

Deep Learning Pipelines for Animal Re-Identification: A Survey from Detection to Population Inference

¹ Serif Oyindamola Oyesiji, ² Toussida Fatah Tanguy Minoungou, ³ Chukwudera Obumneke Anunagba, ⁴Harouna Wendpanga Yann Christian Sankara.

¹ Harrisburg University of Science and Technology, Harrisburg, PA; oyesijiserif@gmail.com

² Independent Researcher, Dallas, TX, USA; tanguyminoungou@gmail.com

³ Barclays, USA; chiokemuo@gmail.com

⁴Independent Researcher, Dallas, TX, USA; christiansankara@gmail.com

Corresponding Author; oyesijiserif@gmail.com

Abstract: Identifying individual animals from images is the quiet infrastructure of modern ecology: abundance estimation, movement analysis, social structure, and long-term monitoring all presuppose knowing which animal is which, and deep learning has transformed the identification from expert craft into a computational pipeline. A rapidly growing literature now spans that pipeline, yet its contributions are published stage by stage, detectors in vision venues, embedding losses in retrieval venues, species systems in ecology venues, and practitioners assembling working systems must synthesize across them. This survey reviews deep learning pipelines for animal re-identification by their anatomy, following the image from acquisition to identity. It first fixes the problem formulations, closed-set classification of known individuals, open-set identification with new-individual discovery, and verification, and the ecological constraints, imbalanced sightings, growing populations, and cross-deployment shift, that distinguish animal re-identification from its person and face antecedents. It then reviews each stage: detection and localization, where camera-trap-scale detectors and species-tuned localizers determine everything downstream; preprocessing, where quality filtering, viewpoint handling, and augmentation govern the effective training distribution; identity representation learning, the survey's core, organizing classification-based, metric-learning, and hybrid objectives, the treatment of patterned coats, fins, and faces, and the growing role of transfer and self-supervised pre-training under scarce labels; and matching, where ranking, gallery management, open-set thresholds, and scalable similarity search convert embeddings into decisions. Datasets and evaluation methodology are reviewed with attention to protocol pitfalls, identity leakage across splits foremost, and the deployment literature, platform systems and the coupling to capture-recapture inference, closes the pipeline. The survey consolidates challenges, open-set calibration, distribution shift, label scarcity, and long-tail sightings, and directions for the field.

1. Introduction

Individual identification predates computer vision in ecology by decades: natural-marking studies established that coat patterns, fin edges, and scars distinguish individuals reliably enough to ground population analysis (Karanth, 1995; Karanth & Nichols, 1998; Jackson et al., 2006), photographic matching was systematized by hand and then by classical algorithms over pattern descriptors (Kelly, 2001; Hiby et al., 2009; Arzoumanian et al., 2005; Town et al., 2013; Crall et al., 2013), and the biometric framing of the enterprise was articulated as a research program in its own right (Kühl & Burghardt, 2013; Awad, 2016), resting on the animal-science tradition of phenotypic characterization that supplies the trait vocabulary such framing formalizes (Ekeocha et al., 2021) and on applied biometric practice for its operational grammar of enrolment, matching, and spoof resistance (Ussher-Eke et al., 2025; Adeniyi et al., 2025). The sensing side scaled in parallel, camera-trap networks becoming standard ecological instrumentation with well-documented survey design and data-volume consequences (O'Connell et al., 2011; Burton et al., 2015; Steenweg et al., 2017; Wearn & Glover-Kapfer, 2019; Ahumada et al., 2020), and citizen-science platforms demonstrating both the scale of accumulating imagery and the labeling bottleneck it created (Swanson et al., 2015); the instrumentation itself rests on power-, bandwidth-, and duty-cycle-constrained sensor networks whose engineering literature governs how much imagery a deployment can actually capture and return (Asiedu & Quainoo, 2023b; Asiedu et al., 2026; Kumuyi et al., 2024a).

Deep learning entered through the general vision stack, convolutional representation learning (Krizhevsky et al., 2012; LeCun et al., 2015; He et al., 2016; Simonyan & Zisserman, 2015; Szegedy et al., 2015) trained at dataset scale (Deng et al., 2009; Lin et al., 2014; Everingham et al., 2010), and reached wildlife imagery first as species-level automation: deep classification of camera-trap images at accuracy and volume that changed what monitoring programs could process (Norouzzadeh et al., 2018; Tabak et al., 2019; Willi et al., 2019; Gomez Villa et al., 2017), with the cross-deployment generalization problem identified early and influentially (Beery et al., 2018). The same convolutional stack was in the same period carried into adjacent applied domains, malware and threat classification at edge scale (Idika et al., 2021; Odejobi et al., 2025), aerial vegetation and corridor classification from UAV imagery (Dagodzo et al., 2022), and clinical image analysis descended from template matching (Ejofodomi et al., 2014), and the cross-domain pattern is informative: the architecture transfers easily and the evaluation discipline does not. Individual-level identification followed on the machinery of person re-identification and face recognition, metric embeddings trained so that images of one individual cluster across views (Chopra et al., 2005; Hadsell et al., 2006; Schroff et al., 2015; Hermans et al., 2017; Zheng et

al., 2015; Ye et al., 2022; Deng et al., 2019), and a species-spanning applied literature demonstrated the transfer, from chimpanzee faces (Freytag et al., 2016; Schofield et al., 2019; Deb et al., 2018) through patterned and marked mammals (Li et al., 2020; Clapham et al., 2020; Nepovinnikh et al., 2020) to marine photo-identification at population scale (Bogucki et al., 2019; Weideman et al., 2017; Cheeseman et al., 2022), with dedicated surveys tracking the area's growth (Schneider et al., 2019; Ravor & Sudarshan, 2020; Vidal et al., 2021) and the broader conservation-technology synthesis situating it among machine learning's ecological roles (Tuia et al., 2022; Christin et al., 2019; Weinstein, 2018; Neethirajan, 2020), alongside the parallel build-out of predictive modeling in agricultural and food systems that shares its data conditions and its decision-support ambitions (Michael & Ogunsola, 2023c).

This survey's organizing claim is that the field's unit of progress is the pipeline, not the model. Working identification systems are compositions, detection, preprocessing, representation, matching, gallery management, and downstream inference, whose stages are studied in different communities and whose end-to-end behavior is governed as much by stage interfaces as by any component's benchmark score. The claim is not special to ecology: the software and data engineering literatures on staged delivery pipelines make the same observation, that composition, versioning, and interface contracts determine deployed behavior more reliably than component quality does (Badmus et al., 2018; Ojukwu et al., 2025; Oyesiji et al., 2024). The survey therefore reviews the literature stage by stage, with formulations and constraints fixed first. Section 2 states scope. Section 3 defines problem formulations and ecological constraints. Sections 4 through 7 review the pipeline stages. Section 8 reviews datasets and evaluation methodology, Section 9 deployment and ecological integration, Section 10 challenges, Section 11 future directions, and Section 12 concludes.

2. Scope and Approach of the Survey

The survey covers deep learning methods for identifying individual animals from images and image sequences, across camera traps, handheld and platform photography, and managed-animal settings, together with the pipeline stages that surround the identification model in deployed systems. Included are the detection and preprocessing stages as they serve identification, the representation-learning core across objective families and body-part modalities, matching and gallery machinery including open-set operation, and the evaluation and deployment literatures specific to the animal setting; the person re-identification and face recognition literatures are included as method sources with their transfer conditions examined rather than assumed (Ye et al., 2022; Zheng et al., 2016; Deng et al., 2019). Excluded are individual identification from non-image modalities, acoustic and genetic, tracking as an end in itself, and species-level classification except as the pipeline stage it constitutes; livestock identification is included where its methods generalize (Andrew et al., 2017; Hansen et al., 2018; Kumar et al., 2018; Qiao et al., 2019), since managed settings often pioneer techniques wild settings later adopt (Wurtz et al., 2019); the husbandry literature surrounding those systems, growth, condition, and physiological monitoring of individually managed animals, is the demand side that makes livestock identification worth automating in the first place (Oluwadele et al., 2024; Tawose et al., 2024).

The synthesis approach follows the pipeline for structure and the surveys for coverage control, using the field's existing reviews as maps (Schneider et al., 2019; Ravor & Sudarshan, 2020; Vidal et al., 2021) while carrying the methodological weight on primary studies. Evaluation claims receive particular scrutiny because the field's protocols vary in ways that change conclusions: split design relative to individual identity, closed-set versus open-set reporting, and per-individual sample imbalance are treated as first-class methodological variables throughout, with the metric-learning reality-check literature imported as a caution against protocol-driven progress claims (Musgrave et al., 2020).

