



An HMM-Based Framework for Identity-Aware Long-Term Multi-Object Tracking From Sparse and Uncertain Identification: Use Case on Long-Term Tracking in Livestock

Anne Marthe Sophie Ngo Bibinbe¹ · Chiron Bang² · Patrick Gagnon³ · Jamie Ahloy-Dallaire¹ · Eric R. Paquet¹

Received: 20 September 2024 / Accepted: 1 September 2025

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2026

Abstract

The need for long-term multi-object tracking (MOT) is growing due to the demand for analyzing individual behaviors in videos that span several minutes. Unfortunately, due to identity switches between objects, the tracking performance of existing MOT approaches decreases over time, making them difficult to apply for long-term tracking. However, in many real-world applications, such as in the livestock sector, it is possible to obtain sporadic identifications for some of the animals from sources like feeders. To address the challenges of long-term MOT, we propose a new framework that combines both uncertain identities and tracking using a Hidden Markov Model (HMM) formulation. In addition to providing real-world identities to animals, our HMM framework improves the F1 score of ByteTrack, a leading MOT approach even with re-identification, on a 10-minute pig tracking dataset with 21 identifications at the pen's feeding station. We also show that our approach is robust to the uncertainty of identifications, with performance increasing as identities are provided more frequently. The improved performance of our HMM framework was also validated on the MOT17 and MOT20 benchmark datasets using both ByteTrack and FairMOT. The code for this new HMM framework and the new 10-minute pig tracking video dataset are available at: <https://github.com/ngobibinbe/uncertain-identity-aware-tracking>

Keywords Multi-object tracking · Uncertain identification · Identity-aware tracking · Long-term tracking · Livestock tracking

1 Introduction

Multi-object tracking (MOT) involves localizing and uniquely identifying objects in a video. While existing approaches succeed in performing MOT for short videos, their performance deteriorates significantly over time on longer videos (Luo et al., 2021; Valmadre et al., 2018). This is primarily due to inherent tracking challenges like object identity switches and loss of tracking identities caused by erratic object movements, occlusions, and the presence of objects with similar appearances (Luo et al., 2021; Valmadre et al., 2018). Existing MOT approaches are usually supervised (Chu et al., 2023; Zhang & Wang, 2021), self-supervised (Bastani et al., 2021; Carl et al., 2018), or unsupervised (Yifu et al., 2021; Bewley et al., 2016; Nicolai et al., 2018). Among these, there is ByteTrack (Yifu et al., 2021) that performs tracking by object detection with the particularity of keeping objects of low detection score in a buffer (eg. occluded objects) to match them, if necessary, with similar tracklets. There is also FairMOT (Zhang & Wang, 2021) which learns at the same time the object detection, tracking, and Re-ID

Communicated by Shangzhe Wu.

✉ Eric R. Paquet
eric.paquet@fsaa.ulaval.ca

Anne Marthe Sophie Ngo Bibinbe
anne-marthe-sophie.ngo-bibinbe.1@ulaval.ca

Chiron Bang
fbangnjenjoc2021@fau.edu

Patrick Gagnon
pgagnon@cdpq.ca

Jamie Ahloy-Dallaire
jamie.ahloy-dallaire.1@ulaval.ca

¹ Animal Science Department, Laval University, Quebec City, QC, Canada

² Department of Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, FL, USA

³ Centre de développement du porc du Québec, Lévis, QC, Canada

features through the detection and Re-ID heads of its network. Most approaches have been tested on benchmarks such as (Voigtlaender et al., 2019) and (Sun et al., 2022) which contain videos lasting usually less than 3 minutes. In many applications of MOT, tracking is often required over long periods, ranging from several minutes to hours. For example, in animal behavior analysis, livestock management, video surveillance, and sports analysis, MOT is often needed for hours-long durations (Psota, 2020; Cui et al., 2023), making 3-minute videos relatively short as a benchmark.

Furthermore, in those MOT applications, it is generally essential not only to detect and track objects but also to assign them to their real-world identities (RWID, e.g., animal IDs, person names, etc.) and to track those RWID long enough for further analysis (e.g. detect health problems or aggressive behaviors) (Psota, 2020; Marsot et al., 2020). Assigning real-world identities to objects is known as identity-aware MOT, and RWID could be provided by external sources like facial recognition algorithms (Dehghan et al., 2015) or automated identification stations as it is the case in the livestock sector (Psota, 2020; Shouu et al., 2016) with feeders and drinkers. Livestock animals could also be equipped with special markers (Psota, 2020) (e.g. tags, QR codes) that are easy to assign to RWID directly from videos, but those markers are rarely used and require extra maintenance. The ideal situation in this sector is to be able to obtain RWID from readily available tools like animal feeders/drinkers where animals are identified through already attached radio frequency identification (RFID) tags used to monitor their feeding/drinking behavior.

2 Related Works

In order to address long-term identity-aware tracking, information on identities is used for re-identification (Shouu et al., 2016; Dehghan et al., 2015). Re-identification consists of correcting the tracking by using identifier (i.e. entity providing RWID identification information) inputs where the identity provided by the identifier is assigned to the object. Here the identifiers could be facial recognition models, active learners, feeding/drinking stations, etc. However, those approaches rely on the accuracy of the identifier to identify the detected objects. Those approaches do not take into account contexts where identifications are sparse over time. (Yu et al., 2016) formulated a solution for identity-aware tracking merging tracking by detection and RWID information from an identifier. They formulated the tracking problem as a quadratic optimization problem with specific constraints ensuring that a detected object can only be associated to one RWID and vice versa. Unfortunately, for this framework, the identities are certain since they are based on face recognition. This

framework is thus not suitable for a context where identities could be assigned to more than one nearby objects due to uncertainty.

While those methods could be suitable for identity-aware MOT, they do not necessarily take into account the uncertainty of the identifier, and the proposed solutions are not exact solutions but approximations. However, uncertain identifications are frequent in various applications, such as in the livestock sector where animals are identified at the feeding stations using RFID tags (Voulodimos, 2010). For example, there could be many animals near the feeding station, as presented in Fig. 1. This frequent situation leads to uncertain identifications of nearby animals. Camilleri (2023) formulated their solution as an integer linear programming (ILP) problem, where they associate RWID with tracklets under the assumption that no identity switch occurs within a tracklet. Their approach accounts for RWID uncertainty by assigning different weights to each RWID-tracklet pair. Since ILP is NP-hard (Aloul et al., 2002; Camilleri, 2023) opted to group frames sharing the same subset of active tracklets, transforming the tracking problem into a set-covering problem. While this problem remains NP-hard, frame grouping improves the method's time efficiency. However, if only a small number of frames can be grouped, the algorithm becomes slower. Consequently, in certain tracking scenarios, the method may experience significant delays, especially when large frame groups cannot be formed. Moreover, it assumes that individual tracklets are perfectly reliable, which is not always the case especially for large groups of frames. Lastly, RWID identifications must be frequent and in some contexts, such as markerless livestock tracking, identifications are sparser, making the approach less applicable.

