



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: V Month of publication: May 2026

DOI: <https://doi.org/10.22214/ijraset.2026.81676>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

SAFAR: A Stray Animal Feeding AI Reward System

Jeevan Singh, Mr. Krishna Garg

Dept. of Artificial Intelligence & Data Science IIMT College of Engineering Greater Noida, India

Abstract: The current paper investigates how stray animals in urban India face challenges each day to obtain enough food. Rather than a lack of empathy from humans, the problem stems from a failure to create a systematic method for transforming empathy into practical actions. In response to this gap, the SAFAR (Stray Animal Feeding Artificial Intelligence Reward) application was created. SAFAR is a mobile tool that uses artificial intelligence and gamification to ensure that urban stray animals are fed.

In terms of implementation, the technology functions in the following way: a user takes a picture using a React Native application, which is then immediately processed by a YOLOv8 computer vision model on a Python FastAPI server. If the AI model detects both an animal and food in the photograph, the user earns credits stored in their Firebase Firestore account. Cooldown intervals are embedded to reduce the risk of cheating. It was found experimentally that image analysis took less than two seconds and worked effectively in daylight conditions. Overall, the application does more than feed stray animals; it creates a dataset that will be useful for local authorities and NGOs in mapping stray animals for vaccination programs

Keywords: Stray animal welfare; YOLOv8; object detection; gamification; React Native; FastAPI; Firebase Firestore; proof-of-work; urban AI; computer vision

I. INTRODUCTION

Early morning observations of medium-sized Indian cities highlight an important yet noticeable problem with respect to stray animals. According to municipal authorities, there are millions of stray dogs in metropolitan cities like Delhi and Bengaluru, which does not include those stray cattle and stray cats. Although some citizens feed stray animals on their own out of concern, these practices are neither systematic nor consistent. This is evidenced by the fact that while a street may have three persons feeding a group of stray dogs, the colony just nearby may remain unfed for several days.

What makes the entire issue difficult to deal with is the absence of formal recognition and feedback processes for the welfare programs conducted through the grassroots. Even though NGOs run campaigns on an occasional basis, daily feeding depends entirely on the actions of the individuals who, without any formal support systems, often perform the task inconsistently. It has been proven through behavioral science that people tend to continue their practices only when there is immediate feedback [1].

SAFAR (Stray Animal Feeding AI Reward) is an AI-based mobile application platform built on two main principles: verifying the feeding event and rewarding the person who feeds the animal. The verification of the feeding event is accomplished by using a YOLOv8 object detector that checks if there is a stray animal in the picture and if there is any kind of feeding activity, such as a bowl, roti, or hand in the picture. The reward will be given in the form of points, which can be redeemed by partners in the future.

II. RELATED WORK

A. Animal Welfare Technology

Most technology applied in the animal welfare space focuses on pet adoption platforms or emergency reporting apps for injured animals. None of these address the daily, routine challenge of structured stray feeding, and none include any verification layer. A user can claim a feeding event without any evidence, which makes digital reward schemes trivially exploitable. SAFAR is designed specifically around that gap.

B. Computer Vision in Animal Monitoring

Models for object detection, specifically those from the YOLO series, have found many uses in wildlife conservation and animal husbandry, such as estimating the number of endangered animals using aerial drones and analyzing the condition of cows in intelligent farms [2]. However, use cases in urban environments characterized by varying lighting conditions and viewpoints, as well as partial occlusions, have not garnered much focus. SAFAR fits this exact use case.

C. GamificationforBehavioralChange

Sweatcoin and Forest apps have proven that small yet meaningful rewards in digital space can motivate people to adopt sustainable behavior in real life [3]. Despite proven resultsoftheimplementationofgamification in the sphereof health care and environment protection, as far as we know, gamification is yet to be used in the field of community animals' welfare.

D. ResearchGap

Currently,thereisnosystemthatprovidesanAI-enabled proof of work for the community-based feeding of strays. Withoutany sortofverifiableproof, reward systems wouldbe untrustworthy.SAFARsolves this problemthroughtheuseof computer vision, not just to identify objects but actions too.

III. PROBLEM STATEMENT

Three main problems are interlinked here. Firstly, there isnoofficialrecognitionoftheroleofindividualfeeders, whichaffectsefficiencynegatively.Secondly,sincethereisno systemofverificationinplace,itpreventstheimplementation of any data-based program aimed at rewarding those people orimprovingtheir lives.Lastly,thepracticeoffeedingstrays doesnotcreateanydatabase,makingitimpossibleforNGOs and local authorities to vaccinate stray dogs.

The problem statement is framed such that efforts are made to build a system that will be able to identify instances of stray animal feeding based on images, assign credit scores to users, and prevent any form of duplication of incidents through considerations of the limitations of consumer-grade mobile applications and cloud systems...

