Does Video Compression Impact Tracking Accuracy?

Takehiro Tanaka, Alon Harell, and Ivan V. Bajić

School of Engineering Science, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada

Abstract—Everyone "knows" that compressing a video will degrade the accuracy of object tracking. Yet, a literature search on this topic reveals that there is very little documented evidence for this presumed fact. Part of the reason is that, until recently, there were no object tracking datasets for uncompressed video, which made studying the effects of compression on tracking accuracy difficult. In this paper, using a recently published dataset that contains tracking annotations for uncompressed videos, we examined the degradation of tracking accuracy due to video compression using rigorous statistical methods. Specifically, we examined the impact of quantization parameter (QP) and motion search range (MSR) on Multiple Object Tracking Accuracy (MOTA). The results show that QP impacts MOTA at the 95% confidence level, while there is insufficient evidence to claim that MSR impacts MOTA. Moreover, regression analysis allows us to derive a quantitative relationship between MOTA and QP for the specific tracker used in the experiments.

Index Terms—Video compression, object tracking, video coding for machines

I. INTRODUCTION

Video compression is ubiquitous in entertainment, monitoring, law enforcement, consumer products, and many other industries. At the same time, object tracking is the cornerstone of many video analysis applications, such as event detection and recognition, visual odometry, navigation, crowd analysis, and so on. Intuitive expectation is that compression degrades tracking accuracy. Yet, there is virtually no evidence in the available literature to support this expectation. The closest available experimental evidence links compression and object detection/classification. The earliest known study on this topic [1] shows that JPEG and Advanced Video Coding (AVC)-based compression impacts pedestrian detection in far infrared (FIR) images, especially at lower bitrates. In [2], the impact of JPEG2000 image compression on classifier accuracy is examined. Dodge and Karam [3] examined the impact of image quality, including JPEG and JPEG2000 compression, on deep neural network (DNN)-based image classification, while [4] examined the impact of image compression on DNN-based object detection. All three studies revealed accuracy degradation at lower bitrates. In [5], the impact of High Efficiency Video Coding (HEVC) [6] on DNN-based based object detection was examined, with similar qualitative conclusions as earlier studies. The lack of annotated tracking data on uncompressed video was highlighted as a major challenge to making further progress in this area.

There are a number of tracking annotations for already-compressed video,[1] so one may wonder why this data cannot be used for studying the impact of further compression on tracking? The reason is that the effect of double compression is fundamentally different from that of single compression. To see this, let X be the original uncompressed signal, and $Y_1 = X + N_1$ be the signal after the first round of compression, where N_1 is the compression (quantization) noise. Another round of compression would produce $Y_2 = Y_1 + N_2 = X + N_1 + N_2$, where N_2 is the noise from the second round of compression. Hence, the net effect on X is the contamination by $N_1 + N_2$. In a simplified case where the two noises are independent, the probability density function (pdf) of their sum is the convolution of their individual pdf's [7]: $f_{N_1+N_2} = f_{N_1} * f_{N_2}$, which will, in general, be different from either f_{N_1} or f_{N_2}. For example, if N_1 and N_2 are uniformly distributed,[2] then $f_{N_1+N_2}$ will be triangular. Hence, its effect on X will not be the same as that of either N_1 or N_2.

In this paper, we make use of the recently released dataset called SFU-HW-Tracks-v1 [8], which contains tracking annotations for a subset of HEVC Common Test Condition (CTC) uncompressed video sequences. Since this dataset is based on uncompressed video, we are able to properly study the effects of compression on object tracking performance. Section II describes data and methods, while Section III presents regression analysis of the influence of compression on tracking performance. Conclusions are presented in Section IV.

II. DATA AND METHODS

A. Dataset

As mentioned earlier, in this study we use the recently released dataset called SFU-HW-Tracks-v1 [8]. The dataset contains object tracking annotations of 13 uncompressed test video sequences. This is an extension of the SFU-HW-Objects-v1 dataset [9], which is currently used for exploration experiments in MPEG Video Coding for Machines (VCM).

