# Efficient Asymmetric Co-Tracking using Uncertainty Sampling

Kourosh Meshgi, Maryam Sadat Mirzaei, Shigeyuki Oba, Shin Ishii Graduate School of Informatics Kyoto University Sakyo-ward, Yoshida-honmachi, Kyoto 606–8501 Email: meshgi-k@sys.i.kyoto-u.ac.jp

Abstract—Adaptive tracking-by-detection approaches are popular for tracking arbitrary objects. They treat the tracking problem as a classification task and use online learning techniques to update the object model. However, these approaches are heavily invested in the efficiency and effectiveness of their detectors. Evaluating a massive number of samples for each frame (e.g., obtained by a sliding window) forces the detector to trade the accuracy in favor of speed. Furthermore, misclassification of borderline samples in the detector introduce accumulating errors in tracking. In this study, we propose a co-tracking based on the efficient cooperation of two detectors: a rapid adaptive exemplar-based detector and another more sophisticated but slower detector with a long-term memory. The sampling labeling and co-learning of the detectors are conducted by an uncertainty sampling unit, which improves the speed and accuracy of the system. We also introduce a budgeting mechanism which prevents the unbounded growth in the number of examples in the first detector to maintain its rapid response. Experiments demonstrate the efficiency and effectiveness of the proposed tracker against its baselines and its superior performance against state-of-the-art trackers on various benchmark videos.

#### I. INTRODUCTION

Nowadays, visual tracking is an inseparable component for high-level visual tasks such as human-computer interface, human behavior analysis, smart appliances, virtual/augmented reality and surveillance. When applied to video sequences in real-life situations, trackers should cope with challenging appearance changes due to illumination variations, motion blur, non-rigid deformations, rotations, mobile imaging platforms and occlusions [1]. While some successful generative trackers [2]-[5] models the target object, they ignore the information hidden in the background. On the other hand discriminative trackers [6]–[12] pose the tracking problem as a classification task. In these models instead of trying to build a complex model of the object, the algorithms seek a decision boundary that best separates the target and background. This re-formulation undermine the inherent issues of generative models like background clutter and model over-simplification [13].

There are many tracking-by-detection visual trackers, which heavily rely on their detector to handle different tracking challenges [9], such as rotations and scale changes. Such schemes treat tracking as a binary classification problem, which separates the object from its local background using a classifier, in which a discriminative classifier is trained with the samples obtained from the tracking sequence, and their performances are affected by their sampling policy. Most trackers only utilize one positive sample, i.e., the tracking result in the current frame [14]. If the tracked location is not accurate, the classifier will be up-



Figure 1: Labeling dilemma in tracking-by-detection. The classifier of the tracker can label the target and the background only when it has high confidence. Labeling the uncertain patches may result in label noise.

dated with the contaminated appearance of the target, leading to a drift over time. To alleviate this problem, multiple samples in the proximity of the estimated target can be used to train the tracker [7], [9]. However, such algorithms are heavily invested in the efficiency and effectiveness of their detectors. Evaluating a massive number of samples for each input frame forces the detector to trade the accuracy in favor of speed to meet the real-time processing requirements. While some trackers aim to enhance the detectors' speed while preserving their accuracy using statistical properties of images (e.g., [12]), generally achieving an adjustable balance between speed and accuracy is desired. Furthermore, dealing with rotations and scale changes challenges such mechanisms. Additionally, misclassification of the borderline input samples in the detector (Figure 1) may introduce accumulating errors in the tracker, degrading its performance significantly [7]. Furthermore, the growth of sample repository in online learning schemes degrades the speed. If not handled properly, the tracker cannot perform long-term tracking [9].

In this study, we propose an efficient co-tracking framework in which an active learning unit orchestrate the information exchange. It consists of a rapid detector with short-term memory, and an uncertainty sampling switcher that query the label of the most uncertain samples of the first detector from an accurate detector with a long-term memory (called the "oracle"). An importance sampling scheme combines the results of the two trackers and handles the scale variation of the target. An exemplar-based detector is employed as the rapid detector and we introduce a budgeting mechanism to prevent the unbounded growth in the number of examples in this detector to maintain its rapid response. In summary we (i) employed active learning in co-tracking framework that leads to increasing the speed and generalization power of the tracker, (ii) actively control the memory of tracker by balancing between short-



Figure 2: The flow of the proposed tracker, UST. After sampling from the region-of-interest determined by optical flow, the samples are mapped onto a low dimension space, and collaboratively labeled by the two classifiers. This collaboration is organized by uncertainty sampling unit. The labeled are used to localize the target through a weighted averaging, and the two classifiers provide training data for each other.

and long-term memories and (*iii*) introduced an intuitive budgeting method for the nearest neighbor classifier.

