DOI: https://doi.org/10.15276/aait.09.2026.41
UDC 004.89:004.932

Abductive identity of small aerial targets for refining neural network classifier labels

Volodymyr Nareznoy[1)]
ORCID: https://orcid.org/0009-0009-1397-5687; shsa1res@gmail.com
Oleg Galchonkov[1)]
ORCID: https://orcid.org/0000-0001-5468-7299; o.n.galchenkov@gmail.com. Scopus Author ID: 56081377900
[1)] Odesa Polytechnic National University, 1, Shevchenko Ave. Odesa, 65044, Ukraine

ABSTRACT

Relevance. Optical channels of counter-drone systems distinguish small unmanned aerial vehicles from birds with neural network classifiers. At long range a target of a few pixels appears in isolated frames, and in a group of objects labels belong to different targets. **Purpose of the article.** To refine classifier labels on a small aerial target by abductive identity, without retraining. **Objectives.** To formalise target identity as an abductive choice and determine its effect on track maintenance, labels and class conclusion. **Methods.** In each frame either the silhouette near the prediction is the target or the target is not visible. An integrity constraint refutes a silhouette deviating from the prediction by over two body lengths, the measured margin of the true target; the minimal explanation is chosen. Labels of one target and the flapping fact feed an abductive logic program. **Scientific novelty.** Identity rests on a physical motion constraint scaled by the measured body length, not on appearance, which tiny silhouettes lack. **Practical significance.** The same classifier held on one target gives more correct labels and explained conclusions. **Results.** For vehicles of nine to twelve pixels frames off the target fell by one quarter to one half without harming bird tracks. The attention window raised correct vehicle labels from one fifth – one third to one third – three fifths. The abductive program removed wrong vehicle conclusions at the cost of fewer decisions; six-pixel vehicles are beyond the method. **Conclusions.** Maintaining the track of one target is a precondition for refining classifier labels; the main limitation is frequent "unknown" answers at low confidence.

Keywords: Abductive reasoning; target identity; integrity constraint; small target; label refinement; flapping flight

For citation: **Nareznoy V. Yu., Galchonkov O. M. "Abductive identity of small aerial targets for refining neural network classifier labels". *Applied Aspects of Information Technology*. 2026; Vol.9 No.4: 632–647. DOI: https://doi.org/10.15276/aait.09.2026.41**

INTRODUCTION

Optical channels of systems for protection against small unmanned aerial vehicles (UAVs) use neural network classifiers to distinguish vehicles from birds [1], [2], [3]. At long range the target occupies a few pixels. The classifier sees it only in isolated frames, and its confidence fluctuates from frame to frame [4], [5]. Systems with several cameras narrow the field of view to enlarge the target [6], and the choice of the classifier architecture improves the quality of object recognition in visible-light camera images [7]. At the limit of the range, however, the target remains a silhouette of a few pixels.

A decision based on one frame of such a target is unreliable, so labels are accumulated along the track. Accumulation, however, makes sense only when all labels belong to one target. In a flock of birds or near other objects the track passes to a neighbouring silhouette, and the labels of different targets are mixed along "one" track. The class conclusion then rests on evidence that does not describe the designated target.

© Nareznoy V., Galchonkov O., 2026

Hence, to refine the conclusion of the classifier, it is first necessary to establish which observations belong to one target. This makes relevant the study of methods for establishing the identity of a small target that raise the share of correct classifier labels without retraining the classifier.

LITERATURE REVIEW AND PROBLEM STATEMENT

In classical target tracking an observation is associated with a track by the nearest neighbour inside a validation gate, a missed detection is allowed for, and the motion is predicted by α-β or Kalman filters [8]. The gate is set either by a fixed radius or statistically, by the covariance of the residual. A fixed radius in pixels has to be tuned to the scale of the target, and a statistical gate narrows on smooth tracks; neither uses the size of the target itself.

In tracking by detection [9], target identity is established by the overlap of bounding boxes in adjacent frames. In [10] an appearance feature obtained by a re-identification network is added. For a silhouette of 5-12 pixels such a feature is not formed, because the object has no texture, and similar birds in a flock look alike.

This is an open access article under the CC BY license (https://creativecommons.org/licenses/by/4.0/deed.uk)

In [11], [12] identity is maintained by associating confident and weak detector boxes with motion taken into account, and in [13] the motion model is corrected by observations after an occlusion. However, identity exists only in frames where the detector returns a box. For a distant target such frames are few, so the track breaks into fragments, and the fragments may belong to different objects.

Classical multiple hypothesis tracking [14], [15] keeps alternative assignments of observations to tracks and selects the most probable one, and track-before-detect methods [16] follow a target without a detector threshold by accumulating likelihood along hypothesised trajectories. These methods do not return the class of the target and do not explain why a particular observation was assigned to it.

Reviews of drone detection [17] compare radar, radio-frequency, acoustic and vision-based sensing, and a review of ground-to-air vision-based methods [18] notes that they mostly detect and track a single UAV, while tracking of several small targets is little explored. Spatio-temporal detection of distant drones in aerial video [19] uses the motion of the target over several frames to separate it from the background. In the challenge on tracking small birds in video [20] the targets occupy a few dozen pixels, so detection and appearance-based association become unreliable, and a metric based on point distance is used instead of box overlap.

Benchmarks for UAV tracking, mostly in the thermal infrared band [1], [21], show that small targets remain the main difficulty: on the benchmark of tiny UAVs in complex scenes [22] the best evaluated tracker reaches a state accuracy of only 35.9 %.

Abductive logic programming [23], [24] defines an explanation of observations as a hypothesis that, together with a background theory, entails the observations and does not violate integrity constraints. Abduction as the choice of the best explanation [25], [26] and neuro-symbolic integration of learned and logical components [27] make it possible to check the conclusion of a neural network against explicit knowledge. In [28] this formalism is applied to refining the class of an aerial object from classifier labels and physical facts. However, it is assumed there that a track of one target already exists.

A generalisation of these works shows that the identity of a small target is established either by appearance, which such a target does not have [10],]20], or only in frames with a detection [11], [12], [13], or by a gate whose size does not follow the target [8] . The unresolved question is how to establish target identity from physical constraints on its motion so that the classifier labels accumulated along the track belong to one target and refine the class conclusion.

RESEARCH AIM AND OBJECTIVES

The aim of the research is to refine the labels of a neural network classifier on a small aerial target by establishing its identity abductively. This makes it possible to take checked and explained decisions on the target class in optical channels of UAV protection systems without retraining the classifier.

To achieve the aim, the following objectives were set:

– to formalise the establishment of target identity as an abductive choice between the hypotheses “the silhouette is the target” and “the target is not visible” with an integrity constraint on the body length;

– to determine the effect of abductive identity on the share of frames off the target (held frames in which the track is not on the designated target) and on the continuity of the observation series;

– to determine how maintaining the track of one target changes the share of correct and wrong classifier labels;

– to determine how the abductive logic program refines the classifier conclusion and how this conclusion depends on the label trust threshold.

MATERIALS AND METHODS

The object of the research is the process of classifying a small aerial target by a neural network classifier in a video stream. The main hypothesis is that classifier labels refine the class conclusion only when all of them belong to one target, and that target identity can be established from physical constraints on motion without appearance features.

Data. Birds are represented by 50 albatross tracks from a video of the colony on Midway Atoll [29] (2641 frames of 1920×1080 pixels, 29.97 Hz, static camera) with a median silhouette size of 4-12 pixels and a length of at least 96 frames. The bird size is the larger side of the dark mask of the silhouette; together with the light parts of the plumage the bird has a median size of 27 pixels. The reference tracks were built by multiple hypothesis tracking [28]: every silhouette starts a track hypothesis that takes the nearest silhouette within a 30-pixel gate around prediction (1), and the hypotheses are ranked by the product of length and straightness. The tracks were selected by a rule fixed before the experiments (length, size and straightness

of at least 0.8, duplicates merged) and checked visually on contact sheets with 8 fragments per track: every fragment shows a bird, but in some frames a neighbouring bird crosses the fragment. The reference uses the same association as basic track maintenance, so it's possible bias favours basic track maintenance rather than the proposed method. UAVs are represented by 24 segments of 12 video sequences of the DUT Anti-UAV dataset [30] against the sky. Small targets of 6, 9 and 12 pixels (the larger side of the annotation box) were obtained by downscaling the frame; the frame rate was taken as 25 Hz.

Classifier. A two-class YOLOv8s detector [31] trained on the birdrone dataset [32] was used. The detector was not retrained: all changes concern only which of its labels belong to the target. Classifier labels in frame t are boxes with a class and a confidence q_t whose centre lies within 25 pixels of the current target position.

Track maintenance. The target is held by an attention window of 120×120 pixels around the prediction. The threshold in the window equals the 99th percentile of the negative contrast of the silhouette relative to the smoothed background, but not less than 8 brightness levels.