IV. SYSTEM ARCHITECTURE

A. Overview

The SAFAR uses the decoupled three-tier architecture design. The Presentation layer uses the React Native mobile application used to authenticate users, take videos using the camera,andshowtheDashboard.TheApplicationlogiclayer consistsofthePythonFastAPIserverthatservestheYOLOv8 model and processes the Smart Validation algorithm. The Data layer uses Google Firebase Firestore that manages the real-time profiles and atomic point updates. All ML Inferences occur in the server-side through the REST APIs without burdening the user's mobile phone.

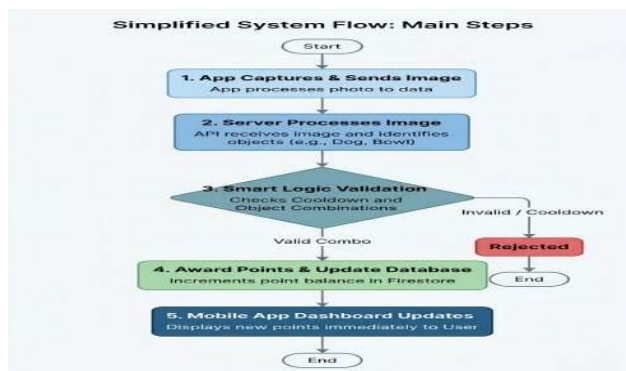


Fig.1.OperationalflowchartoftheSAFARfeeding verification and reward process.

B. MobileClient

React Native and Expo were used to build the front-end forAndroidandiOSfromthesamecodebase.TypeScriptwas extensivelyusedinthefront-endto minimize run-time errors associated with asynchronous API requests and Firestore documents. On initiation of thefeeding session by a user, the front-end will lock itself until the POST request completes to prevent duplicate submissions.The front-end will update the points dashboard using an on-Snapshot listener immediately after the successful completion of the write operation on Firestore documents.

C. AI Backend

Thisbackendexposesoneendpointat/verify-feeding/that takes a multipart POST request containing image bytes and the UID of the user in Firebase. Uvicorn handles the server responsibilitiesofASGI,andtheFastAPIframework supports asynchronous functionality so that the endpoint can handle multiple requests without getting stuck. YOLOv8 parses the inputJPEGandoutputs a flatlistofthedetectedclasslabels, such as ['dog', 'bowl', 'hand'].

D. SmartValidationAlgorithm

Not only does the task involve object recognition but context-based understanding of the situation is required as well. The code uses the ability of the label to correlate between the two given arrays to determine whether or not to award points: valid_animals=["dog","cow","cat","animal"], and valid_food=["hand", "roti", "biscuit", "bowl"]. If either animals or food are detected alone, then different responses will be returned for each case.

E. Anti-Spam and Data Integrity

Before the score can be updated, the backend checks the Firestore database to see if the last_fed value has passed its time cooldown. If the user qualifies, the backend uses the FirestoreIncrement(10) operation to add 10 points to the score along with a SERVER_TIMESTAMP to prevent any race condition issues..

V. TOOLS AND TECHNOLOGIES

Table I summarizes the core technology stack.

Layer	Technology
Frontend	ReactNative, Expo, TypeScript
Backend	Python3.13, FastAPI, Uvicorn
AI/ML	Ultralytics YOLOv8
Database	GoogleCloudFirestore(NoSQL)
Auth	FirebaseAuthentication
Dev Hardware	NVIDIARTX4050(AcerPredator)

TABLE I. Core Technology Stack

VI. RESULTS AND DISCUSSION

A. Test Case 1: Successful Verification

This consisted of an image depicting a person feeding a biscuit to a stray dog. The YOLOv8 model successfully detected the "hand," "dog," and "biscuit," and these specific detections were then verified by the Smart Validation algorithm. This enabled the user to earn 10 points within a span of 1.8 seconds.




Fig. 2. Successful detection of 'dog', 'hand', and 'biscuit' classes by the YOLOv8 model during a valid feeding session

B. Test Case 2: Missing Food Indicator

A picture showing one single cow with no indication of food or beverage was sent using the mobile application. In this regard, while the YOLOv8 algorithm was able to detect the "cow" classification, the Smart Validation system identified that there were no valid_food objects.

Consequently, the response sent by the API was negative, and it automatically prompted an alert pop-up window in the mobile application. This is proof that the front-end application is able to interpret the back-end process and implement the contextual feeding requirements.