B. Methods

Experimental pipeline is shown in Fig. 1. An uncompressed sequence in the YUV420 format is encoded using HEVC HM version 16.20 in the Low Delay format, for various pairs of QP and motion search range (MSR), as shown in Table I. The resulting bitstream is then decoded back into a YUV420

---

[1]https://motchallenge.net

[2]A common model for quanitzation noise.
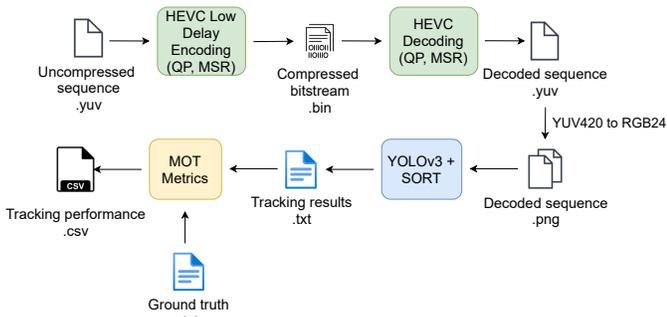
Fig. 1: Experimental pipeline.

TABLE I: QP and MSR values used in the experiment.

HEVC parameter	Values
QP	[18, 22, 26, 30, 34, 38, 42, 46]
MSR	[8, 16, 32, 64]

sequence, and each frame is converted to RGB format using ffmpeg and stored as PNG. Such frames are then input into the tracker (whose details are given below) and the resulting tracking output is compared to the ground truth using various Multiple Object Tracking (MOT) metrics, as described below.

1) Tracker: We use a tracker that follows the tracking-by-detection approach. Specifically, object detection is performed using the Ultralytics implementation[3] of YOLOv3 [11], while the tracking mechanism is based on Simple Online Realtime Tracking (SORT) [12]. According to [12], SORT serves as a baseline method for more sophisticated trackers, as it is constructed from well-known and interpretable components such as Kalman filtering and the Hungarian algorithm. It is worth noting that the current exploration experiments in MPEG VCM use the Joint Detection and Embedding (JDE) tracker [13], whose object detection component is based on the same implementation of YOLOv3 used here. While the tracker used in our experiments is not necessarily state-of-the-art (SOTA) in terms of tracking accuracy, it is based on similar components as SOTA trackers, and provides a solid interpretable baseline for benchmarking tracking performance.

2) Tracking metrics: Tracking accuracy on the uncompressed and compressed videos is measured using the software framework[4] implementing various tracking metrics from [14]–[17]. Among the various metrics available, we chose the Multiple Object Tracking Accuracy (MOTA), defined as

$$\text{MOTA} = 1 - \frac{\sum_t (\text{FN}_t + \text{FP}_t + \text{IDSW}_t)}{\sum_t \text{GT}_t}, \qquad (1)$$

where t is the frame index, FN_t and FP_t are false negatives and false positives in object detection in frame t, IDSW_t represents the number of object ID switches in frame t, and GT_t is the number of ground-truth trajectories in frame t. MOTA is the gold-standard in multiple object tracking accuracy and correlates well with human visual assessment [18].

[3]https://github.com/ultralytics/yolov3, pre-trained on COCO [10].
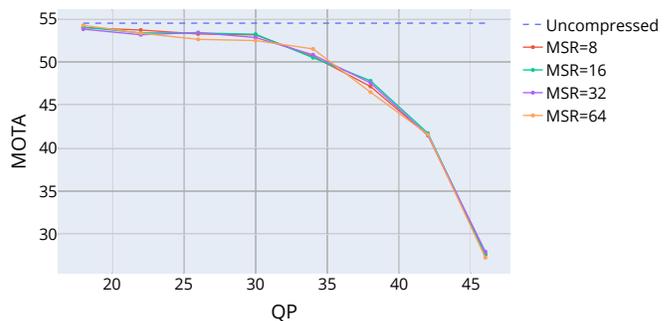[4]https://github.com/cheind/py-motmetrics

Fig. 2: Average MOTA scores across all video sequences for all object classes, for uncompressed sequences (dashed line) and different (QP, MSR) pairs.