The predicted position $\hat{x}_t$ and the velocity v_t are updated as follows:

$$\hat{x}_t = x_{t-1} + v_{t-1},$$
$$v_t = (1-\beta)v_{t-1} + \beta(x_t - x_{t-1}), \quad (1)$$

where x_t is the target position in frame t; $\beta = 0.4$ is the velocity smoothing coefficient. In a frame without a silhouette the velocity is not updated, and the position x_t is taken equal to the prediction $\hat{x}_t$; the target moves along the prediction for up to $K = 5$ frames, after which it is re-acquired using a hint. Three hint strategies were studied: no hint (S0), a hint from classifier detections (SC) and an external target designation at 1 Hz (SE-1), modelled by the reference position with Gaussian noise of 15 pixels.

Abductive identity. For each frame t with candidate silhouettes $b_1, \ldots, b_n$ near the prediction, identity hypotheses are formed. The hypothesis same(b) means that silhouette b is the held target. The hypothesis unseen means that the target is not visible in the frame. An integrity constraint forbids the hypothesis same(b) if the deviation of the silhouette centre c_b from the prediction exceeds two body lengths of the target:

$$same(b) \Rightarrow \|c_b - \hat{x}_t\| \le 2L_t, \quad (2)$$

where L_t is the body length of the target, equal to the median size of the last 15 accepted silhouette masks but not less than 3 pixels; for a bird the mask covers only the dark parts, so L_t is smaller than the visible size of the bird. The size is measured in the same way for both classes, from the accepted silhouette masks. The constraint is applied after five accepted silhouettes so that L_t rests on measured sizes; before that, the 30-pixel gate of basic track maintenance is used. Among the hypotheses *same*(*b*) that are not refuted, the one requiring the smallest deviation from the prediction is chosen:

$$b^* = arg\ \min_b \|c_b - \hat{x}_t\|, \quad (3)$$

where the minimum is taken over the silhouettes b for which *same*(*b*) is not refuted; the deviations in (2) and (3) are Euclidean distances in pixels. If all hypotheses *same*(*b*) are refuted, invisibility of the target is abduced: the track moves along the prediction, and the frame adds neither a silhouette nor a classifier label to the track. Therefore, a neighbouring silhouette lying farther than two body lengths does not enter the track, and its label is not mixed with the labels of the target. The sequence of operations for one frame is shown in Fig. 1 (m – number of consecutive frames without the target; K – maximum number of such frames; wt, ht – width and height of the silhouette).

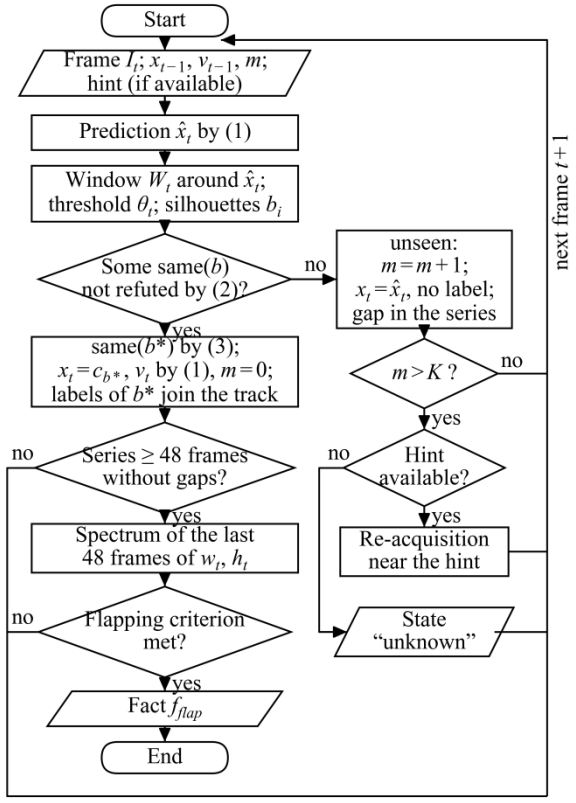


Fig. 1. **Algorithm of track maintenance with abductive identity for one frame**
Source: **compiled by the authors**

The margin in (2) was chosen from the measured deviation of the true target from prediction (1). This deviation is caused not by the dynamics of flight but by the error of the velocity estimate, by the noise of the centre position and, above all, by the break-up of the silhouette into fragments after thresholding. Its distribution in body lengths is shown in Fig. 2.

Fig. 2 shows that the typical deviation is small: 0.07 of the body length for UAVs and 0.13-0.18 for birds (a – birds (reference tracks, 5305 frames, and silhouettes accepted on the target without constraint (2), 17 987 frames); b – UAVs of 6, 9 and 12 pixels (accepted silhouettes, 4351, 10 964 and 21 023 frames); vertical lines – limits Lt and 2Lt; logarithmic ordinate, deviations above 4Lt are counted in the last bin). For birds, however, the distribution has a heavy tail: when the silhouette breaks into body and wings, its centre jumps. One body length is exceeded in 7.6 % of the frames of the reference tracks (4.4 % of the accepted silhouettes), and each such frame breaks a gap-free series of the true target. Two body lengths are exceeded in 2.1 % (0.9 %) of bird frames and in 0.2–0.5 % of UAV frames; on two independent halves of the bird set the share is 2.0 and 2.3 %. The limit $2L_t$ thus covers 98-99 % of the deviations of the true target for both classes and, like the body length itself, scales with the target. The first version of the method used one body length; the coefficient 2 was chosen from this distribution, and the results for the coefficients 1 and 1.5 are given in Table 2. Since the search neighbourhood is 30 pixels, for silhouettes larger than 15 pixels the limit coincides with the gate of basic track maintenance; the constraint acts on smaller silhouettes, in 54 % of bird frames and 94 % of UAV frames.

In form, rule (2)-(3) is nearest-neighbour association in a gate with a missed-detection hypothesis [8]. It differs in three respects: the gate is set in measured body lengths and therefore needs no tuning to range and camera; the hypothesis that the target is not visible is explicit and justified; both hypotheses are written in the same formalism as the class refinement, so the decisions from the assignment of a silhouette to the class conclusion form one explained chain.

The identity theory is written in the logic programming language Prolog in the form of hypotheses, integrity constraints and a minimal explanation. For speed it is replicated in Python; the decisions of the two implementations coincided on 1657 checked frames. Without constraint (2) the method coincides with basic track maintenance, which makes it possible to evaluate the contribution of abductive identity separately. The core of the theory is shown in Fig. 3: the facts f_pred, f_cand and f_body describe the prediction, the candidate silhouettes and the body length; the rule refutes implements constraint (2) with the coefficient body_mult; the rules explain return the minimal explanation (3) or the hypothesis unseen together with its justification.

An example of the abductive choice on a real frame is shown in Fig. 4 (video frame [29], CC BY 2.0).

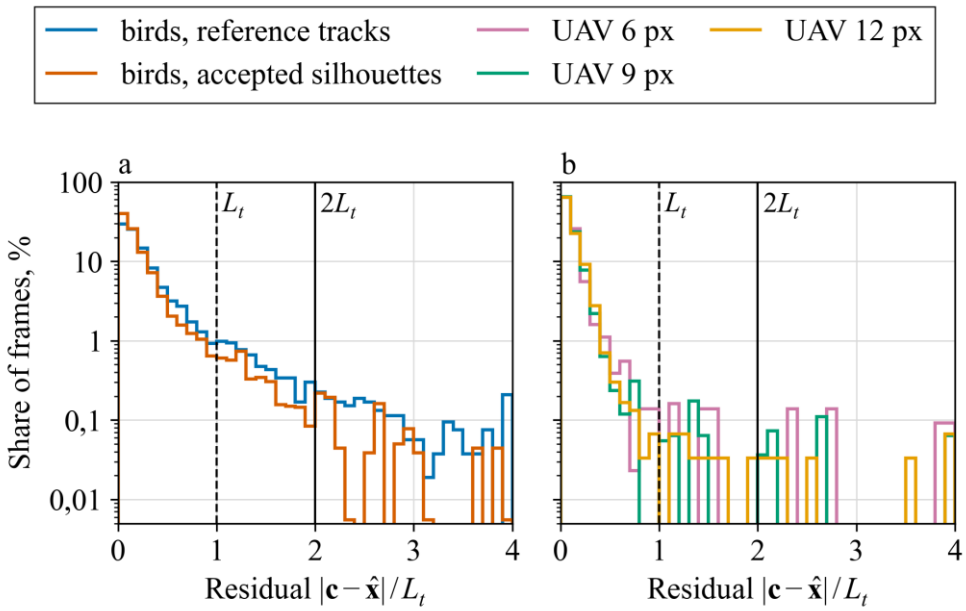


Fig. 2. **Deviation of the silhouette centre from prediction (1) in body lengths of the target: a – birds; b – UAVs of 6, 9 and 12 pixels**