The difference between the proposed co-tracking framework and that of Tang et al. [13] is four-fold: (*i*) the classifiers do not exchange all the data they have problems in labeling, instead, the most informative samples are selected by uncertainty sampling, and exchanged. (*ii*) the update rate of classifiers is different to realize a short and long-term memory mixture, (*iii*) the samples that are labeled for the target localization can be re-used for training and the need for an extra round of sampling and labeling is revoked, (*iv*) since in the proposed asymmetric co-tracking, one of the classifiers scaffolds the other one instead of participating in every labeling process, a more sophisticated classifier with higher computational complexity can be used.

# II. RELATED WORK

Many discriminative models have been adopted in object tracking, where a binary classifier is learned online to separate the target from the background. Numerous supervised or semi-supervised classifiers have been employed for object tracking, such as SVM [15], structured output SVM [9], boosting [14], semi-boosting [16], and online multi-instance boosting [7]. They follow different approaches in tackling foreground-background separation like incorporating a trained SVM into an optical flow tracker [15], using an ensemble of online learned weak classifiers to decides whether a pixel belongs to the target region or background [17], or utilizing online boosting to select discriminative features for separation of target and background [14]. Multiple instance learning (MIL) tracker put all of the ambiguous positive and negative samples into bags to learn a discriminative model [7]. In another stream of studies, the most discriminative feature combination in learned online to build a confidence map for foreground detection [18]. Combining multiple supervised and semisupervised classifiers [6], and governing a learning method using positive and negative constraints [11] are some other of the most successful discriminative approaches for tracking. Such methods are specifically designed to resolve the label noise problem, in which the classifier get confused by even the smallest mistakes in the labeling process. Since the classifier is using a self-learning loop, such mistakes

can accumulate over time and cause the tracker to drift. One of the solutions to this problem is co-tracking [13], in which the self-learning loop is broken and labeling is done collaboratively.

Without model update schemes, trackers accumulate error during run-time (drift) and typically fail if the object disappears temporarily. To address this issue, some online appearance update models have been proposed, e.g., incremental subspace update scheme and adaptive sparse representation update [19]. In the case of tracking-by-detection approaches, this essentially means that the classifier should be re-trained with the relevant samples. In such algorithms, after selecting the samples and labeling them, they are used to update the classifier. The classifier aims for labeling data, while the tracker attempts to localize the target and these two objectives are sometimes contradicting (e.g., in the presence of target-like distractors). Increasing the accuracy of the detector and using unlabeled samples is a typical approach to this problem (e.g., [7], [16]), while trackers such as STRUCK [9] couple these two objectives in a joint learning framework. Another instance was presented in [11] where the recent samples are added to the classifier only if their classifier label is different from the label of constrained classifiers that monitors the performance of the tracker. Combining short and long-term memory to deal with rapid changing targets, occlusions, and environmental change in tracking is another research avenue for model updating [20], however, the proposed schemes are hardly integrated into general trackers.

Active learning techniques build upon such discriminative models and try to improve their convergence speed and generalization power. For instance, Fisher information criterion evaluates the uncertainty of classification model in the MIL tracker [7] to perform active feature selection [21]. Reducing the number of necessary labeled samples [22], unified sample learning and feature selection procedure [23] and reducing the sampling bias by controlling the variance [24] are some of the improvements that active learning provides for the discriminative trackers. Active learner selects the samples that don't know how to label. Uncertainty sampling [25], as one of the most popular forms of active learning, try to query the sample that minimizes a utility function from the oracle. The utility function can be classification confidence [26], margin [27], and entropy [28]. Uncertainty sampling in this sense, optimize the query selection with respect to the utility function. However, this approach is a form of Gibbs sampling and requires probabilistic learning models, it should be treated differently for non-probabilistic classifiers. In this regard, decision trees [26] and nearest neighbors [29] are used with uncertainty sampling with the class label obtained by voting, and for SVMs [30] the proximity to the decision boundary is considered as the utility function.

