

Abstract

Trap-Neuter-Return (TNR) programs mark neutered stray cats by clipping one ear tip. Conventional traps cannot tell a clipped cat from an unclipped one, so where roughly 80% of the population is already neutered, most captures are wasted effort and may teach cats to avoid the trap altogether. We operate an edge-computing trap that runs ear-tip detection on a microcontroller-class NPU and decides autonomously whether to close the gate. Field data from 116 recorded decisions exposes a decision-layer rather than a perception failure: 69% of decisions were reached with zero valid evidence frames, and 24 of 34 captures (71%) were made without any supporting evidence, violating the system's own principle that no capture shall occur under uncertainty. Taking that prototype as a baseline, this project reformulates the capture decision as a utility-based rational agent acting under partial observability, replacing hand-tuned fixed thresholds with explicit belief updating and asymmetric-cost decision making. The perception stack is treated as given; the work proposed here lies entirely in the decision layer. Because every decision record stores per-frame confidences and the thresholds in force at that moment, the new policy can be developed and evaluated by offline replay on already collected data. We expect to eliminate zero-evidence captures while preserving capture recall, and to deliver a policy small enough to run unchanged on the existing microcontroller.

1. Introduction

Motivation. Since Taiwan banned euthanasia for population control in 2017, shelter capacity has been the binding constraint [1], making reduced reproduction the primary management route. TNR is the most studied such intervention and reports the largest effect on population size [2]. It depends on capturing un-neutered individuals, which are distinguished only by the absence of a clipped ear tip [3]. Conventional traps cannot discriminate and capture everything that enters. At our study site roughly 80% of the population is already neutered, so most captures are re-captures: they consume volunteer labour, stress the animal, and, because the chain-driven steel gate is loud, may condition nearby cats to avoid the trap, reducing future access to the individuals that actually matter.

Problem definition. The task is not to detect a cat. It is: given a stream of noisy, intermittent per-frame detections from a fixed camera inside a trap, decide when to capture,

when to release, and when to keep observing, on a device with no network guarantee, within a short observation window, and under strongly asymmetric costs. A false capture is expensive and partially irreversible: animal stress, trap avoidance, wasted labour. A false release is cheap, since the same individual revisits. Formally this is sequential decision making in a partially observable, stochastic, sequential task environment, and the correct objective is expected utility rather than per-frame classification accuracy.

Existing approaches and their gap. Prior smart traps target species-level discrimination; the AI-based box trap of Besala et al. decides whether the animal in frame belongs to the target species [5]. Our task differs in kind: the two classes are the same species in two states, separated only by a few centimetres of ear contour, and the decisive class is severely under-represented. Edge-versus-cloud placement has been studied for latency and energy [4], but that literature assumes the detector output is directly actionable and is silent on what a device should do when its own evidence is insufficient.

Scope. The measurements above come from a prototype trap we operate at the study site, which provides the perception stack, an ear-tip detector quantised onto the device NPU, together with firmware, gate control and per-frame evidence logging. This project takes that perception stack as given and addresses the layer above it: the policy that turns a stream of uncertain detections into a physical action. The decision logic currently in use is a fixed-threshold rule, valuable diagnostically, since it showed that about three quarters of gate rejections are structural rather than threshold-related, pointing at camera placement rather than model accuracy, but brittle as a decision mechanism. Its thresholds are hand-set constants, its accumulator has no notion of how much evidence is enough, and “no clipped ear observed” is silently treated as “not neutered”, which is precisely how 71% of captures came to be made on zero evidence. The belief-and-utility formulation proposed below is not part of the present firmware.

Proposed approach. We replace the fixed-threshold accumulator with an explicit belief-and-utility formulation. Detector confidences are first calibrated into likelihoods, then used to update a posterior belief over {neutered, un-neutered} across the observation window. At each step the agent chooses among CAPTURE, RELEASE and OBSERVE by expected utility under an asymmetric cost matrix, with OBSERVE carrying a time cost bounded by the physical observation window. A third outcome,

UNCERTAIN, becomes representable for the first time. The resulting policy is compiled back into a compact decision table so that the microcontroller firmware remains unchanged in structure.

