

Self-Assessment 3

Numerical Accuracy · Resolution · Convergence · Tolerance · Validation

Ungraded · for your own learning · work at your own pace



AI First Engineering

Figure: AI First Engineering (AIFE) course logo

Taner Onsay, Ph.D. & Sebnem Onsay · © 2026 AI First Engineering

AI TUTOR MODE — read this first

You do not need to complete every question in one sitting. Start with each group. If you can confidently solve the first questions independently, sample the remaining questions. Spend more time in groups where you need AI or human help. In Chapter 3 your AI advances from programming coach to **emerging engineering analyst** — its signature question is “Help me understand what the numerical results are telling me — not just produce more numbers.” If you become stuck, instead of asking for the answer, try:

“I am working on this problem. Do not solve it or compute the answer for me. Ask me one question or give me one small hint at a time. Help me reason through the approximation, the error, the resolution, the convergence evidence, the tolerance, or the validation until I can explain it myself.”

The purpose of this self-assessment is to determine what YOU understand — not what your AI can answer. **This is the working method in action: PREDICT › COMPUTE › INSPECT › EXPLAIN — the judgment stays yours.**

After each group, rate yourself with the four checkboxes. Your pattern of ratings — not a score — is the output: it tells you exactly what to review, practice, or bring to office hours — before Lab-3 asks you to defend an accuracy claim and Homework 3 asks you to certify one. A calculator is all you need — no computer, and no clock: work at your own pace. Every question here is one question in disguise: how accurate is accurate enough — and what evidence justifies that claim?

GROUP 1 · Explain It — Approximation, Reference & the Size of the Gap

One scenario, numbers in view. Answer in your own words — a sentence or a short calculation each. Keep every unit.

A pedestrian footbridge is being checked for how much it flexes under a full crowd. A numerical computation of the kind you built in this chapter reports a midspan deflection of 12.1 mm. A provided analytical solution for the same idealized beam model gives 12.5 mm. The beam’s stiffness and the crowd load are measured and estimated values.

1. One of these two numbers is an approximation; the other will serve as the reference value. Say which is which — and explain why an engineer calls 12.5 mm a reference value rather than “the exact answer,” even though it comes from exact algebra.

2. Compute the absolute error, the relative error, and the percent error of the numerical result — with units where they belong. Then: the same absolute error appears in a computation for a machining fixture whose reference deflection is 2.0 mm. Compute that percent error — and state in one sentence what the comparison shows about which error measure answers which question.

RATE YOURSELF on this group before moving on:

- I can do this independently.
- I can do this with a small AI hint.
- I need substantial AI help.
- I still need human/instructor help.

GROUP 2 · Predict It — Step Size, Resolution & the Grid

Commit each prediction in writing before reading past it. A prediction made after the evidence is a story — PREDICT → COMPUTE → INSPECT → EXPLAIN.

A cooling pump's maximum safe flow is found by this chapter's incremental search: raise the flow from zero in equal steps, test whether the computed pipe pressure at flow + step still respects the pressure limit, and report the last safe flow tested. In this model the pipe pressure rises steadily as the flow rises, so there is exactly one crossing of the limit. With a step of 10 L/min the search reports 140 L/min. A provided analytical solution for the same model gives a reference crossing of 143.7 L/min.

3. PREDICT, then inspect. (a) Before reading further, commit in writing: if the step is cut from 10 to 5 L/min, will the reported flow definitely change — YES / NO — and will the actual error definitely be cut in half — YES / NO? One sentence each. (b) Now the evidence: step 5 reports 140 L/min again; step 2 reports 142 L/min. Using the grid each search can see, explain how two different resolutions produced the identical 140 — and state what that agreement did and did not establish. (c) For this search, what bounds the actual error at each step size — and what is it about this particular search that makes the bound hold? Check the bound against the three errors measured against the reference crossing.

The same search, as code — read it like an engineer; do not run it:

```
q, step = 0.0, 10.0
while pipe_pressure(q + step) <= p_limit:
    q = q + step
print("Max safe flow:", q, "L/min")
```

4. (a) Why must every answer this program can print be a whole multiple of step? (b) State the resolution of the reported answer, in the sense this chapter gave that word. (c) When the program prints 140 at step = 10, between which two flows does the model's crossing necessarily lie — and on which side of the crossing does this search always err? Why is knowing that direction engineering information?

RATE YOURSELF on this group before moving on:

- I can do this independently.
- I can do this with a small AI hint.
- I need substantial AI help.
- I still need human/instructor help.

GROUP 3 · Read the Evidence — Convergence, Cost & Knowing When to Stop

One refinement table, read the way Part 4 taught you. The change column is computed from the answers alone — no reference value needed.

