Data & AI Foundations

Ages 16-17 | Future Pros

Present an AI idea with evidence, limitations, and safeguards.

How to use this guide

Read each section, complete its four lab steps, then answer the questions in the online course. The online progress bar advances only after four correct answers.

Section 1: Analysis

LEARNING GOAL: Analysis. Your finished work should demonstrate: A labeled chart with a careful conclusion.

EXPLAIN THE CONCEPT: Analysis starts with a clear question and suitable data. A chart can show a pattern but may hide missing context. Think about where you might use this skill in a real or imaginary project. Describe the problem it helps you solve and how you would recognize a successful result.

ORIGINAL PROJECT LAB: 1. Choose a small public or fictional dataset. 2. Write one question it can answer. 3. Create a simple chart with labels. 4. Write one observation and one limitation. Finished result: A labeled chart with a careful conclusion.

Understand with an example

The idea. Analysis starts with a question, measurement definitions and suitable denominators. Separate descriptive results from explanations of cause. A changed average does not by itself reveal why behavior changed.

Get ready. Fictional practice scores before a change are 4,6,8; after it they are 6,8,10.

Worked example 1

Follow the reasoning. Means are 6 and 8, a difference of 2 points. This describes the supplied groups. It does not prove that the change caused the increase because group differences or other factors may contribute.

Worked example 2: change one thing

Compare the result. Use only the highest after-score 10 against the full before mean 6. That comparison changes the selection rule and exaggerates the apparent

improvement. Keep comparable inclusion criteria.

Try one action at a time

1. Define the comparison question.
2. Check units and inclusion rules.
3. Calculate both summaries.
4. State a descriptive conclusion and limitation.

A mistake to notice

Pause and check. A numerical difference is not automatically a causal effect.

Your own small challenge

Predict before trying. If both groups are the same learners, does that alone prove causation?

Check your thinking

No. Matching learners removes one difference, but practice, timing and other changes can still affect outcomes.

Show what you know. Point, draw, demonstrate, type, speak or use your usual communication tool. A trusted adult can help without choosing your answer.

GO DEEPER: Identify an assumption, test a safe alternative, and explain the trade-off. Keep a short before-and-after record that someone else could follow.

SAFETY CHECK: Apply identity security, social-engineering defense, secure coding, cloud controls, incident handling, data protection, and professional ethics. Use pretend data, never real passwords or private information. Stop and ask a trusted adult if anything seems unsafe.

REVIEW QUESTIONS: What is the main idea? Why does the first step matter? How would you detect a mistake? What evidence proves the result? What would you improve? Answer in your own words before taking the quiz.

PORTFOLIO REFLECTION: Save a drawing, sanitized screenshot, or short note describing the goal, your approach, your evidence, and one lesson learned. Expected result: A labeled chart with a careful conclusion.

Investigate with pictures and choices

A dashboard combines totals from different time windows.

What should be fixed?

1. Only the color scheme
2. Assume all totals are comparable
3. Align or explicitly label the windows

Worked reasoning

Answer: Align or explicitly label the windows. Comparisons require matching definitions and time populations.

Try a changed example

Write the population and units for each metric.

Picture and meaning matches

Analysis: Checks the data and question

AI model: Produces evaluated predictions

Responsible design: Considers access and effects

Ways to respond

Choose, point, draw, demonstrate, type, speak or use your usual communication tool.

A helper can operate the controls while you choose.

Counting support if useful

Use zero to ten counters or dots. Touch each once. The last count names the quantity. Rearrange the same set and check whether anything was added or removed.

Choose this foundation practice only when it fits the learner goal.

My notes: _____

Hands-on lab

1. Choose a small public or fictional dataset.
2. Write one question it can answer.
3. Create a simple chart with labels.
4. Write one observation and one limitation.

Finished result: A labeled chart with a careful conclusion.

Check your understanding

1. Which idea best describes Analysis?

A. An AI model finds patterns in examples and may make errors on unfamiliar cases.

B. Analysis starts with a clear question and suitable data.

C. A responsible solution protects privacy, checks fairness, and gives people a way

to question a result.

