

Lineage-Aware Markov Partition Refinement for Multi-Species Animal Re-Identification

LifeCLEF at CLEF 2026

Hanif Noer Rofiq^{1,*}

¹Monash University, Wellington Road, Clayton, VIC 3800, Australia

Abstract

Individual animal re-identification is important for wildlife monitoring, population estimation, and behavioural analysis. Unlike species classification, it must distinguish visually similar individuals within the same species. AnimalCLEF 2026 extends this problem to a lookup-and-discovery setting, where test images must either be matched to known individuals or grouped as new individuals. This setting is difficult because the species differ in visual markings, image quality, pose, and identity distribution. Global descriptors provide a useful starting point, but updating the whole partition at once can be unstable because one incorrect merge or split may affect a whole species partition. To address this, we propose Lineage-Aware Markov Partition Refinement (LAMPR), a framework for multi-species open-set animal re-identification. Rather than clustering all test images from scratch at each step, LAMPR treats the current prediction as a partition state and applies constrained, species-specific edits. The method adapts the refinement step to each species: for example, Lynx links use flank-aware geometric verification, Salamander edits use size priors and orientation-normalized matching, Sea Turtle partitions are mostly retained when the descriptor evidence is stable, and Texas Horned Lizard refinement focuses on ventral-pattern evidence. The selected LAMPR configuration achieved 0.34411 public ARI and 0.33830 private ARI, improving over the stronger host benchmark, WildFusion + HDBSCAN, which scored 0.20342 public ARI and 0.21221 private ARI. Component ablations indicate that conservative local-link refinement and species-wise partition replacement provide the clearest gains, while full re-clustering variants degraded performance in several cases despite producing visually plausible results. Code is available at <https://github.com/hanifnoer/animalclef2026-lampr>.

Keywords

animal re-identification, open-set clustering, lookup-and-discovery, partition refinement, AnimalCLEF 2026

1. Introduction

Individual animal re-identification is important for non-invasive wildlife monitoring, population estimation, movement studies, and behavioural analysis [1, 2]. Unlike species classification, individual re-identification must distinguish visually similar individuals within the same species [3, 4, 5]. AnimalCLEF 2026, organized as part of LifeCLEF 2026 [6, 7], extends this problem to a lookup-and-discovery setting. For some species, test images may match known individuals from the provided database, while others belong to unseen individuals that must be discovered and grouped. The task therefore requires identity grouping over the test set rather than prediction over a fixed class list.

This setting is difficult because the four species have different visual cues and image conditions. Lynx images are captured by camera traps and may include night images, back-facing animals, or unclear facial regions, while Salamander images may include human hands or inconsistent orientation [3, 8]. The Lynx subset is based on CzechLynx, a dataset for individual identification and pose estimation of Eurasian Lynx [9]. In our experiments, Sea Turtle images were more stable under descriptor-based matching, likely because head and scute patterns are often visible [10]. Texas Horned Lizard images rely on ventral belly-dot patterns [11], but no labelled training identities are provided for this species in AnimalCLEF 2026 [1].

CLEF 2026 Working Notes, 21 – 24 September 2026, Jena, Germany

*Corresponding author.

✉ hrof0001@student.monash.edu (H. N. Rofiq)

🌐 <https://hanifnoer.github.io/> (H. N. Rofiq)

🆔 0009-0004-0234-3055 (H. N. Rofiq)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Global descriptors such as MegaDescriptor provide a useful initial clustering by embedding whole images into a shared feature space [12]. However, global descriptors can miss fine-grained individual markings when animals are partially visible, strongly occluded, or captured under different viewpoints [13]. Local matchers can add more detailed evidence by checking whether distinctive markings align between image pairs [14, 13]. The main challenge is how to apply this pairwise evidence without changing too much of a partition that is already reliable. A link between two images can connect two larger components, while separating a weakly supported component can also move images that should remain together. This can lead to unstable partitions when one incorrect merge or split propagates through a species partition.

Our main observation is that complete re-clustering was often less stable than local refinement. Several submitted branches used visually plausible alternatives, including SAM3 foreground masking [15], YOLO-based views [16], SIFT graphs [17], DINO patch graphs [18, 19], and anatomy-template registration. However, many full partition rebuilds reduced leaderboard performance because they changed too much structure at once. More reliable improvements came from small, species-specific edits applied to a strong previous partition.