A crash-absorber spring is compressed by 0.5 m, and the energy W it absorbs is computed by summing force $\times$ slice width over small slices Δx — this chapter's stepping pattern, with Δx as the resolution knob. Five runs of the same computation, one knob changed:

Slice Δx (mm)	Computed W (J)	Change from previous (J)	Model evaluations (cost)
20	46.2	—	25
10	47.9	+1.7	50
5	48.7	+0.8	100
2	49.2	+0.5	250
1	49.3	+0.1	500

5. (a) Is the change column evidence consistent with convergence? Say what pattern you are pointing at. (b) Does the table prove that 49.3 J is correct — or even prove that the sequence is converging? State precisely what kind of claim this table can support. (c) What does each refinement cost — and could you decide, from this table alone, whether refining to $\Delta x = 0.5$ mm is justified? Name the one piece of information that is missing.

6. The design review supplies the requirement: before the absorber is certified, W must be known to within ± 1 J. (a) Who owns that ± 1 J — the computer, the AI, the programmer, or someone else — and what justifies it? (b) Based on the available refinement evidence and the ± 1 J requirement, state the Δx at which you would be willing to stop refining. Defend your choice — and name one additional piece of evidence that would strengthen the numerical-accuracy claim. (c) In one sentence: what is the difference between “the changes are shrinking” and a justified stopping criterion?

RATE YOURSELF on this group before moving on:

- I can do this independently.
- I can do this with a small AI hint.
- I need substantial AI help.
- I still need human/instructor help.

GROUP 4 · Judge It — Digits, Budgets & BREAK THE AI

Deliver each judgment with evidence, not vibes. A correct AI claim, verified, is also a finding.

7. Asked for the final answer, a teammate’s AI analyst reports: “ $W = 49.2971486$ J — the exact absorbed energy, to the full precision Python computed.” Using the Group 3 refinement table as your evidence: (a) which digits of that number are defensible, and why those; (b) what exactly is overclaimed — take the words “exact” and “precision” separately; (c) one sentence: why can no fixed decimal-places rule replace the evidence here?

8. A chilled-water cooling loop's flow estimate feeds an efficiency report. Your refinement study shows the numerical answer changing by less than ± 0.01 L/s between the last two refinements. The flow sensor's calibration record shows a measurement uncertainty of about ± 0.4 L/s. Your manager funds ONE improvement: a finer computation, or a recalibrated, better sensor. (a) Which do you fund? Justify it with the one comparison that matters. (b) A teammate objects: "a more accurate simulation is always better engineering." Respond in two sentences, using this case.

9. BREAK THE AI. An AI assistant analyzing a teammate's stepping computation of a safety gate's closing time reports: "I ran the model at two resolutions — step 0.2 s and step 0.1 s — and obtained exactly the same answer, 3.60 s, both times. The solution has therefore converged, and the result is accurate." Deliver your verdict — ACCEPT · MODIFY · REJECT — then critique the claim in three to five sentences: what could the agreement actually be? What evidence would "converged" need? What would "accurate" need beyond convergence? And what question about requirements has nobody asked?

RATE YOURSELF on this group before moving on:

- I can do this independently.
- I can do this with a small AI hint.
- I need substantial AI help.
- I still need human/instructor help.

GROUP 5 · Validate It — CONVERGED ≠ CORRECT

The chapter's second principle earns its keep: a beautifully converged answer can still solve the wrong engineering problem.

10. The footbridge computation from Group 1 is rerun by a teammate with a full refinement study: the changes shrink below 0.001 mm, and the deflection settles to five confident digits. A week later you discover the crowd load was typed in pounds-force where the model expects newtons — a factor of about 4.45. (a) Did the computation

converge anyway? Explain what the convergence evidence was — and was not — evidence about. (b) Name TWO specific pieces of independent validation evidence that could have caught this error, and say what each one checks that convergence cannot.

RATE YOURSELF on this group before moving on:

- I can do this independently.
- I can do this with a small AI hint.
- I need substantial AI help.
- I still need human/instructor help.

Reading your ratings. Mostly “independently” → you own Chapter 3: you are ready for Lab-3’s accuracy claim and Homework 3’s certification. “Small AI hint” clusters → redo just those questions tomorrow without the hint. “Substantial AI help” → rework the matching Chapter 3 lecture part, then retry. Any “human/instructor help” → bring that exact question to your instructor or ask an expert — that is what they are for. If you can answer these small scenarios, you can answer the chapter’s question: how accurate is accurate enough — and what evidence justifies that claim?

AI assists. Humans verify. Engineers decide.