Contribution: In this paper, we propose a long-term identity-aware MOT framework that is able to use sparse and uncertain RWID information from a feeder to improve tracking of a base tracker. The framework is based on the formulation of the tracking problem as a Hidden Markov Model (HMM) problem (Degirmenci, 2014; Sin & Kim, 1995). Formulating the problem like this provides a new identity-aware framework with a solution in $O(TN^3)$ (where T is the number of frames and N the number of RWIDs). The problem formulation provides the ability to combine classic MOT approaches and uncertain sparse identifications. Thus, by using our new HMM framework, approaches like (Shouu et al., 2016; Dehghan et al., 2015) could be used as RWID identity sources and combined with other MOT approaches like (Luo et al., 2021; Yifu et al., 2021; Bewley et al., 2016; Nicolai et al., 2018; Chu et al., 2023; Bastani et al., 2021; Carl et al., 2018).


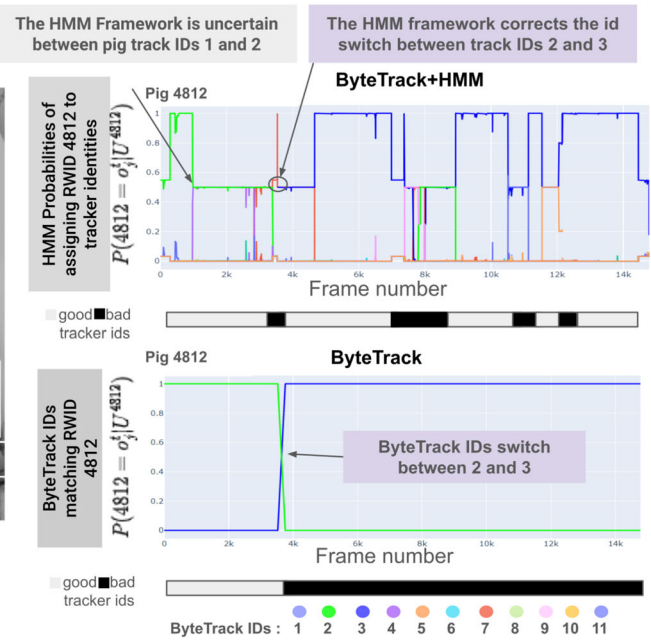


Fig. 1 (Left) On frame 884 of this video, the pig with the RWID 4812 is detected at the feeding station (the red rectangle at the top). At the same time, the tracker ByteTrack detected two pigs with IDs 1 and 2 near the feeding station. In this example, the feeding station is the *identifier* (i.e. the device providing the official RWID). Each pig gets

In addition to the HMM framework, we also provide a new long-term 10-minute tracking dataset of pigs with 21 uncertain RWID identifications obtained from a feeder station.

This approach is the first of its kind in the livestock industry. We tested the framework on the pig dataset with ByteTrack (Yifu et al., 2021) as the base tracker and compared it to ByteTrack and ByteTrack with re-identification. The use of our HMM framework improved the identification F1 score with a minimal number of identifications. Our framework also exhibits stable performance over time unlike the other approaches whose performance deteriorate over time. It also shows a faster convergence in terms of the number of uncertain identifications compared to ByteTrack with re-identification, which was hindered by incorrect identifications. Those findings were validated on six videos from the MOT17 (Dendorfer et al., 2021) dataset, on which the approach was tested due to the lack of open tracking datasets similar to the one we provide. On MOT17, both ByteTrack and FairMOT (Zhang & Wang, 2021) were tested in combination with our HMM framework and we found similar results. Overall, our new HMM framework could provide RWID to objects in a video and improve MOT when sparse and uncertain RWID are provided by an identification source in a scene.



a 50% chance of being RWID 4812. The HMM framework then splits the chance between them (upper right). Later at frame 3666 of the same video, ByteTrack switched IDs 2 and 3 (lower right). Our approach ByteTrack+HMM, using the *identifier* past and future information, corrects this by assigning 4812 to ID 3 after frame 3666 instead of ID 2

3 Method

3.1 HMM Framework

3.1.1 Overview

Let us consider the assignment of RWIDs to detected objects using a combination of multiple HMMs (i.e. one HMM for each RWID) (Degirmenci, 2014; Sin & Kim, 1995). Each HMM is responsible for assigning one given RWID to detected objects (eg. pigs), by taking into account all identifications of that RWID by the *identifier* as presented in fig. 2

The goal is to determine the optimal assignment of RWIDs to detected objects while considering all identifications made by the *identifier*. In our HMM formulation, the RWIDs obtained from the *identifier* serve as observations. Given this formulation, the forward-backward algorithm can efficiently estimate the probability of assigning each RWID to an object at every frame in polynomial time (Degirmenci, 2014). We did not use the Viterbi algorithm in this context, as it would provide the optimal RWID assignment for each individual HMM independently, without accounting for the assignments from the other HMMs. Instead, we combine the forward-backward algorithm with the Hungarian algorithm to merge assignments from individual HMMs, ensuring