# III. PROPOSED METHOD

# A. Architecture of the System

The proposed tracker (Figure 2) consists of two classifiers  $\theta_t^{(1)}$  and  $\theta_t^{(2)}$ , who exchange their information using an uncertainty sampling scheme. The samples are obtained from the region-of-interest (ROI) defined by optical flow with Gaussian probability. Then, the samples are labeled by a collaborative effort of the classifiers. Later, these samples and their labels are employed to update the classifiers. The first classifier ( $\theta_t^{(1)}$ ) is a short-memory highly adaptive exemplar-based classifier that is updated (with respect to a memory budget) with the most informative samples in each tracking episode. The second classifier, the oracle ( $\theta_t^{(2)}$ ), is a long-term memory tracker that is updated with all the samples at fixed intervals to grant the robustness against occlusions and temporal target changes to the tracker.

**Sampling:** In each frame  $F_t, t \in \{1, \ldots, T\}$ , a set of nrandom samples  $\mathbf{p}_t^j$  is generated. These samples are selected from a region-of-interest  $\mathcal{R}_t$  determined by optical flow [31]. To handle still objects, the last known target area is added to the ROI New samples are selected from ROI, based on a Gaussian distribution centered on the last target position  $\mathcal{N}(\mathbf{p}_{t-1}, \Sigma_{search})$ . In each frame, an additional n'samples are selected uniformly from distant locations of the frame (distance larger than  $3 \times \Sigma_{search}$ ), to sample global background, and are automatically labeled as the background. This sampling scheme empowers the foregroundbackground separability in short term by exploiting the locality of the target [9]. Furthermore, it enables the tracker to handle in-scene distractors (e.g., the moving non-target objects in the ROI) and potential occluders using the global sampling for negative exemplars.

**Labeling:** In the proposed asymmetric co-tracking framework, one classifier attempts to label the sample, and it queries the label from the other classifier if a certain condition is met. This is in contrast with using a linear combination of both classifiers based on their general classification accuracy, as adopted in [13]. The proposed tracker can decide for each sample based on the classifier confidence, i.e., for sample  $\mathbf{p}_t^i$  we define a score  $s_t^j$ 

$$s_t^j = h\left(\mathbf{x}_t^{\mathbf{p}_t^j} | \theta_t^{(1)}\right) \tag{1}$$

that reflects the classification score for the image patch  $\mathbf{x}_t^{\mathbf{p}_t^j}$ , with values closer to +1 as possible targets and values closer to -1 as background. Based on uncertainty sampling (elaborated in section III-D), the samples for which the classification score is more uncertain (i.e.,  $s_t^j \to 0$ ), contains more information for the classifier if they are labeled by the other classifier (i.e., the oracle). Therefore, the scores of all samples are sorted, and m samples with the closest values to 0 are selected to be queried from  $\theta_t^{(2)}$ . To handle the