3. Problem Formulations and Ecological Constraints

Three formulations organize the identification stage. Closed-set identification treats individuals as classes and identification as classification, valid where the population is enumerated and stable, small managed groups, habituated study populations, and directly served by softmax training with class-imbalance handling, the same framing applied domains adopt wherever a fixed and enumerable label set makes classification the natural formulation (Akeju et al., 2018). Verification asks whether two images show the same individual, the pairwise formulation on which contrastive training was founded (Chopra et al., 2005; Hadsell et al., 2006). Verification recurs wherever identity claims must be scored rather than classified, and the operational literature there, synthetic-identity detection and real-time risk scoring against calibrated thresholds, is a useful external check on how such scores behave at scale (Elebe et al., 2023; Fadayomi et al., 2024). Open-set identification, the ecologically general case, ranks a query against a gallery of known individuals and must additionally decide whether the query is anyone known at all, since wild populations grow, immigrate, and are never fully enumerated; it is served by metric embeddings with thresholded acceptance, and its calibration, the new-individual decision, is the field's distinctive difficulty, largely absent from the person re-identification setting the methods came from (Ye et al., 2022; Schneider et al., 2019) and better developed in the anomaly and novelty detection literatures, where deciding that an observation belongs to no known class is the primary object rather than a downstream complication (Atakpa & Fobellah, 2023; Adebayo et al., 2023; Adegbite et al., 2023; Ominyi et al., 2024a).

The ecological constraints reshape every stage. Sightings are extremely long-tailed, a few individuals photographed often and many once, so per-class sample assumptions of face-recognition training fail by construction (Cheeseman et al., 2022; Miele et

al., 2021). Appearance evidence is species-dependent, stable patterned pelage and stripes, growing and scarring fin edges, aging faces, and molting or seasonal coats, so the invariances an embedding must learn, and the body parts a pipeline must localize, are biological decisions before they are architectural ones (Kühl & Burghardt, 2013; Vidal et al., 2021). Deployment shift is severe: models trained at some sites and seasons meet new backgrounds, cameras, and populations, the generalization failure documented for camera-trap recognition generally (Beery et al., 2018) and inherited by identification with the added twist that the class set itself changes. And the downstream consumer is statistical inference, capture-recapture estimators whose assumptions, correct identity assignment foremost, convert identification errors into abundance bias, coupling pipeline error rates to scientific conclusions in a way the ecology literature has long quantified for manual matching (Karanth & Nichols, 1998; Royle et al., 2013; Foster & Harmsen, 2012), the same dependence of population-level inference on case-level ascertainment that surveillance modeling contends with in epidemiological settings (Oparah et al., 2022).

4. Stage One: Detection and Localization

Identification begins with finding the animal, and at camera-trap scale detection is the throughput stage. The general detector lineage supplies the components, two-stage precision (Girshick et al., 2014; Girshick, 2015; Ren et al., 2015), single-stage speed in the YOLO and SSD families (Redmon et al., 2016; Redmon & Farhadi, 2017; Redmon & Farhadi, 2018; Bochkovskiy et al., 2020; Ge et al., 2021; Wang et al., 2023; Liu et al., 2016), multi-scale feature pyramids and dense-detection losses (Lin et al., 2017a; Lin et al., 2017b), and transformer formulations (Carion et al., 2020), with efficiency-scaled variants where edge deployment binds (Tan & Le, 2019; Tan et al., 2020; Howard et al., 2017; Sandler et al., 2018); the model-compression literature that has grown around on-device inference, quantization, pruning, and distillation under energy budgets, is the general form of that constraint and increasingly the source of its solutions (Nwakamma et al., 2025; Oyesiji et al., 2025). The wildlife-specific contribution is generalization engineering: generic animal detectors trained across ecosystems to serve as reusable first stages for downstream pipelines, the pattern established for camera-trap review at scale (Beery et al., 2019; Norouzzadeh et al., 2021), motivated directly by the cross-location failures of site-trained models (Beery et al., 2018), and integrated with species classification in the large automation studies (Norouzzadeh et al., 2018; Tabak et al., 2019; Willi et al., 2019). Aerial platforms extend the same stage to survey imagery, where UAV-borne detection has developed its own acquisition, overlap, and georeferencing discipline (Dagodzo, 2018b; Dagodzo & Ahiaeké Patrick, 2021b; Dagodzo & Ahiaeké Patrick, 2022), and remote-sensing and geospatial pipelines supply the georeferenced framing that camera-trap work handles through station metadata (Badmus, 2025; Ahiaeké Patrick et al., 2021).

Identification-oriented localization goes finer than boxes. Species pipelines localize the identity-bearing region, faces for primates and bears (Freytag et al., 2016; Deb et al., 2018; Clapham et al., 2020; Schofield et al., 2019), flanks for striped and spotted species (Li et al., 2020; Nepovinnikh et al., 2020), fins and flukes for marine individuals (Weideman et al., 2017; Bogucki et al., 2019), and the multi-stage pattern, detect animal, localize part, normalize orientation, was established as the pipeline template in the field's system papers (Parham et al., 2018; Crall et al., 2013). Stage-one errors are pipeline errors: missed detections delete sightings from downstream inference, and mislocalized parts feed embeddings noise, which is why the deployment literature treats detector recall on the target fauna, not benchmark mean average precision, as the stage's operative metric (Beery et al., 2019; Schneider et al., 2019), the same recall-weighted logic that governs continuous monitoring systems wherever a missed event is costlier than a reviewed false alarm (Obriki et al., 2023).

5. Stage Two: Preprocessing, Quality, and the Effective Training Distribution

Between localization and representation sits the stage the papers mention least and the deployments tune most. Quality filtering removes blur, occlusion, extreme poses, and truncated crops before they pollute galleries, with sample-quality weighting inside training and matching as the softer alternative, the curation problem that the data-centric training literature treats as a first-class pipeline stage rather than a preprocessing convenience (Oyesiji et al., 2024); viewpoint handling routes left and right flanks, whose patterns are independent evidence for bilaterally asymmetric species, into separate galleries or viewpoint-conditioned models, a biological fact with pipeline consequences the system literature encodes explicitly (Parham et al., 2018; Crall et al., 2013). Orientation and alignment normalization, rotating fins, rectifying flanks, aligning faces, imports the alignment lesson from face recognition, where geometric normalization preceded and enabled the embedding gains (Schroff et al., 2015; Deng et al., 2019), and its animal instantiations range from keypoint-based part rectification to the trailing-edge parameterizations of marine photo-identification (Weideman et al., 2017; Hiby et al., 2009).

Augmentation governs the effective training distribution under label scarcity. The general augmentation repertoire applies (Shorten & Khoshgoftaar, 2019), with the field's specific guidance following its invariance analysis: photometric and background variation simulate the deployment shift that harms generalization (Beery et al., 2018), while transformations that alter the identity-bearing pattern itself, aggressive elastic deformation on patterned coats, are identity-destructive and avoided, the augmentation-as-invariance-declaration view made explicit. Transfer learning is the stage's silent partner: with per-individual labels scarce, pipelines initialize from large-scale pre-training as a matter of course (Deng et al., 2009; Pan & Yang, 2010; Zhuang et al., 2021), increasingly adapting them parameter-efficiently rather than by full fine-tuning, the regime the on-device adaptation literature has characterized

in detail (Oyesiji et al., 2023), and the self-supervised families, contrastive and distillation-based pre-training (Chen et al., 2020; He et al., 2020; Grill et al., 2020; Caron et al., 2021), offer label-free adaptation on the deployment's own unlabeled imagery, an increasingly reported ingredient the survey returns to under future directions. Label production is itself a pipeline stage with its own literature: model-assisted and human-centered annotation workflows, in which a model proposes and an expert disposes, are how most per-individual label sets are actually built, and their design governs both throughput and label noise (Ladapo et al., 2024; Ladapo et al., 2022).