D. It has no purpose in this project.

2. What is the first step in this lab?

A. Write one question it can answer.

B. Create a simple chart with labels.

C. Choose a small public or fictional dataset.

D. Write one observation and one limitation.

3. What should you do after step two?

A. Write one question it can answer.

B. Write one observation and one limitation.

C. Choose a small public or fictional dataset.

D. Create a simple chart with labels.

4. Which result shows that your lab is finished?

A. A labeled chart with a careful conclusion.

B. A copied answer without checking it

C. A skipped lab

D. A blank page with no work

Answer review

1. B. Analysis starts with a clear question and suitable data.

2. C. Start with: Choose a small public or fictional dataset.

3. D. Step three is: Create a simple chart with labels.

4. A. The goal is: A labeled chart with a careful conclusion.

Section 2: AI

LEARNING GOAL: AI. Your finished work should demonstrate: A simple AI evaluation plan.

EXPLAIN THE CONCEPT: An AI model finds patterns in examples and may make errors on unfamiliar cases. Evaluate it with examples it has not seen. Think about where you might use this skill in a real or imaginary project. Describe the problem it helps you solve and how you would recognize a successful result.

ORIGINAL PROJECT LAB: 1. Describe a pretend prediction task. 2. Write two sample inputs and expected outputs. 3. Invent one case where the model might fail. 4. Explain how you would measure its errors. Finished result: A simple AI evaluation plan.

Understand with an example

The idea. Model evaluation needs separate development and evaluation examples,

defined labels and error counts. Accuracy alone may obscure a failure that matters most for the intended use.
Get ready. A fictional classifier predicts Positive for all ten cases. Only two cases are actually Positive.

Worked example 1

Follow the reasoning. There are two true positives and eight false positives, giving 2/10 accuracy. Predicting one label repeatedly is not evidence of understanding the task.

Worked example 2: change one thing

Compare the result. Predict Negative for every case instead. Accuracy becomes 8/10 but both actual Positive cases are missed. A higher overall score can still fail the task's detection goal.

Try one action at a time

1. Define positive and negative labels.
2. Count each outcome type.
3. Calculate overall correctness.
4. Inspect missed cases and consequences.

A mistake to notice

Pause and check. Do not confuse high accuracy on imbalanced data with successful detection.

Your own small challenge

Predict before trying. What is recall for actual Positive cases under the all-Negative rule?

Check your thinking

Zero out of two actual Positives are detected, so recall is 0%. Keep the denominator tied to actual Positive cases.
Show what you know. Point, draw, demonstrate, type, speak or use your usual communication tool. A trusted adult can help without choosing your answer.
GO DEEPER: Identify an assumption, test a safe alternative, and explain the trade-off. Keep a short before-and-after record that someone else could follow.
SAFETY CHECK: Apply identity security, social-engineering defense, secure coding,

cloud controls, incident handling, data protection, and professional ethics. Use pretend data, never real passwords or private information. Stop and ask a trusted adult if anything seems unsafe.

REVIEW QUESTIONS: What is the main idea? Why does the first step matter? How would you detect a mistake? What evidence proves the result? What would you improve? Answer in your own words before taking the quiz.

PORTFOLIO REFLECTION: Save a drawing, sanitized screenshot, or short note describing the goal, your approach, your evidence, and one lesson learned. Expected result: A simple AI evaluation plan.

Investigate with pictures and choices

A model performs well on familiar training data.

What should be tested next?

1. Only the model title
2. Separate representative evaluation data
3. The exact same examples only

Worked reasoning

Answer: Separate representative evaluation data. Held-out evidence helps estimate performance beyond tuning examples.

Try a changed example

Keep acceptance cases separate from development cases.

Picture and meaning matches

Analysis: Checks the data and question

AI model: Produces evaluated predictions

Responsible design: Considers access and effects

Ways to respond

Choose, point, draw, demonstrate, type, speak or use your usual communication tool. A helper can operate the controls while you choose.

Counting support if useful

Use zero to ten counters or dots. Touch each once. The last count names the quantity. Rearrange the same set and check whether anything was added or removed. Choose this foundation practice only when it fits the learner goal.