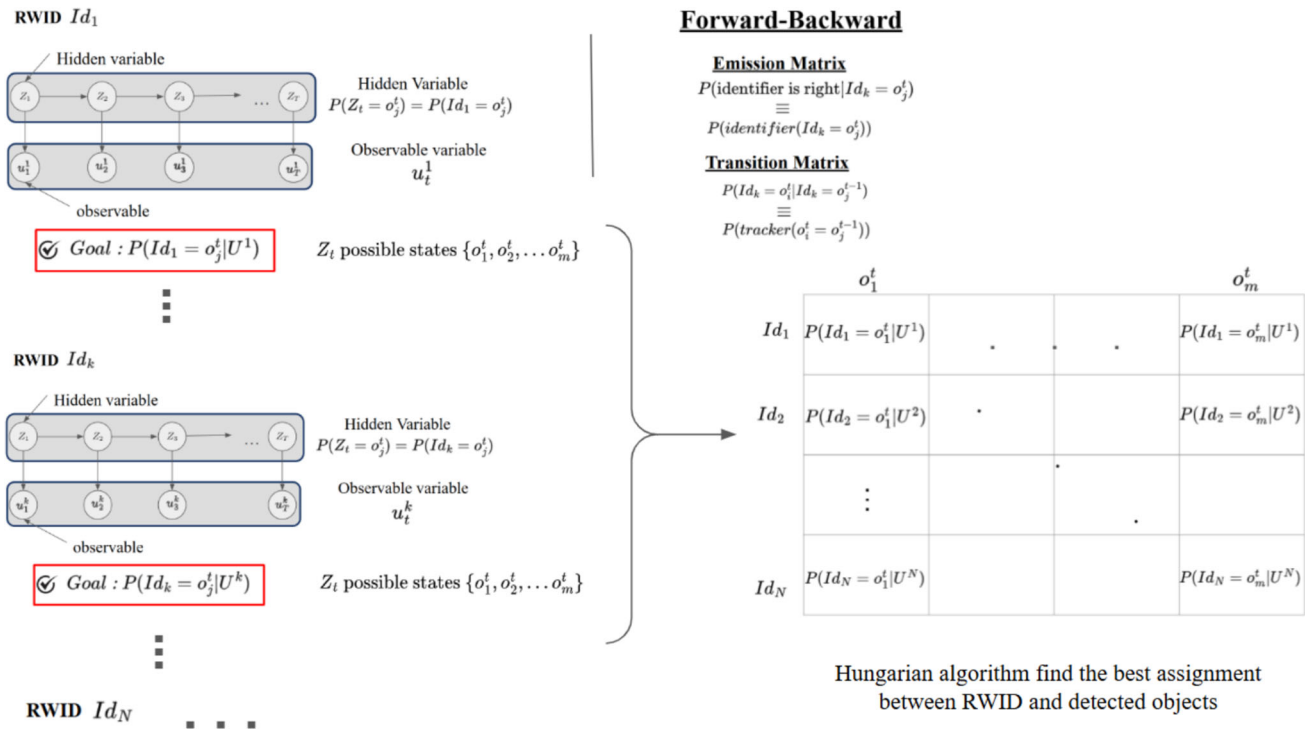


Fig. 2 Our framework defines one HMM for each real-world identity (RWID) of interest (left part of the figure). For each HMM, the hidden state represents the probability of matching the RWID with a detected object (eg. a pig), while the observable state corresponds to the correctness of the *identifier*'s decisions. A forward-backward algorithm is

applied to estimate the probability of matching an RWID to detected objects given all identifications made by the *identifier* of that RWID. Finally, the Hungarian algorithm is used to find the optimal assignment between all RWIDs and detected objects at each frame

a more coherent and comprehensive tracking solution in polynomial time. The complexity of the forward-backward algorithm is in $O(Tm^2)$ for the individual HMM and in $O(NTm^2)$ for all RWIDs. Since the information from the individual HMM is integrated using an Hungarian algorithm at every frame the complexity of this step is in $O(TN^3)$. Since $m \leq N$, we can assume the overall complexity of the approach to be in $O(TN^3)$.

Applying the forward-backward algorithm requires transition and emission matrices, which can be obtained from the *tracker* (e.g., ByteTrack or FairMOT) and the *identifier* (e.g., a feeding station that probabilistically assigns RWIDs to nearby pigs), respectively. In Sect. 3.1.2, we will discuss in detail how the *tracker* and *identifier* generate these matrices.

3.1.2 Problem Formulation

Notations. Let us assume:

- T the total number of frames in a video.
- $O_t = \{o_1^t, \dots, o_m^t\}$ the set of m objects o detected at frame t .
- $ID = \{Id_1, \dots, Id_k, \dots, Id_N\}$ the set of all N RWID.

- $P(Id_k = o_j^t)$ the probability that RWID Id_k is the real identity of the detected object o_j^t at frame t .
- $P(\text{Identifier}(Id_k = o_j^t))$ the probability that o_j^t has the identity Id_k at time t according to the *identifier*.
- $U^k = \{u_1^k, \dots, u_T^k\}$ is the set of observations on identity Id_k (by the *identifier*). Here u_t^k is the event that the *identifier* choice is right (the *identifier* identified well the object with identity Id_k). So we have $P(u_t^k | Id_k = o_j^t) = P(\text{Identifier}(Id_k = o_j^t))$. Unfortunately, in some cases, identifications are highly sparse and not available at every time step t . To address this sparsity, we assume that when no identification is available, the *identifier* assigns an equal probability to all objects using $P(\text{Identifier}(Id_k = o_j^t)) = \frac{1}{m} \forall o_j^t \in O_t$.
- $P(\text{tracker}(o_i^{t-1} = o_j^t))$ is the probability that o_i^{t-1} matches o_j^t in two consecutive frames (i.e. they are the same object). This information is provided by the *tracker*.

The goal of our HMM formulation is to obtain $P(Id_k = o_j^t | U^k)$ (i.e. the probability that the RWID Id_k is really associated to o_j^t knowing the *identifier* observations) for each t, j, k . By considering a HMM for each k having the Id_k

state, and each $o_j^t \in O_t$ as possible values at time t , it is possible to compute $P(Id_k = o_j^t | U^k)$ using the forward-backward algorithm. However, the use of the forward-backward algorithm requires the transition and emission matrix at each t providing respectively equation 1 and 2:

$$P(Id_k = o_j^t | Id_k = o_i^{t-1}) = P(tracker(o_i^{t-1} = o_j^t)) \tag{1}$$

$\forall o_i^{t-1} \in O_{t-1}$ and $o_j^t \in O_t$ which is given by the *tracker*.

$$P(u_i^k | Id_k = o_j^t) = P(Identifier(Id_k = o_j^t)) \tag{2}$$

$\forall o_j^t \in O_t$. which is given by the *identifier*.

3.1.3 Problem-solving: the Forward-Backward Algorithm for our HMM

With the forward-backward algorithm we can process $\alpha_j(t) = P(u_1^k, \dots, u_t^k, Id_k = o_j^t)$, and $\beta_j(t) = P(u_{t+1}^k, \dots, u_T^k | Id_k = o_j^t)$ as in equation 3 and 4 respectively and obtain the marginal probability of a given state (Rabiner, 1989).

$$\alpha_j(t) = \sum_{l=1}^{|O_{t-1}|} \alpha_l(t-1) \times P(Identifier(Id_k = o_j^t)) \times P(tracker(o_l^{t-1} = o_j^t)) \tag{3}$$

with $\alpha_j(0) = 1 \quad \forall j$.

$$\beta_j(t) = \sum_{l=1}^{|O_{t+1}|} \beta_l(t+1) \times P(Identifier(Id_k = o_l^{t+1})) \times P(tracker(o_l^{t+1} = o_j^t)) \tag{4}$$

with $\beta_j(T) = 1 \quad \forall j$.

We are interested in finding $P(Id_k = o_j^t | U^k)$. With the forward-backward algorithm, it is possible to process $P(Id_k = o_j^t | U^k)$ as $\frac{P(U^k, Id_k = o_j^t)}{P(U^k)} = \frac{\alpha_j(t)\beta_j(t)}{P(U^k)}$ (Rabiner, 1989). However $P(U^k)$ is a constant value over time across detections. Knowing that we are interested in finding o_j^t with the highest probability of having identity Id_k , performing the assignment by considering $\alpha_j(t)\beta_j(t)$ as matching value between o_j^t and Id_k will provide the same result as considering $P(Id_k = o_j^t | U^k)$.