This paper proposes Lineage-Aware Markov Partition Refinement (LAMPR). LAMPR treats the current clustering as a partition state and refines it through constrained, species-specific transitions. Each transition modifies only one species partition, or one controlled part of it, while leaving the rest unchanged. A transition is accepted only when it is supported by evidence such as cluster-size plausibility, descriptor similarity, local geometric verification, orientation consistency, or segmentation-based diagnostics. The main contribution of this paper is a partition-refinement framework for composing species-specific corrections without repeatedly rebuilding the whole clustering. Rather than proposing a new backbone model, LAMPR focuses on how different sources of evidence can be used to make controlled edits to an existing partition.

The contributions of this paper are:

- We formalize Lineage-Aware Markov Partition Refinement (LAMPR) for species-wise open-set partition refinement.
- We show that species-local transitions are more reliable than full re-clustering in this heterogeneous setting.
- We introduce evidence constraints based on cluster size, orientation, descriptor similarity, local geometric verification, and segmentation diagnostics.
- We provide component-level ablations that report both useful and unsuccessful branches.

2. Related Work

2.1. Animal re-identification and discovery benchmarks

Recent animal re-identification research has been supported by datasets and toolkits that make individual-level recognition easier to evaluate across species. SeaTurtleID2022 provides a long-span benchmark for sea turtle re-identification [10], while WildlifeDatasets provides a common toolkit for loading and evaluating animal re-identification datasets [12]. WildlifeReID-10k further expands this setting by collecting a large multi-species dataset with many individual identities [20]. For Texas Horned Lizards, prior biological work has shown that ventral spot patterns can support individual identification [11].

AnimalCLEF 2026 is the benchmark used in this work [6, 7, 1]. It differs from earlier closed-set re-identification settings because some test images must be matched to known individuals, while others must be grouped as new individuals. Performance is measured using Adjusted Rand Index (ARI), which compares the predicted identity grouping with the ground truth.

2.2. Global descriptors for animal re-identification

Global descriptor models are widely used in animal re-identification because they represent each image as an embedding that can be compared by similarity search [2, 13]. MegaDescriptor is a common example and provides strong whole-image features for animal re-identification [12]. Such descriptors are useful for building an initial similarity graph or clustering, especially when individual markings are visible and consistent across images.

However, global descriptors can be sensitive to background, viewpoint, illumination, and partial visibility. Prior AnimalCLEF 2025 systems found that descriptor-only matching is often not enough for robust open-set identification, especially when species differ in pose, image quality, and available metadata [21, 3, 8]. This motivates combining global descriptors with more local forms of evidence.

2.3. Local matching and calibrated fusion

Local matching provides pairwise evidence from distinctive markings, which can complement global descriptors when whole-image similarity is ambiguous. This is useful when individual identity depends on fine-grained markings, such as flank spots, skin patterns, or small texture details. ALIKED is a lightweight keypoint and descriptor extraction method that can be used for local image matching [14]. Other local matching methods, such as SIFT, LoFTR, SuperPoint, and LightGlue, have also been used in visual matching pipelines [17, 22, 23, 24].

WildFusion combines global descriptor similarity with local matching evidence using score calibration [13]. AnimalCLEF 2025 solutions further showed that combining MegaDescriptor, ALIKED, EVA02, MiewID, and calibrated score fusion can improve performance in open-set animal re-identification [21, 3, 25, 8]. These works support the use of global descriptors for broad retrieval and local matchers for detailed verification.

2.4. Segmentation and species-specific preprocessing

Segmentation and preprocessing are often used to reduce background noise and make the animal region more consistent before feature extraction. Prior AnimalCLEF 2025 work used segmentation-aware filtering and showed that foreground extraction can help when background clutter or partial occlusion affects the image [21, 25]. However, segmentation must be used carefully because an inaccurate mask can remove useful markings or change the visual structure used for matching.

Several practical strategies have also been used to handle species-specific variation, including orientation normalization, test-time augmentation, and species-specific thresholds [3]. These strategies are relevant for multi-species re-identification because different species may require different evidence such as flank consistency for lynx, orientation-aware pattern matching for salamanders, head and scute patterns for sea turtles, and ventral spot patterns for Texas Horned Lizards.

3. Methods

3.1. Problem formulation

Let $\mathcal{X} = \{x_i\}_{i=1}^N$ be the set of test images, and let $\mathcal{S}$ be the set of species. Each image has a known species label $s_i \in \mathcal{S}$, but the individual identity labels are not fixed in advance. For each species s , let

$$\mathcal{X}_s = \{x_i \in \mathcal{X} : s_i = s\}$$

be the subset of test images belonging to that species. The objective is to assign each image to a predicted individual cluster:

$$y_i = f(x_i), \quad y_i \in \mathcal{Y}_{s_i},$$

where $\mathcal{Y}_{s_i}$ is the set of predicted cluster labels for species s_i . Since new individuals may appear in the test set, $\mathcal{Y}_{s_i}$ is not a fixed class list.