Source: **compiled by the authors**

```
:- dynamic f_pred/2.  % f_pred(X, Y)  predicted position
:- dynamic f_cand/4.  % f_cand(Id, X, Y, Size)  candidate silhouette in the window gate
:- dynamic f_body/1.  % f_body(L)  body length in px (asserted only when grounded)

thr(body_floor, 3.0).
thr(body_mult, 2.0).
thr(window_gate, 30.0).

residual(Id, R) :-
    f_cand(Id, X, Y, _), f_pred(PX, PY),
    R is sqrt((X - PX) ** 2 + (Y - PY) ** 2).

limit(L) :- f_body(B), !, thr(body_floor, F), thr(body_mult, M), L is M * max(B, F).
limit(G) :- thr(window_gate, G).  % constraint not groundable yet

refutes(same(Id), jump(R, L)) :-
    residual(Id, R), limit(L), R > L.

admissible(same(Id), R) :-
    f_cand(Id, _, _, _), \+ refutes(same(Id), _), residual(Id, R).

% explain(-Hypothesis, -Justification)
explain(same(Id), nearest_unrefuted(R)) :-
    findall(R0-I0, admissible(same(I0), R0), L), L \== [], !,
    keysort(L, [R-Id | _]).
explain(unseen, all_refuted(Why)) :-
    findall(I-W, refutes(same(I), W), Why), Why \== [], !.
explain(unseen, no_candidate).
```

Fig. 3. **Core of the identity theory in Prolog (file attention/a4_track_alp.pl of the repository; the header comment is omitted)**

Source: **compiled by the authors**

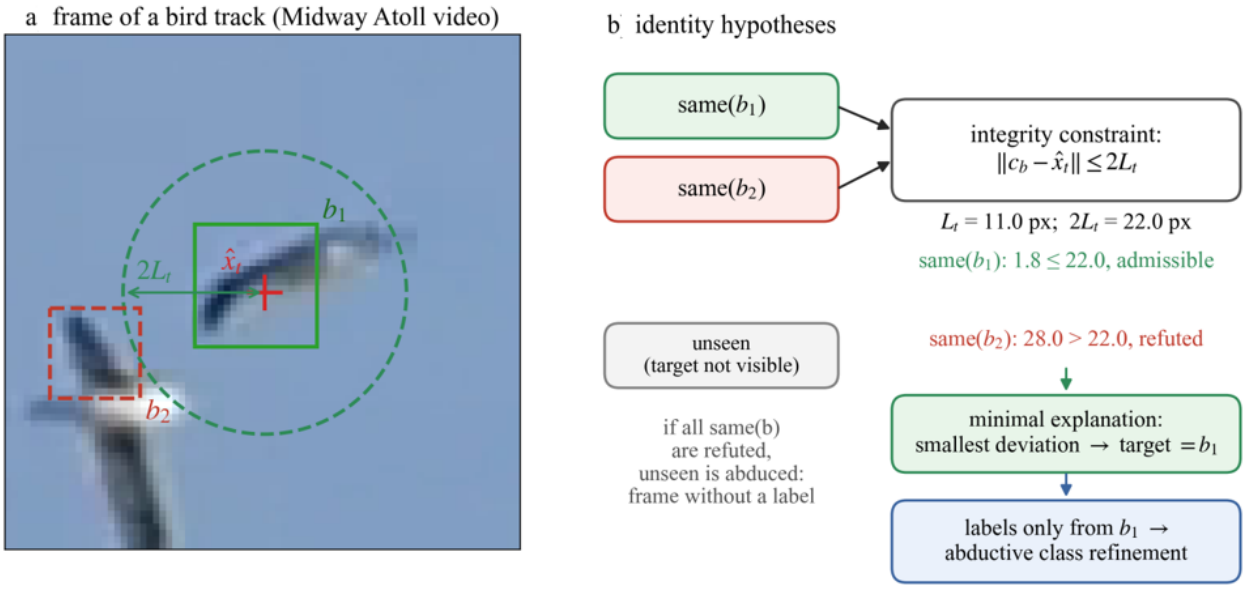


Fig. 4. **Abductive identity of the target on a real frame: a – neighbourhood of the target in a frame of a bird track (cross – predicted position $\hat{x}_t$; dashed circle – admissible deviation of two body lengths, $2L_t$ = 22 pixels; b – identity hypotheses, integrity constraint and minimal *explanation* (b_1 – silhouette accepted as the target; b_2 – refuted silhouette of a neighbouring bird)**

Source: **compiled by the authors**

In the frame of Fig. 4 the target silhouette b_1 lies 1.8 pixels from the prediction, and the wing of a neighbouring bird b_2 lies inside the search neighbourhood 28.0 pixels from it, farther than $2L_t$ = 22 pixels; the hypothesis same(b_2) is refuted, and the minimal explanation chooses b_1. In this frame b_1 would also be chosen without the constraint, as the nearest silhouette. The constraint becomes decisive when the target itself is not visible near the prediction: without it the nearest silhouette, here the wing of the neighbouring bird, would be accepted as the target and its classifier labels mixed with the labels of the target, whereas with it unseen is abduced. On the 50 bird tracks of strategy SE-1 a neighbouring silhouette was refuted while the target was accepted in 305 frames.

The spectral criterion of flapping flight [28] is checked once on the longest continuous series of 48 frames. Its band of 1.5-6 Hz covers the wingbeat frequencies predicted for large birds [33].

Abductive class refinement. If the classifier has a box with a confidence of at least 0.9 on the held target, the class is determined by the classifier. Otherwise, the abductive logic program [28] is used without changes. Only facts measurable at range are passed to it: f_{det} – the class of the majority of classifier labels along the track together with their median confidence, and f_{flap} – fulfilment of the flapping flight criterion. The label f_{det} is taken into account if its confidence is not lower than the trust threshold θ_{tr} = 0.5. Facts about speed, hovering and turning are not asserted without range measurement.

The program returns one of five conclusions, each with its justification: confirmation of the classifier label, revision of the label with the violated constraint indicated, acceptance of the label without objections, a class supported by a physical fact alone when the label is not trusted, or the state “unknown”. A label refuted by a fact is revised: wingbeats are incompatible with the UAV class. The sensitivity of the conclusion was studied by changing the trust threshold θ_{tr} from 0.5 to 0.25.

Metrics. The share of frames off the target is the share of held frames in which the track position is farther from the reference target position than the tolerance r = 2s, where s is the median target size (s is used only in this tolerance and differs from the body length L_t) (the silhouette mask of a bird or the annotation box of a UAV). Frames of extrapolation, in which unseen is abduced, are not held frames and do not enter this share; therefore the share of frames on the target over the whole track is also given, where extrapolated positions are compared with the reference and a lost track counts as off the target. A series is considered continuous on the target if the longest gap-free segment contains at least 48 frames and at least half of its frames lie on the target. The classifier label along the track is determined by a majority of at least 70 % among at least 5 boxes with a confidence of at least 0.25. Tracking by detection, used for comparison, associates the nearest classifier box with a confidence of at least 0.25 within a 45-pixel gate around prediction (1) and extrapolates for up to 5 frames. The limit attainable for the classifier was estimated with ideal association: only boxes near the reference position of the target are assigned to it. To evaluate the start from detection, the track starts in the first frame where the classifier has a box with a confidence of at least 0.5 within 25 pixels of the target. The 95 % confidence intervals were obtained by cluster bootstrap (2000 resamples; UAVs resampled by video sequence, birds by track), and paired conclusions were compared by the exact McNemar test.

Parameters. The coefficients β = 0.4 and K = 5, the 30-pixel gate, the series length of 48 frames and the contrast threshold were inherited without changes from the silhouette channel of [28], where they had been chosen on one key track of video [29]. The 120×120 window and the 99th percentile were taken from the previous stage of the research; 15 masks and 5 silhouettes before the constraint is applied, the label voting parameters and the trust thresholds of 0.5 and 0.9 [28] were fixed before the runs. The DUT Anti-UAV dataset was not used to choose any of these parameters. The coefficient 2 in (2) is the only parameter chosen from data of this study (Fig. 2).

RESEARCH RESULTS

Effect of abductive identity on track maintenance

The share of frames off the target and the share of series on the target without and with abductive identity are given in Table 1.

Table 1 shows that abductive identity reduces the share of frames off the target for both classes (50 bird tracks, 24 UAV segments; for SE-1 – 250 and 120 runs; bold – the better value of a pair; interval – 95 % cluster bootstrap of the difference “with − without” for frames off target). For UAVs of 9 and 12 pixels it decreases from 23-60 % to 11-46 %, and the 95 % intervals of the difference lie below zero for every strategy; for 9 pixels the share of series on the target increases from 33 % to 46-47 %. For birds the share of frames off the target decreases from 48-56 % to 39-53 %, and the share of series on the target does not decrease: 34, 34 and 44 % versus 32, 34 and 40 % for strategies S0, SC and SE-1.

