

REVIEW

Open access

Deep Learning-Driven Machine Vision for Automated Monitoring in Smart Poultry Farms

Ana Rodrigues^{1*}, Tiago Martins¹, Bruno Lopes²

Abstract

The present investigation tackles several pressing difficulties in modern poultry operations, namely pathogen transmission, workforce deficits, inadequate surveillance capacity, and intensifying requirements for superior animal welfare standards. To meet these demands, an elaborate machine vision framework based on deep learning is put forward, wherein the YOLOv11 algorithm is deployed to autonomously observe and administer poultry. A collection of 330 chicken photographs was gathered from Thai agricultural facilities under a range of environmental circumstances and subsequently increased to 1,716 images by means of rotation, along with alterations to brightness and saturation parameters. Training of the YOLOv11 architecture was carried out on the NVIDIA Jetson Orin Nano Developer Kit, with performance gauged through precision, recall, and mean average precision (mAP) metrics. The architecture attained remarkably high fidelity, registering precision of 0.964, recall of 0.938, and mAP of 0.963, correspondingly. The fusion of optical with thermal imaging permitted the apparatus to circumvent limitations imposed by variable illumination and ambient settings. Outcomes from practical deployment evidenced real-time observation and tracing of numerous birds, thereby supporting disease mitigation, feed optimization, and holistic farm governance. This inquiry advances the domain of intelligent agriculture by furnishing a wholly automated, expandable, and exceedingly accurate poultry oversight infrastructure designed to foster ecological responsibility, operational productivity, and elevated animal welfare within contemporary poultry cultivation.

Keywords Deep learning, YOLOv11, Artificial intelligence in agriculture, Smart poultry farms, Machine vision system

*Correspondence:

Ana Rodrigues

ana.rodrigues@ulisboa.pt

¹ Department of Smart Poultry Monitoring and Machine Vision, Faculty of Agriculture, University of Lisbon, Lisbon, Portugal

² Department of Deep Learning for Poultry Health Assessment, Faculty of Agriculture, University of Porto, Porto, Portugal

Introduction

Poultry production stands as a bedrock of worldwide farming, delivering a vital supply of low-cost protein sourced from chicken flesh and eggs. Representing one of the most swiftly accelerating livestock segments, it occupies an indispensable position in sustaining the global populace. Output has expanded markedly throughout recent decades, impelled by climbing consumption—most notably across developing regions. Even so, despite this vigorous enlargement, the domain continues to grapple with a series of formidable impediments that undercut yield, durability, and economic return [1]. Principal among these are the containment of epidemic diseases capable of annihilating entire populations and severing logistical chains, alongside escalating disquiet over humane treatment under high-density rearing regimes. Ecological footprints, especially stemming from refuse handling and resource depletion, pose a further substantial complication, one that is compounded by the imperative for leaner methodologies to shrink the field's carbon output. Concurrently, the sector contends with insufficient labor pools, mediocre functional throughput, and fiscal strain triggered by erratic feed valuations and transnational market rivalry [2].

A pronounced upwelling of technological ingenuity within cultivation has spawned novel approaches to bolster the productivity, durability, and proficiency of farming undertakings. Among such progressions, a focal area is the deployment of mechanization and evidence-based decision processes to refine operational results in intelligent hen husbandry [3]. Time-honored techniques for overseeing hen houses entail heavy reliance on manual input and qualitative judgments, rendering the workflow exertion-heavy, fallible, and inefficient. Machine vision infrastructures steered by deep learning have risen in prominence as a viable countermeasure to these shortcomings, given their capacity to mechanize the surveillance and stewardship of poultry premises with considerable exactness and negligible human involvement [4, 5].

The hen husbandry business is witnessing a swift metamorphosis thanks to automation and cognitive farming, which is opening fresh opportunities to amplify sustainability, output, and creature welfare. As demographic expansion and evolving nutritional leanings propel planetary appetite for fowl meat and eggs, established cultivation conventions are becoming steadily less feasible. Besides being work-intensive and predisposed to human oversight, hand-operated routines are also circumscribed in their ability to furnish the prompt, high-resolution intelligence requisite for sound farm-level decision-making [6]. Mechanization, the Internet of Things (IoT), and artificial intelligence (AI) are among the formidable instruments that cognitive agricultural innovations supply to tackle these conundrums. Hen enterprises are able to survey bird vitality, demeanor, and ambient parameters with greater effectiveness and fidelity through the amalgamation of detectors, mechanized controls, and high-level data analytics. Self-regulating feeding and hydrating setups, atmospheric modulation, and excreta disposal can curtail workforce expenditures, pare down resource draw, and improve creature comfort by engineering optimal surroundings. In addition, sentient farmsteads facilitate contemporaneous information capture and scrutiny, empowering the premature identification of sickness, tension, and anomalous conduct—elements that are crucial for safeguarding flock soundness and forestalling ruinous epizootics [7, 8]. Such a preemptive orientation toward farm oversight not only heightens throughput but also champions conservation-oriented customs by lessening refuse, power utilization, and ecological disturbance.

Welcoming automation and intelligent cultivation instruments is now a necessity for modernized hen output, equipping producers to run more productively, acclimate to marketplace signals, and fulfill the escalating planetary call for premium, fairly priced poultry goods. The present article probes how weaving together these technological strands can sharpen husbandry governance and overall effectiveness, laying a course to a more durable and gainful horizon for the trade. Machine vision allied with deep learning is materializing as a revolutionary toolkit for present-day hen enterprise stewardship, harboring vast capability to magnify throughput, creature wellness, and workflow efficacy. With the hen field beset by rising trials—pathology flare-ups, personnel shortfalls, and heightened ecological accountability—conventional supervisory blueprints habitually fall short of providing the needed exactitude and real-time adaptive faculty [9].

Beyond this, automated visual apparatuses can mechanize repetitive assignments like tallying specimens, gauging corporeal mass and maturation tempo, and tracking alimentation and hydration routines, profoundly lessening dependency on human operators and cutting down on mistakes. Such systems also buttress closer conformity to creature comfort codes—of mounting salience to buyers and oversight entities alike—by furnishing uninterrupted, non-intrusive observation. Blending machine sight with deep learning into hen facility direction signals a pivot toward fact-based governance and sentient cultivation paradigms [10]. This tactic not only sharpens functional throughput and volume but furthers eco-conscious agriculture by making prudent use of inputs and curtailing discards. The following exposition delves into the part that machine vision and deep learning play in remodeling hen enterprise oversight, accenting their potential to channel the trajectory of clever, streamlined, and durable hen output. Engineering sophisticated machine sight installations for intelligent hen farming is intended to overhaul premises administration by wielding frontier tools to settle pivotal sticking points and hoist operational flow. Inherited modes of tracking and running hen production units lean heavily on physical toil, are vulnerable to mistakes, and remain bounded in their capacity to yield granular, instantaneous appraisals of housing circumstances and bird vitality [11]. To rise above these strictures, state-of-the-art machine sight assemblies imbued with deep learning faculties are architected to supply an exhaustive, self-operating remedy that sharpens both throughput and durability.