Algorithm 1: Uncertainty Sampling co-Tracker (UST) **input** : Target position in last frame  $\mathbf{p}_{t-1}$ **output**: Target position in current frame  $\mathbf{p}_t$ ROI  $\mathcal{R}_t \leftarrow OpticalFlow(F_{t-1}, F_t) \cup Area(\mathbf{p}_{t-1})$ for  $j \leftarrow 1$  to n do Make a sample  $\mathbf{p}_t^j \sim \mathcal{N}(\mathbf{p}_{t-1}, \Sigma_{search})$ if  $\mathbf{p}_t^j \in \mathcal{R}_t$  then Calculate  $s_t^j \leftarrow h(\mathbf{x}_t^{\mathbf{p}_t^j}|\theta_t^{(1)})$  (eq(1)) Distinguish uncertain samples  $\mathcal{U}_t$  (eq(2)) if  $\mathbf{p}_t^j \in \mathcal{U}_t$  then  $\theta_t^{(1)}$  is uncertain  $Query \ \theta_t^{(2)} \colon \ell_t^j \leftarrow Sign\left(h(\mathbf{x}_t^{\mathbf{p}_t^j}|\theta_t^{(2)})\right)$ else else Sample not in ROI  $\ell_t^j \leftarrow -1 \ (local \ background)$  $\mathcal{D}_t \leftarrow \mathcal{D}_t \cup \langle \mathbf{x}_t^{\mathbf{p}_t^j}, \ell_t^j \rangle$ Create n' global background samples and add to  $\mathcal{D}_t$ Update  $\theta_t^{(2)}$  with  $\mathcal{D}_{t-\Delta,...,t}$  every  $\Delta$  frames (eq(6)) if  $\sum_{j=1}^n \mathbb{1}(\ell_t^j > 0) > \tau_p$  and  $\sum_{j=1}^n \pi_t^j > \tau_a$  then | Approximate target state  $\hat{\mathbf{p}}_t$  (eq(4)) Update  $\theta_t^{(1)}$  with  $\mathcal{U}_t$  (eq(5)) else target occluded  $\hat{\mathbf{p}}_t \leftarrow \mathbf{p}_{t-1}$ 

situations for which the number of highly uncertain samples are more then m, a range of scores are defined by lower and higher thresholds ( $\tau_l$  and  $\tau_u$ ) and all the samples in this range are considered highly uncertain.

$$\mathcal{U}_t = \{ \mathbf{p}_t^i | \tau_l < s_t^i < \tau_u \text{ or } \left| \{ \exists j \neq i | s_t^j \le s_t^i \} \right| < m \}$$
(2)

in which  $\mathcal{U}_t$  is the list of uncertain samples. The label of the samples  $\ell_t^j \in \mathcal{L}_t$  are then determined by

$$\ell_t^j = \begin{cases} sign\left(h\left(\mathbf{x}_t^{\mathbf{p}_t^j} | \theta_t^{(1)}\right)\right) &, \mathbf{p}_t^j \in \mathcal{U}_t \\ sign\left(h\left(\mathbf{x}_t^{\mathbf{p}_t^j} | \theta_t^{(2)}\right)\right) &, \mathbf{p}_t^j \notin \mathcal{U}_t \end{cases}$$
(3)

and all image patches  $\mathbf{x}_t^{\mathbf{p}_t^j}$  and labels  $\ell_t^j$  are stored in  $\mathcal{D}_t$ .

**Localizing:** To determine the state of the target  $\hat{\mathbf{p}}_t$ , we follow the importance sampling mechanism originally employed by particle filter trackers,

$$\hat{\mathbf{p}}_{t} = \frac{\sum_{j=1}^{n} \pi_{t}^{j} \mathbf{p}_{t}^{j}}{\sum_{j=1}^{1} \pi_{t}^{j}}.$$
(4)

where  $\pi_t^j = s_t^j \mathbb{1}(\ell_t^j > 0)$  and  $\mathbb{1}(.)$  is the indicator function, 1 if true, zero otherwise. This mechanism approximate the state of the target, based on the effect of positive samples, in which samples with higher scores gravitates the final results more toward themselves. Upon the events such as massive occlusion or target loss, this sampling mechanism degenerates [19]. In such cases, the number of positive samples and their corresponding weights shrinks significantly, and the importance sampling is prone to outliers, distractors and occluded patches. To address this issue, two thresholds ( $\tau_p$ and  $\tau_a$ ) are set on the number and average scores of positive samples. If either of thresholds are not exceeded, the target is deemed occluded to avoid tracker degeneracy.



Figure 3: An ever growing example-space of a KNN and the budgeting mechanism to discard the most futile samples. Without budgeting, some samples are added/kept by the classifier that has no effect on the discrimination border.