The prediction can therefore be written as a species-wise partition:

$$P = \{P_s\}_{s \in \mathcal{S}},$$

where each species partition is

$$P_s = \{C_{s,k}\}_{k=1}^{K_s}.$$

Each cluster $C_{s,k}$ is non-empty, clusters within the same species do not overlap, and together they cover all test images of that species:

$$C_{s,k} \neq \emptyset, \quad C_{s,k} \cap C_{s,l} = \emptyset \text{ for } k \neq l, \quad \bigcup_{k=1}^{K_s} C_{s,k} = \mathcal{X}_s.$$

Here, $C_{s,k}$ is a predicted identity cluster for species s , and K_s is the predicted number of individuals for that species. A good partition should group images of the same individual while separating images of different individuals.

3.2. Lineage-Aware Markov Partition Refinement

LAMPR treats the current prediction as a partition state rather than as an output to be discarded after each run. Starting from an initial partition P_0 , the method applies a sequence of refinement actions:

$$P_{t+1} = T(P_t, a_t, E_t),$$

where P_t is the current partition, T is the global transition operator, a_t is the selected refinement action, and E_t is the evidence available at step t . The evidence may include descriptor similarity, local matching scores, orientation metadata, cluster-size statistics, and segmentation diagnostics.

Most transitions are applied to one species at a time:

$$P_{t+1}^{(s)} = T_s(P_t^{(s)}, a_t^{(s)}, E_t^{(s)}), \quad P_{t+1}^{(r)} = P_t^{(r)} \quad \forall r \neq s.$$

Here, T_s denotes the species-local transition operator for species s . This update keeps all other species partitions unchanged while one candidate edit is tested, making each transition a controlled refinement of the previous partition rather than a complete reconstruction of the full clustering.

3.3. Transition operators

LAMPR uses a small set of transition operators. Each operator changes the current partition in a limited way.

A transition is accepted only when it satisfies a set of evidence constraints. These constraints are used to reduce unnecessary partition changes:

- **Format validity:** the submission must contain exactly one row for each official test image identifier, preserve the required `image_id` and `cluster` columns, contain no missing cluster labels, and use cluster labels accepted by the upload system as arbitrary non-empty string identifiers.
- **Species locality:** a transition must only modify the intended species partition.
- **Cluster-size plausibility:** merges or splits should remain plausible under training identity-size statistics when such statistics are available.
- **Similarity evidence:** a proposed link should be supported by descriptor similarity or calibrated matching scores.

Table 1

LAMPR transition operators. Each operator is applied to one species partition or one controlled part of the partition.

Operator	Main target	Purpose
Species partition replacement	All species	Replace one species partition using a stronger candidate state
Component separation	Salamander, Texas Horned Lizard	Separate implausibly large or weakly supported components
Local-link insertion	Lynx, Salamander	Join high-confidence cross-component pairs
Boundary-suppressed refinement	Texas Horned Lizard	Reduce silhouette-driven links and focus on ventral patterns
Retention	Sea Turtle and stable species	Preserve a partition when new edits do not improve evidence

- **Local geometric support:** local links should be supported by ALIKED match quality and geometric consistency [14].
- **Orientation consistency:** side or orientation metadata should be respected when it is relevant.
- **Stability evidence:** a stable partition is retained when candidate edits have weaker descriptor/local support or introduce implausible shape changes, such as excessive singleton creation, a large cluster-count jump, or a maximum component size outside the observed species range. For example, Sea Turtle was retained after the local-link candidate changed the partition without improving the available evidence.

The procedure is summarized in Algorithm 1.

Algorithm 1 Lineage-Aware Markov Partition Refinement (LAMPR)

Require: Initial species-wise partition P_0 , transition set $\mathcal{A}$, evidence bundles $\{E_t\}$, evidence constraints $\mathcal{G}$, number of refinement steps H

- 1: **for** $t = 0, \dots, H - 1$ **do**
 - 2: Select a species s_t and a transition action $a_t \in \mathcal{A}$
 - 3: Construct a candidate partition $\tilde{P}_{t+1}$
 $\tilde{P}_{t+1}^{(s_t)} = T_{s_t}(P_t^{(s_t)}, a_t, E_t^{(s_t)})$
 $\tilde{P}_{t+1}^{(r)} = P_t^{(r)} \quad \forall r \neq s_t$
 - 4: **if** $\tilde{P}_{t+1}$ satisfies all constraints in $\mathcal{G}$ **then**
 - 5: $P_{t+1} \leftarrow \tilde{P}_{t+1}$
 - 6: **else**
 - 7: $P_{t+1} \leftarrow P_t$
 - 8: **end if**
 - 9: **end for**
 - 10: **return** Final partition P_H
-

3.4. Species-specific refinement

LAMPR applies species-specific refinement rules because the identity evidence differs across species. Figures 1 and 2 show representative evidence used for species-specific refinement.