The overriding thrust of the scholarly optics-grounded system is to facilitate perpetual, on-the-spot wellness and behavioral scrutiny of stock by way of computational vision, to flag any pathology, tension, or irregular actions, whilst contending with impediments like line-of-sight blockage, shifting radiance, pigment mutations in plumage, and heterogeneous genetic lines. This anticipatory posture enables swift intercession to tame epidemic threats and curb depletion. Additional capabilities the arrangement sets out to realize encompass the mechanizing of mundane chores—enumerating fowl, charting growth gradients, and logging fodder and liquid uptake—thereby tangibly decreasing payroll outlays and in due course elevating enterprise throughput. Another profoundly weighty target is the bolstering of creature comfort via perpetual non-invasive oversight, which will contribute to upholding decent habitation conditions and satisfying statutory standards fittingly [12]. The apparatus equally aspires to zero in on nutriment, hydration, and energy; through intelligence-fueled understanding, husbandry operators will be positioned to steer more judicious choices and slash squandering, marking one pathway by which the setup pursues its ambition of fine-tuning resource husbandry. Through the materialization of the aims set out above, progressive machine sight frameworks can be harnessed to fabricate a keener, more proficient, and ecologically mindful fowl-rearing praxis. The deliverables relayed in the ensuing text concern the system blueprint and rollout, its tangible utility, and the boons it conveys to hen cultivators vis-à-vis the continuously metamorphosing requisites of the sector.

What remains of this manuscript is laid out in the following fashion: Scholarship germane to both nascent and antecedent contrivances in cognitive husbandry and hen governance is first reviewed. The investigative blueprint, together with the deep-learning-powered machine sight apparatus for intelligent hen rearing, is then delineated. The accumulation and interpretation of data, along with the fidelity of the deep-learning-driven machine sight setup for intelligent poultry operations, are subsequently quantified and deliberated upon. The treatise is finally drawn to a close.

Review of the Literature

This segment furnishes a synopsis of state-of-the-art innovations in intelligent cultivation and flock administration. It opens with a condensed orientation covering existing technologies in precision agriculture, optical machine interpretation frameworks deployed in farming, and deep neural network strategies adopted for agrarian purposes—all presently required for upgrading hen production establishments. Following that, we showcase examinations of the innovations now operational within intelligent farming and poultry oversight. The most recent inquiries into smart flock governance are then laid out.

Survey of innovations in smart farming

Intelligent cultivation, sometimes labelled precision farming, denotes the harnessing of cutting-edge instruments to refine agronomic procedures and customs. A formative exposition of domain breakthroughs was tendered by McBratney *et al.* [13], who expounded the concept of spatially variable field administration as emphasizing the appraisal of disparate topographies bounded by climatic parameters. They revealed how sundry tools—satellite positioning and distant sensing among them—hold the power to transfigure customary husbandry, letting growers boost harvests while trimming resource drawdown. Extending the groundwork laid by these scholars, Zhang *et al.* [14] zoomed in on mechanized contrivances within tillage, especially GPS-steered machinery and aerial drones deployed for seeding, fertilizing, and canopy observation. Their analyses further intimated that such mechanization both diminished manual exertion and heightened the precision of mundane field undertakings while broadly bettering the atmospheric setting for cropping. On the flip side, the investigators underscored how steep start-up capital expenditure habitually hampers broad-based uptake. The treatise even puts forward an IoT-grounded scheme leaning on Deep Reinforcement Learning jointly with ensemble methodology to bolster location-specific cultivation, meeting head-on challenges like input assignment and informational fusion, in the end raising output by 12% and paring down wastage by 10% through state-of-the-art contrivances [15].

Optical machine interpretation frameworks in farming

Attitudes toward deploying automated visual evaluation suites within cultivation are buoyant, given their promise to magnify operational smoothness, yield volume, and ecological durability. These assemblies draw upon imaging hardware and detectors to harvest inputs, subsequently crunched through coded procedures so as to mechanize decision-making [16]. Prime use cases span scrutinizing plant soundness, infestations and sickness, yield forecasting, plus engine-driven collection tasks. Inside picture computation, automated sight can determine crop pathogens, pinpoint growth intervals, and sort reaped botanical specimens, thereby facilitating pinpoint intercessions. A persistent thread woven through agrarian machine sight solutions is the substitution of human action in chores like picking or chemical dispersal with near-total or complete mechanization [17]. This trims workforce outlays while simultaneously elevating steadiness and the quality of farming processes. Two-sided equilibrium likewise emerges as a salient theme here, with spatially aware cultivation via machine perception zeroing in on intelligence fed in a well-timed, place-specific mode so as to achieve sound input husbandry. To cite a case, optical interpretation-steered smart applicators refrain from gratuitous dosing of biocides or soil amendments, thereby capping the piling up of potent synthetic agents and their noxious fallout on ecosystems. That said, such apparatuses face hurdles too—unaffordable price points, the compulsion for instantaneous signal digestion, and the snag of making any single platform transferable across widely dissimilar plot parameters (illumination magnitude, wetness events, etc.) [18]. All the same, contemporary scholarship increasingly homes in on these sticking points, buoyed by nascent detector breakthroughs and more elaborate arithmetic routines. An inexorably mounting appetite for resilient, heavy-yielding cultivation blueprints guarantees that farm-oriented visual computing gear will stay indispensable as an emerging technological pillar in the trade [19].

Architectures grounded in deep learning

Algorithmic strategies rooted in deep neural nets for item spotting can be broadly sorted into dual pathways: unified-stage and bifurcated-stage identifiers [20]. Bifurcated-stage schemata, represented by Mask RCNN, initially cook up zones of potential relevance (ROIs) in a lead-off round via a Region Proposal Network (RPN) [21]. To yield pinpoint object geolocation and labelling, those candidate patches get further winnowed and polished in the follow-on round. The heightened exactness of this route traces back to the focused honing of detected entities. Contrariwise, by divining enclosure quadrilaterals and entity tags directly within a single forward traversal of the network, unified-stage identifiers of the You Only Look Once (YOLO) lineage telescope this pipeline, trading off a sliver of fidelity for pronounced acceleration dividends. That trait aside, because unified-stage constructs avoid segmenting recognition into discrete phases of region nomination and refinement, they manage speedier execution. As such, they fit hand-in-glove with applications mandating live-stream data throughput. In the most recent period, both orientations have marched forward to grapple with the swiftness-versus-correctness compromise. Without question, throughout the last dozen years, a profusion of unified-stage and bifurcated-stage identifiers has been purpose-built around targeted performance profiles touching the equilibrium between exactness and tempo [22]. Among the most routinely fielded bifurcated-stage identifiers stand R-FCN and Cascade R-CNN, which sharpen both positioning and category assignment using bespoke sub-networks; Fast R-CNN, which amplifies economy in feature exploitation; and Faster R-CNN, which bakes in a region nomination pipeline. On the unified-stage flank, the Single Shot MultiBox Detector (SSD), RetinaNet—which pioneered focal loss to redress category skew—and the YOLOv5, along with YOLOv6 model clans, which emphasize inference fleetness, count as the mightiest exemplars in this regard. YOLOv8 and Mask R-CNN represent the unified-stage and bifurcated-stage frameworks that have seen the heaviest deployment in farm-related scenarios in the latest wave, with eminently striking results. Indeed, across multiple head-to-head assessments, they have unmistakably evidenced superior prowess at locating and taxonomizing agrarian targets under heterogeneous settings, as catalogued in Table 1 of this write-up. These constructs are distinctively matched to the fluid tableaux typical of husbandry backdrops because they reconcile swiftness, exactitude, and sturdiness. Mask R-CNN (owing to its knack for fine-grained data segmentation) and YOLOv8 (due to its extraordinary throughput speed) have been singled out for the present exploration, grounded on these coalescing antecedent contributions [23–26].