6. Stage Three: Identity Representation Learning

The representation stage is the survey's core, and its literature organizes by objective family. Classification-based training treats individuals as classes over a shared backbone (He et al., 2016; Krizhevsky et al., 2012), with margin-based softmax variants importing the face-recognition refinement that inter-class separation can be enforced in the loss geometry itself (Deng et al., 2019; Wen et al., 2016); it is simple, strong where classes are sampled adequately, and structurally awkward for open populations, since new individuals require head changes, which is why its dominant animal use is as a training objective whose penultimate features serve as embeddings rather than as a deployed classifier (Schofield et al., 2019; Freytag et al., 2016). Metric-learning objectives train the embedding directly: contrastive and triplet families with hard mining (Chopra et al., 2005; Hadsell et al., 2006; Schroff et al., 2015; Hermans et al., 2017), the strong-baseline training recipes person re-identification consolidated (Luo et al., 2019; Sun et al., 2018; Wang et al., 2018), and hybrid classification-plus-metric objectives that the animal literature adopts as its default composite (Schneider et al., 2020; Miele et al., 2021; Nepovinnikh et al., 2020). The metric-learning evaluation caution applies with force: controlled comparisons found much of the reported loss-function progress attributable to protocol variance rather than objectives (Musgrave et al., 2020), and the animal literature's heterogeneous protocols make cross-paper loss rankings weak evidence, a methodological point Section 8 develops.

Species morphology stratifies the design space more than architecture does. Patterned species, stripes, spots, and pelage, reward texture-sensitive representations and tolerate strong geometric normalization, the regime where deep embeddings displaced the classical descriptor matchers that founded the area (Crall et al., 2013; Lowe, 2004; Bay et al., 2008; Dalal & Triggs, 2005); edge-and-shape species, fins and flukes, carry identity in contours where alignment and curvature representation dominate (Weideman et al., 2017; Bogucki et al., 2019; Cheeseman et al., 2022); face-bearing species import the face stack most directly, detection, alignment, margin-trained embeddings (Deb et al., 2018; Clapham et al., 2020; Schofield et al., 2019; Deng et al., 2019); and weakly-marked species remain the hard case, where identity evidence is subtle, per-image reliability is low, and multi-image aggregation across a sighting becomes part of the representation itself (Miele et al., 2021; Vidal et al., 2021). Backbone choice tracks the general vision timeline, residual and efficiency-scaled convolutional networks (He et al., 2016; Tan & Le, 2019) with transformer backbones entering the recent literature (Dosovitskiy et al., 2021; Liu et al., 2021), and interpretability tooling, activation attribution and embedding visualization, serves the field's distinctive validation need, confirming that models attend to biological identity evidence rather than background shortcuts (Selvaraju et al., 2017; van der Maaten & Hinton, 2008; Beery et al., 2018); the wider interpretable machine learning literature makes the same argument for operational systems, that attribution is a precondition of trust wherever a model's output feeds a consequential decision (Komi & Adamolekun, 2021), and visualization is the reporting layer through which that scrutiny is usually exercised (Ogbole et al., 2021).

7. Stage Four: Matching, Galleries, and Open-Set Decisions

Embeddings become identifications through a matching layer whose design the deployment literature, more than the methods literature, has specified. Ranking against the gallery is nearest-neighbor retrieval in embedding space, exact at small scale and approximate at population scale, where graph-based and quantized indexes make million-image galleries interactive (Malkov & Yashunin, 2020; Johnson et al., 2021), with the ordinary systems constraints of served retrieval infrastructure, index build cost, memory residency, and elastic scaling under bursty query load (Ahmed et al., 2020; Mayo et al., 2021); the person re-identification evaluation stack, cumulative match characteristics and mean average precision over ranked galleries, transfers as the stage's metric vocabulary (Zheng et al., 2015; Ye et al., 2022). Gallery composition is a managed policy in working systems: multiple exemplars per individual across viewpoints and years, quality-weighted prototype aggregation, and curation of the confusable, the operational layer the platform systems built and the methods papers inherit silently (Parham et al., 2018; Berger-Wolf et al., 2017; Crall et al., 2013).

The open-set decision is the stage's scientific crux. Accepting a match wrongly merges individuals; rejecting wrongly fragments them; and both propagate into capture-recapture inference with opposite biases the ecological literature has quantified since the manual era (Foster & Harmsen, 2012; Karanth & Nichols, 1998). Deployed practice therefore favors human-in-the-loop operation at calibrated thresholds, the system presents ranked candidates with similarity scores, experts confirm, and abstention routes to review, the workflow the platform literature institutionalized (Berger-Wolf et al., 2017; Parham et al., 2018), a specific instance of the human-in-the-loop machine learning pattern whose general treatment covers exactly the design questions the platforms face, when to query the expert, how to route abstentions, and how to treat reviewer disagreement as signal rather than

noise (Ladapo et al., 2022; Ladapo et al., 2026) and the large marine campaigns ran at population scale (Cheeseman et al., 2022). Threshold calibration against verification score distributions, per species and per deployment, is the reported practice, the same operating-point discipline that real-time risk scoring systems formalize as explicit score-to-action policy (Fadayomi et al., 2024; Oluwo et al., 2024); principled new-individual detection, open-set recognition with calibrated novelty probabilities rather than tuned cutoffs, remains thinner in the animal literature than the problem's centrality warrants, a gap Section 10 records. Sequence and multi-image evidence strengthens the stage where it exists: sightings are usually bursts, and aggregating ranks or embeddings across a burst before deciding is the cheap ensemble deployments use, with video tracking supplying the association where imagery is continuous (Bewley et al., 2016; Wojke et al., 2017; Zhang et al., 2022; Luiten et al., 2021).

8. Datasets and Evaluation Methodology

The field's evaluation substrate is a patchwork the survey treats critically. Beyond comparability, that substrate raises governance questions, provenance, sensitivity classification, and traceable licensing of imagery contributed by many parties, which the data-governance literature treats as design requirements rather than administrative afterthoughts (Aliliele et al., 2024a; Orise & Niniola, 2025). Species datasets with individual labels anchor method comparison, striped and spotted mammal benchmarks (Li et al., 2020), seal and primate collections (Nepovinnikh et al., 2020; Freytag et al., 2016; Deb et al., 2018), livestock sets (Andrew et al., 2017; Hansen et al., 2018; Kumar et al., 2018), and the population-scale marine competitions whose long-tailed identity distributions mirror deployment reality (Cheeseman et al., 2022; Bogucki et al., 2019), atop the camera-trap corpora that serve the detection stage (Swanson et al., 2015; Yu et al., 2013; Chen et al., 2014). Three protocol variables decide what results mean. Split design relative to identity: closed-set splits share individuals across train and test and measure classification; disjoint-identity splits measure the transferable embedding, the setting deployment actually needs; and mixed or unreported designs, common in the applied literature, blur the two, the field's version of the leakage failure the general methodology literature warns against and the metric-learning reality check demonstrated empirically (Musgrave et al., 2020). Open-set reporting: gallery-and-query evaluation with novel-identity queries, and the resulting acceptance operating curves, is the deployment-faithful protocol and remains the minority practice. Imbalance handling: per-individual sample counts span orders of magnitude, so macro-averaged and per-individual reporting diverge from pooled accuracy, and the long-tail literature's cautions apply directly (Cheeseman et al., 2022; Miele et al., 2021). The surveys document the resulting comparability problem across the applied literature (Schneider et al., 2019; Ravoor & Sudarshan, 2020; Vidal et al., 2021), and this survey's synthesis position is procedural: disjoint-identity splits, open-set protocols, and per-individual reporting constitute the minimal standard under which the field's numbers compose into knowledge. Release-gating practice from the model-governance literature offers a complementary mechanism, evaluation treated as a versioned artifact with pre-registered acceptance criteria rather than a table assembled after the result is known (Nwakamma et al., 2024a).

9. Deployment Pipelines and Ecological Integration

The deployed literature closes the loop the methods papers open. Platform systems assembled the full pipeline as infrastructure, detection, part localization, matching, expert review, and gallery curation as a maintained service across species and institutions (Berger-Wolf et al., 2017; Parham et al., 2018), with lineage back to the algorithmic matchers that first made photo-identification computational (Crall et al., 2013; Hiby et al., 2009; Arzoumanian et al., 2005; Town et al., 2013). Population-scale campaigns demonstrated the pipeline at its intended magnitude, marine photo-identification across decades of imagery and thousands of individuals with hybrid machine-expert matching (Cheeseman et al., 2022; Bogucki et al., 2019), and study systems from chimpanzee video archives (Schofield et al., 2019) to camera-trap networks (Norouzzadeh et al., 2021; Beery et al., 2019) established the automation's scientific throughput. Adjacent operational settings run a similar stack for management rather than science: airport wildlife-hazard programs use detection and identification to quantify strike risk and target mitigation on a decision loop measured in hours rather than seasons (Adeyelu & Dagodzo, 2024; Adeyelu, 2020; Adeyelu, 2026), and health surveillance programs pair individual-level case ascertainment with population reporting in the same two-layer arrangement (Oparah et al., 2023; Atakpa et al., 2023).

