PIERRE BOCKNER · EMBEDDED SYSTEMS & EDGE AI

TinyML Feasibility Scorecard

Is your AI idea feasible on a microcontroller? Ten criteria, honest scores · bockner.com

> Score each criterion **0, 1 or 2** using the anchors — an unknown is a 0, not a 1. Sum the ten scores (max 20) and read the results band at the end. The scorecard complements, not replaces, a feasibility spike on your real target hardware. Version 2026-11.

The ten criteria

1 · Problem fit

Is this a pattern-recognition problem where explicit rules or thresholds genuinely fail?

0 — unclear problem, or ML for its own sake

1 — plausible ML benefit, unproven

2 — rules demonstrably insufficient

Score: ☐ 0 ☐ 1 ☐ 2

2 · Data availability

Do you have representative, labeled data from the target environment — not just lab or synthetic data?

0 — none / synthetic only

1 — partial or lab-only data

2 — field data + labeling process

Score: ☐ 0 ☐ 1 ☐ 2

3 · Memory budget

Does the model (weights + activations + tensor arena) fit flash/RAM of the target with headroom?

0 — unknown or clearly too big

1 — borderline, needs quantization/pruning

2 — fits with ≥30% headroom

Score: ☐ 0 ☐ 1 ☐ 2

4 · Inference latency

Does inference meet your deadline on the actual target hardware — measured, not estimated?

0 — not measured

1 — measured on dev board only

2 — on target, meets deadline with margin

Score: ☐ 0 ☐ 1 ☐ 2

5 · Sensor pipeline

Are sampling, filtering and feature extraction deterministic and validated across operating conditions?

0 — ad hoc

1 — works in the lab

2 — validated across conditions

Score: ☐ 0 ☐ 1 ☐ 2

6 · Power budget

Does the inference duty cycle fit your energy budget (battery lifetime, energy harvesting)?

0 — not considered

1 — estimated only

2 — measured on target

Score: ☐ 0 ☐ 1 ☐ 2

7 · Hardware fit

Is the compute class right (MCU vs. MPU/accelerator), with long-term supply availability?

0 — undecided or unsuitable **1** — candidate selected, unverified **2** — verified + supply longevity confirmed

Score: ☐ 0 ☐ 1 ☐ 2

8 · Deployment & updates

Is there a path for model conversion (e.g. TFLite Micro/CMSIS-NN), versioning and safe OTA updates?

0 — none **1** — conversion works, updates manual **2** — automated pipeline + OTA concept

Score: ☐ 0 ☐ 1 ☐ 2

9 · Validation & drift

Do you have a held-out test set, field-accuracy monitoring and a trigger for retraining?

0 — none **1** — offline test set only **2** — field validation + drift plan

Score: ☐ 0 ☐ 1 ☐ 2

10 · Ownership

Who maintains the model lifecycle (data, retraining, releases) after launch — named, with a process?

0 — nobody **1** — implicit / one person **2** — named owner + process

Score: ☐ 0 ☐ 1 ☐ 2

Your result

16–20 · Green	**Feasible.** Start with a 2–4 week feasibility spike on the target hardware to confirm latency, memory and accuracy before committing to a product plan.
10–15 · Yellow	**Feasible with constraints.** Resolve the lowest-scoring criteria first — usually data, memory or latency. A scoped spike de-risks exactly these unknowns.
0–9 · Red	**Reconsider.** Revisit the problem definition, data situation or hardware class before investing in model development. An honest "not yet" saves months.

Rule of thumb: a single 0 in criteria 3, 4 or 6 (memory, latency, power) can sink an otherwise green project — these are physics, not effort. Validate them on real hardware first.

Want a candid feasibility assessment? Book a free, confidential 15–20 minute private technical review: bockner.com/review. Edge AI development & TinyML services: bockner.com/edgeai.

Pierre Bockner · Senior Embedded Software Engineer · Nürnberg, Germany · bockner.com · info.bockner@posteo.de · General guidance — validate on your target hardware.