

Robust Timestamp Correction for Multi-Camera Trajectory Reconstruction in Road Safety Analysis

Nimra Jabeen^{1,2,*}, Cédric Vandelaer², Tim Jacobs², Evelien Polders¹, Wim Ectors¹, Tom Brijs¹

¹UHasselt, Transportation Research Institute (IMOB), Martelarenlaan 42, 3500 Hasselt, Belgium

*Nimra.jabeen@uhasselt.be**

²Cegeka EngineeringSTUDIO, Kempische Steenweg 307, 3500 Hasselt, Belgium

Abstract:

Multi-camera video systems offer rich data for proactive road safety analysis, but their reliability can be severely degraded by timestamp anomalies such as freezes and duplications. Using real-world traffic video from the [MIA](#) (Mobiliteit Innovatief Aanpakken) STRIKE project, we identify these faults and systematically evaluate several correction strategies, including first-timestamp filtering, group-by-timestamp merging, OCR-based verification, and timestamp interpolation. The results demonstrate that OCR-validated timestamp interpolation most effectively restores a strictly increasing and physically plausible time axis, yielding longer and more stable trajectories. Although interpolation assumes short-term linear temporal progression, it proved reliable for brief timestamp freezes (<1 s) observed in the dataset. These findings provide practical guidance for agencies operating multi-camera systems, highlighting timestamp validation as a prerequisite for reliable surrogate safety analysis.

Keywords: Multi-Camera Systems, Timestamp Correction, Trajectory Reconstruction, Surrogate Safety Analysis, Sensor Fusion

Introduction:

Multi-camera video systems are increasingly employed in automated road safety analysis (e.g., traffic conflict detection) because they provide wider coverage and enable 3D localization of road users by fusing detections from different viewpoints (Zhou et al., 2024). This fusion enhances the reconstruction of vehicle and pedestrian trajectories, which is crucial for identifying potential conflicts or near-misses. However, the effectiveness of multi-camera applications critically depends on precise time synchronization between camera streams (Zhou et al., 2024). Even small temporal misalignments can degrade perception and tracking performance (Shahriar et al., 2025). In practice, sensor clocks may drift or freeze, causing timestamp inconsistencies that compromise cross-camera data fusion (Teh, Kempa-Liehr, & Wang, 2020). Invalid timestamps disrupt the temporal ordering of frames, fragment trajectories, and generate false or missing detections, ultimately undermining the reliability of safety analyses (Shahriar et al., 2025).

In our study, we observed uncommon timestamp anomalies in real traffic video collected during the [MIA-STRIKE](#) road safety project (Ectors et al., 2024). Under normal operation, each frame (i.e., detection dictionary) should be uniquely indexed in time, with timestamps t_i strictly increasing with frame index i . However, initial inspection revealed recurring anomalies:

- Multiple consecutive frames shared the same timestamp t .
- In many cases, the metadata x_i was identical, suggesting repeated detections.
- In other cases, metadata varied slightly (e.g., new object classes appearing), raising questions about whether the fault originated in the detection pipeline or the source video.

Such anomalies violate the expected monotonic time progression and can cause stalled trajectories, ambiguous cross-camera mappings, and unreliable surrogate safety indicators such as Time-to-Collision (TTC) and Post-Encroachment Time (PET), which assume a consistent temporal base.

To address this issue, we focus on a two-camera roadside setup in an urban environment and systematically compare strategies for handling timestamp faults. These include filtering, grouping, and an optical character recognition (OCR) procedure to extract visually embedded frame timecodes as an independent verification of recorded timestamps. The OCR-based validation reveals hidden clock freezes and helps localize the source of the anomaly. By integrating these approaches, we develop a processing framework that restores temporal coherence, ensuring that reconstructed trajectories maintain physically consistent timing for dynamic and safety-related analysis.

For agencies deploying roadside multi-camera systems, undetected timestamp faults may silently compromise safety indicators such as TTC and PET, underscoring the need for systematic temporal validation procedures. The following sections detail our proposed processing strategies.

Methodology:

We developed multiple complementary strategies to process the detection streams. Below, we describe each strategy in detail, including its formal definition and how it handles the case of stagnant timestamps. We also discuss the potential drawbacks of each approach, which motivated the subsequent refinements.

First-Timestamp Filtering

The first-timestamp strategy aims to enforce a strictly monotonic timestamp sequence by retaining only the first occurrence of each timestamp value. Formally, let $D = \{(t_i, x_i)\}_{i=1}^N$ be the sequence of detections. We define the filtered set D_{FT} as:

$$D_{FT} = (t_j, x_j) \mid t_j \neq t_{j-1} \text{ or } j = 1$$

This process ensures that each timestamp appears exactly once in the final output. It is particularly effective at removing redundant entries when timestamp stalling causes the same frame to be reprocessed multiple times. However, it risks discarding frames that contain new or distinct metadata, such as newly detected road users in subsequent frames with the same t_j .