Table 1. Result augmentation techniques of the model.

Model	mAP	Recall	Precision
RoboFlow	0.961	0.950	0.954
YOLOv11	0.963	0.983	0.964

You only look once (YOLO)

Across the span of the last handful of years, the YOLO dynasty of object spotting and instance demarcation models has undergone blistering transformation, each descendant arriving with boons in either inference pace, correctness, or both. The newest unified-stage design, YOLOv8, was sculpted atop the scaffolding bequeathed by forerunner YOLO editions, YOLOv3 and YOLOv5 among them. Standing apart from bifurcated-stage architectures, YOLOv8 collapses object recognition into a single step by outrightly inferring enclosure quadrilaterals and class confidences sans any stand-alone region proposal network.

The embrace of an anchor-free, centerpoint-rooted scheme for item pinpointing stands as a momentous leap within YOLOv8, conferring multiple edges over more old-school anchor-reliant tactics typified by YOLOv5, YOLOv6, and YOLOv7. To churn out a richer mix of projections and stiffen the terminal prediction's exactness alongside toughness, YOLOv8 wields Pseudo Ensemble, alternatively christened Pseudo Supervision (PS)—a ploy that entails schooling numerous model variants fitted with dissimilar parametrizations on the selfsame body of examples.

YOLOv8 further leans on the Darknet-53 skeleton, a deeply stacked convolutional neural construct spanning 53 strata that is purpose-calibrated for entity identification and trait extraction. The ousting of the C3 building block in favor of the C2F unit marks one eye-catching design shake-up in YOLOv8. The C3 segment, habitually dubbed the convolutional module, carries out a chain of convolution-driven operations to chew over incoming signals. When set beside its earlier incarnations, the C2F unit—a souped-up reworking of the C3 component—bolsters inference throughput alongside correctness [27].

Besides this, YOLOv8 banks on the Darknet-53 structural plan, a deeply layered convolutional neural edifice outfitted with 53 tiers attuned for entity spotting and feature abstraction. The displacement of the C3 module with the C2F module ranks as a salient architectural shuffle in YOLOv8. By wielding a sequence of convolution-based operations, the C3 piece—also called the convolutional module—makes sense of input streams. The C2F piece, an ameliorated take on the C3 module, ups the ante on computation tempo and exactitude relative to prior age cohorts [28].

To satisfy heterogeneous mandates for computational cadence and exactness, YOLOv8 comes in an assortment of profiles [29–31]. For streaming applications on hardware-pinched apparatus, YOLOv8-Tiny dispenses brisk inference at the expense of a measure of exactness; YOLOv8-Small hits a sweet spot between pace and finer-grained recognition scope; YOLOv8-Standard proffers dependable muscle under diverse operating conditions; and YOLOv8-Large accords primacy to high exactness for mission-critical settings where fidelity and granularity are overriding.

Fresh headway in YOLOv8 has smoothed its pathway into a sweeping array of farming-related usage, bearing out its potency in untangling idiosyncratic snags that surface amid the heterogeneous cultivation tableaux previously profiled. Tuning the manipulation of raw, low-tier signals renders YOLOv8 a formidable means for flagging agrarian creepy-crawly infestations and plant maladies at a nascent stage—something pivotal to preserving crop soundness [32]. To guarantee early spotting and remedial action against botanical ailments, for one, enhanced offshoots of YOLOv8 have been put into service to pinpoint disorders in glasshouse vegetables [33]. The infusion of attention-oriented contrivances into YOLOv8 so as to augment object-flagging capacity marks yet another novelty. Yang *et al.* [34] put this gambit to the test in order to lift tomato-identification exactness within cluttered, packed horticultural tableaux.

Related works

Mechanized observation alongside data amplification procedures is progressing as well. An image manipulation protocol for delineating the outlines of dressed poultry was contrived by Chen and Wang [7], targeting heightened operational tempo in automated dressing workflows. Beyond this, Paszke *et al.* [35] impart understanding into the PyTorch programming ecosystem, a platform that has grown into a linchpin for deploying intricate architectures throughout agricultural artificial intelligence settings. In a parallel vein, Singh *et al.* [36] put forward an all-encompassing scrutiny of AI and IoT rollouts targeted at flock surveillance, concentrating with particular intensity on visual-spectrum wellness evaluation and acoustic-driven demeanor scrutiny. A swelling wave of inquiry zeros in on sanitary oversight. Exemplifying this, Ahmed *et al.* [37] unveiled a prognostic engine for premature pathology flagging among hens, leveraging kinetic and biological metrics harvested from on-body IoT gadgetry, netting fidelity as high as 97%.

In the year 2022, Cakic *et al.* [38] probed the fashion in which deep neural networks together with high-caliber computing resources could be enlisted within intelligent hen holdings to execute target pinpointing. The endeavor set out to sculpt object detection blueprints operable on peripheral nodes outfitted with visual sensors. Such blueprints augment mechanization in premises supervision by rendering workable undertakings like on-the-spot carcass identification and bird tallying [39]. Through harnessing Faster R-CNN alongside the Detectron2 suite, the research unit lands on tolerable exactness benchmarks, although supplementary iterative testing proves necessary to polish recognition frameworks for deceased specimens, gesturing toward conceivable future expandability within agricultural automation answers. Meanwhile, Kim and Lee [40] debuted an RNN-anchored arrangement to catch breakdowns in airflow apparatus, a function indispensable to upholding atmospheric fitness and sickness prophylaxis inside enclosed hen enclosures.