3.1.4 Numerical Stability

In the context of long-term tracking, the HMM is applied on a large numbers of frames. Since α and β are estimated from successive multiplication of small probabilities ≤ 1 , this could rapidly lead to numerical instability (Bishop et

al., 2006). To solve this issue, for each t , we normalize $\alpha(t)$ and $\beta(t)$ by $\sum_j \alpha_j(t)$ and $\sum_j \beta_j(t)$. Instead of using $\alpha(t)$ and $\beta(t)$ on the forward-backward algorithm, we use $\hat{\alpha}_j(t) = \frac{\alpha_j(t)}{\sum_j \alpha_j(t)}$ and $\hat{\beta}_j(t) = \frac{\beta_j(t)}{\sum_j \beta_j(t)}$. By this process, the probability propagation is maintained and we avoid going under machine epsilon. However, by doing so, the exact probability $P(Id_k = o_j^t | U^k)$ is lost but for each t and j , according to (Bishop et al., 2006) there exists c^t such that $\hat{P}(Id_k = o_j^t | U^k) = P(Id_k = o_j^t | U^k) \times c^t$. Therefore they can be used equivalently at each t Fig 3 .

3.2 Dataset

3.2.1 Experimental Setup

The dataset was collected from a pen of growing pigs of $12.7 m^2$ ($4.87m \times 2.6m$) shortly after they were transferred from the nursery. A camera (4MP resolution with a $2.7 \sim 13.5$ mm varifocal lens) recording at 25 fps was mounted on the ceiling to capture the entire area. A feeding station (i.e. the *identifier*) is positioned in the top part of the pen (Fig 1)

3.2.2 Animal Subjects

In the pen, there was a fixed number of 15 pigs of similar age, weight, and appearance, as is common in the livestock sector. The behavior of the 15 pigs was recorded during 10 minutes from 9:00 to 9:10 am (Eastern Time). During this 10-minute video, pigs were energetic and displayed various behaviors, including feeding (drinking and eating), playing, running across the pen, and engaging in interactions with other pigs including minor aggressions and play (Fig. 4).

3.2.3 Data Annotation

The goal of the dataset was to provide a comprehensive tracking resource suitable for various types of trackers, including those that perform tracking through pose estimation. For each pig, key points (the two ears, tail, neck, and nose) were manually annotated along with their identities. Based on these key points, a bounding box was created as the smallest enclosing rectangle around them. In rare cases, only two keypoints (e.g., the tail and the neck) are available, making it impossible to construct a bounding box directly. This occurs when other keypoints are occluded by other pigs or due to the pig's position within the pen. In such cases, a rectangle is generated by centering it on the line connecting the two keypoints, with a fixed width equal to the average width of all pig bounding boxes.

The 10-minute video was manually annotated with the key points and identities of the 15 animals, approximately

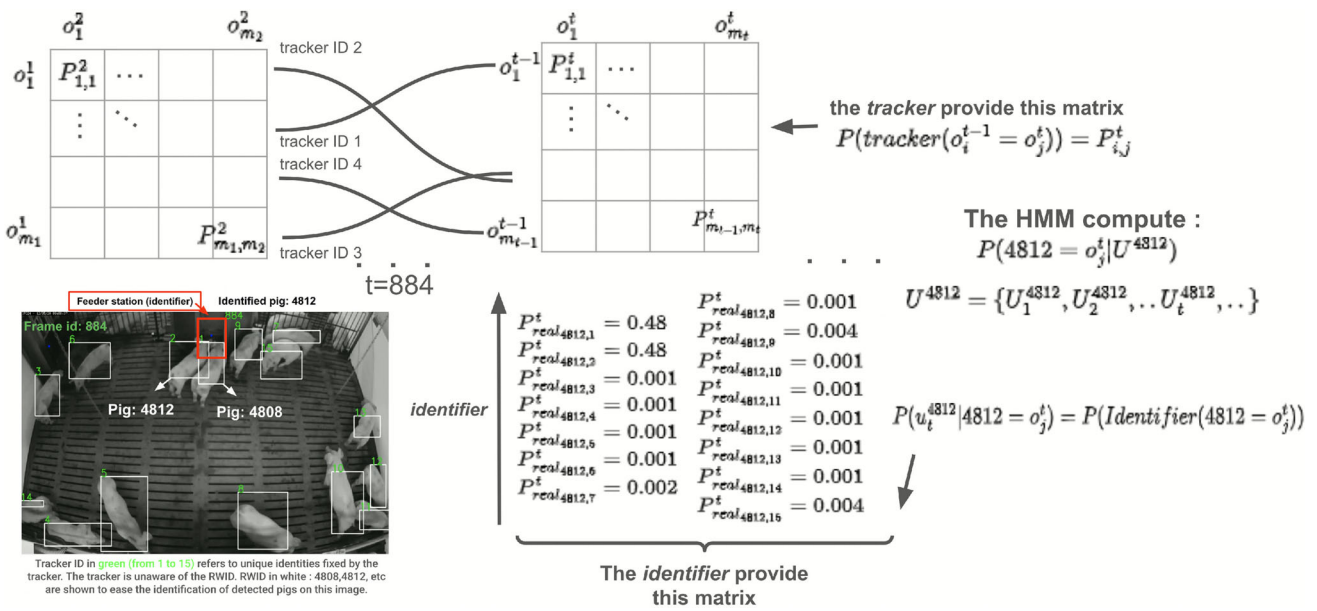


Fig. 3 HMM tracking reformulation explained. (Bottom left) Here, at frame number $t=884$, the pig 4812 is identified at the feeder, but there are two pigs near the feeder. In this situation, pigs with tracker IDs 1 and 2 will have the highest probabilities of being assigned to RWID 4812 since they are close to the feeder. We retrieve this set of probabilities as a line of the emission matrix at frame number 884. The

process is similar for all frames. At the top we could see the transition matrix provided by the tracker at each frame number. These sets of emissions and transition matrices will be used together in the HMM framework to propose probabilities of assigning a RWID to a detected object given all identifications at the feeder. In this figure, $p^{real}_{k,j} = p(Identifier(Id_k = o_j^t))$




Fig. 4 Snapshots of pig activities in the park, extracted from our 10-minute video, are presented. (a) Two distinct frames captured 2 seconds apart ((a-1) at 09:01:53 and (a-2) at 09:01:55) highlight the changes in the pigs' positions over this short interval, illustrating their high levels of activity. (b) Various behaviors are depicted, including pigs

moving and interacting with each other. These activities often result in occlusions and erratic movements, compounded by the fact that all pigs share similar appearances. The combination of high activity levels and visual similarity among the pigs significantly increases the complexity of tracking in this scenario




Fig. 5 F1 score over time

every second, resulting in 780 annotated frames. In addition to these tracking annotations, 21 sporadic identifications from the feeder are provided, distributed throughout the video as shown in Fig. 5. The 21 identifications were recorded from 11 pigs that passed by the feeder. The full dataset and annotations are available at : <https://github.com/ngobibibnbe/uncertain-identity-aware-tracking/tree/main/dataset>.