2. Method

2.1. Task environment specification. We characterise the agent in PEAS terms [6]. Performance: expected cost over captures and releases, not raw accuracy. Environment: trap interior, one cat at a time, other cats visible through the mesh. Actuators: gate motor and remote notification. Sensors: camera and two infrared beams. The environment is partially observable, stochastic, sequential and single-agent. This classification determines the admissible solution family and is the reason a per-frame classifier is structurally insufficient.

2.2. Confidence calibration. Raw detector confidences are not probabilities. We fit one-parameter temperature scaling [7] on a held-out split so that the reported confidence of the clipped class becomes a usable likelihood, evaluated by expected calibration error and reliability diagrams. This step is necessary because the downstream belief update multiplies these values; miscalibration propagates directly into wrong actions.

2.3. Belief update over the observation window. The window is treated as a sequence of observations. Frames that fail the structural gate are treated as missing rather than as negative evidence; this single change removes the “absence of evidence equals evidence of absence” defect. Valid frames update the posterior. We compare an independent-likelihood update against a two-state hidden Markov formulation that models the fact that the underlying state is constant within a visit while observability varies.

2.4. Utility-based action selection. We define a cost matrix with asymmetric entries reflecting the field reality above and select the action maximising expected utility. Because OBSERVE is available but costly and the window is bounded at 6 s, about nine frames at the measured 703 ms median interval, this is an optimal-stopping problem, solved by backward induction over the remaining window.

2.5. Offline policy evaluation by replay. Every decision record stores per-frame confidences, per-condition gate rejection counts and the thresholds in force, so any candidate policy can be scored on the 116 recorded decisions and 1,125 recorded frames without new field data collection. This makes rapid iteration possible within one semester and is the main reason the project is feasible.

2.6. Deployment constraint. The final policy is quantised into a lookup table over (accumulated belief, frames remaining) so that it fits the existing RTL8735B firmware. Any policy that cannot be so compiled is rejected. This mirrors the constraint that already governed our architecture choice, where RT-DETR-L lost its entire decision-layer coverage after quantisation despite a mAP

drop of only 0.087.

3. Expected Results



Figure 1. One recorded observation window (nine frames, 0.6–6.3 s). Only f02 became a valid evidence frame; the remainder failed the structural gate. Under the current rule this window ends in a capture on effectively no evidence. The proposed policy is expected to end it as UNCERTAIN.

Metric	Baseline: fixed-threshold rule	Target: proposed layer
Captures on zero valid evidence	24 / 34 (71%)	0
Decisions with zero valid evidence	69%	reported as UNCERTAIN
Un-neutered recall	not computable (no ground truth)	per class, with CI
Policy on device	hand-set constants	lookup table

Table 1. Targets against the current rule.

We report per-class recall rather than overall accuracy: where 80% of individuals are already neutered, a degenerate “always neutered” policy would score about 80%. The primary deliverable is a decision-cost curve comparing the proposed policy against the current rule across a range of cost ratios. A negative result remains possible: if calibrated confidences on the clipped class are too weak, no policy can recover them, and the conclusion becomes a quantified statement that the bottleneck is perception rather than decision logic. Separating those two causes cleanly is itself a result.

References

- [1] Yan, T.-Y., and Teng, K. T.-y. Trends in animal shelter management, adoption, and animal death in Taiwan from 2012 to 2020. *Animals*, 13:1451, 2023.
- [2] Smith, L. M., Hartmann, S., Munteanu, A. M., Dalla Villa, P., Quinell, R. J., and Collins, L. M. The effectiveness of dog population management: A systematic review. *Animals*, 9:1020, 2019.
- [3] Agriculture Knowledge Portal, Taiwan. Introduction to TNR. 2020.
- [4] Koubaa, A., Ammar, A., Alahdab, M., Kanhouh, A., and Azar, A. T. DeepBrain: Experimental evaluation of cloud-based computation offloading and edge computing in the Internet-of-Drones for deep learning applications. *Sensors*, 20:5240, 2020.
- [5] Besala, F. I., Niimoto, R., Lee, J. H., and Okamoto, S. Development of AI-based smart box trap system for capturing a harmful wild boar. *ROBOMECH Journal*, 12:4, 2025.
- [6] Russell, S., and Norvig, P. *Artificial Intelligence: A Modern Approach*. Pearson, 4th edition, 2020.
- [7] Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *ICML*, 2017.