourself. If your run-through takes 7 minutes, cut 40% of content.

Presentation Structure



Hook → Problem → Data → Model → Demo → Lessons. Memorize this.

Slide Design Principles

Less Is More

- Maximum 3–5 slides (not counting title and demo)
- One idea per slide, one chart per slide
- Font size ≥ 24 pt (if your audience squints, the font is too small)
- Dark text on light background (not the other way around)

What to Show:

- One EDA chart that reveals something interesting
- Confusion matrix or metrics comparison table
- Screenshot of your deployed app (backup for live demo)

What NOT to Show:

- Code (ever – show results, not implementation)
- Walls of text (if you must read your slides, you have too much text)
- Every chart you made (pick the one that tells the story)

Live Demo Best Practices

Making Your Demo Work

- Test your demo 5 minutes BEFORE your presentation
- Have the app/API already running when you start
- Use pre-filled example inputs (do NOT type during demo)
- Show 2–3 different predictions to demonstrate flexibility

Streamlit Demo Flow:

1. Show the dashboard landing page
2. Adjust input sliders and explain each one
3. Show prediction updating in real time

FastAPI Demo Flow:

1. Open Swagger docs page (/docs)
2. Send a sample request with example data
3. Show the JSON response and explain the prediction

Backup plan: If demo crashes, switch to screenshots and say “Here is what it does”

Handling Q&A

Common Questions You WILL Be Asked

- “Why did you choose this model over alternatives?”
- “How did you handle missing data / class imbalance?”
- “What is your model’s biggest weakness?”
- “Could this model be biased? How did you check?” (L47)
- “What would you do differently with more time?”

How to Handle Questions Well:

- Repeat the question (buys thinking time, ensures everyone heard)
- Answer honestly – “I did not test that” beats making something up
- Connect to course: “As we learned in L27, accuracy is misleading for imbalanced data”
- “Great question, I would explore that as future work” (when genuinely unsure)

Good Q&A demonstrates understanding better than any slide can

Checkpoint: The #1 Presentation Mistake

What Is the Single Biggest Mistake Presenters Make?

It is NOT:

- Going over time (that is #2)
- Demo failure (that is #3)
- Bad slides (that is #4)

The #1 mistake: Not explaining WHY anyone should care.

- Bad: “I built a credit scoring model with 87% F1 score”
- Good: “Every year, banks lose 2B to defaults. My model catches 87% of them before they happen, potentially saving 1.7B”

Before your first slide, answer: “So what? Why does this matter?”

Start with impact, not implementation. Your audience is not a Python interpreter.

Telling the Data Story

The Narrative Arc of a Data Science Project

Every good story has:

1. **Setting:** “In the lending industry, 5% of loans default...”
2. **Conflict:** “Traditional methods miss 40% of defaults...”
3. **Journey:** “We explored 10,000 historical loans and found...”
4. **Resolution:** “Our model catches 89% of defaults”
5. **Implication:** “Saving 1.5M annually”

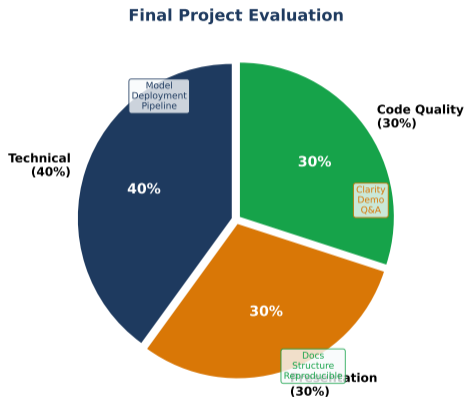
Sound like: “I noticed X, so I tried Y, and discovered Z”

NOT like: “First I did A, then B, then C, then D”

Data science tells stories with evidence. Make yours compelling.

Visualizing Results Effectively

One Chart That Tells Your Story



Pick the ONE visualization that best summarizes your key finding

Audience Adaptation

Your Classmates Are Your Audience Today. But Think Bigger.

Classmates (Today):

- Know ML terminology – use jargon freely
- Will ask about model comparison and evaluation

Job Interview (Tomorrow):

- May not know ML – explain from first principles
- Will ask “How would this scale?” and “What is the business case?”

Executive (Future):

- Has 2 minutes of attention – lead with the punchline
- Will ask “How confident are you?” and “What could go wrong?”

Practice explaining your project to a non-technical friend. If they get it, you are ready.

Time Management

5 Minutes Goes Faster Than You Think

Common Time Traps:

- Spending 2 minutes on the problem (should be 30 seconds)
- Walking through code (never do this)
- Explaining every chart you made (pick one)
- Starting the demo after 4 minutes (too late)

Time Budget:

- 0:00–0:30 **Hook:** Why this matters (one sentence)
- 0:30–1:00 **Data:** Source, size, one key insight
- 1:00–2:00 **Model:** What, why, key result
- 2:00–3:30 **Demo:** Live application
- 3:30–4:30 **Results:** Key metric, comparison
- 4:30–5:00 **Close:** Lessons learned, one sentence

Practice with a timer. If you run over in rehearsal, you will run over in the presentation.