4 Experiments

4.1 Testing the HMM Framework on a New 10-minute Video

To test our framework, we used ByteTrack (Yifu et al., 2021) as the base tracker. We evaluated three different tracking approaches in our experiments on the 10-minute video: ByteTrack alone (ByteTrack), ByteTrack with re-identification (ByteTrack+Re-ID), and ByteTrack in combination with our HMM framework (ByteTrack+HMM).

ByteTrack is a tracking-by-detection model and it requires a detection model. We trained a YOLOX (Zheng et al., 2021) model from 11700 annotations of pigs (80% training, 20% validation) using 500 epochs and got an mAP of 96.5% for the training set and 95% for the validation set. We note that using the same dataset to train the YOLOX model and to evaluate our new approach could influence the evaluation results. However, since all methods rely on the same model, this effect should have minimal impact on the comparison.

For the ByteTrack+HMM approach, ByteTrack provides the transition matrix. As for the emission matrix, when an RWID Id_k is detected at a given frame t , the probability, according to the *identifier*, that Id_k corresponds to a detected animal o_j at time t is computed as a function of the pig's distance to the feeding station. The closer the pig is to the station, the higher the probability :

$$P(\text{Identifier}(Id_k = o_j^t)) = \frac{1/\text{distance}(o_j^t, \text{station})}{\sum_l 1/\text{distance}(o_l^t, \text{station})}$$

When no identification is available at a given time t , the identification probability is set to :

$$P(\text{Identifier}(Id_k = o_j^t)) = \frac{1}{N} = \frac{1}{15},$$

since there is no indication of which detected animal could be associated with a given RWID. Here, 15 corresponds to the total number of pigs in the video. From this, individual HMMs are constructed for the RWIDs identified at least once at the feeder, and matching probabilities between detected animals and RWIDs are computed using the forward-backward algorithm. The Hungarian algorithm is then used to determine the optimal assignment between RWIDs and detected animals at each frame t . In each assignment, if the matching probability between an RWID Id and a detected animal o is lower than $1/N$, the match is disregarded due to high uncertainty. As a result, the animal is considered undetected rather than misidentified.

For the ByteTrack and ByteTrack+Re-ID approaches, ByteTrack's internal object IDs are matched with the closest RWIDs at the very first frame. For the ByteTrack approach the RWIDs are associated to the same internal object throughout the entire video. For the ByteTrack+Re-ID approach, if a pig is detected at the feeder station but does not have the expected RWID, an identity switch is assumed and the identities of the detected object at the feeder and the object that previously held that RWID are swapped. This correction is then applied to future frames.

Since we have access to real RWIDs and their corresponding pig bounding boxes as ground truth approximately every second throughout the 10-minute video, we can effectively evaluate the performance of different approaches by comparing the RWIDs assigned by each approach to those in our dataset.

As noted by (Yu et al., 2016; Camilleri, 2023), traditional MOT metrics do not fully address the requirements of identity-aware tracking, as they do not account for the correct assignment of RWIDs. Therefore, new metrics are needed that consider both the detection of ground truth objects and the accurate assignment of RWIDs.

To evaluate the performance of different approaches, we used I-TP (Identity-aware True Positive: ground truth object detected with the correct RWID), I-FP (Identity-aware False Positive: either an unmatched detected object OR a ground truth object detected with the wrong RWID), and I-FN (Identity-aware False Negative: either a ground truth object not detected OR a ground truth object detected without an assigned RWID). Additionally, we computed Micro-Precision, Micro-Recall, and Micro-F1 Score, following the methodology in Yu et al. (2016). These metrics were chosen because they provide a comprehensive measure of how well

Table 1 Tracking performance of the different approaches using Identity-aware True Positive (I-TP), Identity-aware False Positive (I-FP), Identity-aware False Negative (I-FN), Micro-Precision, Micro-Recall, and Micro-F1 score (see the text for their definitions) Bold values indicate the best-performing method in each column

Dataset	I-TP	I-FP	I-FN	Micro-Precision	Micro-Recall	Micro-F1 score
ByteTrack	4125	4252	291	0.49	0.48	0.48
ByteTrack+Re-ID	4222	4093	291	0.51	0.49	0.50
ByteTrack+HMM	5264	1751	1642	0.75	0.61	0.67

each approach assigns the correct RWID to detected individuals throughout the video.

Unlike (Yu et al., 2016), which applies a fixed 1-meter cutoff, we consider a detected object matched to a ground truth object if $\text{IOU} > 0$. Based on this criterion, Micro-Precision, Micro-Recall, and Micro-F1 Score are computed as follows:

$$\text{Micro-Precision} = \frac{\text{I-TP}}{\text{I-TP} + \text{I-FP}}$$

$$\text{Micro-Recall} = \frac{\text{I-TP}}{\text{I-TP} + \text{I-FN}}$$

$$\text{Micro-F1 score} = \frac{2 \cdot \text{Micro-Precision} \cdot \text{Micro-Recall}}{\text{Micro-Precision} + \text{Micro-Recall}}$$

As shown in Table 1, our new HMM framework improves identity-aware tracking in comparison to ByteTrack alone and ByteTrack+Re-ID for most metrics. We notice higher I-TP, lower I-FN, higher Micro-Precision, Recall, and F1 score. This improvement can be attributed to the HMM framework's ability to propagate identifications effectively, as well as its probabilistic nature, which facilitates the recognition of low-probability assignments. However, this probabilistic approach also accounts for the observed increase in I-FN, as it inherently considers a broader range of potential identifications, including those with lower confidence.

When analyzing the evolution of the performance of the three approaches over time as a function of an increasing number of identifications (Fig. 5), we can observe that the performance of ByteTrack+HMM exhibits greater stability compared to ByteTrack alone or ByteTrack+Re-ID, both of which show a decline in performance as time progresses. Additionally, we can see that as the number of identifications increases (after 6K frames), the HMM approach becomes more precise.