III. ANALYSIS

The goal of our analysis is to evaluate and quantify the nature of the relationship between compression and tracking. More specifically, we are interested in the effects of the values of MSR and QP on MOTA, as explained before. As described in section II, for each pair of values of QP and MSR we have the MOTA score for multiple video sequences. We being by evaluating the average effect of QP and MSR on MOTA as can be seen in Fig. 2. The average MOTA score is computed across all video sequences and all object classes[5] at different QP and MSR values. As shown in the figure, MOTA scores are highest, on average, on uncompressed sequences, and degrade as the sequence is compressed, i.e., as QP increases. This seems to agree with intuition. On the other hand, MOTA scores do not seem to depend much on MSR, at least for the tracker employed in our experiments.[6] In the remainder of this section we employ regression analysis [20] to statistically verify observations from Fig. 2. The same type of analysis can be applied to other tracking metrics, other parameters (besides QP and MSR) and other trackers.

A. Regression model

Because the relationship between MOTA and QP shows clear non-linearity in Fig. 2, we first transform QP so that MOTA and the transformed predictor variable QP' form an approximately linear relationship. To find the appropriate transformation for QP, we experimented with various forms and eventually decided on the following one:

$$\text{QP}' = \frac{1}{\frac{\text{QP}}{52} - 1}, \qquad (2)$$

where 52 was chosen to allow the full range of QP values $[0, 51]$ [6], without the danger of a zero denominator.

The scatter plots of MOTA vs. QP and MOTA vs. QP' on BasketballPass are shown in Fig. 3 where we can see the transformation indeed makes the relationship between QP' and MOTA approximately linear. On the other hand, from Fig. 2

[5]Here, "all" object classes means all classes that exist in the ground truth for each sequence.
[6]There are trackers, such as MV-YOLO [19], that directly use motion vectors from the video bitstream, and these may be more sensitive to MSR.

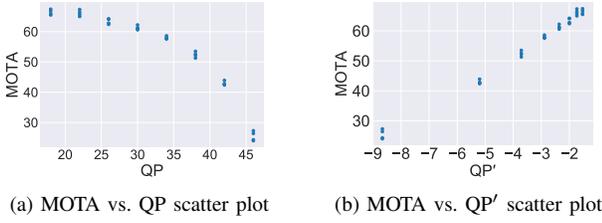(a) MOTA vs. QP scatter plot   (b) MOTA vs. QP' scatter plot

Fig. 3: Scatter plots of MOTA vs. QP before and after QP transformation on the BasketbalPass sequence; the four MOTA values at each QP are produced by different MSR values.

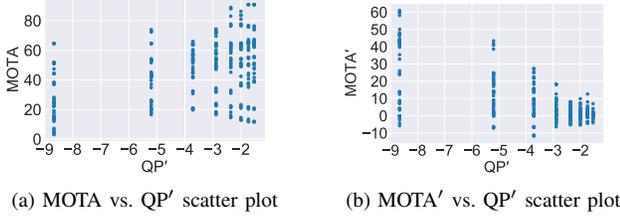(a) MOTA vs. QP' scatter plot   (b) MOTA' vs. QP' scatter plot

Fig. 4: Scatter plots of MOTA vs. QP' before and after MOTA transformation (4) across all video sequences.

(as well as Fig. 3), MOTA is approximately constant with MSR, and thus already approximately linear, and appropriate for linear regression.

Because the MOTA scores, even for uncompressed videos, vary greatly between the sequences (as seen in Fig. 4a), we need to isolate the effect of compression on tracking before performing regression. We do this by modeling the effects of compression using a subtractive term on MOTA:

$$\text{MOTA}_i(\mathbf{Q}) = \text{MOTA}_i^0 - \text{MOTA}_i'(\mathbf{Q}), \tag{3}$$

where i is the sequence index, MOTA^0 is the MOTA of an uncompressed sequence, and $\text{MOTA}'(\mathbf{Q})$ is the degradation in MOTA due to compression under parameters $\mathbf{Q}$ (in our case, $\mathbf{Q} = (\text{QP}', \text{MSR})$). Re-arranging (3), we obtain:

$$\text{MOTA}_i'(\mathbf{Q}) = \text{MOTA}_i^0 - \text{MOTA}_i(\mathbf{Q}), \tag{4}$$

Fig. 4 shows the scatter plots of MOTA and MOTA' at each QP', i.e., before and after MOTA transformation. We will demonstrate below that the deviations of MOTA' from the regression model are approximately Gaussian. This will validate the use of statistical tests such as t-test, which rely on the Gaussianity (normality) of regression errors [20].