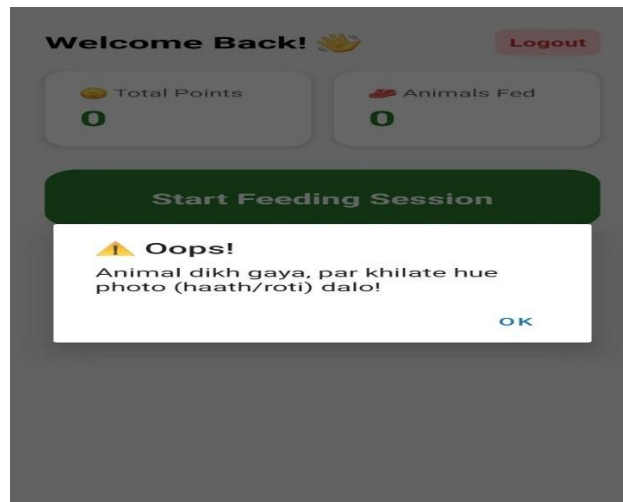


Fig. 3. Mobile application interface displaying a contextual validation alert when food objects are missing from the captured image.

C. Test Case 3: Spam Prevention

Both images were valid and submitted by the same user one after another. However, the second image was submitted during the time period when the cooldown timer was still active, which resulted in the rejection of the second image at the timestamp check step before the YOLO model inference process.

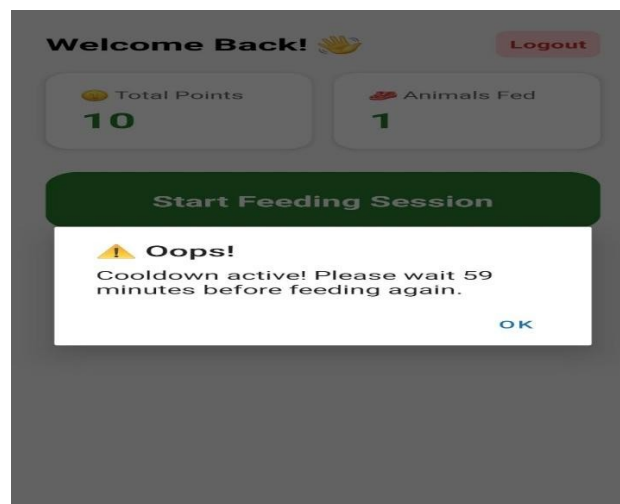


Fig. 4. Anti-spam mechanism triggering a 60-minute cooldown alert on the user dashboard to prevent duplicate or rapid submissions

D. Overall Performance

From around 40 trials conducted in different light settings, the end-to-end latency was observed to stay below 2 seconds when tested locally on a network that had the RTX 4050 server. This was consistent when conducted in favorable day light. In darker indoor settings, the accuracy became poorer; details follow below.

VII. CURRENT LIMITATIONS

There are three major limitations associated with this build of the algorithm. Firstly, low light conditions: The accuracy of YOLOv8 goes down during the night, thus making the night time scenarios virtually non-existent in terms of implementation of an incentive system. This can be solved via pre-processing of the video (histogram equalization and CLAHE), as well as via creating a special training dataset for use during low light conditions. Secondly, reliance on a dedicated GPU:

The back-end needs a dedicated GPU in order to have sub-second latency; cloud computing without a GPU will significantly increase latency, making the instant-reward system inefficient

VIII. FUTURE SCOPE

Future advancements will bring more benefits to the system. For instance, the application of geospatial heat maps based on GPS coordinates attached to each verified feed post will help track the total number of strays in the city. Using leaderboards might make the application social in nature, which can be supported by the empirical evidence showing that using leaderboards ensures continued participation from users over a long time period. Partnering with other companies and NGOs that provide incentives like discounted pet food based on users' performance can be another area of interest for future studies.

IX. CONCLUSION

In the SAFAR model, there is a difference between the well-meaning behavior of people and their sustained and contributing involvement with social good. This, however, can be addressed through the deployment of a feedback-loop system. Using computer vision to verify and gamification as motivation makes it easier for people to engage in occasional acts of social good and make themselves more visible in the process. The scope of the project suggests another important idea that is conveyed in it, which is that there are multiple uses of artificial intelligence other than those that may seem obvious at first glance.

X. ACKNOWLEDGMENT

The author is grateful to Mr. Krishna Garg for his continued advice and constructive criticism throughout the course of this work, to the Department of AI&DS, IIMT College of Engineering, for permitting access to their facilities, and to the family members who endured the debugging sessions with patience.

REFERENCES

- [1] B.F. Skinner, *Science and Human Behavior*. New York: Macmillan, 1953.
- [2] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," arXiv:1804.02767, 2018.
- [3] S. Nicholson, "A user-centered theoretical framework for meaningful gamification," Proc. Games+Learning+Society Conf., Madison, WI, 2012.
- [4] React Native Documentation. [Online]. Available: <https://reactnative.dev/> (Accessed: Jan. 2026).
- [5] Ultralytics YOLOv8 Documentation. [Online]. Available: <https://docs.ultralytics.com/> (Accessed: Jan. 2026).
- [6] FastAPI Documentation. [Online]. Available: <https://fastapi.tiangolo.com/> (Accessed: Mar. 2026).
- [7] Google Firebase Documentation. [Online]. Available: <https://firebase.google.com/docs> (Accessed: Mar. 2026).



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)