4.2 Testing the Impact of Uncertainty and the Number of Identifications on the HMM Framework

To assess the performance of ByteTrack+Re-ID and ByteTrack+HMM as a function of the level of uncertain identification information, and their respective robustness to




Fig. 6 F1 score as a function of the number of identifications from the *identifier*. Error bars represent standard deviation from 20 random simulations where 25% of observations have been shuffled

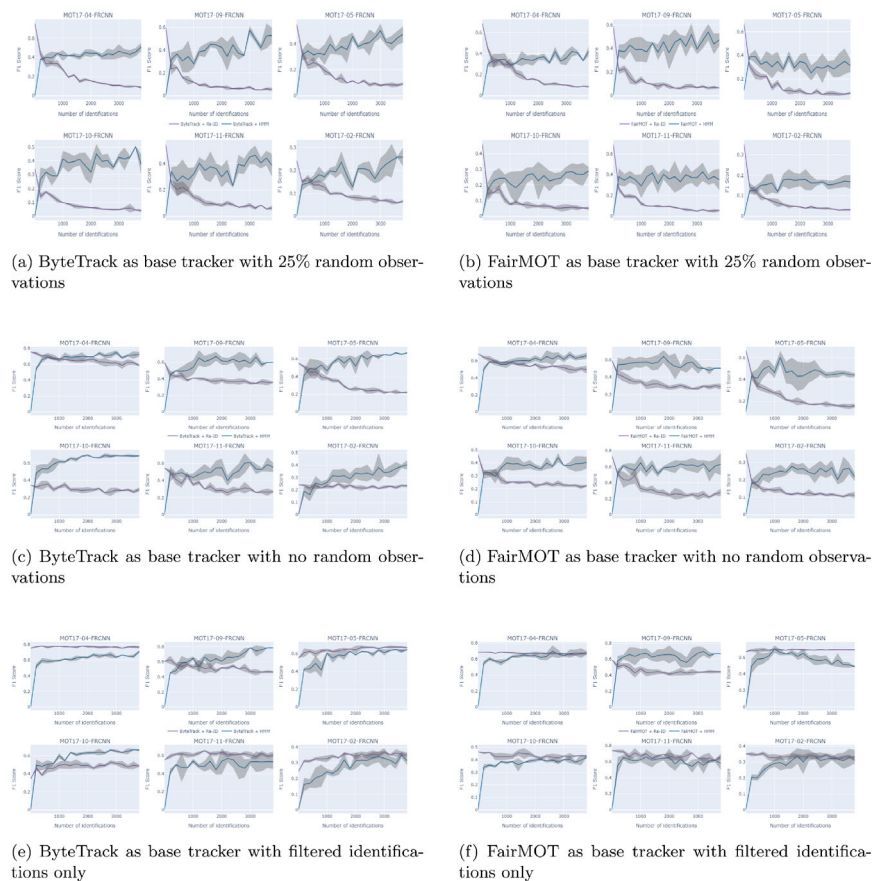
identification errors, we generated artificial identifications over the annotated 10-minute video. Here, 25% of the artificial identification information was randomly assigned to all animals, while 75% was assigned with higher probability to the animal with the real identity based on the distance to the ground truth from which we extract the identity. To ensure that our random simulations are representative this procedure was repeated 20 times.

On Fig. 6 we see that ByteTrack+HMM tends to have higher F1 score compared to ByteTrack+Re-ID when the number of identifications is over 8. This is due to the fact that the ByteTrack+Re-ID model obtain all the RWID information at the beginning of the video while ByteTrack+HMM needs to obtain them from the identifications. For RWID that are never identified by the *identifier* (e.g. pigs not visiting the feeder in a given video), the matching probabilities between all detected objects and those RWID will be the same. For those animals, this will lead to lower performance for ByteTrack+HMM.

4.3 Testing the HMM Framework with Another MOT Approach (FairMOT) on MOT17 and MOT20

To further test our HMM framework in videos outside the livestock sector, we used videos of MOT17 and MOT20 benchmarks (Dendorfer et al., 2021). These datasets feature videos of pedestrians filmed from different points of view (e.g. moving camera, elevated viewpoint, nighttime, daytime,

Fig. 7 F1 score as a function of the number of identifications from the artificial identifier in the MOT17 dataset. Error bars present standard deviation from 20 random simulations. In 7a and 7b, observations are generated such that 25% of them are randomly distributed. These do not include identification errors caused by crowded scenes during the artificial observation generation process, so there are more than 25% of misleading observations generated. In 7c and 7d none of the observations are randomly distributed, but rather are all intended to be correct. In 7e and 7f, observations are filtered to remove misleading identification candidates, and no random identifications are added to limit the number of misleading observations



etc.). The videos are relatively short, ranging from 18 seconds to 1 minute, with frame rates varying from 14 to 30 fps for MOT17. For MOT20 videos are still short ranging from 17 seconds to 2 minutes and 13 seconds, with frame rates of 25 fps. The number of people also varies between videos, from 26 to 133 in total for MO17, and from 90 to 1211 for MOT20. In all those videos, there are many individuals entering and exiting the scene throughout the video. The MOT17 and MOT20 dataset does not provide RWID identifications so we decided to simulate artificial identifications based on the ground truth provided in the MOT17 dataset. This is similar to what was presented in the previous section, but for the MOT17 and MOT20 videos.

We were also interested to test if the performance of our new HMM framework could also be extended to another MOT approach. For this experiment we tested both ByteTrack and FairMOT. We evaluated the performance of ByteTrack and FairMOT combined with Re-ID or our HMM framework as a function of the number of artificial random identifications evenly distributed throughout the videos. For our MOT17 and MOT20 experiment, we used the object detection provided by the model provided on the github of the authors (Yifu et al., 2021; Zhang & Wang, 2021) for both ByteTrack and FairMOT respectively. We conducted the tests

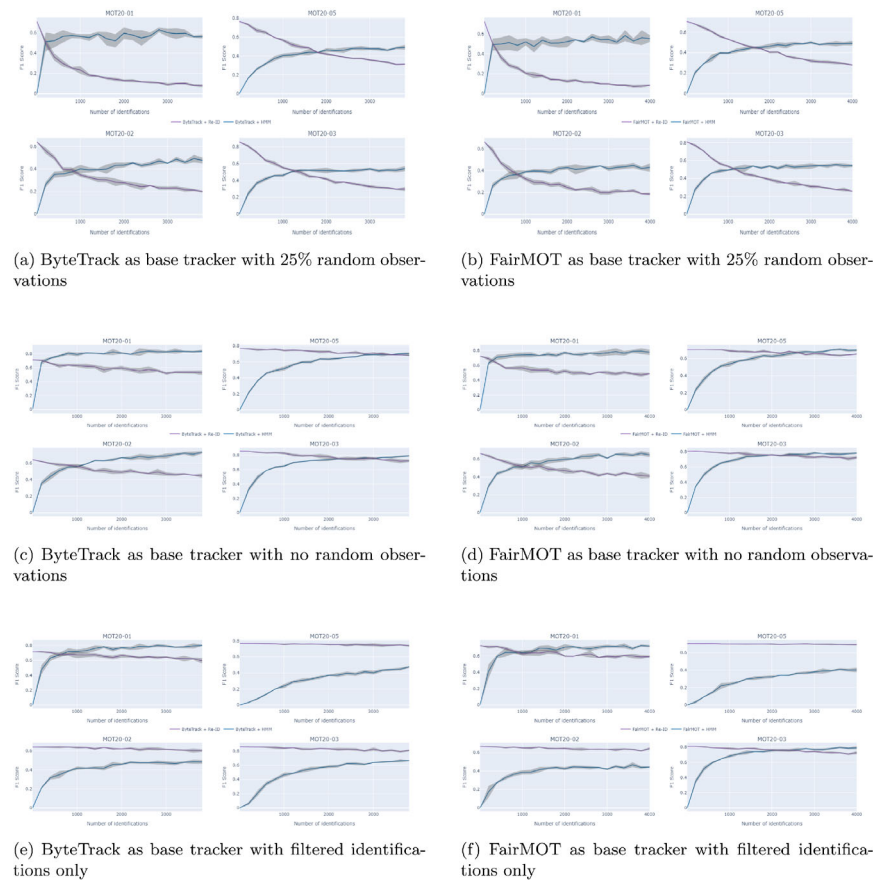
on the validation sets of the benchmarks, which represent the second half of each video.