Group-by-Timestamp Merge

To mitigate the loss of valid detections caused by the first-timestamp strategy, we developed a group-by-timestamp merging approach. Instead of removing all but the first occurrence of each timestamp, this method clusters all detections with the same timestamp T_k and merges their metadata while filtering duplicates based on object class and bounding box similarity.

Let $\mathcal{T} = T_k$ be the set of unique timestamps. For each T_k , we define a group G_k as the set of all object detections associated with that timestamp:

$$G_k = \{x_j \mid t_j = T_k\}$$

represent the group of detections at timestamp T_k . The merged dataset is initially:

$$D_{Group} = \{(T_k, X_k) \mid X_k = \bigcup_{x_j \in G_k} x_j\}$$

To reduce redundancy, detections with the same object class and similar bounding boxes ($\|B(x_{j_1}) - B(x_{j_2})\| < \epsilon$) are deduplicated. Distinct detections are retained:

$$X_k^* = \text{prune}(G_k), D_{Group}^* = (T_k, X_k^*)$$

This method preserves spatial diversity while collapsing duplicate frames. However, it does not resolve timestamp errors, making it less effective when time metadata is corrupted.

OCR-Based Frame Verification

To determine whether the root cause of the timestamp duplication was in the detection pipeline or in the original video itself, we implemented an OCR-based diagnostic procedure. The goal was to extract visually embedded timestamps from each frame and compare them against the system-recorded timestamps t_j .

Let each frame be denoted by image f_j , and let the OCR-extracted timestamp be:

$$o_j = \text{OCR}(f_j)$$

We computed differences $\Delta_j = o_j - o_{j-1}$ across frames and investigated instances where:

$$t_j = t_{j+1} = \dots = t_{j+n}, \text{ and } o_j = o_{j+1} = \dots = o_{j+n}, \text{ yet } f_j \neq f_{j+1}$$

That is, even though frames are visually different both the system and embedded timestamps remain constant. This verified that the issue was not in the detection logic but in the camera's internal clock or encoding process which embedded static timestamps across multiple frames despite ongoing motion. These segments required correction beyond what filtering or grouping could achieve.

Timestamp Interpolation

Following the identification of timestamp stalling through OCR, we developed an interpolation-based strategy to reconstruct a physically consistent timeline. For each sequence of consecutive frames $\{d_j = (t_j, x_j)\}_{j=a}^b$ where:

$$t_j = t_a \text{ for all } j \in [a, b]$$

and where the next valid timestamp $t_{b+1} > t_a$ is available, we replace the constant timestamps with linearly spaced interpolated values \hat{t}_j defined as:

$$\hat{t}_j = t_a + \frac{j - a + 1}{b - a + 2} \cdot (t_{b+1} - t_a), \text{ for } j = a, \dots, b$$

This yields a strictly increasing timestamp sequence:

$$\hat{t}_a < \hat{t}_{a+1} < \dots < \hat{t}_b$$

evenly distributed between t_a and t_{b+1} . In the analysed dataset, timestamp stagnation intervals typically ranged between 100–400 ms. Interpolation was therefore restricted to short intervals (<1 s) to avoid distortion of acceleration estimates and other high-frequency dynamic measures. If no subsequent timestamp exists (i.e., if $b = N$), we define $t_{b+1} = t_a + \delta$, where δ is a default time increment (e.g., 50 milliseconds).

The corrected dataset becomes:

$$D_{\text{Interp}} = \{(\hat{t}_j, x_j) \text{ for interpolated frames, and } (t_j, x_j) \text{ otherwise}\}$$

This approach maintains all metadata while restoring a consistent temporal order. It is particularly effective in preserving the fidelity of velocity estimates, trajectory continuity, and downstream safety metrics such as TTC and PET.

Results:

The comparative analysis of the four timestamp-reconstruction strategies (Raw, First-Timestamp, Group-by-Timestamp, and Interpolated) revealed substantial differences in trajectory continuity, stability, and completeness, emphasizing the critical role of temporal correction in multi-camera trajectory reconstruction. Among all methods, the Interpolated approach achieved the most reliable and temporally coherent trajectory reconstruction, producing the highest number of valid trajectories (906) and the longest average trajectory length (42.5 frames per object), while the Raw dataset contained only 283 trajectories with an average of 24.6 frames. These values were obtained by computing the number of unique trajectory IDs and the average number of frames per ID within each processed dataset, supplemented by visual inspection of sample trajectories and