Table 1. **Effect of abductive identity (tolerance r = 2s), %**

Target, strategy	Frames off target, without / with identity	95 % interval of the difference, p.p.	Series ≥ 48 frames on target, without / with	Frames on target over the whole track, without / with
Birds, S0	48 / **39**	−17…−2	32 / **34**	**39** / 37
Birds, SC	56 / **53**	−11…+5	34 / 34	**40** / 37
Birds, SE-1	49 / **41**	−13…−3	40 / **44**	46 / **48**
UAV 6 px, S0	80 / **75**	−14…−1	13 / 13	15 / **17**
UAV 6 px, SC	81 / **73**	−20…−1	13 / 13	18 / **23**
UAV 6 px, SE-1	83 / **75**	−22…−1	13 / **15**	18 / **24**
UAV 9 px, S0	54 / **33**	−39…−6	33 / **46**	38 / **43**
UAV 9 px, SC	59 / **43**	−30…−4	33 / **46**	40 / **45**
UAV 9 px, SE-1	60 / **46**	−25…−3	33 / **47**	40 / **47**
UAV 12 px, S0	23 / **11**	−24…−2	83 / 83	72 / **77**
UAV 12 px, SC	27 / **15**	−25…−2	83 / **88**	74 / **83**
UAV 12 px, SE-1	28 / **17**	−23…−2	83 / 83	73 / **80**

Source: **compiled by the authors**

The share of frames on the target over the whole track, which counts extrapolated frames, also does not decrease: for UAVs of 9 and 12 pixels it grows from 38-40 % to 43-47 % and from 72-74 % to 77-83 %, and for birds it stays at 37-48 % (the differences are not significant).The gain is therefore not caused by more frequent abduction of unseen. For UAVs of 6 pixels the indicators hardly change: the silhouette mask of such a vehicle has only 2×2 pixels, and the attention window mostly holds something other than the target itself (73-83 % of frames off the target), so this size lies outside the working range of the method. No wrong decisions "bird" for UAVs was obtained with abductive identity. For birds, decisions about another object became fewer (3, 3 and 21 instead of 4, 6 and 23), while decisions about the target itself remained the same (2, 2 and 19).

Comparison with other gates and sensitivity to the parameters

To compare abductive identity with classical association, constraint (2) was replaced by a fixed gate of 10 and 20 pixels and by a statistical gate of 3.03 $\hat{\sigma}$ (χ^2 with two degrees of freedom, 0.99), where $\hat{\sigma}$ is estimated from the deviations of the track; the coefficient of the limit was also varied (Table 2).

According to Table 2 (strategy S0; bold – the best value in a column), only the limit $2L_t$ is not worse than basic track maintenance in the share of series for any class and size. The fixed gate of 10 pixels gives the fewest frames off the target for UAVs but breaks bird tracks (2 % of series). The gate of 20 pixels is close to $2L_t$ only because the targets here are 9-12 pixels ($2L_t$ = 14-30 pixels): for a target twice as far away the same 20 pixels would be four body lengths and would admit neighbours, and for a target twice as close one body length, which rejects the target itself. Even at this favourable scale it gives more UAV frames off the target than $2L_t$ for every strategy (40-51 % versus 33-46 % for 9 pixels and 15-24 % versus 11-17 % for 12 pixels). The statistical gate narrows on smooth UAV tracks and loses the target (29 % of series for 12 pixels instead of 83 %). The limit of one body length gives the fewest bird frames off the target (28 %) but halves the bird series (16 %), because it rejects the true silhouette in 7.6 % of

Table 2. **Association rules without a hint: frames off target / series on target, %**

Association rule	Birds: frames off target / series	UAV 9 px: frames off target / series	UAV 12 px: frames off target / series
Basic track maintenance, gate 30 px	48 / 32	54 / 33	23 / **83**
Fixed gate 10 px	**25** / 2	**5** / **50**	**2** / 79
Fixed gate 20 px	39 / 32	40 / 46	15 / 79
Statistical gate χ^2 (0.99)	27 / 6	49 / 25	17 / 29
Constraint (2), limit L_t	28 / 16	32 / 46	11 / **83**
Constraint (2), limit 1.5 L_t	35 / 30	32 / 46	11 / **83**
Constraint (2), limit 2 L_t (proposed)	39 / **34**	33 / 46	11 / **83**

Source: **compiled by the authors**

frames; a lower bound of L_t of 4 or 6 pixels does not change this (16-18 % of series). The results are stable with respect to β and K: with β = 0.2-0.6 and K = 3-8 the share of frames off the target is 33-38 % for 9-pixel and 10-13 % for 12-pixel UAVs, and the number of series changes by at most one of 24; for birds only β = 0.6 reduces the series to 26 %.

The operation of abductive identity on a bird track is shown in Fig. 5a – frame fragments (dashed line – attention window; cross – predicted position; green box – silhouette accepted as the target; yellow boxes – classifier detections with qt ≥ 0.25); Fig. 5b – classifier confidence qt, track maintenance state and silhouette height ht; Fig.5c – spectrum of the series ht (coloured background – band 1.5-6 Hz)

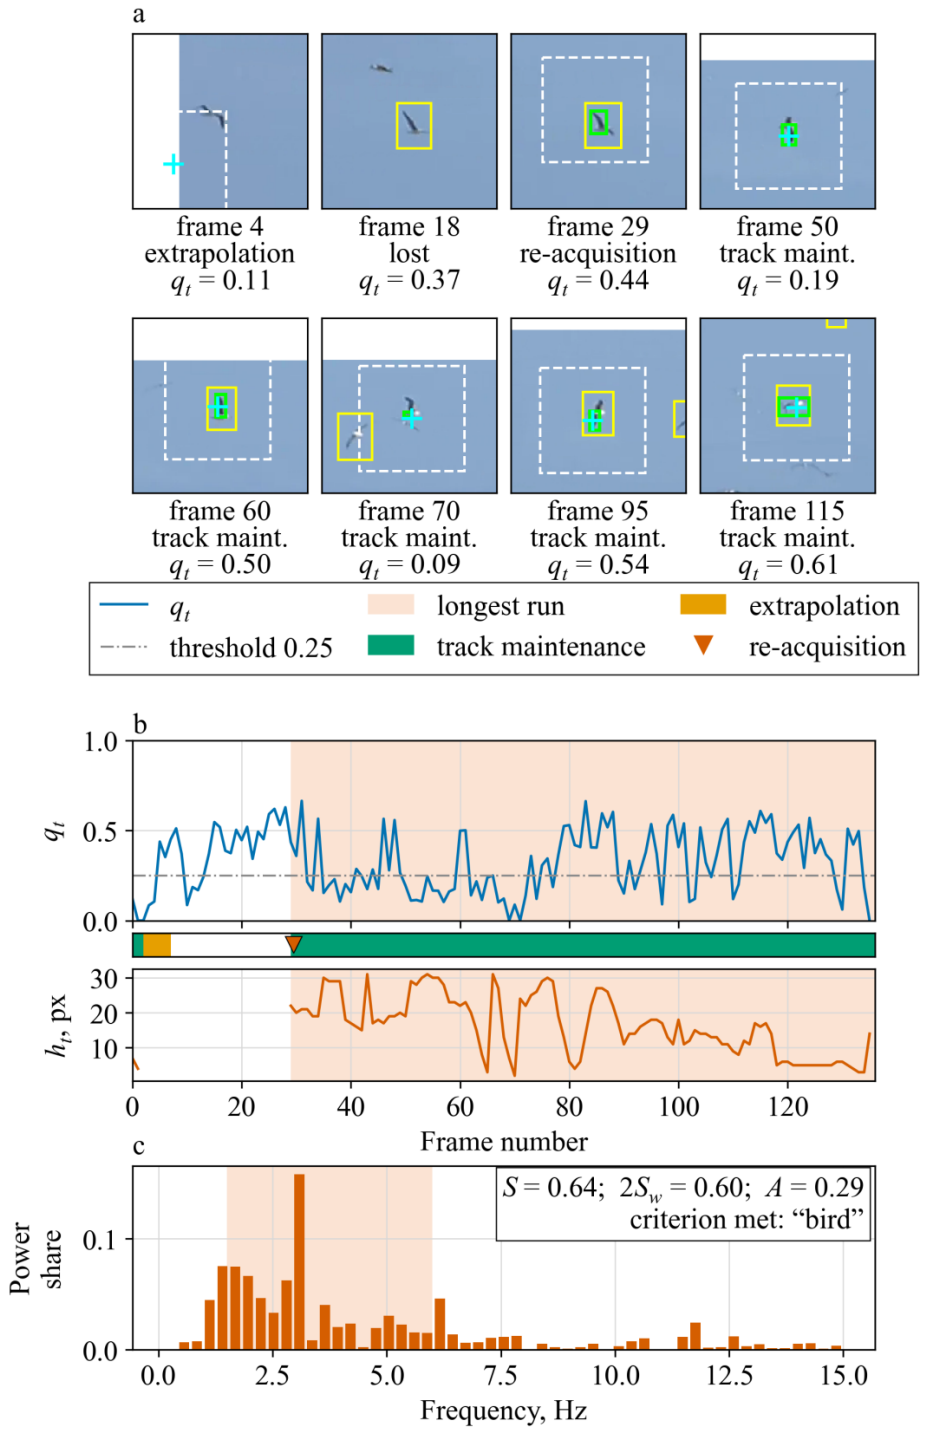


Fig. 5. **Track maintenance on a bird track with abductive identity (video [29] , strategy SE-1): a – frame fragments; b – classifier confidence q_t ; c – spectrum of the series ht**

Source: **compiled by the authors**

In frame 70 of Fig. 5 the classifier returns a box on a neighbouring bird, but the track remains on the designated target, and the label of the neighbouring bird does not enter the track. The continuous series is 107 frames long, and the power share in the wingbeat band is 0.64 at a threshold of 0.60, so the target is classified as a bird.