Integration with inference is the pipeline's purpose and its accountability. Identification outputs feed capture-recapture and spatially explicit abundance models whose estimators assume identity assignments of known reliability (Karanth & Nichols, 1998; Royle et al., 2013; O'Connell et al., 2011), and the misidentification-bias literature translates matching error rates into estimator bias, making the pipeline's precision-recall operating point a parameter of the science rather than an engineering detail (Foster & Harmsen, 2012; Bolger et al., 2012). Identification also feeds health and disease work on marked individuals, where longitudinal identity is what allows infection, exposure, and condition histories to be attached to an animal rather than to a site, the individual-level substrate One Health monitoring assumes (Ayinde, 2025; Ayinde et al., 2024; Olorunkosebi et al., 2023). The deployment synthesis across this literature is consistent: automation multiplies expert attention rather than replacing it, calibrated abstention protects the inference, and the gallery, versioned, curated, and quality-weighted, is the asset on which long-term monitoring value compounds (Berger-Wolf et al., 2017; Cheeseman et al., 2022; Tuia et al., 2022). Where galleries span institutions that cannot pool imagery outright, federated training and lifecycle-aware data management offer the architecture under which collaboration remains possible (Sanni et

al., 2023; Ladapo et al., 2026), and the operational reporting layer that comparable monitoring programs treat as standard, dashboards over throughput, queue depth, reviewer agreement, and model drift, remains underdeveloped in this one (Filani et al., 2022; Shittu et al., 2023; Ilodigwe & Adesemoye, 2024a).

10. Challenges and Open Problems

Five challenges consolidate. Open-set calibration: the new-individual decision lacks the principled treatment its inferential weight demands, with tuned thresholds standing in for calibrated novelty probabilities and little animal-specific research on open-set recognition despite its centrality (Schneider et al., 2019; Vidal et al., 2021). Distribution shift: cross-site, cross-camera, and cross-season transfer degrades identification as it does detection (Beery et al., 2018), and the field lacks shift-explicit benchmarks at the individual level, so robustness claims rest on anecdote. Robustness of this kind is also a distributional question, since models that perform unevenly across sites distribute their errors unevenly across the programs that depend on them, an argument made explicit in another applied forecasting domain (Komi, 2023); and where identification carries enforcement or commercial consequence, deliberate manipulation joins incidental shift as a threat model the biometric and adversarial machine learning literatures already treat seriously (Adeniyi et al., 2025; Adebayo et al., 2022). Label economics: per-individual annotation is expert work, the long tail guarantees most individuals are barely sampled, and the self-supervised and transfer machinery (Chen et al., 2020; He et al., 2020; Grill et al., 2020; Zhuang et al., 2021) is promising but unevenly evaluated under the field's protocols, and the annotation workflows through which scarce expert attention is actually spent receive less design attention than their cost warrants (Ladapo et al., 2024). Appearance change over time: growth, scarring, molt, and aging erode gallery validity at species-specific rates, making temporal robustness and gallery-refresh policy an empirical question the long-running platforms are positioned to answer and the methods literature has barely posed (Cheeseman et al., 2022; Berger-Wolf et al., 2017). And evaluation validity: protocol heterogeneity, identity leakage in splits, and closed-set reporting inflate the apparent state of the art relative to deployment reality, the field-specific instance of the general metric-learning and benchmarking cautions (Musgrave et al., 2020), and the single reform with the highest ratio of value to cost.

11. Future Directions

Five directions follow from the synthesis. Protocol standardization first: community adoption of disjoint-identity splits, open-set evaluation with acceptance curves, per-individual reporting, and shift-structured test sets would make the literature's rapid output cumulative, the reform the surveys converge toward (Schneider et al., 2019; Ravor & Sudarshan, 2020; Musgrave et al., 2020). Foundation-style pre-training for animal identity: large-scale self-supervised pre-training on aggregated wildlife imagery, evaluated for transfer across species and deployments, would attack label economics at the root (Caron et al., 2021; Chen et al., 2020; Tuia et al., 2022), with parameter-efficient adaptation the likely mechanism by which such a backbone is specialized per species and per deployment (Oyesiji et al., 2023). Uncertainty-first pipelines: calibrated match and novelty probabilities propagated from embedding to estimator, so that capture-recapture inference consumes identification uncertainty rather than assuming it away, joining the pipeline to its statistical consumer formally (Royle et al., 2013; Foster & Harmsen, 2012); the forecasting literature that treats predictive uncertainty as a first-class output rather than a diagnostic afterthought is the closest available methodological model (Wedraogo & Sanni, 2024). And longitudinal science from the platforms: the deployed systems accumulate exactly the temporal, cross-site, expert-adjudicated data the open problems need, appearance drift, threshold stability, reviewer-model agreement, and their publication as research datasets would let the field study its own deployments with the rigor it applies to benchmarks (Berger-Wolf et al., 2017; Cheeseman et al., 2022). And edge-resident identification: on-camera detection and embedding extraction would cut transmission volumes to a fraction and make near-real-time alerting feasible in bandwidth-poor deployments, which is what the compression, on-device inference, and edge-intelligence literatures have been building toward (Nwakamma et al., 2025; Oyesiji et al., 2025; Komi & Ganiu, 2023; Asiedu et al., 2026), with the governance question, who is accountable for an automated identification acted on without review, arriving with it (Akintola et al., 2025). Capacity is the quiet precondition for all five, since the practitioners who will run these pipelines are trained by curricula that are themselves only beginning to absorb the methods (Joseph & Uzodu, 2024a; Orise & Niniola, 2025).

12. Conclusion

This survey has reviewed deep learning for animal re-identification as a pipeline: detection and localization determining what enters; preprocessing and augmentation shaping the effective distribution; representation learning, organized by objective family and species morphology, producing the identity embedding; and matching, gallery management, and open-set decision-making converting similarity into science, with datasets, evaluation protocols, and deployed platforms reviewed as the infrastructure that makes the stages a field. The synthesis yields a consistent picture. The representation core has been transformed by the metric-learning stack and is no longer the binding constraint for well-marked species; the binding constraints are at the edges, open-set

calibration, distribution shift, label economics, temporal appearance change, and in the middle of the field's own methodology, where protocol heterogeneity and identity leakage blur what has actually been achieved.

The direction of the field follows its consumers. Animal re-identification exists to serve population inference and long-term monitoring, and the pipeline view makes the obligations explicit: operating points are scientific parameters, abstention protects estimators, galleries are longitudinal assets, and evaluation must mirror the open, shifting, long-tailed populations the science addresses. The literature reviewed here shows a field with its components proven and its composition maturing; standardized protocols, uncertainty-carrying pipelines, and the research use of its own deployments are the steps by which its rapid progress becomes reliable ecological instrumentation.