Finance: Presenting to Stakeholders

Rules for Finance Presentations

Always Include:

- Monetary impact (“saves X per year” or “reduces risk by Y%”)
- Benchmark comparison (beat the S&P 500? Beat manual review?)
- Risk acknowledgment (what happens when the model is wrong?)

Never Say:

- “My model predicts the stock market” (overconfident)
- “100% accurate on test data” (overfitting or data leakage)
- “No limitations” (every model has limitations)

Power Phrase: “Based on historical data from [period], our model achieves [metric], suggesting [business implication].”

In finance, credibility comes from honesty about uncertainty, not from big promises

Peer Evaluation

Rate Each Presentation (Anonymous)

For each presenter, score 1–5 on:

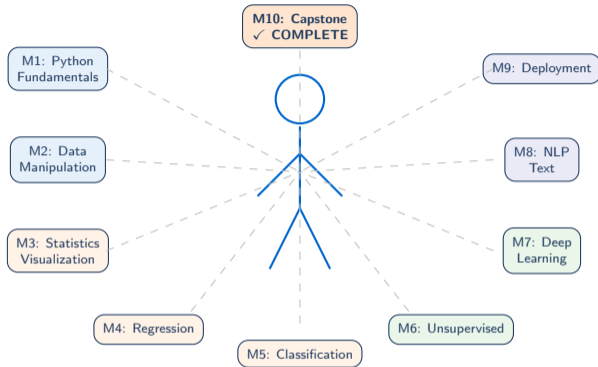
1. **Problem clarity:** Did you understand what they were solving?
2. **Technical quality:** Appropriate model and evaluation?
3. **Visualization:** Clear, informative charts?
4. **Demo:** Working deployment that adds value?
5. **Presentation style:** Clear, confident, within time?

Constructive Feedback:

- Write one thing they did well (be specific)
- Write one suggestion for improvement (be constructive)
- “Your confusion matrix clearly showed the model’s strength” > “Good job”

Give the kind of feedback you would want to receive

48 Lessons. One Data Scientist.



You are a data scientist now.

48 lessons. 10 modules. From “Hello World” to deployed ML. Congratulations.

Course Retrospective

From Zero to Deployed ML in 48 Lessons

Module 1–2 (L01–L12): Python & Data Manipulation

- Python basics, pandas DataFrames, data cleaning, time series

Module 3 (L13–L20): Statistics & Visualization

- Descriptive stats, distributions, hypothesis testing, matplotlib/seaborn

Module 4–5 (L21–L28): Supervised Learning

- Linear/logistic regression, trees, evaluation, class imbalance

Module 6 (L29–L32): Unsupervised Learning

- K-Means, hierarchical clustering, PCA, ML pipelines

Module 7–8 (L33–L40): Deep Learning & NLP

- Neural networks, backprop, text processing, sentiment analysis

Module 9–10 (L41–L48): Deployment, Ethics, Capstone

48 lessons, one integrated skill set. The knowledge compounds.

What Comes Next After BSc

Where to Go From Here

Deepen Your Skills:

- Kaggle competitions (real problems with leaderboards)
- Fast.ai courses (practical deep learning, free)
- *Hands-On ML* (Geron), *Storytelling with Data* (Knaflic)

Build Your Portfolio:

- GitHub profile with 3–5 polished projects (README, clean code, deployed demo)
- Blog posts explaining your approach (shows communication skills)
- LinkedIn – post your projects, engage with the community

Finance-Specific:

- CFA Institute's certificate in ESG investing
- QuantConnect for algorithmic trading backtests
- Bloomberg Terminal skills (if available at your institution)

The best way to learn data science is to **DO** data science – keep building

The Goal Was Never to Memorize

Memorize formulas

Pass the exam

Get the grade

Learn how to learn.

Build things that matter.

Keep going.

Thank you for a great semester. Now go build something the world needs.

Final Takeaways

Congratulations!

You have completed a full data science curriculum:

- 48 lessons from Python basics to deployed ML models
- Hands-on experience with real finance data
- A portfolio project you can show to employers
- Understanding of ML ethics and responsible AI

Final Advice:

- Data science is a practice, not a body of knowledge – keep practicing
- Communication matters as much as technical skill
- Always ask “So what?” – models without business impact are just math
- Ethics first: Just because you CAN build something does not mean you SHOULD

Thank you for a great semester!

Memory: The goal was never to memorize – it was to learn how to learn. Keep going.