A failure of abductive identity is shown in Fig. 6.

In Fig. 6 a large neighbouring bird crosses the target (video [29], strategy S0; a – cross – reference position of the target; x – track position; dashed circle – limit 2Lt = 18 pixels around the prediction; b – dashed line – tolerance 2s). In frame 58 its silhouette lies inside the limit $2L_t$ and nearer to the prediction than the target, so the minimal explanation accepts it, and the track passes to the neighbour; without a hint it does not return, and the distance to the target grows to 14-20 target sizes. Constraint (2) does not protect against a neighbour that is nearer to the prediction than the target itself; such crossings require a hint or features other than the position. A persistent transition of this kind after a long stretch on the target occurred in 5 of the 50 bird tracks without a hint; in one of them it is the reference that passes to a crossing bird.

Classifier labels on one target

The labels accumulated by the classifier along the track with different track maintenance methods are given in Table 3 (classifier [32] , strategy SC; a wrong UAV label is “bird”, a wrong bird label is “UAV”; bold – the best share of correct labels without ideal association). For comparison, the limit attainable with ideal association of classifier boxes with the target is given.

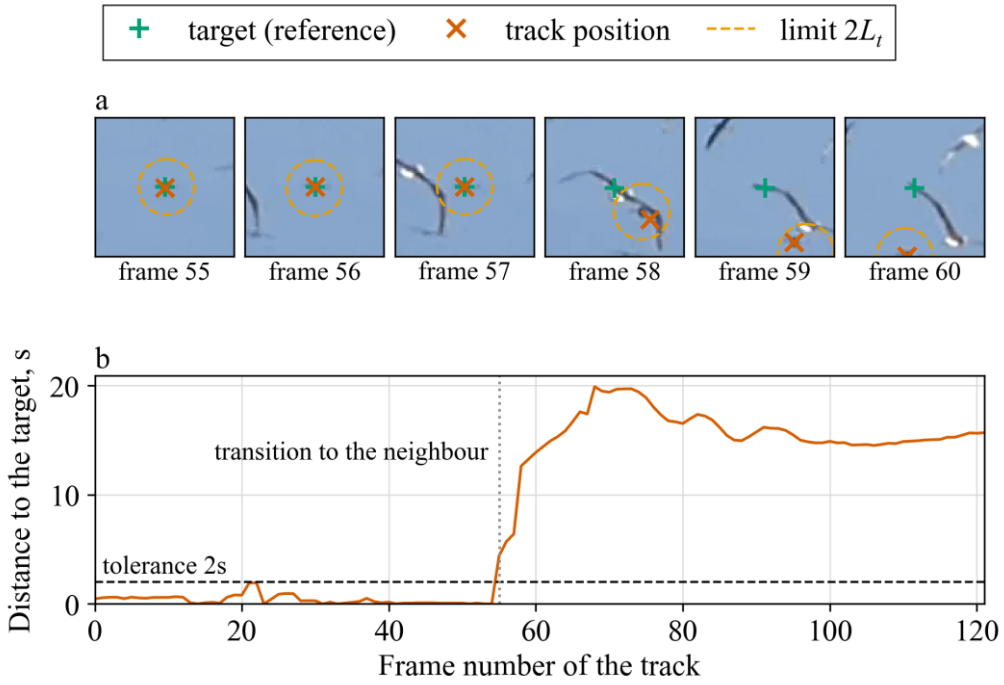


Fig. 6. **Failure of abductive identity on a bird track: a – frames before and after a neighbouring bird crosses the target ; b – distance from the track to the reference in target sizes *s***

Source: **compiled by the authors**

Table 3. **Classifier labels along the track: correct / wrong, %**

Target	Tracking by detection	Attention window	Attention window with abductive identity	Attention window, crop 120 → 640 px	Ideal association
UAV 6 px	21 / 4	33 / 4	33 / 4	**50** / 4	50 / 4
UAV 9 px	25 / 4	**54** / 8	**54** / 8	42 / 17	67 / 8
UAV 12 px	33 / 8	62 / 12	**67** / 12	54 / 33	67 / 12
Birds	4 / 8	**22** / 34	18 / 30	4 / 46	28 / 34

Source: **compiled by the authors**

Table 3 shows that the share of correct UAV labels of the same classifier increases from 21-33 % with tracking by detection to 33-62 % with the attention window and approaches the limit of 50-67 % attainable with ideal association; for 9 and 12 pixels the increase of 29 percentage points is significant (95 % intervals +4…+63 and +13…+42 p.p.). The classifier was not changed: only the share of its boxes that belong to the target increased. This gain is due to the attention window. Abductive identity does not change the share of labels significantly: 33-67 % correct UAV labels, and 30 % instead of 34 % wrong "UAV" labels for birds (p = 0.63). Its contribution is that the labels and facts are collected from the designated target (Table 1). A crop of the window enlarged to the classifier input size raises the visibility of 6-pixel UAVs from 12 % to 64 % of frames, but for 9 and 12 pixels it increases the share of wrong labels. For birds, wrong "UAV" labels amount to 34 % even with ideal association.

Abductive refinement of the classifier conclusion

The conclusions of the classifier without and with the abductive program are given in Table 4.

Table 4 (label trust threshold 0.5; precision – share of correct among the decisions taken; coverage – share of targets with a decision; bold – the better value) shows that with the abductive program there are no wrong conclusions for UAVs (0 of 72 versus 4 for the classifier) and one for birds (versus 4); the precision of the decisions taken rises from 83 % to 100 % for UAVs and from 33 % to 83 % for birds. The price is coverage: the program takes a decision for 18 % of UAV series versus 32 % for the classifier. On this sample the difference in wrong conclusions is not significant (exact McNemar test, p = 0.125). At range the classifier never had a box with a confidence of 0.9, and the median confidence of the labels is mostly below the trust threshold of 0.5, so the answer "unknown" prevails.

Each conclusion is returned with its justification. For a bird track, the majority of classifier labels is "UAV" with a median confidence of 0.46, below the trust threshold, and flapping flight is found on a series of 63 frames; the program returns "bird" with the justification kinematic_only(wing_beat_oscillation): the label is not trusted, and the class is supported by a physical fact. At a trust threshold of 0.4 the same track gives revised (drone, refuted_by(flapping_flight), supported_by(wing_beat_oscillation)): the label "UAV" is refuted by flapping flight. For a UAV segment with the label "UAV" of confidence 0.55 and without flapping flight the program returns accepted (drone, not_refuted): the label is accepted because nothing refutes it, but it is not physically confirmed. When the label is below the trust threshold and no fact supports a class, the program returns no_admissible_hypothesis, that is, "unknown". For the operator, a revision indicates an error of the classifier, and acceptance without confirmation calls for further observation. A conclusion of the type "confirmed" did not occur: only the label and the flapping fact are passed to the program, and confirmation of the UAV class requires a positive fact, such as hovering or a long straight flight, which was not measured.

The results of the start from a classifier detection are given in Table 5 (detection with a confidence of at least 0.5 within 25 pixels of the target; label trust threshold 0.5).

The classifier found the target with a confidence of at least 0.5 only for 21 of 50 birds and for 6, 14 and 15 of 24 UAV segments, after a median of 24-42 frames. After such a start the share of frames on the target is 23-78 % versus 22-83 % with the reference start, and there are no wrong decisions "bird" for UAVs. For birds the start from detection gives more frames on the target than the reference start (60 % versus 38 %). On the same 21 tracks the reference start gives 47 %, and on the same number of first frames 56 %; the difference is thus explained by the composition of the subsample and by the age of the track, which drifts off the target with time. The dependence of the conclusion on the trust threshold is shown in Fig. 7.