3.4.1. Lynx

Lynx images contain strong illumination changes and left/right flank asymmetry. The Lynx refinement keeps the original camera-trap images and uses segmentation output only as a foreground guide for

subject-focused enhancement. Candidate links are first selected using EVA02 descriptor similarity [26], then checked with ALIKED [14] geometric verification. Cross-component links are restricted by flank orientation and local geometric support, so visually similar but opposite-side views are not joined too easily.




Figure 1: Lynx same-flank local-link examples. The examples show candidate pairs with consistent flank orientation, which is used as a constraint before local-link refinement.

3.4.2. Salamander

Salamander images contain large orientation variation and many partial-body views. The Salamander refinement uses neutral grey padding and evaluates both original-orientation and orientation-normalized views. It combines descriptor similarity with ALIKED [14] local verification. It also uses training identity-size priors to avoid creating implausibly large components.

3.4.3. Sea Turtle

The Sea Turtle refinement is conservative. Since the descriptor-based partition was already stable in our experiments, LAMPR retains this species partition unless there is clear evidence that an additional edit improves it.

3.4.4. Texas Horned Lizard

Texas Horned Lizard images rely on ventral spot patterns for individual identification [11]. Since no labelled training identities were provided for this species [1], the refinement does not use training identity-size priors. The selected rule preserves the original orientation and avoids vertical flipping. The refinement uses SAM3 foreground views [15] together with boundary-suppressed ventral-pattern evidence, so the partition is guided more by internal spot structure than by silhouette or crop-boundary effects.

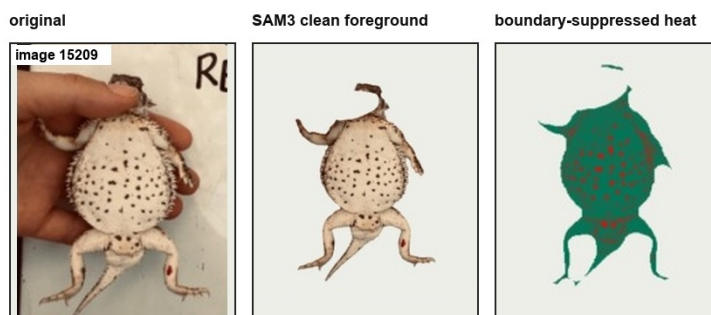


Figure 2: Texas Horned Lizard ventral-pattern evidence. The figure shows the original image, the SAM3 foreground view, and the boundary-suppressed heatmap used to emphasize internal ventral spot patterns while reducing silhouette and crop-boundary cues.

Table 2
Pipeline overview for the selected LAMPR submission.

Pipeline stage	Updated species	Role in the selected system
Initial candidate partitions	All species	Generate descriptor-, segmentation-, and shape-aware partitions that serve as candidate states for later refinement.
Species-wise replacement	Salamander, Sea Turtle	Construct a species-wise hybrid partition by retaining the selected Lynx and Texas partitions while replacing Salamander and Sea Turtle with stronger descriptor-hybrid candidates.
Texas refinement	Texas Horned Lizard	Apply orientation-preserving ventral-pattern refinement followed by boundary-suppressed split-only refinement. Other species remain unchanged.
Local-link refinement	Lynx, Salamander	Add conservative cross-cluster links only when supported by EVA02 similarity, ALIKED geometry, orientation constraints, and component-size guards. This defines the final partition.

3.5. Selected pipeline overview

Table 2 gives a compact overview of the selected LAMPR pipeline. The system first constructs candidate partitions, then applies species-local refinements to the incumbent partition state.

3.6. Implementation

The descriptor stages used MegaDescriptor-L-384 [12] and MiewID embeddings from the conservationlabs/miewid-msv3 checkpoint [27]. Descriptor similarities were computed with cosine similarity after row normalization. For Lynx and Salamander, local-link refinement used EVA02-B-16 descriptors initialized with the merged2b_s8b_b131k pretrained weights [26], together with ALIKED-N16 local features loaded as aliked-n16 [14]. ALIKED extraction used at most 512 keypoints per image. Images were resized to 512×512 after neutral-grey square padding. EVA02 scores were computed by averaging embeddings from five 86% crops. ALIKED matches were checked by estimating a homography and a partial affine transform with RANSAC, then using the larger inlier count. The projection threshold was 14 pixels for Lynx and 12 pixels for Salamander.