References

- Adebayo, A., Adegbite, M. P., & Ahmed, M. O. (2022). Adversarial machine learning in critical infrastructure: A conceptual framework for threat modeling AI enabled OT systems. *World Journal of Innovation and Modern Technology*, 6(1), 184-234. <https://doi.org/10.56201/wjimt.v6.no1.2022.pg184.234>
- Adebayo, A., Adegbite, M. P., & Ahmed, M. O. (2023). AI augmented threat detection in industrial control systems: A systematic review of machine learning approaches for ICS anomaly detection. *International Journal of Engineering and Modern Technology*, 9(3), 287-340. <https://doi.org/10.56201/ijcsmt.v9.no3.2023.pg287.340>
- Adegbite, M. P., Adebayo, A., & Ahmed, M. O. (2023). ICS and SCADA threat detection architectures in energy sector networks: A systematic review of SIEM, NDR, and anomaly detection approaches. *World Journal of Innovation and Modern Technology*, 7(2), 121-182. <https://doi.org/10.56201/wjimt.v7.no2.2023.pg121.182>
- Adeniyi, A. I., Odejebi, O., & Taiwo, T. (2025). Countermeasures against bias and spoofing in modern facial recognition systems. *World Journal of Advanced Research and Reviews*, 25(1), 1914-1930.
- Adeyelu, O. O. (2020). An artificial intelligence-driven conceptual model for wildlife strike risk quantification and aircraft airworthiness impact assessment at high-traffic African airports. *Iconic Research and Engineering Journals*, 4(1), 339-368. <https://doi.org/10.64388/IREV4I1-1718482>
- Adeyelu, O. O. (2026). A root-cause investigation framework for recurring runway wildlife hazards and standardized corrective action protocols at Nigerian international airports. *International Journal of Social Sciences and Management Research*, 12(5), 842-880. <https://doi.org/10.67163/ijssmr.vol.12no5.2026.pg842.880>
- Adeyelu, O. O., & Dagodzo, D. (2024). Advances in artificial intelligence and data-driven wildlife hazard monitoring and incident reduction at international airports in West Africa. *International Journal of Scientific Research in Science and Technology*, 11(5), 872-910. <https://doi.org/10.32628/IJSRST52310285>
- Ahiaeke Patrick, M. C., Okonkwo, C. S., Mayo, W., & Okeke, O. T. (2021). Model for data driven vendor evaluation and bid selection using geospatial intelligence. Shodhshauryam, *International Scientific Refereed Research Journal*, 4(4), 426-443. <https://doi.org/10.32628/SHISRRJ214461>
- Ahmed, K. S., Odejebi, O. D., & Oshoba, T. O. (2020). Predictive model for cloud resource scaling using machine learning techniques. *Journal of Frontiers in Multidisciplinary Research*, 1(1), 173-183.
- Ahumada, J. A., Fegraus, E., Birch, T., Flores, N., Kays, R., O'Brien, T. G., Palmer, J., Schuttler, S., Zhao, J. Y., Jetz, W., Kinnaird, M., Kulkarni, S., Lyet, A., Thau, D., Duong, M., Oliver, R., & Dancer, A. (2020). Wildlife Insights: A platform to maximize the potential of camera trap and other passive sensor wildlife data for the planet. *Environmental Conservation*, 47(1), 1-6.
- Akeju, B., Edivri, J., Ogbale, J. I., Okoruwa, P. O., Fadayomi, O., & Abolaji, T. O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. *Iconic Research and Engineering Journals*, 1(9), 476-492. <https://doi.org/10.64388/IREV1I9-1713778>
- Akintola, A. S., Fawehinmi, Y. O., Aiyenitaju, O., Chinemerem, B., & Emmanuel, O. (2025). Digital transformation in telehealth: A systematic review of user trust, privacy, and regulatory governance in AI-powered remote monitoring systems. *Journal of Scientific Research and Reports*, 31(11), 163-182. <https://doi.org/10.9734/jsrr/2025/v31i113658>
- Aliliele, C., Mbonu, I. S., & Iwuanyanwu, U. (2024a). A conceptual framework for enterprise data sensitivity classification and regulatory traceability mechanisms. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 3103-3124. <https://doi.org/10.62225/2583049X.2024.4.6.5991>
- Andrew, W., Greatwood, C., & Burghardt, T. (2017). Visual localisation and individual identification of Holstein Friesian cattle via deep learning. *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2850-2859.
- Arzoumanian, Z., Holmberg, J., & Norman, B. (2005). An astronomical pattern-matching algorithm for computer-aided identification of whale sharks *Rhincodon typus*. *Journal of Applied Ecology*, 42(6), 999-1011.
- Asiedu, W., & Quainoo, R. (2023b). Toward maintenance free wireless infrastructure: Simulating performance and reliability in large scale intermittently powered IoT networks. *Gyanshauryam, International Scientific Refereed Research Journal*, 6(1), 489-510.
- Asiedu, W., Quainoo, R., & Asiedu, A. (2026). Predictive power management for intermittent IoT devices using lightweight machine learning under uncertain energy harvest. *Gulf Journal of Engineering and Technology*, 2(4), 108-118.
- Atakpa, M. I., & Fobellah, A. N. (2023). Anomaly detection in financial time-series data: A conceptual model for healthcare and banking applications. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(2), 967-997. <https://doi.org/10.32628/CSEIT2342441>
- Atakpa, M. I., Abolaji, T. O., & Akeju, B. (2023). A review of natural language processing for social media-based public health surveillance and analytics. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(6), 1030-1060. <https://doi.org/10.32628/CSEIT23906783>

- Awad, A. I. (2016). From classical methods to animal biometrics: A review on cattle identification and tracking. *Computers and Electronics in Agriculture*, 123, 423–435.
- Ayinde, M. K. (2025). West Nile virus dynamics in American robins (*Turdus migratorius*; AMRO) of Illinois [Unpublished thesis]. Western Illinois University.
- Ayinde, M. K., Penddah, A. L., Mensah, G. K., & Olorunkosebi, M. T. (2024). Environmental microplastics as vectors for avian virus transmission and their implications for food safety within a one health framework. *Journal of Current Research in Food Science*, 5(2), 210-219. <https://doi.org/10.22271/foodsci.2024.v5.i2c.288>
- Badmus, O. (2025). Integrated predictive and remote-sensing framework for early warning and regulatory compliance in environmentally sensitive urban zones. *IIARD International Journal of Geography and Environmental Management*, 11(12), 212-239.
- Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2018). A systematic review of CI/CD pipeline strategies in Salesforce DevOps: Tools, practices, and deployment outcomes. *Iconic Research and Engineering Journals*, 2(6).
- Bay, H., Ess, A., Tuytelaars, T., & Van Gool, L. (2008). Speeded-up robust features (SURF). *Computer Vision and Image Understanding*, 110(3), 346–359.
- Beery, S., Morris, D., & Yang, S. (2019). Efficient pipeline for camera trap image review. *arXiv preprint arXiv:1907.06772*.
- Beery, S., Van Horn, G., & Perona, P. (2018). Recognition in terra incognita. *Proceedings of the European Conference on Computer Vision*, 456–473.
- Berger-Wolf, T. Y., Rubenstein, D. I., Stewart, C. V., Holmberg, J. A., Parham, J., Menon, S., Crall, J., Van Oast, J., Kiciman, E., & Joppa, L. (2017). Wildbook: Crowdsourcing, computer vision, and data science for conservation. *arXiv preprint arXiv:1710.08880*.
- Bewley, A., Ge, Z., Ott, L., Ramos, F., & Upcroft, B. (2016). Simple online and realtime tracking. *Proceedings of the IEEE International Conference on Image Processing*, 3464–3468.
- Bochkovskiy, A., Wang, C.-Y., & Liao, H.-Y. M. (2020). YOLOv4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*.
- Bogucki, R., Cygan, M., Khan, C. B., Klimek, M., Milczek, J. K., & Mucha, M. (2019). Applying deep learning to right whale photo identification. *Conservation Biology*, 33(3), 676–684.
- Bolger, D. T., Morrison, T. A., Vance, B., Lee, D., & Farid, H. (2012). A computer-assisted system for photographic mark-recapture analysis. *Methods in Ecology and Evolution*, 3(5), 813–822.
- Burton, A. C., Neilson, E., Moreira, D., Ladle, A., Steenweg, R., Fisher, J. T., Bayne, E., & Boutin, S. (2015). Wildlife camera trapping: A review and recommendations for linking surveys to ecological processes. *Journal of Applied Ecology*, 52(3), 675–685.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., & Zagoruyko, S. (2020). End-to-end object detection with transformers. *Proceedings of the European Conference on Computer Vision*, 213–229.
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., & Joulin, A. (2021). Emerging properties in self-supervised vision transformers. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9650–9660.
- Cheeseman, T., Southerland, K., Park, J., et al. (2022). Advanced image recognition: A fully automated, high-accuracy photo-identification matching system for humpback whales. *Mammalian Biology*, 102, 915–929.
- Chen, G., Han, T. X., He, Z., Kays, R., & Forrester, T. (2014). Deep convolutional neural network based species recognition for wild animal monitoring. *Proceedings of the IEEE International Conference on Image Processing*, 858–862.
- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. *Proceedings of the International Conference on Machine Learning*, 1597–1607.
- Chopra, S., Hadsell, R., & LeCun, Y. (2005). Learning a similarity metric discriminatively, with application to face verification. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 539–546.
- Christin, S., Hervet, É., & Lecomte, N. (2019). Applications for deep learning in ecology. *Methods in Ecology and Evolution*, 10(10), 1632–1644.
- Clapham, M., Miller, E., Nguyen, M., & Darimont, C. T. (2020). Automated facial recognition for wild bears. *Ecology and Evolution*, 10(23), 12883–12892.
- Crall, J. P., Stewart, C. V., Berger-Wolf, T. Y., Rubenstein, D. I., & Sundaresan, S. R. (2013). HotSpotter: Patterned species instance recognition. *Proceedings of the IEEE Workshop on Applications of Computer Vision*, 230–237.
- Dagodzo, D. (2018b). A review of UAV applications in electrical transmission line inspection: Methods, technologies, and challenges. *Iconic Research and Engineering Journals*, 2(6), 234-254. <https://doi.org/10.64388/IREV2I6-1716083>
- Dagodzo, D., & Ahiaek Patrick, M. C. (2021b). An integrated framework for UAV, LiDAR, and GIS in infrastructure corridor management. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(5), 497-524. <https://doi.org/10.32628/CSEIT217566>
- Dagodzo, D., & Ahiaek Patrick, M. C. (2022). A review of right-of-way encroachment detection methods using geospatial technologies. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 8(1), 638-667. <https://doi.org/10.32628/CSEIT2281226>
- Dagodzo, D., Ahiaek Patrick, M. C., & Aliliele, C. (2022). AI and deep learning for vegetation classification in power corridor management: A review. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 8(1), 668-698. <https://doi.org/10.32628/CSEIT2281227>
- Dalal, N., & Triggs, B. (2005). Histograms of oriented gradients for human detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 886–893.
- Deb, D., Wiper, S., Gong, S., Shi, Y., Tymoszek, C., Fletcher, A., & Jain, A. K. (2018). Face recognition: Primates in the wild. *Proceedings of the IEEE International Conference on Biometrics Theory, Applications and Systems*, 1–10.

- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Deng, J., Guo, J., Xue, N., & Zafeiriou, S. (2019). ArcFace: Additive angular margin loss for deep face recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4690–4699.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale. *Proceedings of the International Conference on Learning Representations*.
- Ejofodomi, O. A., Gideon, E. N., Oladipo, G. O., & Oshomah, E. R. (2014). Automated detection of architectural detection in mammograms using template matching. *International Journal of Biomedical Science and Engineering*, 2(1), 1-6. <https://doi.org/10.11648/j.ijbse.20140201.11>
- Ekeocha, A. H., Aganga, A. A., Adejoro, F. A., Oyebojani, A., Oluwadele, J. F., & Tawose, O. M. (2021). Phenotypic characteristics of indigenous chickens in selected regions of Nigeria. *Journal of World's Poultry Research*, 11(3), 352-358. <https://doi.org/10.36380/jwpr.2021.42>
- Elebe, O., Hammed, N. I., Omoegun, G. O., Fadayomi, O., & Bello, A. D. (2023). An advanced machine learning model for detecting synthetic identity fraud in e-commerce platforms [Preprint]. ResearchGate.
- Everingham, M., Van Gool, L., Williams, C. K. I., Winn, J., & Zisserman, A. (2010). The PASCAL visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2), 303–338.
- Fadayomi, O., Bello, A. D., Elebe, O., Hammed, N. I., & Omoegun, G. O. (2024). An adaptive fraud risk scoring model for real-time transaction monitoring at scale. *IIARD International Journal of Banking and Finance Research*, 10(10), 212-230.
- Filani, O. M., Nnabueze, S. B., Ike, P. N., & Wedraogo, L. (2022). Real-time risk assessment dashboards using machine learning in hospital supply chain management systems. *International Journal of Multidisciplinary Evolutionary Research*, 3(1), 65-76. <https://doi.org/10.54660/IJMER.2022.3.1.65-76>
- Foster, R. J., & Harmsen, B. J. (2012). A critique of density estimation from camera-trap data. *Journal of Wildlife Management*, 76(2), 224–236.
- Freytag, A., Rodner, E., Simon, M., Loos, A., Köhl, H. S., & Denzler, J. (2016). Chimpanzee faces in the wild: Log-Euclidean CNNs for predicting identities and attributes of primates. *Proceedings of the German Conference on Pattern Recognition*, 51–63.
- Ge, Z., Liu, S., Wang, F., Li, Z., & Sun, J. (2021). YOLOX: Exceeding YOLO series in 2021. *arXiv preprint arXiv:2107.08430*.
- Girshick, R. (2015). Fast R-CNN. *Proceedings of the IEEE International Conference on Computer Vision*, 1440–1448.
- Girshick, R., Donahue, J., Darrell, T., & Malik, J. (2014). Rich feature hierarchies for accurate object detection and semantic segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 580–587.
- Gomez Villa, A., Salazar, A., & Vargas, F. (2017). Towards automatic wild animal monitoring: Identification of animal species in camera-trap images using very deep convolutional neural networks. *Ecological Informatics*, 41, 24–32.
- Grill, J.-B., Strub, F., Altché, F., et al. (2020). Bootstrap your own latent: A new approach to self-supervised learning. *Advances in Neural Information Processing Systems*, 33, 21271–21284.
- Hadsell, R., Chopra, S., & LeCun, Y. (2006). Dimensionality reduction by learning an invariant mapping. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1735–1742.
- Hansen, M. F., Smith, M. L., Smith, L. N., Salter, M. G., Baxter, E. M., Farish, M., & Grieve, B. (2018). Towards on-farm pig face recognition using convolutional neural networks. *Computers in Industry*, 98, 145–152.
- He, K., Fan, H., Wu, Y., Xie, S., & Girshick, R. (2020). Momentum contrast for unsupervised visual representation learning. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9729–9738.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hermans, A., Beyer, L., & Leibe, B. (2017). In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*.
- Hiby, L., Lovell, P., Patil, N., Kumar, N. S., Gopalaswamy, A. M., & Karanth, K. U. (2009). A tiger cannot change its stripes: Using a three-dimensional model to match images of living tigers and tiger skins. *Biology Letters*, 5(3), 383–386.
- Howard, A. G., Zhu, M., Chen, B., et al. (2017). MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.
- Idika, C. N., Salami, E. O., Ijiga, O. M., & Enyejo, L. A. (2021). Deep learning driven malware classification for cloud-native microservices in edge computing architectures. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(4).
- Ilodigwe, L., & Adesemoye, A. C. (2024a). Advances in real-world evidence, predictive analytics, and data-driven decision making for improving health system performance and patient outcomes. *International Journal of Health and Pharmaceutical Research*, 9(5), 174-197. <https://doi.org/10.56201/ijhpr.v9.no5.2024.pg174.197>
- Jackson, R. M., Roe, J. D., Wangchuk, R., & Hunter, D. O. (2006). Estimating snow leopard population abundance using photography and capture-recapture techniques. *Wildlife Society Bulletin*, 34(3), 772–781.
- Johnson, J., Douze, M., & Jégou, H. (2021). Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3), 535–547.
- Joseph, O. B., & Uzundu, N. C. (2024a). Integrating AI and machine learning in STEM education: Challenges and opportunities. *Computer Science & IT Research Journal*, 5(8), 1732-1750. <https://doi.org/10.51594/csitrj.v5i8.1379>