Table 4. **Conclusions of the classifier and of the abductive program**

Target (number)	Classifier: correct / wrong / unknown	Classifier: precision / coverage, %	With abductive program: correct / wrong / unknown	With abductive program: precision / coverage, %
UAV 6 px (24)	5 / 1 / 18	83 / 25	2 / **0** / 22	**100** / 8
UAV 9 px (24)	6 / 1 / 17	86 / 29	5 / **0** / 19	**100** / 21
UAV 12 px (24)	8 / 2 / 14	80 / 42	6 / **0** / 18	**100** / 25
UAV, all (72)	19 / 4 / 49	83 / 32	13 / **0** / 59	**100** / 18
Birds (50)	2 / 4 / 44	33 / 12	5 / **1** / 44	**83** / 12

Source: **compiled by the authors**

Table 5. **Start from a classifier detection**

Target (number)	Targets found	Median delay, frames	Frames on target, reference start: all targets / same targets / same first frames, %	Frames on target, start from detection, %	Conclusions after the start from detection: correct / wrong / unknown
Birds (50)	21	42	38 / 47 / 56	60	8 / 2 / 11
UAV 6 px (24)	6	37	22 / 16 / 15	23	2 / 0 / 4
UAV 9 px (24)	14	24	47 / 51 / 57	54	6 / 0 / 8
UAV 12 px (24)	15	28	83 / 84 / 89	78	8 / 0 / 7

Source: **compiled by the authors**

As Fig. 7 shows (a – UAVs, 72 series: 24 segments × 3 sizes; series of one segment are not independent); b – birds (50 tracks); dotted line – classifier without the abductive program), at a threshold of 0.4 the share of correct conclusions for UAVs is 29 % versus 26 % for the classifier, and wrong decisions “bird” amount to 3 % versus 6 %; this difference of two series is not significant. Lowering the threshold to 0.3 increases the share of correct conclusions for UAVs to 49 % ($p = 0.002$, 95 % interval of the difference +9…+40 p.p.), while wrong decisions “bird” grow to 8 %. For birds both correct (2 → 12) and wrong (4 → 15) conclusions grow, so the threshold of 0.5 is recommended as the main one.

DISCUSSION OF THE RESULTS

The decrease in the share of frames off the target for UAVs (Table 1) is explained by the fact that constraint (2) rejects neighbouring silhouettes lying farther than two body lengths from the prediction. Without the constraint, such a silhouette, being the nearest to the prediction, was accepted as the target, and the track passed to another object. For 9-pixel UAVs this also increased the share of continuous series on the target, because the track stays on one object longer.

The limit of one body length broke bird series not because of the noise of the centre position ($\sigma \approx 1.4$ pixels on the reference tracks) but because of the heavy tail of the deviation (Fig. 2): when the thresholded silhouette breaks into body and wings, its centre jumps beyond one body length in 7.6 % of frames, and each such frame breaks a gap-free series of 48 frames. That is why a lower bound of L_t related to the noise did not help. The limit of two body lengths covers this tail, keeps bird series and leaves the UAV results practically unchanged, since the UAV deviation rarely exceeds even one body length. For birds the coefficient sets a compromise: one body length reduces wrong “UAV” labels along the track more strongly (from 17 to 10 of 50 versus 15 with $2L_t$) but halves the series on the target. Larger silhouettes are not affected, since for them the limit reaches the search gate.

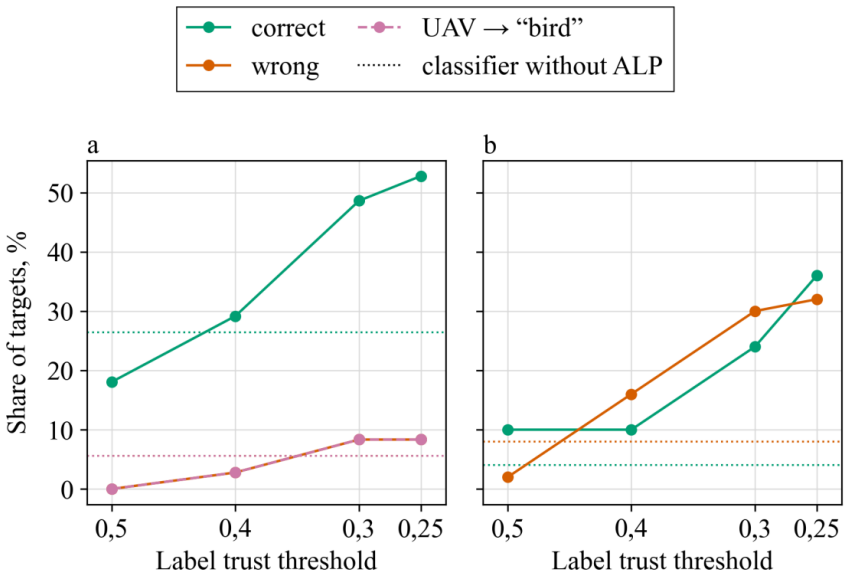


Fig. 7. **Dependence of the conclusion of the abductive logic program (ALP) on the classifier label trust threshold, which decreases from left to right: a – UAVs; b – birds**

Source: **compiled by the authors**

An identity check based on the size and contrast of the silhouette was considered and rejected. It halved the share of series on the target for birds, because the change of the silhouette size during a wingbeat is exactly the feature checked by the flapping flight criterion. A kinematic constraint applied to tracking by detection without the attention window did not change the results: in such tracking there are almost no transitions to other objects, and the limitation is missing boxes.

The increase in the share of correct classifier labels (Table 3) is explained by the attention window, which keeps classifier boxes near the target in frames where tracking by detection has lost it. Track maintenance does not correct the systematic error of the classifier: with ideal association it assigns 34 % of bird tracks to UAVs, and a larger number of labels only make this error more stable. It is exactly this error that the abductive program revises when the fact of flapping flight refutes the label “UAV”.

The results of Table 4 and Fig. 7 are explained by the fact that the program [28] takes the classifier label into account only when its confidence is not lower than the trust threshold. At range the label confidence is mostly below 0.5, and flapping flight of a UAV is not observed, so the program more often answers “unknown”. Lowering the threshold brings back wrong decisions “bird”, because the absence of wingbeats does not refute the class “bird”: soaring birds do not flap their wings.

The method is fast enough for real time. On one thread of a laptop the identity step in Python takes a median of 0.49 ms (95th percentile 0.87 ms) per target and frame, and the Prolog query 0.08 ms (0.24 ms), 30–70 times less than the frame period of 33–40 ms; the time of the detector is not included. A target of 9 pixels corresponds to a range of about 600 m for a UAV of 0.5 m and a camera of 1920 pixels with a 10° field of view, $D = l \cdot W/(n \cdot FOV)$, and to about 1.2 km at 5°. On a pan–tilt platform prediction (1) has to be transferred to the current frame by the platform angles: a pure rotation of the camera maps the image by the homography $H = K \cdot R \cdot K^{-1}$, which does not depend on the scene depth [34], and such compensation is used for pan–tilt cameras [35]. The sky has little texture, so the platform angles must be synchronised with the frame [36]. A residual compensation error modelled as a random walk of the frame degrades any track maintenance in image coordinates: for 12-pixel UAVs the share of frames off the target without the constraint grows from 23 % to 41 % at 1 pixel per frame. Constraint (2) keeps a smaller share at any error (38-45 % versus 58-62 % for 9 pixels and 14–32 % versus 37-41 % for 12 pixels at 1-4 pixels per frame), and long series require a residual error of at most 1-2 pixels per frame, that is, angles synchronised with the frame better than 1-2 ms at a rotation rate of 5°/s. An external target designation (SE-1) can be provided by radar or another camera of the system.

The limitations of the study are as follows. Small UAVs were obtained by downscaling the frame; atmospheric effects and the motion of the camera were not modelled, apart from the random walk of the frame. The frame rate of the dataset [30] was taken as 25 Hz; formulas (1)-(3) are written in frames and pixels and do not depend on it, and taking 30 Hz changes only the flapping criterion (wrong decisions “bird” 1.6 % instead of 0 %). Birds are represented by one large species from one video. The band of 1.5-6 Hz covers large birds; small birds flapping faster than 10 Hz are close to or above the Nyquist frequency at 25-30 frames per second, and soaring birds do not flap at all. The absence of wingbeats does not refute the class “bird”, so for such birds the program answers “unknown” or keeps the classifier label. For UAVs the reference is the annotated boxes, so the deviation of the silhouette centre was compared with a reference directly only for birds. The coefficient of the limit was chosen on the same data, and the sample is small and dependent (24 segments of 12 videos), so the conclusions are given with cluster bootstrap intervals and the McNemar test. The start from a classifier detection occurred only for 25-62 % of distant targets. For 6-pixel UAVs the silhouette mask has only 2×2 pixels, and track maintenance based on it does not work.

It is advisable to develop the research in two directions. The first is physical facts that positively indicate the UAV class at range, since in their absence the program often answers “unknown” and never confirms the UAV class. The second is verification on several bird species, with a moving camera and with real compensation of the platform motion.

CONCLUSIONS

1. The establishment of the identity of a small target is formalised as an abductive choice between the hypotheses “the silhouette is the target” and “the target is not visible” with an integrity constraint: the deviation of the silhouette from the prediction cannot exceed two body lengths of the target, a

margin that covers 98-99 % of the measured deviation of the true target. The constraint does not require appearance features, and the decisions of the logic and program implementations coincided on 1657 frames.