My notes: _____

Hands-on lab

1. Describe a pretend prediction task.
2. Write two sample inputs and expected outputs.
3. Invent one case where the model might fail.
4. Explain how you would measure its errors.

Finished result: A simple AI evaluation plan.

Check your understanding

1. Which idea best describes AI?

A. A responsible solution protects privacy, checks fairness, and gives people a way to question a result.

B. It has no purpose in this project.

C. Analysis starts with a clear question and suitable data.

D. An AI model finds patterns in examples and may make errors on unfamiliar cases.

2. What is the first step in this lab?

A. Describe a pretend prediction task.

B. Explain how you would measure its errors.

C. Write two sample inputs and expected outputs.

D. Invent one case where the model might fail.

3. What should you do after step two?

A. Describe a pretend prediction task.

B. Invent one case where the model might fail.

C. Write two sample inputs and expected outputs.

D. Explain how you would measure its errors.

4. Which result shows that your lab is finished?

A. A skipped lab

B. A blank page with no work

C. A simple AI evaluation plan.

D. A copied answer without checking it

Answer review

1. D. An AI model finds patterns in examples and may make errors on unfamiliar cases.
2. A. Start with: Describe a pretend prediction task.
3. B. Step three is: Invent one case where the model might fail.
4. C. The goal is: A simple AI evaluation plan.

Section 3: Responsible design

LEARNING GOAL: Responsible design. Your finished work should demonstrate: An ethical AI proposal slide.

EXPLAIN THE CONCEPT: A responsible solution protects privacy, checks fairness, and gives people a way to question a result. Think about where you might use this skill in a real or imaginary project. Describe the problem it helps you solve and how you would recognize a successful result.

ORIGINAL PROJECT LAB: 1. List who may be affected by your pretend model. 2. Remove private fields the task does not need. 3. Write a way to report a wrong result. 4. Present the idea and its limitations on one slide. Finished result: An ethical AI proposal slide.

Understand with an example

The idea. A responsible design links the intended use to error costs, privacy choices and human review. Avoid using a toy model for consequential real decisions about people. State where the system must stop or defer.

Get ready. A fictional study-helper labels notes by topic and sometimes cannot decide.

Worked example 1

Follow the reasoning. An Uncertain result can ask the learner to choose a topic instead of presenting an unsupported label as certain. Keep only the data needed for this practice task.

Worked example 2: change one thing

Compare the result. Force every uncertain note into a topic to make the interface look complete. The apparent completeness hides unresolved errors; preserve a review path.

Try one action at a time

1. Define the low-stakes task.
2. List likely errors and effects.
3. Minimize input data.
4. Design a visible uncertainty path.

A mistake to notice

Pause and check. Do not describe a planned review process as an operational safeguard without evidence.

Your own small challenge

Predict before trying. Does adding a human-review button guarantee responsible operation?

Check your thinking

No. Define who reviews, what evidence they see and what happens when no review occurs. A button alone does not establish an effective process.

Show what you know. Point, draw, demonstrate, type, speak or use your usual communication tool. A trusted adult can help without choosing your answer.

GO DEEPER: Identify an assumption, test a safe alternative, and explain the trade-off. Keep a short before-and-after record that someone else could follow.

SAFETY CHECK: Apply identity security, social-engineering defense, secure coding, cloud controls, incident handling, data protection, and professional ethics. Use pretend data, never real passwords or private information. Stop and ask a trusted adult if anything seems unsafe.

REVIEW QUESTIONS: What is the main idea? Why does the first step matter? How would you detect a mistake? What evidence proves the result? What would you improve? Answer in your own words before taking the quiz.

PORTFOLIO REFLECTION: Save a drawing, sanitized screenshot, or short note describing the goal, your approach, your evidence, and one lesson learned. Expected result: An ethical AI proposal slide.

Investigate with pictures and choices

An AI tool recommends an action on a private record.

What must still happen?

1. Independent resource authorization
2. Automatic access because AI asked
3. Only a spelling check

Worked reasoning

Answer: Independent resource authorization. A model suggestion does not establish permission to access or change a record.

Try a changed example

Compare a permitted and cross-tenant request with synthetic data.