- Karanth, K. U. (1995). Estimating tiger *Panthera tigris* populations from camera-trap data using capture-recapture models. *Biological Conservation*, 71(3), 333–338.
- Karanth, K. U., & Nichols, J. D. (1998). Estimation of tiger densities in India using photographic captures and recaptures. *Ecology*, 79(8), 2852–2862.
- Kelly, M. J. (2001). Computer-aided photograph matching in studies using individual identification: An example from Serengeti cheetahs. *Journal of Mammalogy*, 82(2), 440–449.
- Komi, N. M. (2023). When accurate forecasts are not fair forecasts: Deep learning load prediction and equity across districts. *Shodhshauryam, International Scientific Refereed Research Journal*, 6(6), 476-546. <https://doi.org/10.32628/SHISRRJ247102>
- Komi, N. M., & Adamolekun, A. (2021). Interpretable machine learning for early failure prediction in distributed renewable energy assets. *International Journal of Multidisciplinary Research and Growth Evaluation*, 2(6), 1015-1038. <https://doi.org/10.54660/IJMRGE.2021.2.6.1015-1038>
- Komi, N. M., & Ganiu, O. S. (2023). Edge intelligence for resilient microgrid control: Advances, energy sovereignty, and open challenges. *Gyanshauryam, International Scientific Refereed Research Journal*, 6(3), 525-586. <https://doi.org/10.32628/GISRRJ236339>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097–1105.
- Kühl, H. S., & Burghardt, T. (2013). Animal biometrics: Quantifying and detecting phenotypic appearance. *Trends in Ecology & Evolution*, 28(7), 432–441.
- Kumar, S., Singh, S. K., Singh, R., & Singh, A. K. (2018). Deep learning framework for recognition of cattle using muzzle point image pattern. *Measurement*, 116, 1–17.
- Kumuyi, O., Uzoka, E., Akeju, B., & Ozowara, D. E. (2024a). Architecture for machine learning-enabled predictive energy management using IoT sensor networks. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 3138-3149. <https://doi.org/10.62225/2583049X.2024.4.6.6004>
- Ladapo, O. O., Dosunmu, A. A., Jooda, D., & Abolaji, T. O. (2022). Human-in-the-loop machine learning: A state of the art. *Journal of Frontiers in Multidisciplinary Research*, 3(1), 656-669. <https://doi.org/10.54660/JFMR.2022.3.1.656-669>
- Ladapo, O. O., Jooda, D., Dosunmu, A. A., & Abolaji, T. O. (2024). Keeping humans in the loop: Human-centered automated annotation with generative AI. *International Journal of Multidisciplinary Futuristic Development*, 5(1), 81-95. <https://doi.org/10.54660/IJMF.2024.5.1.81-95>
- Ladapo, O. O., Jooda, D., Dosunmu, A. A., & Abolaji, T. O. (2026). On the disagreement problem in human-in-the-loop federated machine learning. *International Journal of Multidisciplinary Research and Growth Evaluation*, 7(3), 178-192. <https://doi.org/10.54660/IJMRGE.2026.7.3.178-192>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Li, S., Li, J., Tang, H., Qian, R., & Lin, W. (2020). ATRW: A benchmark for Amur tiger re-identification in the wild. *Proceedings of the 28th ACM International Conference on Multimedia*, 2590–2598.
- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., & Belongie, S. (2017a). Feature pyramid networks for object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2117–2125.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017b). Focal loss for dense object detection. *Proceedings of the IEEE International Conference on Computer Vision*, 2980–2988.
- Lin, T.-Y., Maire, M., Belongie, S., et al. (2014). Microsoft COCO: Common objects in context. *Proceedings of the European Conference on Computer Vision*, 740–755.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., & Berg, A. C. (2016). SSD: Single shot multibox detector. *Proceedings of the European Conference on Computer Vision*, 21–37.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., & Guo, B. (2021). Swin transformer: Hierarchical vision transformer using shifted windows. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2), 91–110.
- Luiten, J., Osep, A., Dendorfer, P., Torr, P., Geiger, A., Leal-Taixé, L., & Leibe, B. (2021). HOTA: A higher order metric for evaluating multi-object tracking. *International Journal of Computer Vision*, 129(2), 548–578.
- Luo, H., Gu, Y., Liao, X., Lai, S., & Jiang, W. (2019). Bag of tricks and a strong baseline for deep person re-identification. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 1487–1495.
- Malkov, Y. A., & Yashunin, D. A. (2020). Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(4), 824–836.
- Mayo, W., Ogbale, J. I., Okoruwa, P. O., & Babatope, O. M. (2021). Designing an AI-predictive maintenance model for e-commerce systems using machine learning and cloud analytics. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(5), 416-440. <https://doi.org/10.32628/IJSCSEIT>
- Michael, O. N., & Ogunsola, O. E. (2023c). Development of predictive data models for enhancing crop resilience and climate adaptation in agrarian economies. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(1), 1510-1527. <https://doi.org/10.62225/2583049X.2023.3.1.15143>
- Miele, V., Dussert, G., Spataro, B., Chamaillé-Jammes, S., Allainé, D., & Bonenfant, C. (2021). Revisiting animal photo-identification using deep metric learning and network analysis. *Methods in Ecology and Evolution*, 12(4), 690–702.
- Musgrave, K., Belongie, S., & Lim, S.-N. (2020). A metric learning reality check. *Proceedings of the European Conference on Computer Vision*, 681–699.

- Neethirajan, S. (2020). The role of sensors, big data and machine learning in modern animal farming. *Sensing and Bio-Sensing Research*, 29, 100367.
- Nepovinnyykh, E., Eerola, T., & Kälviäinen, H. (2020). Siamese network based pelage pattern matching for ringed seal re-identification. *Proceedings of the IEEE Winter Conference on Applications of Computer Vision Workshops*, 25–34.
- Norouzzadeh, M. S., Morris, D., Beery, S., Joshi, N., Jojic, N., & Clune, J. (2021). A deep active learning system for species identification and counting in camera trap images. *Methods in Ecology and Evolution*, 12(1), 150–161.
- Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C., & Clune, J. (2018). Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences*, 115(25), E5716–E5725.
- Nwakamma, S., Ojukwu, J., & Oyesiji, S. O. (2024a). LLM-as-a-judge for automated model governance: A review of methods, biases, and release-gating practices. *World Journal of Innovation and Modern Technology*, 8(6), 185-216. <https://doi.org/10.56201/wjimt.v8.no6.2024.pg185.216>
- Nwakamma, S., Ojukwu, J., & Oyesiji, S. O. (2025). On-device multimodal small language models for privacy-preserving edge intelligence through quantization, pruning, distillation, and energy-efficient inference. *World Journal of Innovation and Modern Technology*, 9(12), 271-302. <https://doi.org/10.56201/wjimt.v9.no12.2025.pg271.302>
- O'Connell, A. F., Nichols, J. D., & Karanth, K. U. (Eds.). (2011). *Camera traps in animal ecology: Methods and analyses*. Springer.
- Obriki, O. D., Arumosoye, O. M., & Obogo, S. F. (2023). Advances in continuous hazard monitoring systems for workplace safety. *International Journal of Advanced Multidisciplinary Research and Studies*, 3(6), 2742-2759. <https://doi.org/10.62225/2583049X.2023.3.6.6075>
- Odejobi, O. D., Okonkwo, C. S., Ahiaekwe Patrick, M. C., Okeke, O. T., & Mayo, W. (2025). AI-augmented secure software engineering: Leveraging deep learning for autonomous threat detection and mitigation. *International Journal of Engineering and Modern Technology*, 11(12), 101-121. <https://doi.org/10.56201/ijemt.vol.11.no12.2025.pg101.121>
- Ogbole, J. I., Okoruwa, P. O., Babatope, O. M., & Oyewole, T. (2021). Developing an integrated data visualization model for continuous business performance monitoring and optimization. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(5), 441-467. <https://doi.org/10.32628/IJSRCSEIT>
- Ojukwu, J., Oyesiji, S. O., & Nwakamma, S. (2025). WebAssembly component model and WASI for portable cloud-to-edge software: Language interoperability, security, performance, and container comparisons. *International Journal of Computer Science and Mathematical Theory*, 11(12), 232-263. <https://doi.org/10.56201/ijcsmt.vol.11.no12.2025.pg232.263>
- Olorunkosebi, M. T., Ayinde, M. K., & Mensah, G. K. (2023). Integrating artificial intelligence and advanced biosensors for rapid and sensitive detection of antimicrobial resistant foodborne pathogens. *Journal of Advances in Microbiology Research*, 4(2), 175-182. <https://doi.org/10.22271/micro.2023.v4.i2b.293>
- Oluwadele, J. F., Tawose, O. M., & Adetumbi, T. (2024). Evaluation of feed intake, growth performance, and carcass characteristics of West African dwarf rams fed mango leaves, Napier grass, neem seed cake, and concentrate for fattening. *EUREKA: Life Sciences*, (4), 11-19. <https://doi.org/10.21303/2504-5695.2024.003603>
- Oluwo, K., Dada, T., & Isiekwu, C. P. (2024). Transforming insurance underwriting with machine learning: A review and application cases. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 3008-3019.
- Ominyi, M., Anichukwueze, C. C., & Uzougbo, N. S. (2024a). Developing a conceptual framework for AI-assisted compliance monitoring and anomaly detection in high-volume transaction environments. *World Journal of Innovation and Modern Technology*, 8(6), 204-228. <https://doi.org/10.56201/wjimt.v8.no6.2024.pg204.228>
- Oparah, O. S., Ezeh, F. E., Olatunji, G. I., & Ajayi, O. O. (2022). Big data-enabled predictive models for anticipating infectious disease outbreaks at population and regional levels. *International Journal of Advanced Multidisciplinary Research and Studies*.
- Oparah, O. S., Ezeh, F. E., Olatunji, G. I., & Ajayi, O. O. (2023). Framework for designing national real-time disease surveillance dashboards for public health stakeholders. *Shodhshauryam, International Scientific Refereed Research Journal*, 6(1), 208-227.
- Orise, C., & Niniola, S. (2025). Data governance in digital health learning systems: A systematic review of implementation frameworks and best practices. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 11(5), 520-538.
- Oyesiji, S. O., Nwakamma, S., & Ojukwu, J. (2023). Parameter-efficient multilingual adaptation of on-device language models: A survey. *International Journal of Engineering and Modern Technology*, 9(3), 287-320. <https://doi.org/10.56201/ijcsmt.v9.no3.2023.pg287.320>
- Oyesiji, S. O., Nwakamma, S., & Ojukwu, J. (2024). A conceptual framework for context-aware data curation in fine-tuning pipelines. *International Journal of Engineering and Modern Technology*, 10(11), 197-228. <https://doi.org/10.56201/ijemt.v10.no11.2024.pg197.228>
- Oyesiji, S. O., Nwakamma, S., & Ojukwu, J. (2025). Compressing recommender and ranking models for edge deployment: A survey. *International Journal of Engineering and Modern Technology*, 11(12), 205-237. <https://doi.org/10.56201/ijemt.vol.11.no12.2025.pg205.237>
- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359.
- Parham, J., Stewart, C., Crall, J., Rubenstein, D., Holmberg, J., & Berger-Wolf, T. (2018). An animal detection pipeline for identification. *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 1075–1083.
- Qiao, Y., Su, D., Kong, H., Sukkarieh, S., Lomax, S., & Clark, C. (2019). Individual cattle identification using a deep learning based framework. *IFAC-PapersOnLine*, 52(30), 318–323.
- Ravoor, P. C., & Sudarshan, T. S. B. (2020). Deep learning methods for multi-species animal re-identification and tracking: A survey. *Computer Science Review*, 38, 100289.