2. It is established that abductive identity reduces the share of frames off the target for UAVs of 9-12 pixels from 23-60 % to 11-46 % and for 9-pixel UAVs increases the share of continuous series on the target from 33 % to 46-47 %, without reducing the share of series for birds (34-44 % versus 32-40 %); for 6 pixels (a mask of 2×2 pixels) the method does not work. Unlike fixed and statistical gates, the limit in body lengths is not worse than basic track maintenance for any class and size.

3. It is established that the attention window raises the share of correct UAV labels of the same classifier from 21-33 % to 33-62 % without retraining, while abductive identity does not change the share of labels but ensures that they are collected from the designated target. Track maintenance does not correct the systematic error of the classifier on birds.

4. It is established that the abductive logic program makes the classifier conclusion checked and explained: at a trust threshold of 0.5 there were no wrong conclusions for UAVs (0 of 72 versus 4) and one for birds (versus 4), at the cost of decisions for only 18 % of UAV series; on this sample the difference is not significant. Lowering the threshold to 0.3 raises the share of correct UAV conclusions from 26 % to 49 % ($p = 0.002$), with 8 % wrong decisions “bird”.

Conflict of interests: The authors declare that there is no conflict of interests, in particular financial, personal, authorship or any other, which could have influenced the research and the results published in this article

Financing: Research was carried out without financial support

Accessibility data: The program code, the identity theory in Prolog and the scripts that reproduce the results are available in the open repository https://github.com/vnareznoy-dev/abductive_class_refinement/tree/v1.2-article4 (directory attention); archived copy of this version in Zenodo: https://doi.org/10.5281/zenodo.24104995. The open videos and datasets used are cited in the section Materials and Methods

Use of artificial intelligence tools: The authors used artificial intelligence technologies within acceptable limits to provide their own verified data, which is described in the research methodology section

REFERENCES

1. Jiang, N., Wang, K., Peng, X., Yu, X., Wang, Q., Xing, J., et al. “Anti-UAV: A large-scale benchmark for vision-based UAV tracking”. *IEEE Transactions on Multimedia*. 2023; 25: 486–500, https://www.scopus.com/pages/publications/85120075300.
DOI: https://doi.org/10.1109/TMM.2021.3128047.

2. Coluccia, A., Fascista, A., Schumann, A., Sommer, L., Dimou, A., Zarpalas, D., et al. “Drone vs. bird detection: Deep learning algorithms and results from a grand challenge”. *Sensors*. 2021; 21 (8): 2824, https://www.scopus.com/pages/publications/85104311184. DOI: https://doi.org/10.3390/s21082824.

3. Taha, B. & Shoufan, A. “Machine learning-based drone detection and classification: State-of-the-art in research”. *IEEE Access*. 2019; 7: 138669–138682, https://www.scopus.com/pages/publications/85078236502. DOI: https://doi.org/10.1109/ACCESS.2019.2942944.

4. Coluccia, A., Fascista, A., Dimou, A., Zarpalas, D., Sommer, L., Schumann, A., et al. “The Drone-vs-Bird detection grand challenge at IJCNN 2025”. *Proceedings of the 2025 International Joint Conference on Neural Networks (IJCNN)*. 2025. p. 1–8, https://www.scopus.com/pages/publications/105023962744. DOI: https://doi.org/10.1109/IJCNN64981.2025.11228314.

5. Kondo, Y., Ukita, N., Yamaguchi, T., Hou, H.-Y., Shen, M.-Y., Hsu, C.-C., et al. “MVA2023 small object detection challenge for spotting birds: Dataset, methods, and results”. *Proceedings of the 18th International Conference on Machine Vision and Applications (MVA).* 2023. p. 1–11, https://www.scopus.com/pages/publications/85168094609.
DOI: https://doi.org/10.23919/MVA57639.2023.10215935.

6. Unlu, E., Zenou, E., Riviere, N. & Dupouy, P.-E. “Deep learning-based strategies for the detection and tracking of drones using several cameras”. *IPSJ Transactions on Computer Vision and Applications*. 2019; 11: 7, https://www.scopus.com/pages/publications/85069679620.

DOI: https://doi.org/10.1186/s41074-019-0059-x.

7. Galchonkov, O. M., Baranov, A. M. & Baskov, I. O. “Improving the quality of mine detection by a mine-detecting drone by selecting an architecture for classifying objects in visible-light camera images”. *Applied Aspects of Information Technology*. 2026; 9 (1): 9–24. DOI: https://doi.org/10.15276/aait.09.2026.01.

8. Bar-Shalom, Y., Li, X. R. & Kirubarajan, T. “Estimation with applications to tracking and navigation: Theory, algorithms and software”. *John Wiley & Sons*. 2001. DOI: https://doi.org/10.1002/0471221279.

9. Bewley, A., Ge, Z., Ott, L., Ramos, F. & Upcroft, B. “Simple online and realtime tracking”. *Proceedings of the 2016 IEEE International Conference on Image Processing (ICIP)*. 2016. p. 3464–3468, https://www.scopus.com/pages/publications/85006782971. DOI: https://doi.org/10.1109/ICIP.2016.7533003.

10. Wojke, N., Bewley, A. & Paulus, D. “Simple online and realtime tracking with a deep association metric”. *Proceedings of the 2017 IEEE International Conference on Image Processing (ICIP)*. 2017. p. 3645–3649, https://www.scopus.com/pages/publications/85045318361.
DOI: https://doi.org/10.1109/ICIP.2017.8296962.

11. Zhang, Y., Sun, P., Jiang, Y., Yu, D., Weng, F., Yuan, Z., et al. “ByteTrack: Multi-object tracking by associating every detection box”. *Computer Vision – ECCV 2022. Lecture Notes in Computer Science*. 2022. p. 1–21, https://www.scopus.com/pages/publications/85142725518. DOI: https://doi.org/10.1007/978-3-031-20047-2_1.

12. Aharon, N., Orfaig, R. & Bobrovsky, B.-Z. “BoT-SORT: Robust associations multi-pedestrian tracking”. *arXiv*. 2022. DOI: https://doi.org/10.48550/arXiv.2206.14651.

13. Cao, J., Pang, J., Weng, X., Khirodkar, R. & Kitani, K. “Observation-centric SORT: Rethinking SORT for robust multi-object tracking”. *Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2023. p. 9686–9696, https://www.scopus.com/pages/publications/85199960057. DOI: https://doi.org/10.1109/CVPR52729.2023.00934.

14. Reid, D. B. “An algorithm for tracking multiple targets”. *IEEE Transactions on Automatic Control*. 1979; 24 (6): 843–854, https://www.scopus.com/pages/publications/0018689571. DOI: https://doi.org/10.1109/TAC.1979.1102177.

15. Blackman, S. S. “Multiple hypothesis tracking for multiple target tracking”. *IEEE Aerospace and Electronic Systems Magazine*. 2004; 19 (1): 5–18, https://www.scopus.com/pages/publications/1442313954. DOI: https://doi.org/10.1109/MAES.2004.1263228.

16. Davey, S. J. & Rutten, M. G. “A comparison of three algorithms for tracking dim targets”. *Proceedings of the 2007 Information, Decision and Control Conference*. 2007. p. 342–347, https://www.scopus.com/pages/publications/34548862864. DOI: https://doi.org/10.1109/IDC.2007.374574.

17. Seidaliyeva, U., Ilipbayeva, L., Taissariyeva, K., Smailov, N. & Matson, E. T. “Advances and challenges in drone detection and classification techniques: A state-of-the-art review”. *Sensors*. 2024; 24 (1): 125. DOI: https://doi.org/10.3390/s24010125.

18. Yasmeen, A. & Daescu, O. “Recent research progress on ground-to-air vision-based anti-UAV detection and tracking methodologies: A review”. *Drones*. 2025; 9 (1): 58, https://www.scopus.com/pages/publications/85215940942. DOI: https://doi.org/10.3390/drones9010058.

19. Sangam, T., Dave, I. R., Sultani, W. & Shah, M. “TransVisDrone: Spatio-temporal transformer for vision-based drone-to-drone detection in aerial videos”. *Proceedings of the 2023 IEEE International Conference on Robotics and Automation (ICRA)*. 2023. p. 6006–6013, https://www.scopus.com/pages/publications/85168680108. DOI: https://doi.org/10.1109/ICRA48891.2023.10161433.

20. Kondo, Y., Ukita, N., Kanayama, R., Yoshida, Y., Yamaguchi, T., Yu, X., et al. “MVA 2025 small multi-object tracking for spotting birds challenge: Dataset, methods, and results”. *Proceedings of the 19th International Conference on Machine Vision and Applications (MVA)*. 2025. p. 1–15, https://www.scopus.com/pages/publications/105019059110.
DOI: https://doi.org/10.23919/MVA65244.2025.11175061.