Similar observations can be made on all plots of Figs. 7 and 8 where base trackers (ByteTrack and FairMOT) + HMM initially show lower performance but improve over the increasing number of identifications, eventually outperforming ByteTrack+Re-ID and FairMOT+Re-ID as more identifications are available, even with a significant number of misleading identifications.

However, compared to our 10-minute video, where the scene is less crowded and animals are less occluded, we noticed that many of the highest probabilities obtained from our artificial identifications were not assigned to the right person. Additionally, we observe a decline in the performance of FairMOT+Re-ID and ByteTrack+Re-ID when we increase the number of uncertain identifications. The consequence of this problem is the difficulty for FairMOT+Re-ID and ByteTrack+Re-ID to improve the performance of their base trackers, as shown in Figs. 7a, 7b, 7c, and 7d.

To reduce the effect of this problem we performed another experiment where we removed some of those misleading identifications (Fig. 7e, Fig. 7f, Fig. 8e, and Fig. 8f). Briefly, we added filters that require a minimal IOU of 0.7 for the detection associated to the ground truth identification and a

Fig. 8 F1 score as a function of the number of identifications from the artificial identifier in the MOT20 dataset. Error bars present standard deviation from 20 random simulations. In 8a and 8b, observations are generated such that 25% of them are randomly distributed. These do not include identification errors caused by crowded scenes during the artificial observation generation process, so there are more than 25% of misleading observations generated. In 8c and 8d none of the observations are randomly distributed, but rather are all intended to be correct. In 8e and 8f, observations are filtered to remove misleading identification candidates, and no random identifications are added to limit the number of misleading observations



probability higher than 0.5 for at least one detected object. This procedure does not guarantee the removal of all misleading observations, but helps mitigate at least some of them. With these filters, fewer identifications were provided to ByteTrack+HMM, FairMOT+HMM, ByteTrack+Re-ID, and FairMOT+Re-ID as we can see in Figs. 7f and 7e. In that experiment, no random observations were introduced, and we can observe that Re-ID is able to improve the base trackers on most datasets, similarly to our HMM approach. However, we also notice that in this context, our HMM is not that effective since we are in the scenario without filtering. This is likely due to the reduced number of identifications, which confirms the importance of the number of identifications for the HMM framework.

4.4 Discussion

The goal of our work was to provide a solution for identity-aware long-term tracking with sparse, uncertain identity information. As shown in Table 1 our HMM framework improves upon the basic tracker and the classic re-identification. This means that, at a specific time t , by using past and future identification information, an HMM-based tracker can make better decisions on identity assignment.

When examining the F1 score over time for the three tracking approaches presented in Fig. 5, we observe a decrease in performance over time for ByteTrack and ByteTrack+Re-ID. Meanwhile, ByteTrack+HMM remains stable and becomes more accurate over time, due to the increased number of identifications after 6K frames. We expect the performance of the HMM-based tracker to improve with the number of identifications provided by the *identifier* (as observed in Fig. 6). In our evaluation with a variable number of identifications (Figs. 6, 7 and 8), Re-ID-based trackers tend to stagnate with no significant improvement due to the identification errors whereas HMM-based trackers exhibit continuous improvements in the F1 score as the number of identifications increases. This is because the HMM framework makes better use of well-distributed identification information to make correct assignments despite the presence of unreliable identifications.

Figures 7 and 8 illustrate that the number of identifications is a crucial factor for the HMM framework, for both ByteTrack and FairMOT. The HMM framework demonstrates its ability to extract more information from a large number of uncertain identifications compared to a smaller number of more certain identifications, outperforming Re-ID in this regard. This indicates that the HMM framework is

robust to identification errors, suggesting that it is a promising approach for uncertain long-term identity-aware tracking.

5 Conclusion

In this work, we propose a framework based on an HMM reformulation of the tracking in order to tackle the uncertain identity-aware long-term tracking problem where sparse identity information is available. This framework takes advantage of the forward-backward algorithm for information propagation to provide a solution for this problem. The HMM modeling requires two key pieces of information: probabilities linking objects to identities and a tracking matrix matching objects between consecutive frames. The experiments showed that the framework is better than a basic re-identification, can handle identity information error, and is more accurate if more identifications are provided, even with some mistakes, which are desirable properties for accurate long-term tracking.

While our work demonstrates that integrating sparse and uncertain external identifications can improve tracking performance, the proposed method relies on an offline approach, which limits its applicability to real-time tracking. While implementing a time window could allow for real-time use, developing a fully online version remains an important direction for future research. Although our experiments cover multiple domains, the challenges associated with sporadic identification observed in the livestock sector may similarly affect other fields. We believe this HMM framework holds promise for applications such as human video surveillance, player tracking, and robotic vision.

Acknowledgements This work was supported by the MAPAQ (Quebec Ministry of Agriculture, Fisheries and Food) Innov'Action program. Thanks to the CDPQ (Centre de développement du porc du Québec) employees for gathering video data and Andréane Bélanger-Roy for annotations.

Data Availability Data are publicly available at this link : <https://github.com/ngobibibnbe/uncertain-identity-aware-tracking>.

Declarations

Competing interests The authors declare no competing interests.