The exploration conducted by Hewawasam *et al.* [41] probes the fabrication of a fused, smart-algorithm-powered suite purposed toward refining poultry husbandry operations. The inquiry homes in on robotizing pivotal poultry-rearing sequences—specifically embryonic detection within ova, fowl flesh grade scoring, ancestry typing, and ailment recognition. Exploiting machine cognition, IoT, and computational sight methodologies, the assemblage yields on-demand transparency into flock workflows. The evidence underscores the latent capacity for sharpened exactitude and swiftness throughout every modular component, cumulatively feeding into amplified output and monetary return for the domain. Pain points like hand-operated hemoglobin speck flagging and impressionistic meat decay scoring are countered through pictorial processing and neural net formulations. The suite's punch is spotlighted by way of the fruitful deployment of sophisticated recipes such as CNN and YOLOv5 spanning disparate functions—ancestry recognition and malady flagging among them. The scholarship wraps up with advisories for subsequent refinement, most notably in scaling up ancestry and sickness taxonomization breadth. Vinod *et al.* [42] furnish an autonomous contraption for ovum enumeration inside poultry operations, wielding computational sight and AI stratagems to stretch exactitude while compressing manual fallibility. Through the incorporation of formulae like YOLO and Faster R-CNN, the assembly makes feasible contemporaneous spotting and tallying of eggs, steeply slashing workforce and temporal overhead while strengthening dependability. The paper also casts light on the unit's capability to register shell imperfections such as fissures, feeding into better warehousing oversight and quality benchmarking. This inquiry interlocks with sustained momentum in AI and foreshadows weighty dividends for the hen-raising trade via bolstered throughput and ecological durability. The current paper, authored by Natho *et al.* [43], records the construction of a concept-stage rig for mechanized hen sustenance dispensing through an onboard logic controller. It calls out the feebleness ingrained in hand-dispensed protocols, which balloon payroll expenditures and yield fodder loss. The suggested setup operates through Arduino microcontroller circuitry to orchestrate the dispensation cadence at set time lags, ensuring pinpoint nutriment delivery and expenditure compression. The investigation equally spotlights how mechanization lifts yield, slashes physical toil, and betters flock condition by way of more sanitary feeding customs.

Materials and Methods

Within the following paragraphs, we delineate a bolstered machine sight apparatus for administering hen facilities through the deep learning conceptual schema exhibited in **Figure 1**, commencing with the gathering of RGB stills from poultry establishments across Thailand. The drilling and vetting corpora were produced by manually labeling these captures,

which were gathered under an assortment of meteorological backdrops, embracing both bright and heavily overcast intervals. The previously cited deep learning formulations were subsequently fitted on the drilling corpus, and the validation corpus was drawn upon to measure their competence at instance demarcation. Once manipulated to warehouse the fowl records, the information was deposited inside the cloud ecosystem, where it might be exploited as inputs for figuring hen sustenance allotments and for transmitting instructions to the bulk ingredient hopper to emancipate raw substances commensurate with the verified hen needs.




Figure 1. A system model.

Data acquisition in the study

Over the course of this endeavor, we harvested information about hens housed in Thai facilities, tracing development from the neonatal phase right up until the birds satisfied market specifications, typically around week nine. An average bird carries roughly 1,228 grams of mass, and the mechanized sustenance allocation is anchored in the theoretical underpinnings of earlier scholarship [43]. Capturing of fowl datasets may be undertaken across dissimilar surroundings—variable radiance and atmospheric regimes being representative examples—and spatial concentration of stock can additionally be photographed, as portrayed in **Figure 2**. In the arena of machine learning, intelligence caliber is of critical weight. Data harvesting, categorization, and manipulation sit among the non-negotiable functions enacted during the interval preceding ingestion by the DNN procedure. Photographic stills comprise the raw substrate for this conduit. Each frame was extracted from motion captures recorded within the hen-rearing compound. The central aim of the mechanism advanced in this treatise is to locate targets inside the harvested photographs—in the immediate instance, hens. This hurdle originated as an offshoot of autonomous, mechanized spotting within the visual computation discipline. Hence, it stands apparent that certain AI procedures might be deployed throughout this conduit. The process calls for generous pools of examples to execute effectively. Toward realizing this goal, we compiled and classified 500 hen images, making use of a bounding-box annotation utility within the RoboFlow toolkit [44] to label items inside these captures, a toolkit endowed with the power to automatically adjust and classify stills, as visualized in **Figure 2**.



Figure 2. Data collection and annotation.

Our elected path here hinges on the ability to leverage the bounding-box annotation workspace, which affords the option to house the labels in multiple encoding conventions. Under the Common Objects in Context (COCO) specification, it becomes obligatory to preserve every finalized label in a suitable file architecture. This arrangement permits later viewing and the coaching sequence directed at constructing a forecast framework [45]. The nominated coding suites for DNN exploration and model assembly are congruent with this convention. An associated JSON digest incorporating a skeleton for recording item taxonomy and spatial coordinates for every label accompanies each group of tagged captures in the COCO specification. The COCO yardstick additionally supplies the cardinal benchmarks required to gauge framework soundness, as expanded upon subsequently within the outcomes passage.

Data preparation

At this juncture, pivotal stages involve ascertaining that adequate intelligence stands ready for drilling the framework, while informational pretreatment for structuring and enrichment gets executed via the RoboFlow instrument [44]. A RoboFlow screen grab summarizing the data pool and pretreatment enacted for this examination is reproduced in Figure 3.




Figure 3. Data preparation.

Pretreatment, enrichment, classification, ordering, architecture drilling, and architecture fielding represent just a sampling of the wide-ranging data orchestration functionalities that the RoboFlow workspace makes available. The pretreatment segment of the present examination solely necessitated rescaling the raw captures to conform with the actual dataset dimensions. The enrichment sequence involved pivoting the frame; captures were pivoted at stochastic degrees bounded between -45% and +45%. This conditions the framework to spot targets irrespective of how they are turned inside the visual field. The luminosity level of captures was shifted across a spectrum from -15% to +15%. This replicates dissimilar radiance milieus and equips the framework to function effectively within both muted and intense luminance settings. The color vividness was amplified across a spectrum from -15% to +15%. This modulation alters the depth of pigmentation in the captures, hardening the architecture's resistance to swings in chromatic display; frames underwent trimming of as much as 20% relative to their baseline dimensions. This maneuver conditions the framework to perceive fragmentary items and elevates its execution when targets are only fractionally exposed inside the visual field. We arrived at a pool comprising 2,600 captures following the enrichment conduit, springing from the seed batch of 500. It warrants mention that the unpaid version of RoboFlow imposes a boundary on dataset volume subsequent to enrichment; even so, a tally of 2,600 frames constituted a becoming magnitude of visual material. Of note, tagged captures segmented into more granular groupings and steered via RoboFlow proved arrangeable through RoboFlow. This additionally yielded a dataset constituting a unified, expansive corpus of decipherable, marked, and refined captures. As required for the machine learning pipeline, the pool was additionally partitioned into coaching, tuning, and assessment subsets employing RoboFlow. A screen view of the RoboFlow workspace featuring a summary of the corpus and pretreatment wielded in this examination is reproduced in **Figure 3**. The intelligence sits ready for deployment in machine cognition probes and the fabrication of forecasting blueprints for target identification after it has undergone harvesting, annotation, pretreatment, and enrichment.