The final transition guards were deliberately conservative. Lynx links required a calibrated probability of at least 0.98, local score at least 0.96, at least 50 inliers, no opposite-flank relation, and merged component size at most 84. Salamander links required calibrated probability at least 0.90, local score at least 0.40, at least 4 inliers, EVA score at least 0.86, merged component size at most 5, and at most 10 accepted links. Texas refinement used a split-only boundary-suppressed rule. An intra-cluster edge was retained only when the ventral-pattern score was at least 0.310 and either the point score was at least 0.24 or the stack-gain score was at least 0.42. Sea Turtle labels were retained from the species-hybrid state.

4. Results

4.1. Evaluation setting and benchmark comparison

Submissions were evaluated using Adjusted Rand Index (ARI) [28, 29], the official AnimalCLEF 2026 leaderboard metric. Let n_{ij} be the contingency-table count between predicted cluster i and ground-truth

identity j , and let $a_i = \sum_j n_{ij}$, $b_j = \sum_i n_{ij}$, and N be the number of evaluated images. ARI is

$$\text{ARI} = \frac{\sum_{ij} \binom{n_{ij}}{2} - \frac{\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}}{\binom{N}{2}}}{\frac{1}{2} \left[\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2} \right] - \frac{\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}}{\binom{N}{2}}}.$$

The metric measures agreement between the predicted partition and the ground-truth identity grouping while adjusting for chance. It penalizes both over-clustering, where images of the same individual are split into different clusters, and under-clustering, where different individuals are merged into the same cluster.

The official leaderboard team name was **Hanif Noer Rofiq**. On the final leaderboard, the selected LAMPR configuration achieved rank 43 with 0.34411 public ARI and 0.33830 private ARI. The public and private ARI values in this section are taken from the official Kaggle AnimalCLEF 2026 leaderboard and submission records [1].

Table 3 compares the selected LAMPR configuration with the two host benchmark submissions. Compared with the higher-scoring host benchmark, WildFusion + HDBSCAN, LAMPR increased public ARI by 0.14069 and private ARI by 0.12609. Another evaluated configuration achieved a slightly higher public ARI of 0.34459, but its private ARI decreased to 0.33543. Therefore, the 0.34411/0.33830 configuration is used as the selected result.

Table 3

Overall comparison with the host benchmark submissions.

Configuration	Public ARI	Private ARI
MegaDescriptor + MiewID + DBSCAN benchmark	0.19401	0.20646
WildFusion + HDBSCAN benchmark	0.20342	0.21221
Selected LAMPR configuration	0.34411	0.33830
Highest-public LAMPR configuration	0.34459	0.33543

During development, the public leaderboard score was used to compare submitted candidate states such as those listed in Tables 4 and 5. It was not used to train the pairwise calibrators, which used only labelled training identities, and it was not used as an input to any image-pair transition rule. Private leaderboard scores were unavailable during development and are reported only as final evaluation outcomes. This distinction is important because the highest-public LAMPR state was not the best private-scoring state.

4.2. Ablation and diagnostic results

Table 4 summarizes the main submitted states that led to the selected LAMPR configuration. These states show how the partition improved through foreground descriptor evidence, species-wise replacement, and species-local refinement.

Table 5 reports diagnostic submissions used to evaluate alternative refinement strategies that were not retained in the selected configuration. These diagnostic results provide negative evidence for the final design. Several alternative branches reduced ARI because they changed too much of the partition at once. The selected LAMPR configuration avoided this by keeping the stronger existing partition and applying only constrained species-specific edits.

The results in Table 4 show that the main gains came from preserving useful partition structure and applying species-specific edits. The species-hybrid state improved private ARI from 0.29066 to 0.32483, while the Lynx and Salamander local-link refinements further improved the private score. The selected configuration used the Salamander EVA02/ALIKED refinement because it achieved the best private ARI.

The diagnostic branches in Table 5 support the observation that full re-clustering was often less stable than local refinement. SAM3+SIFT, SAM3+DINO, SAM+YOLO, and anatomy-template branches

Table 4

Main submitted states leading to the selected LAMPR configuration.