References

- Aloul, F.A., Ramani, A., Markov, I.L., & Sakallah, K.A. (2002). "Generic ILP versus specialized 0-1 ILP: an update". In: *IEEE/ACM International Conference on Computer Aided Design, 2002. ICCAD 2002*. Pp. 450–457. <https://doi.org/10.1109/ICCAD.2002.1167571>.
- Bastani, F., He, S., & Madden, S. (2021). Self-Supervised Multi-Object Tracking with Cross-Input Consistency. *Advances in Neural Information Processing Systems, 17*, 13695–13706. <https://doi.org/10.48550/arxiv.2111.05943>
- Bewley, A., Ge, Z., Ott, L., Ramos, F., & Upcroft, B. (2016). "Simple online and realtime tracking". In: *2016 IEEE International Conference on Image Processing (ICIP)*. IEEE, pp. 3464–3468. <https://doi.org/10.1109/ICIP.2016.7533003>.
- Bishop, Christopher, M., Nasser, M., & Nasrabadi (2006). *Pattern recognition and machine learning*. Vol. 4. 4. Springer. <https://doi.org/10.1117/1.2819119>.
- Camilleri, M. P., Zhang, L., Bains, R. S., Zisserman, A., & Williams, C. K. (2023). Persistent animal identification leveraging non-visual markers. *Machine Vision and Applications, 34*(4). <https://doi.org/10.1007/s00138-023-01414-1>. issn: 1432-1769.
- Chu, P., Wang, J., You, Q., Ling, H., & Liu, Z. (2023). "TransMOT: Spatial-Temporal Graph Transformer for Multiple Object Tracking". In: *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 4859–4869. <https://doi.org/10.1109/WACV56688.2023.00485>.
- Cui, Y., Zeng, C., Zhao, X., Yang, Y., Wu, G., & Wang, L. (2023). "Sportsmot: A large multi-object tracking dataset in multiple sports scenes". *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9921–9931.
- Degirmenci, Alperen (2014). *Introduction to Hidden Markov Models*. https://scholar.harvard.edu/files/adeqirmenci/files/hmm_adeqirmenci_2014.pdf.
- Dehghan, A., Tian, Y., Torr, P. H., & Shah, M. (2015). "Target Identity-aware Network Flow for online multiple target tracking". In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1146–1154. <https://doi.org/10.1109/CVPR.2015.7298718>.
- Dendorfer, P., Ošep, A., Milan, A., Schindler, K., Cremers, D., Reid, I., & Leal-Taixé, L. (2021). "Motchallenge: A benchmark for single-camera multiple target tracking". In: *International Journal of Computer Vision* 129, pp. 845–881. <https://doi.org/10.1007/s11263-020-01393-0>.
- Ge, Z., Liu, S., Wang, F., Li, Z., & Sun, J. (2021). "YOLOX: Exceeding YOLO Series in 2021". *arXiv* 5, p. 12. <https://doi.org/10.48550/arxiv.2107.08430>
- Luo, W., Xing, J., Milan, A., Zhang, X., Liu, W., & Kim, T. K. (2021). "Multiple object tracking: A literature review". In: *Artificial Intelligence* 293, p. 103448. issn: 0004-3702. <https://doi.org/10.1016/j.artint.2020.103448>.
- Marsot, M., Mei, J., Shan, X., Ye, L., Feng, P., Yan, X., & Zhao, Y. (2020). "An adaptive pig face recognition approach using Convolutional Neural Networks". *computers and Electronics in Agriculture* 173, p. 105386. <https://doi.org/10.1016/j.compag.2020.105386>.
- T. Psota, E., Schmidt, T., Mote, B., & C. Pérez, L. (2020). Long-Term Tracking of Group-Housed Livestock Using Keypoint Detection and MAP Estimation for Individual Animal Identification. *Sensors, 20*(13). <https://doi.org/10.3390/s20133670>
- Rabiner, Lawrence R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE, 77*(2), 257–286. <https://doi.org/10.1109/5.18626>
- Sin, Bongkee, & Kim, Jin H. (1995). Nonstationary hidden Markov model. *Signal Processing, 46*(1), 31–46. [https://doi.org/10.1016/0165-1684\(95\)00070-T](https://doi.org/10.1016/0165-1684(95)00070-T). issn: 0165-1684.
- Sun, P., Cao, J., Jiang, Y., Yuan, Z., Bai, S., Kitani, K., & Luo, P. (2022). "DanceTrack: Multi-Object Tracking in Uniform Appearance and Diverse Motion". In: *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20961–20970. <https://doi.org/10.1109/CVPR52688.2022.02032>.
- Valmadre, J., Bertinetto, L., Henriques, J. F., Tao, R., Vedaldi, A., Smeulders, A. W., & Gavves, E. (2018). "Long-Term Tracking in the Wild: A Benchmark". In: *Computer Vision – ECCV 2018: 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part III*. Munich, Germany: Springer-Verlag,

- pp. 692–707. isbn: 978-3-030-01218-2. https://doi.org/10.1007/978-3-030-01219-9_41.
- Voigtlaender, P., Krause, M., Osep, A., Luiten, J., Sekar, B. B. G., Geiger, A., & Leibe, B. (2019). “MOTS: Multi-Object Tracking and Segmentation”. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7934–7943. <https://doi.org/10.1109/CVPR.2019.00813>.
- Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., & Murphy, K. (2018). “Tracking Emerges by Colorizing Videos”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 11217 LNCS, pp. 402–419. issn: 16113349. https://doi.org/10.1007/978-3-030-01261-8_24.
- Voulodimos, Athanasios S., et al. (2010). A complete farm management system based on animal identification using RFID technology. *Computers and electronics in agriculture*, 70(2), 380–388. <https://doi.org/10.1016/j.compag.2009.07.009>
- Wojke, N., Bewley, A., & Paulus, D. (2018). “Simple online and real-time tracking with a deep association metric”. In: *Proceedings - International Conference on Image Processing, ICIP 2017-September*, pp. 3645–3649. issn: 15224880. <https://doi.org/10.1109/ICIP.2017.8296962>.
- Yu, S. I., Meng, D., Zuo, W., & Hauptmann, A. (2016). “The Solution Path Algorithm for Identity-Aware Multi-object Tracking”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3871–3879. <https://doi.org/10.1109/CVPR.2016.420>.
- Yu, S. I., Yang, Y., Li, X., & Hauptmann, A. G. (2016). “Long-Term Identity-Aware Multi-Person Tracking for Surveillance Video Summarization”. In: <https://doi.org/10.48550/arXiv.1604.07468>.
- Zhang, Y., Sun, P., Jiang, Y., Yu, D., Weng, F., Yuan, Z., & Wang, X. (2021). “ByteTrack: Multi-Object Tracking by Associating Every Detection Box”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 13682 LNCS, pp. 1–21. issn: 16113349. <https://doi.org/10.48550/arxiv.2110.06864>.
- Zhang, Yifu, Wang, Chunyu, et al. (2021). Fairmot: On the fairness of detection and re-identification in multiple object tracking. *International journal of computer vision*, 129, 3069–3087. <https://doi.org/10.1007/s11263-021-01513-4>

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.