Application of deep learning

The YOLOv11 architecture that powered the setup underwent training on a workstation running Ubuntu 18.04 LTS 64-bit fitted with an NVIDIA Jetson Orin Nano Developer Kit packing 8 GB of RAM, a selection driven by the fact that this particular hardware configuration is utilized as an on-site evaluation platform within a hen production facility situated in Thailand—a setting where economical energy draw paired with adequate computational grunt for such an agricultural context is paramount. YOLOv11 outstrips every earlier release in terms of both fidelity and resource economy, YOLOv8 included. The nano variant YOLOv11n achieves an improved mAP score (39.5% mAP), a leaner parameter footprint, and dramatically swifter CPU inference turnaround (56.1 ms CTT). As such, it emerges as the prime candidate for fielding in processor-constrained solutions. The extra-large variant YOLOv11x advances beyond its forerunner (YOLOv8x) by registering both a loftier mAP metric (54.7% mAP) and a reduced BFLOPS toll (194.9) per individual frame. The YOLOv11x configuration discovers a near-ideal reconciliation of compute frugality against spotting precision within its structural blueprint. For these reasons, the YOLOv11 lineage constitutes a markedly stronger option for live-stream item recognition, supplying models able to deliver high accuracy alongside compute-aware design wherever dual performance benchmarks must be met simultaneously. In sum, the YOLOv11 model family typifies some of the strongest available baselines for anyone embarking on building present-day visual intelligence pipelines. At its core, the architecture leans on PyTorch, a server-side computational scaffold executing inside Linux ecosystems. As regards the model-fitting regimen, the learning rate = 0.001, mini-batch cardinality = 32, and dropout likelihood = 0.5 were adopted, all while restraining the possibility of overfitting. Over one thousand fitting cycles were conducted. The fitting routine was halted once the architecture's metrics showed zero forward movement throughout an interval of 20 uninterrupted epochs computed against the held-out tuning portion, prior to the architecture reaching the 232-epoch mark. This intervention boosted the architecture's capacity to generalize and suppressed over-specialization on the coaching sample. The momentum coefficient and weight penalty term for both architectures were pegged at 0.937 and 0.0005, respectively, with a foundational learning rate of 0.01 maintained right through fitting. These precise hyperparameter choices were calibrated toward extracting maximum fitting efficacy while simultaneously shrinking the danger that the architecture would mold itself too tightly around the coaching distribution. To introduce steadiness into the model optimization pathway and to shrink the odds of the process becoming mired in a suboptimal local basin, a gradual ramp-up phase was instituted over the initial three epochs, setting momentum to 0.8 and clamping the bias learning rate at 0.1. The instrument chain

required for fashioning sound forecasting architectures receives treatment in this passage. Python acted as the implementation tongue, counting among the most universally embraced programming languages within the machine cognition community. Python's gentle learning curve, the extensive ecosystem of ready-made libraries that streamline efficient problem-solving, along with the sizable developer collective, supply the chief motives undergirding this preference [46]. A sweeping assortment of packages stands at the ready, each extending utility toward constructing deep neural nets whenever Python serves as the implementation substrate. The dominant orientation of the present inquiry homes in on object spotting difficulties, with a specific concentration on fowl identification. For this express purpose, a handful of genuinely formidable Python packages are available. Our election settled on Detectron2, a library that principally leans upon the PyTorch numerical engine. The PyTorch library commands extensive deployment inside scholarly environments and bends readily to a vast spectrum of machine cognition problems. Its notable virtues envelop approachability and reconfigurability. PyTorch is recurrently set side-by-side against other tooling that fills a comparable niche, centrally TensorFlow. Both platforms see intensive utilization in the crafting of AI systems, given that declaring either unambiguously superior proves untenable [47]. Contemporary item spotting and boundary demarcation formulations constitute the Detectron2 library's stock in trade [48]. Setting up and tuning the Detectron2 bundle can, on occasion, introduce friction, the degree of which hinges upon the computing landscape at hand. Our hands-on practice indicates that exploiting the Google Colab sandbox represents a sound route for academic and pilot-stage trial runs, on account of its frictionless expandability [49]. Google Colab brings to the table facilities that quicken and streamline environment configuration and toolkit installation workflows. Scores of AI investigators lean upon Google Colab, which is offered under both zero-cost and subscription-based/premium licensing. The bedrock of this service centers on furnishing end users with leased virtualized compute slices for a finite window of access. This temporal allotment fluctuates under the zero-cost tier, and coders are furnished no anticipatory signal indicating for how long their assigned virtual node will hold up before the session is severed. Under the premium tier, this permitted interval stretches considerably longer. Over the course of architecture fitting, a real possibility exists that a session might suddenly lapse, with outputs failing to be persisted. This pitfall can be circumvented by inserting a modest extra block of scripting that periodically persists architecture snapshots at staggered points throughout the fitting operation. A further nuisance arises from the purge of a coder's deposited materials once the connection to the Google Colab virtual node is torn down. One workaround for this obstacle involves lodging datasets on Google Drive. The Achilles' heel of this tack lies in the unavoidable step of pulling materials back from Google Drive at each fresh connection or whenever one means to wield them in tandem with this virtual node—an operation that can devour significant clock cycles. Accordingly, as long as one enters with awareness of the possible shortcomings tied to operating Google Colab in its zero-cost guise, the platform can behave as a tool that concretely facilitates bringing the task to completion.

Performance evaluation

Five non-overlapping assessment yardsticks were pressed into service while gauging the instance demarcation prowess exhibited by the Roboflow and YOLOv11 frameworks, namely: precision, recall, Area Under the Curve (AUC), inference rapidity, and mean Average Precision (mAP) at an intersection over union threshold of 0.5 (mAP@ 0.5 IOU). The quotient formed by correctly forecasting positive happenings against the universe of all forecast positive happenings corresponds to the precision measure that Eq. 1 unpacks. Following a parallel logic, recall, as formulated in Eq. 2, captured the share of accurately flagged positive instances set beside the overall count of target item occurrences actually present within the data pool. The mean average precision (mAP), where mAP is understood as the arithmetic mean taken across AP quantities spanning k item groupings, and AP figures were computed under a cutoff test stipulating 50% overlap between forecast and real-world entity perimeters or bounding enclosures in Eq. 4. The AUC engulfs every plausible decision boundary, and the AUC figure tied to each framework was derived computationally via Eq. 5. The tempo of inference, which quantifies the temporal expenditure required to process individual frames, likewise factored into evaluating the alacrity with which the framework delivered forecast answers.