State or transition	Public ARI	Private ARI	Finding
Initial graph baseline	0.14150	0.18585	Descriptor graph clustering alone gave a weak starting partition.
SAM3 masked ReID fusion	0.20535	0.25546	SAM3 foreground evidence improved descriptor-based clustering over the initial graph baseline.
Shape-constrained fusion partition	0.27632	0.29066	Shape-based constraints produced a stronger private-score partition.
Species hybrid	0.29758	0.32483	Replacing selected species partitions gave a clear private-score gain.
Texas ventral-pattern refinement	0.31811	0.32206	Orientation-preserving ventral-pattern refinement improved the Texas partition.
Texas boundary-suppressed refinement	0.32684	0.32140	Boundary-suppressed ventral-pattern evidence improved public ARI, but did not improve private ARI over the previous Texas branch.
Lynx EVA02/ALIKED local-link refinement	0.32903	0.33630	Flank-aware local verification improved private ARI.
Salamander EVA02/ALIKED local-link refinement	0.34411	0.33830	Conservative Salamander local links produced the selected result.
Extended Lynx-Salamander refinement	0.34459	0.33543	Relaxed link constraints improved public ARI but reduced private ARI.

Table 5

Diagnostic branches that were not retained in the selected configuration.

Branch	Public ARI	Private ARI	Finding
SAM3 + SIFT graph	0.08609	0.13020	Building a full local graph from SAM3 foreground views and SIFT matches reduced performance.
SAM3 + DINO patch graph	0.06852	0.11812	Patch-level graph reconstruction did not produce stable species partitions.
SAM+YOLO species-view fusion	0.06303	0.04912	Detector-based view fusion degraded the partition when used as an independent rebuild.
Anatomy-template registration	0.01136	0.00219	Template alignment did not handle partial visibility and pose variation well.
Sea Turtle EVA02/ALIKED local-link refinement	0.33609	0.33518	Additional local links reduced the Sea Turtle partition quality compared with the selected configuration.

all reduced both public and private ARI. These results indicate that visually plausible preprocessing or alignment steps did not necessarily improve the partition when they were used to rebuild the clustering too broadly. The Sea Turtle local-link result also supports retaining the descriptor-based Sea Turtle partition rather than adding extra local links.

5. Discussion

LAMPR improved over both host benchmarks, including WildFusion + HDBSCAN. Since this benchmark already combines similarity fusion with density-based clustering, the improvement suggests that the main gain did not come only from using stronger visual features. The results instead show that how partition edits are composed matters. LAMPR keeps a useful previous partition and changes only the

species or component with enough supporting evidence.

This is also reflected in the species-wise refinements. The four species required different strategies: Lynx benefited from flank-aware local verification, Salamander from orientation-aware matching and size constraints, Sea Turtle from retaining a stable descriptor-based partition, and Texas Horned Lizard from ventral-pattern refinement. This supports the use of species-specific evidence rather than one uniform clustering rule for all species.

The diagnostic branches show that more processing did not always improve the partition. SAM3+SIFT, SAM3+DINO, SAM+YOLO, and anatomy-template registration produced visually plausible alternatives, but reduced ARI when they were used to rebuild the clustering too broadly. This suggests that full partition reconstruction can introduce new merge or split errors even when some local visual evidence is reasonable. In contrast, constrained species-local refinement limited the size of each change and was more reliable.

The contrast between the selected configuration and the highest-public configuration further supports this conservative design. Relaxed local-link rules improved the visible public split but reduced robustness on the hidden split, suggesting that conservative link acceptance is preferable to maximizing the public score alone.

LAMPR is also designed to be more scalable than repeatedly rebuilding the full partition. Each transition is local to one species or component, so new evidence can be added without recomputing all clusters from the beginning. This is useful in a lookup-and-discovery setting where identity databases may be updated over time. However, the current implementation still depends on manually designed species-specific guards and empirically selected transition settings. Scaling the framework to more species would require more automatic guard calibration, candidate-link selection, and transition acceptance.

The main limitation is that LAMPR depends on the quality of the initial partition. If the starting clustering is weak, small species-local edits may not be enough to recover the correct identity structure. Some transition rules were also selected through empirical testing rather than a fully automatic policy. Future work could focus on learning when to accept or reject partition edits, especially for species with limited or no labelled training identities.

6. Conclusion

This paper presented Lineage-Aware Markov Partition Refinement (LAMPR) for multi-species open-set animal re-identification in AnimalCLEF 2026. The task combines lookup and discovery, requiring images to be grouped into identity partitions rather than assigned only to a fixed class list. LAMPR addresses this setting by treating the current clustering as a partition state and applying constrained, species-specific refinements instead of repeatedly rebuilding the full clustering.