Picture and meaning matches

Analysis: Checks the data and question

AI model: Produces evaluated predictions

Responsible design: Considers access and effects

Ways to respond

Choose, point, draw, demonstrate, type, speak or use your usual communication tool.

A helper can operate the controls while you choose.

Counting support if useful

Use zero to ten counters or dots. Touch each once. The last count names the quantity. Rearrange the same set and check whether anything was added or removed.

Choose this foundation practice only when it fits the learner goal.

My notes: _____

Hands-on lab

1. List who may be affected by your pretend model.
2. Remove private fields the task does not need.
3. Write a way to report a wrong result.
4. Present the idea and its limitations on one slide.

Finished result: An ethical AI proposal slide.

Check your understanding

1. Which idea best describes Responsible design?

A. An AI model finds patterns in examples and may make errors on unfamiliar cases.

B. It has no purpose in this project.

C. Analysis starts with a clear question and suitable data.

D. A responsible solution protects privacy, checks fairness, and gives people a way to question a result.

2. What is the first step in this lab?

A. List who may be affected by your pretend model.

B. Present the idea and its limitations on one slide.

C. Remove private fields the task does not need.

D. Write a way to report a wrong result.

3. What should you do after step two?

A. List who may be affected by your pretend model.

B. Write a way to report a wrong result.

C. Remove private fields the task does not need.

D. Present the idea and its limitations on one slide.

4. Which result shows that your lab is finished?

A. A skipped lab

B. A blank page with no work

C. An ethical AI proposal slide.

D. A copied answer without checking it

Answer review

1. D. A responsible solution protects privacy, checks fairness, and gives people a way to question a result.

2. A. Start with: List who may be affected by your pretend model.

3. B. Step three is: Write a way to report a wrong result.

4. C. The goal is: An ethical AI proposal slide.

Section 4: Guided Project & Worked Examples

COURSE INVESTIGATION: Thresholds trade false alarms against misses

YOUR PRACTICE INPUT: A fictional classifier reviews ten cases, of which four truly need attention. A threshold produces three true positives, two false positives, one false negative, and four true negatives.

WHY IT WORKS: Precision is 3/5 = 60%; recall is 3/4 = 75%. Accuracy is 7/10 = 70%. Different measures answer different questions.

PLAN BEFORE YOU ACT: Draw or write the starting state, the question you need to answer, and the evidence you expect to see. A paper trace is acceptable; label it as a prediction unless you actually run a test.

WORK THE CASE: Build both confusion matrices and write a human-review recommendation based on the cost of misses and false alarms.

COMPARE WITH THE EXPECTED RESULT: The comparison states counts and denominators, retains the lost true positive, and does not claim a universally best threshold.

EXPLAIN YOUR REASONING: Show the connection between the supplied input, each decision, and your result. If a number, label, or observation differs, return to the first differing step rather than changing several things together.

KEEP YOUR EVIDENCE: Save the initial plan, a drawing or sanitized output, and one explanation of what the result does and does not prove. Use pointing, speech,

drawing, typing, or your usual communication method.
LEARNING GOAL: Guided Project & Worked Examples. Your finished work should demonstrate: A completed guided project with documented evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.
EXPLAIN THE CONCEPT: Review the course concepts through two applied examples. Compare training, validation, and real-world data roles conceptually and identify leakage risk. Evaluate precision/recall tradeoffs in a simple fictional safety classifier without claiming one metric is always best. Guided project: Define an ethical AI solution using synthetic/public data, baseline analysis, model concept, evaluation measures, human review, privacy, and failure handling. Work with synthetic or fictional information and ask an adult for help when appropriate. Think about where you might use this skill in a real or imaginary project. Describe the problem it helps you solve and how you would recognize a successful result.
ORIGINAL PROJECT LAB: 1. Read both worked examples and explain the reasoning. 2. Define an ethical AI solution using synthetic/public data, baseline analysis, model concept, evaluation measures, human review, privacy, and failure handling. 3. Record the steps you completed and the result you observed. 4. Compare your result to the goal and correct one issue. Finished result: A completed guided project with documented evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Understand with an example

The idea. An evaluation report should connect a fixed test set, model outputs and a limited conclusion. Preserve inconvenient failures and identify changes to data or rules.
Get ready. Use ten fictional cases: two Positive and eight Negative. Compare always-Positive and always-Negative rules.