With QP and MOTA transformed, we can now represent a multiple linear regression model as:

$$\text{MOTA}_{i,j}' = \beta_0 + \beta_1 \cdot \text{QP}_j' + \beta_2 \cdot \text{MSR}_j + \beta_3 \cdot \text{QP}_j' \cdot \text{MSR}_j + \epsilon_{i,j} \tag{5}$$

where j is the index of the compression parameter combination, β_0, β_1, β_2, and β_3 are regression parameters, and $\epsilon_{i,j}$ is assumed to be an independent normally distributed (Gaussian) error with mean 0 and unknown variance σ_j^2 (note that the variance only changes with j, not i). The normality of errors will be verified below.

TABLE II: Regression model parameter estimates across all object classes and all video sequences.

Parameter	Estimated value	p-value
β_0	-4.32	$< 10^{-6}$
β_1	-2.97	$< 10^{-9}$
β_2	$< 10^{-2}$	0.83
β_3	$< 10^{-2}$	0.77

Allowing σ_j^2 to vary with j, we apply a weighted least squares method [20] to estimate β's, with the weight $w(\mathbf{Q}_j)$ given by

$$w(\mathbf{Q}_j) = \frac{1}{s^2(\mathbf{Q}_j)}, \tag{6}$$

where $s^2(\mathbf{Q}_j)$ is the sample variance of $\text{MOTA}_{i,j}'$ at a given set of parameters $\mathbf{Q}_j$ (taken over the different video sequences).

B. Results

Utilizing Statsmodels [21], we obtain regression results shown in Table II. As seen in the table, estimated regression parameters β_2 and β_3 are very small, suggesting that MOTA' is not influenced by MSR or the product of MSR and QP'. To test the significance of each parameter, we performed the following individual hypothesis test for each β_c as a t-test:

$$\begin{aligned} H_0 &: \beta_c = 0 \\ H_1 &: \beta_c \neq 0 \end{aligned} \tag{7}$$

where $c \in \{0, 1, 2, 3\}$. The null hypothesis H_0 is $\beta_c = 0$, meaning that the corresponding predictor variable does not impact the response variable MOTA'. The alternative hypothesis H_1 is $\beta_c \neq 0$, meaning that the corresponding predictor variable impacts MOTA'. The p-value for each hypothesis test is given in Table II.

The results show that H_0 can be rejected for β_0 and β_1 (p-values are very small), meaning that these values significantly differ from 0. In other words, QP' has a significant impact on MOTA'. On the other hand, p-values for β_2 and β_3 are fairly large. Since they are larger than 0.05, this means there is insufficient evidence to reject H_0 for these parameters at the 95% confidence level. We presented individual hypothesis tests for simplicity, but a joint hypothesis test for (β_2, β_3) gives similar results. Hence, from the available data, we cannot conclude that MSR impacts MOTA'.

To validate our assumption regarding the normality of regression errors, we plot the sorted studentized residuals in Fig. 5, following [20]. The studentized residual $r_{i,j}$ can be represented as

$$r_{i,j} = \frac{\text{MOTA}_{i,j}' - \widehat{\text{MOTA}'}(\mathbf{Q}_j)}{s(\mathbf{Q}_j)} \tag{8}$$

where $\widehat{\text{MOTA}'}(\mathbf{Q}_j)$ is MOTA' predicted by the regression model at a given $\mathbf{Q}_j$ and $s(\mathbf{Q}_j)$ is the sample standard deviation of MOTA' scores at a given $\mathbf{Q}_j$. The horizontal axis represents sorted expected values of studentized residuals under normality, which can be computed as [20, p. 111]

$$\mathbb{E}[r_k] = z\left(\frac{k - 0.375}{n + 0.25}\right), \tag{9}$$

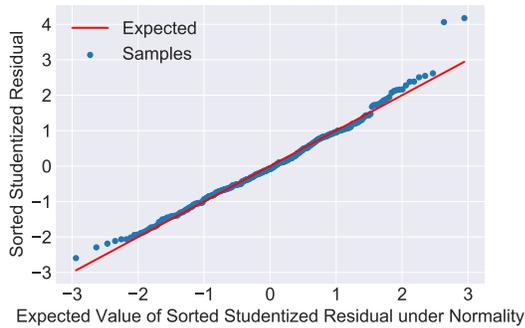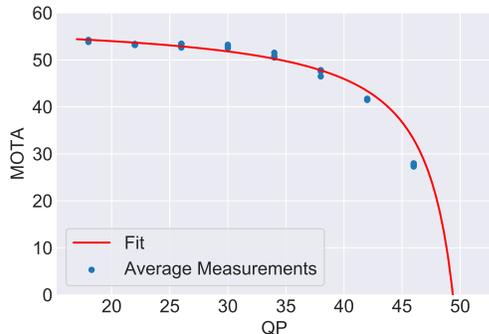Fig. 5: Approximate normality of regression errors.