21. Huang, B., Li, J., Chen, J., Wang, G., Zhao, J. & Xu, T. “Anti-UAV410: A thermal infrared benchmark and customized scheme for tracking drones in the wild”. *IEEE Transactions on Pattern Analysis*

and Machine Intelligence. 2024; 46 (5): 2852–2865, https://www.scopus.com/pages/publications/85178014259. DOI: https://doi.org/10.1109/TPAMI.2023.3335338.

22. Xie, B., Zhang, C., Wang, F., Liu, P., Lu, F., Chen, Z. & Hu, W. "CST Anti-UAV: A thermal infrared benchmark for tiny UAV tracking in complex scenes". *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. 2025. p. 6216–6225, https://www.scopus.com/pages/publications/105035142029.
DOI: https://doi.org/10.1109/ICCVW69036.2025.00647.

23. Kakas, A. C., Kowalski, R. A. & Toni, F. "Abductive logic programming". *Journal of Logic and Computation*. 1992; 2 (6): 719–770, https://www.scopus.com/pages/publications/77957187070. DOI: https://doi.org/10.1093/logcom/2.6.719.

24. Denecker, M. & Kakas, A. "Abduction in logic programming". In *Kakas, A. C. & Sadri, F. (eds.) Computational Logic: Logic Programming and Beyond*. *Lecture Notes in Computer Science*. 2002; 2407: 402–436, https://www.scopus.com/pages/publications/84867643604.
DOI: https://doi.org/10.1007/3-540-45628-7_16.

25. Poole, D. "Explanation and prediction: An architecture for default and abductive reasoning". *Computational Intelligence*. 1989; 5 (2): 97–110, https://www.scopus.com/pages/publications/84990556284. DOI: https://doi.org/10.1111/j.1467-8640.1989.tb00319.x.

26. Josephson, J. R. & Josephson, S. G. (eds.) "Abductive inference: Computation, philosophy, technology". *Cambridge University Press*. 1994. DOI: https://doi.org/10.1017/CBO9780511530128.

27. Garcez, A. d'A. & Lamb, L. C. "Neurosymbolic AI: The 3rd wave". *Artificial Intelligence Review*. 2023; 56 (11): 12387–12406, https://www.scopus.com/pages/publications/85150173157. DOI: https://doi.org/10.1007/s10462-023-10448-w.

28. Nareznoy, V. "Abductive class refinement for airborne objects: Code and reports (version v1.0-article3)". *Zenodo*. 2026. DOI: https://doi.org/10.5281/zenodo.23730247.

29. Starr, F. & Starr, K. "Midway Atoll – Bird sightings – Mar–Apr 2015" [Video]. *Wikimedia Commons, CC BY 2.0*. – Available from: https://commons.wikimedia.org/wiki/File:Midway_Atoll_-_Bird_Sightings_-_Mar-Apr_2015.webm. – [Accessed: Apr, 2026].

30. Zhao, J., Zhang, J., Li, D. & Wang, D. "Vision-based anti-UAV detection and tracking". *IEEE Transactions on Intelligent Transportation Systems*. 2022; 23 (12): 25323–25334, https://www.scopus.com/pages/publications/85131721334. DOI: https://doi.org/10.1109/TITS.2022.3177627.

31. Jocher, G., Chaurasia, A. & Qiu, J. "Ultralytics YOLOv8" [Software], AGPL-3.0. – Available from: https://github.com/ultralytics/ultralytics. – [Accessed: Apr, 2026].

32. "Birdrone dataset, version 10". *Roboflow Universe, CC BY 4.0*. – Available from: https://universe.roboflow.com/fatin-zamri/birdrone-dataset/dataset/10. – [Accessed: Nov, 2026].

33. Pennycuick, C. J. "Predicting wingbeat frequency and wavelength of birds". *Journal of Experimental Biology*. 1990; 150 (1): 171–185. DOI: https://doi.org/10.1242/jeb.150.1.171.

34. Hartley, R. & Zisserman, A. "Multiple view geometry in computer vision". *2nd ed. Cambridge University Press*. 2004. DOI: https://doi.org/10.1017/CBO9780511811685.

35. Allebosch, G., Van Hamme, D., Veelaert, P. & Philips, W. "Robust pan/tilt compensation for foreground–background segmentation". *Sensors*. 2019; 19 (12): 2668, https://www.scopus.com/pages/publications/85068141838. DOI: https://doi.org/10.3390/s19122668.

36. Lisanti, G., Masi, I., Pernici, F. & Del Bimbo, A. "Continuous localization and mapping of a pan–tilt–zoom camera for wide area tracking". *Machine Vision and Applications*. 2016; 27 (7): 1071–1085, https://www.scopus.com/pages/publications/84979673485. DOI: https://doi.org/10.1007/s00138-016-0799-x.

Received 30.07.2026
Received after revision 10.09.2026
Accepted 18.09.2026

DOI: https://doi.org/10.15276/aait.09.2026.41
УДК 004.89:004.932

Абдуктивне встановлення тотожності малорозмірної повітряної цілі для уточнення міток нейромережевого класифікатора

Нарєзной Володимир Юрійович[1)]
ORCID: https://orcid.org/0009-0009-1397-5687; shsa1res@gmail.com
Галчонков Олег Миколайович[1)]
ORCID: https://orcid.org/0000-0001-5468-7299; o.n.galchenkov@gmail.com. Scopus Author ID: 56081377900
[1)] Національний університет "Одеська політехніка", пр. Шевченка, 1. Одеса, 65044, Україна

АНОТАЦІЯ

Актуальність. Оптичні канали систем захисту від малих безпілотних літальних апаратів розрізняють апарати і птахів нейромережевими класифікаторами. На великій дальності ціль у кілька пікселів видно лише в окремих кадрах, а в групі об'єктів мітки належать різним цілям. **Мета статті.** Уточнення міток класифікатора на малорозмірній повітряній цілі абдуктивним встановленням тотожності без перенавчання. **Задачі.** Формалізувати тотожність цілі як абдуктивний вибір і визначити її вплив на утримання, мітки та висновок про клас. **Методи.** Для кожного кадру розглядаються дві гіпотези: силует біля прогнозу є ціллю, або ціль невидима. Обмеження цілісності спростовує силует, відхилення якого від прогнозу перевищує дві довжини тіла; цей запас задано за виміряним відхиленням справжньої цілі. Обирається мінімальне пояснення. Мітки однієї цілі та факт махового польоту передаються абдуктивній логічній програмі. **Наукова новизна.** Тотожність спирається на фізичне обмеження руху, масштабоване за виміряною довжиною тіла, замість ознак зовнішнього вигляду, яких малі силуети не мають. **Практична значущість.** Той самий класифікатор, утримуваний на одній цілі, дає більше правильних міток і пояснені висновки без перенавчання. **Результати.** Для апаратів розміром від дев'яти до дванадцяти пікселів частка кадрів поза ціллю зменшилась на чверть – половину без погіршення треків птахів. Вікно уваги підвищило частку правильних міток апарата з п'ятої – третьої частини до третини – трьох п'ятих, а тотожність її не змінила. Абдуктивна програма усунула хибні висновки щодо апаратів ціною меншої кількості рішень; для апаратів розміром шість пікселів метод не працює. **Висновки.** Утримання однієї цілі є передумовою уточнення міток класифікатора; основним обмеженням є часті відповіді «невідомо» за низької впевненості.

Ключові слова: абдуктивне міркування; тотожність цілі; обмеження цілісності; малорозмірна ціль; уточнення міток; маховий політ

ABOUT THE AUTHORS



Volodymyr Yu. Nareznoy - PhD student, Department of Information Systems. Odesa Polytechnic National University, 1, Shevchenko Ave. Odesa, 65044, Ukraine
ORCID: https://orcid.org/0009-0009-1397-5687; shsa1res@gmail.com
Research field: Abductive logic programming; computer vision; protection against unmanned aerial vehicles

Нарєзной Володимир Юрійович - аспірант кафедри Інформаційних систем. Національний університет "Одеська політехніка", пр. Шевченка, 1. Одеса, 65044, Україна
Галузь досліджень: абдуктивне логічне програмування; комп'ютерний зір; захист від безпілотних літальних апаратів



Oleg M. Galchonkov - Candidate of Engineering Sciences, Associate Professor, Department of Information Systems. Odesa Polytechnic National University, 1, Shevchenko Ave. Odesa, 65044, Ukraine
ORCID: https://orcid.org/0000-0001-5468-7299; o.n.galchenkov@gmail.com. Scopus Author ID: 56081377900
Research field: Artificial intelligence; evolution systems; machine learning

Галчонков Олег Миколайович - кандидат технічних наук, доцент кафедри Інформаційних систем. Національний університет "Одеська політехніка", пр. Шевченка, 1. Одеса, 65044, Україна
Галузь досліджень: штучний інтелект; еволюційні системи; машинне навчання