Worked example 1

Follow the reasoning. The first rule has 20% accuracy and detects both Positives. The second has 80% accuracy and detects neither. Show the counts alongside percentages so readers can evaluate the trade-off.

Worked example 2: change one thing

Compare the result. Change the test set to five Positive and five Negative cases.

Each constant rule now has 50% accuracy, showing how prevalence affects this metric. The rules themselves did not improve.

Try one action at a time

1. Define labels and test cases.
2. Build outcome counts.
3. Compare accuracy and missed Positives.
4. Change composition and explain the metric.

A mistake to notice

Pause and check. Do not quietly change the test population while attributing the score change to a better model.

Your own small challenge

Predict before trying. Can these ten invented cases establish a real model's production performance?

Check your thinking

No. They illustrate evaluation logic only. Real performance requires suitable representative data and an appropriate evaluation process.

Show what you know. Point, draw, demonstrate, type, speak or use your usual communication tool. A trusted adult can help without choosing your answer.

GO DEEPER: Identify an assumption, test a safe alternative, and explain the trade-off. Keep a short before-and-after record that someone else could follow.

SAFETY CHECK: Apply identity security, social-engineering defense, secure coding, cloud controls, incident handling, data protection, and professional ethics. Use pretend data, never real passwords or private information. Stop and ask a trusted adult if anything seems unsafe.

REVIEW QUESTIONS: What is the main idea? Why does the first step matter? How would you detect a mistake? What evidence proves the result? What would you improve? Answer in your own words before taking the quiz.

PORTFOLIO REFLECTION: Save a drawing, sanitized screenshot, or short note describing the goal, your approach, your evidence, and one lesson learned. Expected result: A completed guided project with documented evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Investigate with pictures and choices

A model is accurate overall but unreliable for one important subset.

What should the report show?

1. Only the overall score
2. A promise of fairness
3. Subset results and their uncertainty

Worked reasoning

Answer: Subset results and their uncertainty. Aggregate metrics can hide consequential differences and small samples.

Try a changed example

State the evidence required before using the model for the intended task.

Picture and meaning matches

Analysis: Checks the data and question

AI model: Produces evaluated predictions

Responsible design: Considers access and effects

Ways to respond

Choose, point, draw, demonstrate, type, speak or use your usual communication tool.

A helper can operate the controls while you choose.

Counting support if useful

Use zero to ten counters or dots. Touch each once. The last count names the quantity. Rearrange the same set and check whether anything was added or removed.

Choose this foundation practice only when it fits the learner goal.

My notes: _____

Hands-on lab

1. Read both worked examples and explain the reasoning.
2. Define an ethical AI solution using synthetic/public data, baseline analysis, model concept, evaluation measures, human review, privacy, and failure handling.
3. Record the steps you completed and the result you observed.
4. Compare your result to the goal and correct one issue.

Finished result: A completed guided project with documented evidence: Problem/data

brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Check your understanding

1. Which idea best describes Guided Project & Worked Examples?

A. Analysis starts with a clear question and suitable data.

B. Review the course concepts through two applied examples.

C. An AI model finds patterns in examples and may make errors on unfamiliar cases.

D. It has no purpose in this project.

2. What is the first step in this lab?

A. Define an ethical AI solution using synthetic/public data, baseline analysis, model concept, evaluation measures, human review, privacy, and failure handling.

B. Record the steps you completed and the result you observed.

C. Read both worked examples and explain the reasoning.

D. Compare your result to the goal and correct one issue.

3. What should you do after step two?

A. Define an ethical AI solution using synthetic/public data, baseline analysis, model concept, evaluation measures, human review, privacy, and failure handling.