The subsequent formula serves as the basis for deriving these yardsticks:

$$Precision = \frac{TP}{TP+FP} \quad (1)$$

$$Recall = \frac{TP}{TP+FN} \quad (2)$$

$$IoU = \frac{AreaOverlap}{AreaUnion} = \frac{TP}{FP+TP+FN} \quad (3)$$

$$mAP = \left(\frac{1}{K}\right) \sum_{i=0}^k (AP)_i \quad (4)$$

$$AUC = \int_0^1 TPR(FPR)^{-1}(u) du \quad (5)$$

Wherein the tokens TP, FP, and FN denote true positive, false positive, and false negative entity instances, in that sequence. $(AP)_i$ designates the mean accuracy calculated pertaining to the i^{th} grouping within the ensemble of k categories, the character k standing for the total cardinality of entity groupings. The region caught beneath the precision-recall graph belonging to a given grouping receives the designation AP. FPR refers to the false positive proportion, TPR to the true positive proportion, and t to the framework's wall-clock expenditure for inferring outputs tied to a designated (individual) frame.

Results and Discussion

Prediction model evaluation

The full compilation of RGB stills comprises 2,600 captures sourced from hen production sites across Thailand. These undergo partitioning into three splits: 10% set aside for examination, 10% kept for tuning, and 80% devoted to coaching. The architecture consumed close to 4.6 hours in fitting mode, burning through 232 epochs atop the full corpus for fowl spotting. The trendlines belonging to the coaching and tuning splits reproduced in **Figures 4 and 5** lay bare the manner in which the architecture's fitness evolved, as traced by a selection of indicators—box loss (box_loss), segmenttness (seg_loss), classification (cls_loss), and focal spread (dfl_loss). These tools evaluate the architecture's skill at geolocating broiler chicks by type via a specially tailored loss expression that reins in category skew as fitting proceeds.




Figure 4. RoboFlow plots training and validation sets showing epochs.




Figure 5. YOLOv11 plots training and validation sets showing epochs.

Table 1 makes plain the striking exactness and retention capacity of the YOLOv11-powered thermal imaging framework for spotting, clocking values of 0.964 and 0.938, respectively. **Figures 4 and 5** render visible two divergent construals of mAP. The mean accuracy determined at an IoU threshold of 0.5 is captured by the foremost one, mAP50. Feeding from the provided stills compilation, the RoboFlow-engineered framework returns a mean of 0.95 (**Figure 4**). The other, mAP50-95, is computed in increments of 0.5 spanning IoU benchmarks from 0.5 through 0.95. The YOLOv11-based framework executes a straightforward operation of singling out hens across an assortment of radiance backdrops, a claim upheld by readings clearing 0.97. Weakened effectivity readings manifest throughout the performance yardsticks and losses belonging to the YOLOv11-powered framework—train/box_loss, train/seg_loss, train/cls_loss, and train/dfl_loss exhibited in (**Figure 5**)—hinting that a copious and well-matched compilation becomes indispensable for exact and trustworthy spotting if this obstacle is to be cleared. Hence, in consonance with the debut of the new YOLOv11 generation, two varieties of image enrichment protocols for the architecture—the mosaic recipe alongside the established methodology—were harnessed in the course of this inquiry (**Table 1**). The categorization loss path traced across architecture fitting (train/cls_loss) undergoes steady contraction as each successive epoch elapses, setting off from 4.5 and bottoming out at 0.21, as rendered in **Figure 6**. By way of contrast, the tuning categorization loss (val/cls_loss)

surpasses the train/cls_loss by a multiplier of three, igniting at 3.1 and tapering to 0.6. This stark gulf hints that the enrichment pipeline stands as the underlying wellspring of the gap separating the two loss trajectories.




Figure 6. Plots of classification loss of the YOLOv11 model.

Discussion

An RGB camera yields a detection framework for pinpointing disease episodes in flocks that is markedly more faithful and sharply tuned than the counterpart assembled through YOLOv11. The YOLOv11-derived framework lends itself to tracing the behavioural routines and movements of hens during sunlit intervals. Where lens-based surveillance hardware is concerned, the strength of ambient light directly shapes the grade of the stills harvested. The struggle faced by antecedent architectures revolved around the problem of having to pick out targets under heterogeneous radiance regimes. A broadened field of capture—one that enables the crisper imaging of a much wider array of subjects—is

secured by fitting an RGB camera with the capacity to record frames across a sweep of brightness levels and meteorological states. What follows is that the tuning of the algorithm is itself shaped by the caliber of the imaging device. The set of stills, when pooled, compiles the image-centric corpora that the present investigation leans on to assemble a concurrent image-centric YOLOv11 architecture, thereby supplying alternate windows into fowl condition. The main thrust behind opting for an RGB camera springs from the ambition to obtain shots possessing heightened exactitude. For densely stocked hen enterprises where round-the-clock flock observation is an absolute necessity, such a consideration becomes acutely salient. This strategy attunes the advanced architecture toward more dependable hen tailing under heterogeneous localized or microclimatic bands. Grainy depictions of birds, harvested from the full compass of orientations and drawn from manifold stances or lines of sight, are thereby acquired.

Embedding frontier deep-learning-powered machine sight assemblies into hen-rearing practice has delivered on its ambitions, and the data marshalled in this section offer testimony to that. From an operational standpoint, the YOLOv11 construct—which capitalizes on spotting and tailing of hens within AI-augmented visual conduits to drive high-resolution fowl husbandry—has garnered reinforcement through lofty scores on both exactitude and recall axes. The coupling of lens-driven observation with infrared heat mapping gave a notable lift to the rig's competence at performing robustly inside non-uniform radiance and atmospheric backdrops, effectively sweeping aside a long-standing bottleneck that had dogged conventional oversight setups.

This appraisal, alongside scores of companion studies, chimes with the pre-existing documentary trail surrounding mechanized livestock cultivation. The work has gestured toward the feasibility of attaining true on-the-fly target spotting and trajectory tracing inside agrarian contexts through assorted deep neural formulations, with particular weight falling on the YOLO family. The demonstration of calibrated exactitude by the rig in singling out specimens, coupled with its flair for tracking and disaggregating parameters such as somatic gain, alimentation conduct, and wellness status, has underscored the practical utility of cognitive infrastructures within present-day hen output systems.

Despite the fact that the reported findings represented virgin territory, several friction points were logged. The weightiest vulnerability consisted of lean reliance on scrupulously tagged, high-grade corpora. While enrichment protocols—first and foremost the mosaic regimen—did afford some help in stretching the architecture's generalizing power, a sticking point lingers concerning the sheer bulk of tagged data collections that needs resolving before rollout can become genuinely viable. In addition, while the architecture demonstrated an ability to function throughout a continuum of brightness conditions, supplementary refinement may well be called for to lift its delivery inside crowded tableaux where bird bodies occlude one another.