The selected LAMPR configuration achieved 0.34411 public ARI and 0.33830 private ARI, outperforming both host benchmarks, including WildFusion + HDBSCAN. The ablation and diagnostic results show that the main improvements came from species-wise partition replacement and conservative local-link refinement. They also show that visually plausible full reconstruction branches, such as segmentation-based views, patch graphs, and template alignment, did not necessarily improve the final partition.

These results suggest that preserving reliable partition structure is important in heterogeneous multi-species discovery settings. Because LAMPR applies local transitions to selected species or components, it provides a scalable direction for updating identity partitions without recomputing the whole clustering from the beginning. Future work could make this process more automatic by learning when to accept or reject partition edits, especially when scaling to more species or to species with limited labelled training data.

Acknowledgments

The author thanks the AnimalCLEF 2026 organizers for providing the challenge, datasets and benchmark submissions. The author also acknowledges the dataset contributors and open-source tool maintainers whose resources supported this work.

Declaration on Generative AI

During the preparation of this work, the author used ChatGPT (OpenAI, GPT-5.5 Thinking) in order to: grammar and spelling check, paraphrase, reword and formatting assistance. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

- [1] L. Adam, L. Pícek, K. Papafitsoros, D. Williams, D. Biffi, Animalclef26 @ cvpr & clef, <https://kaggle.com/competitions/animal-clef-2026>, 2026. Kaggle.
- [2] P. C. Ravoora, S. T.S.B., Deep learning methods for multi-species animal re-identification and tracking - a survey, *Computer Science Review* 38 (2020) 100289. URL: <https://www.sciencedirect.com/science/article/pii/S1574013720303890>. doi:10.1016/j.cosrev.2020.100289.
- [3] D. Kim, B. Park, H. Bae, S. Lee, C. Lee, Fusion of global and local descriptors with feature calibration for multi-species animal re-identification: Animalclef 2025, Working Notes of CLEF - Conference and Labs of the Evaluation Forum (2025).
- [4] S. Schneider, G. W. Taylor, S. C. Kremer, Similarity learning networks for animal individual re-identification-beyond the capabilities of a human observer, in: *Proceedings of the IEEE/CVF winter conference on applications of computer vision workshops*, 2020, pp. 44–52.
- [5] C. Beyan, A. Osman Tur, E. Karimi, From species-specific models to universal re-id: a survey of animal re-identification, *Information Fusion* 133 (2026) 104323. URL: <https://www.sciencedirect.com/science/article/pii/S1566253526002022>. doi:10.1016/j.inffus.2026.104323.
- [6] L. Pícek, L. Adam, S. Kahl, R. Bossy, L. Chrobak, H. Goëau, K. Papafitsoros, H. Klinck, W.-P. Vellinga, R. Planqué, T. Denton, K. Barnard, C. Nédellec, L. Deléger, M. Courtin, G. Martellucci, I. Moummad, F. Vinatier, P. Bonnet, A. Joly, Overview of LifeCLEF 2026: Ai challenges for biodiversity understanding and ecosystem management, in: *International Conference of the Cross-Language Evaluation Forum for European Languages (CLEF)*, Springer, 2026.
- [7] L. Adam, K. Papafitsoros, D. A. Williams, D. Biffi, L. Pícek, Overview of AnimalCLEF 2026: Discovery and re-identification of individual animals, in: *Working Notes of CLEF 2026 - Conference and Labs of the Evaluation Forum*, 2026.
- [8] N. Semenova, Meta-algorithm for open-set animal re-id: Wildfusion, xgboost, and dual-backbone arcface, in: *CEUR Workshop Proceedings*, volume 4038, 2025, pp. 3154–3165.
- [9] L. Pícek, J. Straka, M. Jirik, E. Belotti, M. Dul'a, J. Krausová, M. Bojda, V. Cermak, L. Bufka, R. Dvořák, et al., Czechlynx: A dataset for individual identification and pose estimation of the eurasian lynx, *Scientific Data* 13 (2026) 511.
- [10] L. Adam, V. Čermák, K. Papafitsoros, L. Pícek, Seaturtleid2022: A long-span dataset for reliable sea turtle re-identification, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024, pp. 7146–7156.
- [11] D. Biffi, M. R. Tucker, A. Ackel, D. A. Williams, Identification of individual texas horned lizards (*phrynosoma cornutum*) using genotypes and ventral spot patterns, *Ecology and Evolution* 15 (2025) e71167.
- [12] V. Čermák, L. Pícek, L. Adam, K. Papafitsoros, Wildlifedatasets: An open-source toolkit for animal re-identification, in: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2024, pp. 5953–5963.