B. Compare your result to the goal and correct one issue.

C. Read both worked examples and explain the reasoning.

D. Record the steps you completed and the result you observed.

4. Which result shows that your lab is finished?

A. A completed guided project with documented evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

B. A copied answer without checking it

C. A skipped lab

D. A blank page with no work

Answer review

1. B. Review the course concepts through two applied examples.

2. C. Start with: Read both worked examples and explain the reasoning.

3. D. Step three is: Record the steps you completed and the result you observed.

4. A. The goal is: A completed guided project with documented evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Section 5: Independent Challenge & Mastery

INDEPENDENT TRANSFER CASE: Thresholds trade false alarms against misses

KEEP THE ORIGINAL BASELINE: A fictional classifier reviews ten cases, of which four truly need attention. A threshold produces three true positives, two false positives, one false negative, and four true negatives.

CHANGE ONE CONDITION: Raising the threshold removes one false positive but also one true positive. Precision becomes 2/3, while recall falls to 2/4.
YOUR TASK: Build both confusion matrices and write a human-review recommendation based on the cost of misses and false alarms.
PREDICT, THEN CHECK: State which part of your original reasoning should remain true and which part must be recalculated or retested. Keep the original and changed results side by side.
REVIEW CRITERIA: The comparison states counts and denominators, retains the lost true positive, and does not claim a universally best threshold.
MENTOR DISCUSSION: Explain why the changed condition matters. Identify a tempting wrong conclusion, show the observation that would challenge it, and choose one next test. A statement of uncertainty is better than inventing a result.
PORTFOLIO NOTE: Keep your baseline, changed input, prediction, observed or paper-traced result, explanation, and one improvement. These practice prompts do not replace the course’s existing quiz or final assessment.
LEARNING GOAL: Independent Challenge & Mastery. Your finished work should demonstrate: An independent demonstration with evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.
EXPLAIN THE CONCEPT: Apply what you learned without copying the guided solution. Independent challenge: Perform a lightweight model-card review: intended use, excluded use, data limits, bias, errors, monitoring, and human escalation. Safety and responsibility: Never use private personal data without proper authority; validate AI output; document limitations; keep humans accountable for consequential decisions. Reflect on your work: What did you build or discover? What did you test or change after feedback? What safe or responsible choice did you make? Think about where you might use this skill in a real or imaginary project. Describe the problem it helps you solve and how you would recognize a successful result.
ORIGINAL PROJECT LAB: 1. Explain the challenge and make a short plan. 2. Perform a lightweight model-card review: intended use, excluded use, data limits, bias, errors, monitoring, and human escalation. 3. Test the result and record evidence of your work. 4. Review the safety scenario and describe one improvement. Finished result: An independent demonstration with evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Understand with an example

The idea. A data and AI portfolio case needs a bounded question, reproducible counts and explicit limits. Distinguish a manually applied rule from a trained model or deployed system.

Get ready. Invent eight fictional cases with three Positive labels and five Negative labels, then propose a simple rule.

Worked example 1

Follow the reasoning. An always-Negative baseline is correct on five of eight cases, or 62.5%, and misses all three Positives. Use this baseline to explain why correctness alone is insufficient.

Worked example 2: change one thing

Compare the result. A second rule correctly detects two Positives and all five Negatives. It has 7/8 accuracy, or 87.5%, while one Positive remains missed. Report the miss rather than only the improvement.

Try one action at a time

1. Publish fictional labels and predictions.
2. Compute outcome counts.
3. Compare against a baseline.
4. Explain the remaining error and limits.

A mistake to notice

Pause and check. A toy evaluation is not permission to deploy an automated decision about real people.

Your own small challenge

Predict before trying. Give Positive recall for the second rule as a fraction.

Check your thinking

It detects 2 of the 3 actual Positive cases, so recall is 2/3, approximately 66.7%.

Show what you know. Point, draw, demonstrate, type, speak or use your usual communication tool. A trusted adult can help without choosing your answer.

GO DEEPER: Identify an assumption, test a safe alternative, and explain the trade-off. Keep a short before-and-after record that someone else could follow.