Fig. 6: MOTA scores averaged over all sequences vs. QP.

where z is a percentile function [20], k is the index of expected values of studentized residuals under normality, from the smallest value at $k = 1$ to the largest value at $k = n$, and n is the total number of measurements.[7] As shown in Fig. 5, regression residuals (blue) follow closely the red line, which shows the expected residual values under normality. In other words, the residuals are approximately normal. The coefficient of determination is $R^2 = 0.985$, indicating close fit to the red line. This validates the regression analysis presented above.

Using the values of β_0 and β_1 from Table II in (5), and then substituting transformations from (2), (4), we obtain the following relationship between the average MOTA across all video sequences (denoted $\overline{\text{MOTA}}$) and QP:

$$\overline{\text{MOTA}} = 58.81 + 2.97 \cdot \frac{1}{\frac{\text{QP}}{52} - 1}. \tag{10}$$

Fig. 6 shows $\overline{\text{MOTA}}$ vs. QP from (10), as well as the average measured MOTA across all sequences at different QP values. As seen in the figure, the agreement with measurements is fairly good. The values obtained from (10) for QP > 46 should be taken with a grain of salt since there are no data points for QP > 46, but at lower QP the agreement is good.

C. When does MOTA "drop"?

Using the available data, we can also answer the following question: *At which QP value does MOTA start to significantly deviate from uncompressed performance?* To answer this question, we look at the difference between uncompressed MOTA

---

[7] $n = 384$ in our case: 8 QP values $\times$ 4 MSR values $\times$ 12 sequences.

TABLE III: One-tailed t-test results for (11).

QP Value	p-value
18	0.03
22	$< 10^{-2}$
26	$< 10^{-3}$
30	$< 10^{-3}$
34	$< 10^{-4}$
38	$< 10^{-5}$
42	$< 10^{-8}$
46	$< 10^{-11}$

and compressed MOTA, in other words, MOTA$'$ defined in (4). We want to determine when the average MOTA$'$ becomes significantly larger than 0, so we formulate the following hypothesis test at each QP (or, equivalently, each QP$'$):

$$\begin{aligned} H_0 &: \mu_{\text{MOTA}'} \leq 0 \\ H_1 &: \mu_{\text{MOTA}'} > 0 \end{aligned} \tag{11}$$

where $\mu_{\text{MOTA}'}$ is the average MOTA$'$ at a given QP across all sequences. Note that we formulate the null hypothesis H_0 as $\mu_{\text{MOTA}'} \leq 0$ instead of $\mu_{\text{MOTA}'} = 0$, because a one-tailed t-test is more powerful than a two-tailed t-test when testing for the difference in one direction.

One-tailed t-test results are shown in Table III. As seen in the table, all p-values are less than 0.05, so we can reject H_0 for all QP at the 95% confidence level. This means that even at the lowest QP value of 18, MOTA has significantly departed (at the 95% confidence level) from its uncompressed value. If we want to be more strict and ask for 99% confidence, then the lowest QP for which we can reject H_0 is 22, because p-value for QP $= 22$ is less than 10^{-2}. Nevertheless, both QP $= 18$ and QP $= 22$ are synonymous with fairly high quality compression, so we conclude that even high-quality compression may impact tracking accuracy.

IV. CONCLUSIONS

In this paper, we analyzed the impact of video compression on object tracking accuracy using rigorous statistics. We made use of the recently released dataset of tracking annotations on uncompressed video sequences, and examined the behavior of Multiple Object Tracking Accuracy (MOTA) on HEVC-compressed sequences with different QP values and motion search range (MSR). Tracking was performed using a combination of YOLOv3 object detection and SORT tracking. Our regression analysis indicates that QP has a significant impact on MOTA, while there is insufficient evidence for MSR impact on MOTA. As a result of regression analysis, we derived a relationship between MOTA and QP for our chosen tracker, which can serve as a guideline for estimating the impact of compression on tracking. Moreover, we showed that even low to moderate compression can have a statistically significant impact on tracking accuracy. While numerical values presented here are specific to the chosen tracker, similar analysis can be performed on other trackers as well.