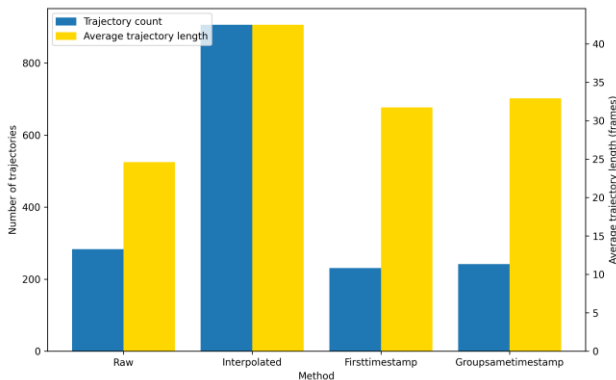


Figure 1: Number of Trajectories VS Average Trajectory Length speed-time plots. This indicates that interpolation not only reconstructs missing timestamps but also prevents premature trajectory fragmentation caused by frozen or duplicated time entries (see Figure 1: Number of Trajectories Vs Average Trajectory Length). In terms of motion stability, the Interpolated trajectories demonstrated the lowest variability in speed, with a coefficient of variation (CV) of approximately 1.43 compared to 1.75 for Raw, 1.96 for First-Timestamp, and 1.91 for Group-by-Timestamp, confirming smoother velocity transitions and fewer discontinuities (see Figure 2: Speed Stability CV vs Acceleration Stability CV).




Figure 3: Speed vs Time for all Methods

Although mean speed values were relatively similar across all datasets (around 0.7 m/s), the Interpolated trajectories exhibited more realistic and continuous temporal evolution, whereas the Raw trajectories showed abrupt spikes and drops in speed, and both First-Timestamp and Group-by-Timestamp methods retained small oscillations linked to repeated or clustered timestamps (see Figures 3: Speed vs. Time for all methods). The time-series visualizations further highlighted that interpolation restored continuous and progressive motion, eliminating resets caused by stalled timestamps, while the other methods either lost information through filtering or failed to correct temporal stagnation. Overall, the Interpolated method provided the optimal balance between trajectory completeness and motion smoothness, ensuring

the optimal balance between trajectory completeness and motion smoothness, ensuring

temporally consistent and physically plausible trajectories suitable for reliable multi-camera road safety analysis.

Discussion:

Temporal recalibration through timestamp interpolation substantially enhances the consistency and reliability of multi-camera trajectory data for road safety analysis by restoring a physically plausible time sequence and improving motion stability, which is essential for surrogate safety measures such as TTC and PET that rely on continuous velocity and acceleration estimates. Compared to first-timestamp filtering and group-based merging, interpolation uniquely preserves both temporal and kinematic continuity, resulting in longer and more stable trajectories. From an operational perspective, these findings underscore the importance of integrating timestamp validation into routine data-quality assessment for multi-camera systems, with OCR-based verification serving as an effective diagnostic tool to detect embedded clock freezes or encoding faults. Interpolation is recommended only when timestamp freeze intervals are short (<1 s), object motion remains continuous, and no abrupt scene transitions occur; for extended freezes, exclusion or re-synchronization strategies may be preferable to avoid distortion of high-frequency dynamic indicators. While the approach assumes approximately linear temporal progression between valid timestamps and may introduce bias if applied indiscriminately, it provides a robust and practical preprocessing solution when applied within clearly defined temporal-gap thresholds relative to camera frame rate and motion dynamics.

Conclusion:

The comparative analysis confirms that timestamp interpolation provides the most robust correction for frozen or duplicated timestamps in multi-camera video systems. It enhances trajectory continuity, smoothness, and stability, enabling more reliable road safety analytics. These results highlight the importance of temporal recalibration as a preprocessing step in any multi-sensor fusion framework, ensuring that synchronized video data can be used for automated traffic conflict detection and safety assessment. Future work should establish formal constraints on the application of interpolation such as acceptable temporal gaps to ensure its validity for high-precision conflict analysis.

Acknowledgements:

This project has received funding from the European Union's Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101119590.

References

- Zhou, X., Dai, Y., Qin, H., Qiu, S., Liu, X., Dai, Y., Li, J., & Yang, T. (2024). Subframe-Level Synchronization in Multi-Camera System Using Time-Calibrated Video. *Sensors*, 24(21), 6975. <https://doi.org/10.3390/s24216975>
- Shahriar, M. H., Barat, M. M. A., Sundar, H., Zhang, N., Ramakrishnan, N., Hou, Y. T., & Lou, W. (2025). Temporal Misalignment Attacks against Multimodal Perception in Autonomous Driving. *arXiv preprint arXiv:2507.09095*.
- Teh, H. Y., Kempa-Liehr, A. W., & Wang, K. I. K. (2020). Sensor data quality: A systematic review. *Journal of Big Data*, 7(1), 11.
- Ectors, W., Tran, T., Brijs, T., Janssens, D., & Wets, G. (2024). *Verkeersveiligheidsonderzoek met privacy-by-design slimme camera's: Studie's in Pelt, Hasselt (Kiewit), Lummen en Maasmechelen* [Unpublished manuscript]. UHasselt, Instituut voor Mobiliteit. https://assets.vlaanderen.be/image/upload/v1726582353/MIA_Rapport_Verkeersveiligheidsonderzoek_met_slimme_cameras_eyvghw.pdf