SAFETY CHECK: Apply identity security, social-engineering defense, secure coding, cloud controls, incident handling, data protection, and professional ethics. Use pretend data, never real passwords or private information. Stop and ask a trusted adult if anything seems unsafe.
REVIEW QUESTIONS: What is the main idea? Why does the first step matter? How would you detect a mistake? What evidence proves the result? What would you improve? Answer in your own words before taking the quiz.
PORTFOLIO REFLECTION: Save a drawing, sanitized screenshot, or short note describing the goal, your approach, your evidence, and one lesson learned. Expected result: An independent demonstration with evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Investigate with pictures and choices

An agent calls a tool that changes data.
What belongs outside the model?
1. A longer personality description
2. Enforced permissions and exact-action controls
3. Only a request to behave

Worked reasoning

Answer: Enforced permissions and exact-action controls. Executable authorization and action controls cannot depend only on generated text.

Try a changed example

Use an advisory-only design when the action boundary is unverified.

Picture and meaning matches

Analysis: Checks the data and question
AI model: Produces evaluated predictions
Responsible design: Considers access and effects

Ways to respond

Choose, point, draw, demonstrate, type, speak or use your usual communication tool.
A helper can operate the controls while you choose.

Counting support if useful

Use zero to ten counters or dots. Touch each once. The last count names the

quantity. Rearrange the same set and check whether anything was added or removed.

Choose this foundation practice only when it fits the learner goal.

My notes: _____

Hands-on lab

1. Explain the challenge and make a short plan.
2. Perform a lightweight model-card review: intended use, excluded use, data limits, bias, errors, monitoring, and human escalation.
3. Test the result and record evidence of your work.
4. Review the safety scenario and describe one improvement.

Finished result: An independent demonstration with evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Check your understanding

1. Which idea best describes Independent Challenge & Mastery?

A. It has no purpose in this project.

B. Analysis starts with a clear question and suitable data.

C. Apply what you learned without copying the guided solution.

D. An AI model finds patterns in examples and may make errors on unfamiliar cases.

2. What is the first step in this lab?

A. Review the safety scenario and describe one improvement.

B. Perform a lightweight model-card review: intended use, excluded use, data limits, bias, errors, monitoring, and human escalation.

C. Test the result and record evidence of your work.

D. Explain the challenge and make a short plan.

3. What should you do after step two?

A. Test the result and record evidence of your work.

B. Perform a lightweight model-card review: intended use, excluded use, data limits, bias, errors, monitoring, and human escalation.

C. Review the safety scenario and describe one improvement.

D. Explain the challenge and make a short plan.

4. Which result shows that your lab is finished?

A. A blank page with no work

B. An independent demonstration with evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

C. A copied answer without checking it

D. A skipped lab

Answer review

1. C. Apply what you learned without copying the guided solution.
2. D. Start with: Explain the challenge and make a short plan.
3. A. Step three is: Test the result and record evidence of your work.
4. B. The goal is: An independent demonstration with evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Mastery & safe-use plan

Ages 16-17

Digital safety: Apply identity security, social-engineering defense, secure coding, cloud controls, incident handling, data protection, and professional ethics.

AI literacy: Use AI as an accountable engineering assistant: validate outputs, protect data, document use, test security, and keep humans responsible.

Challenge: Work through a realistic ticket, incident, architecture, data, or application scenario with professional evidence.

Portfolio: Produce a career-ready artifact such as a repository, incident report, architecture diagram, dashboard, automation, or technical brief.

Mastery check: Present the work as if to a mentor or interviewer: problem, approach, evidence, tradeoffs, security, and next improvement.

Project worksheet

My project: Present an ethical AI solution

Who will use or review it? _____

My plan: _____

Checkpoint 1: A labeled chart with a careful conclusion.

Checkpoint 2: A simple AI evaluation plan.

Checkpoint 3: An ethical AI proposal slide.

Checkpoint 4: A completed guided project with documented evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

Checkpoint 5: An independent demonstration with evidence: Problem/data brief, Evaluation table, Model-card style report, Responsible-AI presentation.

One thing I tested: _____

One thing I changed after testing: _____

What I learned: _____

Trusted adult or mentor feedback: _____

Family review

Show your finished work to a trusted adult. Explain one thing you learned and one thing you would improve.

Beat IT Kids | For learning at home. Do not put real passwords or private details in practice projects.