REFERENCES

[1] T. Hase, W. Hintermaier, A. Frey, T. Strobel, U. Baumgarten, and E. Steinbach, "Influence of image/video compression on night vision based pedestrian detection in an automotive application," in *Proc. IEEE Vehicular Technology Conference (VTC Spring)*, May 2011, pp. 1–5.

[2] R. Wagner, M. Thom, R. Schweiger, M. Gabb, A. Röhlig, and A. Rothermel, "Influence of image compression on cascade classifier components," in *Proc. IEEE 10th Jubilee International Symposium on Intelligent Systems and Informatics*, 2012, pp. 367–371.

[3] S. Dodge and L. Karam, "Understanding how image quality affects deep neural networks," in *Proc. QoMEX*, 2016.

[4] H. Hsiang, K.-C. Chen, P.-Y. Li, and Y.-Y. Chen, "Analysis of the effect of automotive ethernet camera image quality on object detection models," in *Proc. Int. Conf. Artificial Intelligence in Information and Communication (ICAIIC)*, 2020, pp. 21–26.

[5] K. Kajak, "Impact of video compression on the performance of object detection algorithms in automotive applications," M.S. thesis, KTH Royal Institute of Technology, Sweden, 2020.

[6] G. J. Sullivan, J. Ohm, W. Han, and T. Wiegand, "Overview of the high efficiency video coding (HEVC) standard," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 22, no. 12, pp. 1649–1668, Dec. 2012.

[7] H. Stark and J. W. Woods, *Probability, Statistics, and Random Processes for Engineers*, Pearson Prentice Hall, 4th edition, 2012.

[8] T. Tanaka, H. Choi, and I. V. Bajić, "SFU-HW-Tracks-v1: Object tracking dataset on raw video sequences," *arXiv preprint arXiv:2112.14934*, Dec. 2021.

[9] H. Choi, E. Hosseini, S. R. Alvar, R. A. Cohen, and I. V. Bajić, "A dataset of labelled objects on raw video sequences," *Data in Brief*, vol. 34, Feb. 2021.

[10] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár, "Microsoft COCO: Common Objects in Context," *arXiv preprint arXiv:1405.0312*, 2015.

[11] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," *arXiv preprint arXiv:1804.02767*, 2018.

[12] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, "Simple online and realtime tracking," in *Proc. IEEE ICIP*, Sep. 2016, pp. 3464–3468.

[13] Z. Wang, L. Zheng, Y. Liu, and S. Wang, "Towards real-time multi-object tracking," in *Proc. ECCV*, 2020.

[14] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: The CLEAR MOT metrics," *EURASIP J. Image and Video Processing*, May 2008.

[15] A. Milan, L. Leal-Taixé, I. D. Reid, S. Roth, and K. Schindler, "MOT16: A benchmark for multi-object tracking," *arXiv preprint arXiv:1603.00831*, 2016.

[16] Y. Li, C. Huang, and R. Nevatia, "Learning to associate: Hybridboosted multi-target tracker for crowded scene," in *Proc. IEEE CVPR*, 2009, pp. 2953–2960.

[17] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, "Performance measures and a data set for multi-target, multi-camera tracking," in *Proc. ECCV 2016 Workshop on Benchmarking Multi-Target Tracking*, 2016, pp. 17–35.

[18] L. Leal-Taixé, A. Milan, K. Schindler, D. Cremers, I. D. Reid, and S. Roth, "Tracking the trackers: An analysis of the state of the art in multiple object tracking," *arXiv preprint arXiv:1704.02781*, 2017.

[19] S. R. Alvar and I. V. Bajić, "MV-YOLO: motion vector-aided tracking by semantic object detection," in *Proc. IEEE MMSP*, 2018, pp. 1–5.

[20] M. Kutner, *Applied Linear Statistical Models*, McGraw-Hill/Irwin, 5th edition, 2005.

[21] S. Seabold and J. Perktold, "Statsmodels: Econometric and statistical modeling with Python," in *Proc. 9th Python in Science Conference (SCIPY)*, 2010.