- [13] V. Cermak, L. Pícek, L. Adam, L. Neumann, J. Matas, Wildfusion: Individual animal identification with calibrated similarity fusion, in: A. Del Bue, C. Canton, J. Pont-Tuset, T. Tommasi (Eds.), *Computer Vision – ECCV 2024 Workshops*, Springer Nature Switzerland, Cham, 2025, pp. 18–36.
- [14] X. Zhao, X. Wu, W. Chen, P. C. Y. Chen, Q. Xu, Z. Li, Aliked: A lighter keypoint and descriptor extraction network via deformable transformation, *IEEE Transactions on Instrumentation and Measurement* 72 (2023) 1–16. doi:10.1109/TIM.2023.3271000.
- [15] N. Carion, L. Gustafson, Y.-T. Hu, S. Debnath, R. Hu, D. Suris, C. Ryali, K. V. Alwala, H. Khedr, A. Huang, J. Lei, T. Ma, B. Guo, A. Kalla, M. Marks, J. Greer, M. Wang, P. Sun, R. Rädle, T. Afouras, E. Mavroudi, K. Xu, T.-H. Wu, Y. Zhou, L. Momeni, R. Hazra, S. Ding, S. Vaze, F. Porcher, F. Li, S. Li, A. Kamath, H. K. Cheng, P. Dollár, N. Ravi, K. Saenko, P. Zhang, C. Feichtenhofer, Sam 3: Segment anything with concepts, 2026. URL: <https://arxiv.org/abs/2511.16719>. arXiv:2511.16719.
- [16] R. Khanam, M. Hussain, Yolov11: An overview of the key architectural enhancements, 2024. URL: <https://arxiv.org/abs/2410.17725>. arXiv:2410.17725.
- [17] D. G. Lowe, Distinctive image features from scale-invariant keypoints, *International Journal of Computer Vision* 60 (2004) 91–110.
- [18] M. Caron, H. Touvron, I. Misra, H. Jegou, J. Mairal, P. Bojanowski, A. Joulin, Emerging properties in self-supervised vision transformers, in: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 9630–9640. doi:10.1109/ICCV48922.2021.00951.
- [19] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. HAZIZA, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, P. Bojanowski, DINOv2: Learning robust visual features without supervision, *Transactions on Machine Learning Research* (2024). URL: <https://openreview.net/forum?id=a68SUt6zFt>, featured Certification.
- [20] L. Adam, V. Cermák, K. Papafitsoros, L. Pícek, Wildlifereid-10k: Wildlife re-identification dataset with 10k individual animals, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025, pp. 2124–2134.
- [21] L. Adam, K. Papafitsoros, R. Kovář, V. Čermák, L. Pícek, Overview of animalclef 2025: recognizing individual animals in images, *Working Notes of CLEF - Conference and Labs of the Evaluation Forum* (2025).
- [22] J. Sun, Z. Shen, Y. Wang, H. Bao, X. Zhou, Loftr: Detector-free local feature matching with transformers, in: *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 8918–8927. doi:10.1109/CVPR46437.2021.00881.
- [23] D. DeTone, T. Malisiewicz, A. Rabinovich, Superpoint: Self-supervised interest point detection and description, in: *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2018, pp. 224–236. doi:10.1109/CVPRW.2018.00060.
- [24] P. Lindenberger, P.-E. Sarlin, M. Pollefeys, Lightglue: Local feature matching at light speed, in: *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023, pp. 17581–17592. doi:10.1109/ICCV51070.2023.01616.
- [25] R. Pakhomov, G. Demidov, K. Bogdan, S. Lanskikh, D. Dinmuhametov, A. Khlopotnykh, Individual wildlife recognition via hybrid global-local matching and segmentation-aware filtering, *Working Notes of CLEF - Conference and Labs of the Evaluation Forum* (2025).
- [26] Y. Fang, Q. Sun, X. Wang, T. Huang, X. Wang, Y. Cao, Eva-02: A visual representation for neon genesis, *Image and Vision Computing* 149 (2024) 105171. URL: <https://www.sciencedirect.com/science/article/pii/S0262885624002762>. doi:10.1016/j.imavis.2024.105171.
- [27] L. Otarashvili, T. Subramanian, J. Holmberg, J. J. Levenson, C. V. Stewart, Multispecies animal re-id using a large community-curated dataset, 2024. URL: <https://arxiv.org/abs/2412.05602>. arXiv:2412.05602.
- [28] L. Hubert, P. Arabie, Comparing partitions, *Journal of classification* 2 (1985) 193–218.
- [29] W. M. Rand, Objective criteria for the evaluation of clustering methods, *Journal of the American Statistical association* 66 (1971) 846–850.