A parallel complication stems from the habitual freight of computation hunger that frameworks built atop deep neural computation bear. The YOLOv11 construct was fitted leveraging the NVIDIA Jetson Orin Nano Developer Kit; however, modifications at the silicon level would, with near certainty, become a prerequisite for satisfactory execution inside production-scale hen facilities. A hefty quantum of work still stretches ahead for fielding such architectures in concert with peripheral AI engines, where the preoccupation shifts toward stingy energy budgets without trading away detection fidelity. Weaving IoT threads into the rig will, beyond that, ease the flow of data capture and adaptive steering. In consequence, flows of this nature can unlock auto-adjustments to feeding rhythms, ambient thermal regulation, and illness avoidance countermeasures by resting upon forward-projecting analytics mediated via live data interchange with farm administration suites.

The contribution made by this study settles onto an already weighty accumulation of scholarship devoted to AI-infused cognitive cultivation. The indications drawn from the data are that deep neural computation fused with machine sight opens expansive possibilities for raising ecological soundness, creature comfort, and operational tempo within hen husbandry. The road ahead should arc toward variegating the corpora, paring back compute footprints, and drawing in multimodal detector streams so as to stretch both exactitude and pragmatic utility for the governance of hen enterprises.

Conclusion

An efficacious, upgraded perceptual computation apparatus molded specially for cognitive hen installations was architected and operationalized by tapping deep neural computation, in particular the recently introduced YOLOv11 formulation. The rig's scorecard of performance indicators clocked an applaudable bracket: 0.964, 0.938, with a mean average precision (mAP) of 0.963, thereby receiving validation for the task of on-the-spot flock observation. The wedding of lens-driven and heat-based sensing pathways allowed the neat circumvention of bygone observational hurdles rooted in wavering brightness and ambient variables. Data enrichment regimens, above all the mosaic protocol, occupied a pivotal niche in fortifying the construct's generalizing edge and its holistic delivery. The architecture possesses the wherewithal to deliver competently inside dissimilar contextual backdrops while keeping multiple trajectories locked onto parallel specimens, thereby resetting what counts as achievable in the mechanized administration of hen production units. A milestone of such magnitude emerges as the fruition of a longstanding, pressing hunger among producers for yield-maximizing, durable, and welfare-attentive rearing modalities. Beyond this, the additional snags catalogued over the span of the inquiry included the hefty compute bill demanded for digesting streaming live intelligence and the requirement of scrupulously labeled corpora. Tracking this line of argument, the constraints sketch out yet more ground awaiting future investigative efforts, most pointedly touching on parsimonious architecture retuning for peripheral computation settings and on denser protocols for intelligence gathering. Peering past the technical reverberations, the work inclines toward a defined, foreseeable tomorrow wherein AI-governed ecologies stand to fundamentally recast the everyday texture of hen facility labor. It stands to substantially assist in nipping disease in the bud, fine-tuning feed-to-mass conversion factors, and steering of premises through its command of consumption tracking and prompt diagnostic output. This contribution knits together disparate threads of comprehension within cognitive cultivation and stands as a single demonstration of how deep neural computation might be brought to bear on the next wave of modernization flowing through hen husbandry praxis. The fluid absorption and steering of this operation will, in truth, stand as a meaningful signpost on the trajectory toward ecologically attuned, lean, and ethically anchored fowl bioproduction frameworks. Forthcoming endeavors would do well to home in on wellness and conduct-tracing functionality. This must likewise spill over into the fashioning of upgraded analytics engines that undergird prognostic supervision over physiological condition and growth trajectories.

Acknowledgements

We would like to express our heartfelt gratitude to the Research and Development Institute at Rajamangala University of Technology Suvarnabhumi (<http://rdi.rmutsb.ac.th>) for their generous financial support, which made this research possible. The funding provided by the institute was crucial in covering the costs associated with data collection, laboratory experiments, and other essential resources. Without this support, the successful completion of this project would not have been achievable.

Conflict of interest

None

Financial support

None

Ethics statement

None

Published online: 15 March 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Abadula TA, Jilo SA, Hussein JA, Abadura SZ. Poultry production status, major constraints, and future prospective. *J World Poult Sci*. 2022;1(1):22-8.
- Chowdhury MAH, Ashrafudoulla M, Mevo SIU, Mizan MFR, Park SH, Ha S. Current and future interventions for improving poultry health and poultry food safety and security: a comprehensive review. *Compr Rev Food Sci Food Saf*. 2023;22(3):1555-96.
- Choukidar GA, Dawande N. Smart poultry farm automation and monitoring system. In: 2017 International Conference on Computing, Communication, Control and Automation (ICCUBEA). IEEE; 2017. p. 1-5.
- Berrahal FZ. Intelligent Solutions for Modern Agriculture: Leveraging Artificial Intelligence in Smart Farming Practices. 2024; p. 269-76.
- Choudhary R, Mantri SP, Barse VA, Chitnis S. Leveraging AI in smart agro-informatics: a review of data science applications. *Int Res J Adv Eng Manag*. 2024;2(6):1964-75.
- Yang D, Cui D, Ying Y. Development and trends of chicken farming robots in chicken farming tasks: a review. *Comput Electron Agric*. 2024;221:108916.
- Chen Y, Wang S. Poultry carcass visceral contour recognition method using image processing. *J Appl Poult Res*. 2018;27(3):316-24.
- Landge P, Kiplinger G, Sakhare V, Rajankar O, Shinde B. Smart Poultry Farm Automation. 2024; p. 198-204.
- Taleb HM, Mahrose K, Abdel-Halim AA, Kasem H, Ramadan GS, Fouad AM, et al. Using artificial intelligence to improve poultry productivity – a review. *Ann Anim Sci*. 2025;25(1):23-33.
- Goyal V, Yadav A, Mukherjee R. A literature review on the role of internet of things, computer vision, and sound analysis in a smart poultry farm. *ACS Agric Sci Technol*. 2024;4(4):368-88.
- Hamadani H, Hamadani A, Shabir S. Artificial intelligence in animal farms for management and breeding. In: *A Biologist's Guide to Artificial Intelligence*. Elsevier; 2024. p. 167-82.
- Egon K, Karl L, Eugene R. Animal Welfare and Ethical Considerations in Precision Livestock Farming (PLF). 2023.
- McBratney A, Whelan B, Ancev T, Bouma J. Future directions of precision agriculture. *Precis Agric*. 2005;6(1):7-23.
- Zhang N, Wang M, Wang N. Precision agriculture—a worldwide overview. *Comput Electron Agric*. 2002;36(2):113-32.
- Choudhari A, Bhoyar D, Badole W. Design of an IoT based monitoring system utilizing deep reinforcement learning and ensemble learning for precision agriculture. In: 2024 2nd International Conference on Sustainable Computing and Smart Systems (ICSCSS).

IEEE; 2024. p. 845-52.