- Redmon, J., & Farhadi, A. (2017). YOLO9000: Better, faster, stronger. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7263–7271.
- Redmon, J., & Farhadi, A. (2018). YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.
- Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 779–788.
- Ren, S., He, K., Girshick, R., & Sun, J. (2015). Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 28, 91–99.
- Royle, J. A., Chandler, R. B., Sollmann, R., & Gardner, B. (2013). *Spatial capture-recapture*. Academic Press.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4510–4520.
- Sanni, J. O., Iwuanyanwu, U. A., Essien, M. A., & Attah, A. (2023). Lifecycle-aware marketing automation using federated learning for secure cross-organizational data management. *Gyanshauryam, International Scientific Refereed Research Journal*, 6(6), 337-364. <https://doi.org/10.32628/GISRRJ236642>
- Schneider, S., Taylor, G. W., & Kremer, S. C. (2020). Similarity learning networks for animal individual re-identification: Beyond the capabilities of a human observer. *Proceedings of the IEEE Winter Conference on Applications of Computer Vision Workshops*, 44–52.
- Schneider, S., Taylor, G. W., Linquist, S., & Kremer, S. C. (2019). Past, present and future approaches using computer vision for animal re-identification from camera trap data. *Methods in Ecology and Evolution*, 10(4), 461–470.
- Schofield, D., Nagrani, A., Zisserman, A., Hayashi, M., Matsuzawa, T., Biro, D., & Carvalho, S. (2019). Chimpanzee face recognition from videos in the wild using deep learning. *Science Advances*, 5(9), eaaw0736.
- Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 815–823.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. *Proceedings of the IEEE International Conference on Computer Vision*, 618–626.
- Shittu, M. A., Shittu, H. A., Adeleke, O. J., & Adedokun, O. J. (2023). Digital twin modeling for real time monitoring and fault detection in smart substations. *International Journal of Industrial Engineering Research and Development*, 14(2), 25-44. https://doi.org/10.34218/IJIERD_14_02_003
- Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, 6, 60.
- Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *Proceedings of the International Conference on Learning Representations*.
- Steenweg, R., Hebblewhite, M., Kays, R., et al. (2017). Scaling-up camera traps: Monitoring the planet's biodiversity with networks of remote sensors. *Frontiers in Ecology and the Environment*, 15(1), 26–34.
- Sun, Y., Zheng, L., Yang, Y., Tian, Q., & Wang, S. (2018). Beyond part models: Person retrieval with refined part pooling. *Proceedings of the European Conference on Computer Vision*, 480–496.
- Swanson, A., Kosmala, M., Lintott, C., Simpson, R., Smith, A., & Packer, C. (2015). Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian species in an African savanna. *Scientific Data*, 2, 150026.
- Szegedy, C., Liu, W., Jia, Y., et al. (2015). Going deeper with convolutions. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1–9.
- Tabak, M. A., Norouzzadeh, M. S., Wolfson, D. W., et al. (2019). Machine learning to classify animal species in camera trap images: Applications in ecology. *Methods in Ecology and Evolution*, 10(4), 585–590.
- Tan, M., & Le, Q. (2019). EfficientNet: Rethinking model scaling for convolutional neural networks. *Proceedings of the International Conference on Machine Learning*, 6105–6114.
- Tan, M., Pang, R., & Le, Q. V. (2020). EfficientDet: Scalable and efficient object detection. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10781–10790.
- Tawose, O. M., Ekeocha, A. H., & Oluwadele, J. F. (2024). Haematological and biochemical responses of West African dwarf goats fed water hyacinth-cassava peels silage. *Nigerian Journal of Animal Production*, 50(3). <https://doi.org/10.51791/njap.v50i3.4030>
- Town, C., Marshall, A., & Sethasathien, N. (2013). Manta Matcher: Automated photographic identification of manta rays using keypoint features. *Ecology and Evolution*, 3(7), 1902–1914.
- Tuia, D., Kellenberger, B., Beery, S., et al. (2022). Perspectives in machine learning for wildlife conservation. *Nature Communications*, 13, 792.
- Ussher-Eke, D., Okoh, O. F., & Ijiga, O. M. (2025). The role of biometric and IoT-based attendance systems in streamlining HR administrative functions, enhancing workforce accountability, and reducing labor inefficiencies. *International Journal for Multidisciplinary Research*, 7(4).
- van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9, 2579–2605.
- Vidal, M., Wolf, N., Rosenberg, B., Harris, B. P., & Mathis, A. (2021). Perspectives on individual animal identification from biology and computer vision. *Integrative and Comparative Biology*, 61(3), 900–916.
- Wang, C.-Y., Bochkovskiy, A., & Liao, H.-Y. M. (2023). YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7464–7475.

- Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., & Liu, W. (2018). CosFace: Large margin cosine loss for deep face recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5265–5274.
- Wearn, O. R., & Glover-Kapfer, P. (2019). Snap happy: Camera traps are an effective sampling tool when compared with alternative methods. *Royal Society Open Science*, 6(3), 181748.
- Wedraogo, L., & Sanni, J. O. (2024). Machine learning models addressing uncertainty in cross-channel campaign performance forecasting accuracy. *International Journal of Advanced Multidisciplinary Research and Studies*, 4(6), 2875-2890. <https://doi.org/10.62225/2583049X.2024.4.6.5648>
- Weideman, H. J., Jablons, Z. M., Holmberg, J., et al. (2017). Integral curvature representation and matching algorithms for identification of dolphins and whales. *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2831–2839.
- Weinstein, B. G. (2018). A computer vision for animal ecology. *Journal of Animal Ecology*, 87(3), 533–545.
- Wen, Y., Zhang, K., Li, Z., & Qiao, Y. (2016). A discriminative feature learning approach for deep face recognition. *Proceedings of the European Conference on Computer Vision*, 499–515.
- Willi, M., Pitman, R. T., Cardoso, A. W., et al. (2019). Identifying animal species in camera trap images using deep learning and citizen science. *Methods in Ecology and Evolution*, 10(1), 80–91.
- Wojke, N., Bewley, A., & Paulus, D. (2017). Simple online and realtime tracking with a deep association metric. *Proceedings of the IEEE International Conference on Image Processing*, 3645–3649.
- Wurtz, K., Camerlink, I., D'Eath, R. B., et al. (2019). Recording behaviour of indoor-housed farm animals automatically using machine vision technology: A systematic review. *PLoS ONE*, 14(12), e0226669.
- Ye, M., Shen, J., Lin, G., Xiang, T., Shao, L., & Hoi, S. C. H. (2022). Deep learning for person re-identification: A survey and outlook. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6), 2872–2893.
- Yu, X., Wang, J., Kays, R., Jansen, P. A., Wang, T., & Huang, T. (2013). Automated identification of animal species in camera trap images. *EURASIP Journal on Image and Video Processing*, 2013, 52.
- Zhang, Y., Sun, P., Jiang, Y., et al. (2022). ByteTrack: Multi-object tracking by associating every detection box. *Proceedings of the European Conference on Computer Vision*, 1–21.
- Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., & Tian, Q. (2015). Scalable person re-identification: A benchmark. *Proceedings of the IEEE International Conference on Computer Vision*, 1116–1124.
- Zheng, L., Yang, Y., & Hauptmann, A. G. (2016). Person re-identification: Past, present and future. *arXiv preprint arXiv:1610.02984*.
- Zhuang, F., Qi, Z., Duan, K., et al. (2021). A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1), 43–76.