Kumar A, Karn N, Sharma H. IoT, AI, and robotics applications in the agriculture sector. In: Satapathy S, Muduli K, editors. *Advances in Business Information Systems and Analytics*. IGI Global; 2024. p. 243-72.

Kim K, Park J. A smart farm seedling automation system using machine vision. *Asia-Pac J Converge Res Interchange*. 2024;10(5):45-54.

Mavridou E, Vrochidou E, Papakostas GA, Pachidis T, Kaburlasos VG. Machine vision systems in precision agriculture for crop farming. *J Imaging*. 2019;5(12):89.

Ghazal S, Munir A, Qureshi WS. Computer vision in smart agriculture and precision farming: techniques and applications. *Artif Intell Agric*. 2024;13:64-83.

Carranza-García M, Torres-Mateo J, Lara-Benítez P, García-Gutiérrez J. On the performance of one-stage and two-stage object detectors in autonomous vehicles using camera data. *Remote Sens*. 2020;13(1):89.

He K, Gkioxari G, Dollar P, Girshick R. Mask r-CNN. *IEEE Trans Pattern Anal Mach Intell*. 2020;42(2):386-97.

Soviany P, Ionescu RT. Optimizing the trade-off between single-stage and two-stage deep object detectors using image difficulty prediction. In: 2018 20th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC). IEEE; 2018. p. 209-14.

Xu B, Wang W, Falzon G, Kwan P, Guo L, Sun Z, et al. Livestock classification and counting in quadcopter aerial images using Mask R-CNN. *Int J Remote Sens*. 2020;41(21):8121-42.

Jabir B, Moutaouakil KE, Falih N. Developing an Efficient System with Mask R-CNN for Agricultural Applications. 2023.

Mu X, He L, Heinemann P, Schupp J, Karkee M. Mask R-CNN based apple flower detection and king flower identification for precision pollination. *Smart Agric Technol*. 2023;4:100151.

Duong-Trung H, Duong-Trung N. Integrating YOLOv8-agri and DeepSORT for advanced motion detection in agriculture and fisheries. *EAI Endorsed Trans Indust Netw Intell Syst*. 2024;11(1).

Ye J, Wu Y, Rong W. Based on the optimization and performance evaluation of YOLOv8 object detection model with multi-backbone network fusion. In: 2024 IEEE International Conference on Mechatronics and Automation (ICMA). IEEE; 2024. p. 269-74.

Thanh HV, Thu NT, Phuong TM, Vuong ND, Jo KH. Inverted bottleneck convolution module for YOLOv8. In: 2024 IEEE 33rd International Symposium on Industrial Electronics (ISIE). IEEE; 2024. p. 1-6.

Wang G, Chen Y, An P, Hong H, Hu J, Huang T, et al. UAV-YOLOv8: a small-object-detection model based on improved YOLOv8 for UAV aerial photography scenarios. *Sensors*. 2023;23(16):7190.

Wang T, Zhang K, Zhang W, Wang R, Wan S, Rao Y, et al. Tea picking point detection and location based on mask-RCNN. *Inf Process Agric*. 2023;10(2):267-75.

Li H, Guo C, Yang Z, Chai J, Shi Y, Liu J, et al. Design of field real-time target spraying system based on improved YOLOv5. *Front Plant Sci*. 2022;13:1072631.

Zhang L, Ding G, Li C, Li D. DCF-yolov8: an improved algorithm for aggregating low-level features to detect agricultural pests and diseases. *Agronomy*. 2023;13(8):2012.

Wang X, Liu J. Vegetable disease detection using an improved YOLOv8 algorithm in the greenhouse plant environment. *Sci Rep*. 2024;14(1):4261.

Yang G, Wang J, Nie Z, Yang H, Yu S. A lightweight YOLOv8 tomato detection algorithm combining feature enhancement and attention. *Agronomy*. 2023;13(7):1824.

Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. Version 1. 2019.

Singh M, Kumar R, Tandon D, Sood P, Sharma M. Artificial intelligence and IoT based monitoring of poultry health: a review. In: 2020 IEEE International Conference on Communication, Networks and Satellite (Comnetsat). IEEE; 2020. p. 50-4.

Ahmed G, Malick RAS, Akhunzada A, Zahid S, Sagri MR, Gani A, et al. An approach towards IoT-based predictive service for early detection of diseases in poultry chickens. *Sustainability*. 2021;13(23):13396.

Cacic S, Popovic T, Krco S, Nedic D, Babic D, Jovovic I. Developing edge AI computer vision for smart poultry farms using deep learning and HPC. *Sensors*. 2023;23(6).

Islam MM, Sourov Tonmoy S, Quayum S, Sarker AR, Umme Hani S, Mannan MA. Design and implementation of automated poultry farm with distinguish features. In: 2019 International Conference on Robotics, Electrical and Signal Processing Techniques (ICREST). IEEE; 2019. p. 273-6.

Kim SJ, Lee MH. Design and implementation of a malfunction detection system for livestock ventilation devices in smart poultry farms. *Agriculture*. 2022;12(12):2150.

Hewawasam T, Gunasekara J, Jayarathna N, Indrawansa Y, Rathnayake S, Panduwawala P. Integrated AI-based system for comprehensive poultry management. In: 2023 5th International Conference on Advancements in Computing (ICAC). IEEE; 2023. p. 382-7.

Vinod A, Mohanty DC, John A, Depuru BK. Application of Artificial Intelligence in Poultry Farming - Advancing Efficiency in Poultry Farming by Automating the Egg Counting Using Computer Vision System. 2023.

Natho P, Chamuthai L, Boonying S, Tantidontanet N. A prototype platform for automated chicken feeding control with an embedded system. *Int J Adv Sci Comput Eng*. 2023;5(1):66-74.

Ciaglia F, Zuppichini FS, Guerrie P, McQuade M, Solawetz J. Roboflow 100: A Rich, Multi-Domain Object Detection Benchmark. Version 3. 2022.

Lin TY, Maire M, Belongie S, Bourdev L, Girshick R, Hays J, et al. Microsoft COCO: Common Objects in Context. Version 3. 2014.

Raschka S, Patterson J, Nolet C. Machine learning in Python: main developments and technology trends in data science, machine learning, and artificial intelligence. *Information*. 2020;11(4):193.

Florencio F, Valenc T, Moreno ED, Junior MC. Performance analysis of deep learning libraries: TensorFlow and PyTorch. *J Comput Sci*. 2019;15(6):785-99.

Wen H, Huang C, Guo S. The application of convolutional neural networks (CNNs) to recognize defects in 3d-printed parts. *Materials*. 2021;14(10):2575.

Carneiro T, Medeiros Da Nobrega RV, Nepomuceno T, Bian GB, De Albuquerque VHC, Filho PPR, et al. Performance analysis of Google Colaboratory as a tool for accelerating deep learning applications. *IEEE Access*. 2018;6:61677-85.