**Updating:** In the proposed tracker, the first classifier,  $\theta_t^{(1)}$  is updated using the samples that it queried from  $\theta_t^{(2)}$ , i.e, with the those it was uncertain about and had oracle label them  $(\mathcal{U}_t)$ . This uncertainty is either due to the model shortcomings of the classifier (e.g., simple observation model) or because of intrinsic ambiguity of the sample. Using a sophisticated classifier as the oracle alleviate these issues, and by providing the label back to the first classifier, it scaffolds it for better classification in similar circumstances, potentially improving the speed of future sample evaluation as well as the generalization.

$$\theta_{t+1}^{(1)} = \psi(\theta_t^{(1)}, \mathcal{U}_t, \mathcal{L}_t) \tag{5}$$

in which  $\psi(.)$  is the update function (r.t. III-C). On the other hand, to realize a dual-memory scheme to handle temporal target changes and occlusions, the oracle is equipped with a non-volatile memory and updated less frequently (every  $\Delta$  frames) with all the data sampled during this period,

$$\theta_{t+1}^{(2)} = \begin{cases} u(\theta_t^{(2)}, \mathcal{D}_{t-\Delta, \dots, t}) & \text{, if } t \neq k\Delta \\ \theta_t^{(2)} & \text{, if } t = k\Delta \end{cases}$$
(6)

in which u(.) is a classifier re-train function. In summary, the first detector rapidly adapts to the target in order to estimate the target location considering its recent changes. The second detector, however, cares for best object detection performance, has a long-term memory and is robust to noise and temporal occlusions. The co-training of these two detectors balances the desired level of speed and accuracy for the tracker. This differs from the dual memory scheme of MUSTer [20], in which the short-term memory is based on the information obtained from the current frame, and the long-term memory that has exponential forgetting curve and get updated only when no occlusion is detected.

# B. Realization

In this study, the short-term memory classifier is implemented using a k-nearest neighbor classifier in which all the samples have a short lifetime to realize the budgeting mechanism. The histogram of colors and bag of visual words (of SIFT) forms the feature vector of every patch  $\mathbf{x}_t^{\mathbf{p}_t^j}$ , and its dimensionality is reduced to 20 using PCA. The KDtree-based KNN classifier, its budgeted memory, the lazy update behavior of KNNs, and the reduced dimensionality of the feature space render the KNN suitable for real-time tracking. The oracle in this study is a part-based detector [32]. The features, part-base detector dictionary, and the parameters (k of KNN, thresholds  $\tau_l, \tau_u, \tau_a, \tau_p$ , number of samples n, n', search radius  $\Sigma_{search}$ , and update latency  $\Delta$ ) are trained/tuned via a cross-validation approach.

# C. Budgeting

Online learning of discriminative trackers has its own challenges. The sample size of most of the adaptive clas-



Figure 4: KNN local minima and the budgeting method that help solving it. The budgeting prevent the distractor to reinforce itself by influencing the neighbors and spread throughout the sample space.

sifiers is constantly growing, making them slower by the time. For a KNN classifier, even with a robust KD-tree architecture, the computation cost rapidly increase over time. This is similar to the curse of kernelization, in which the number of support vectors increases with the amount of training data in kernel classifiers. To allow for real-time operation, there is a need to control the number of support vectors. Recently, approaches have been proposed for online learning of classification SVMs on a fixed budget, meaning that the number of support vectors is constrained to remain within a specified limit as it is employed in [9]).

Reducing the dataset for KNN classification has been studied in the literature (e.g., the condensed nearest neighbor [33]), yet it is not suitable for tracking in which the distribution of target and background is non-stationary, and there is a need to keep/remove the samples based on the temporal properties of tracking task. We propose a simple accounting method for a sample x with the nearest neighborhood  $\eta(x)$  as KNN budgeting rules:

- 1) Discard the new sample x for which all  $\eta(x)$  have similar labels (*absorbed*);
- 2) Attach a timer  $\alpha$  to each new x that counts down upon processing each new frame. If the timer goes off  $(\alpha \rightarrow 0)$ , the sample is flagged;
- 3) Mark the sample for which all neighbors have opposite labels as *outlier*;
- For each added sample x, increment the timer of all η(x) whose labels differ with the new sample;
- 5) For each flagged sample if all of  $\eta(\mathbf{x})$  has similar labels to  $\mathbf{x}$  it is discarded (*absorbed*), if  $\mathbf{x}$  is an outlier and none of  $\eta(\mathbf{x})$  has the same label, it should also be discarded. The remaining flagged samples are called *prototypes*.

This scheme tends to discard the most futile samples from the sample pool while preserving the recent or essential ones. Figure 3 depicts the sample 2D feature space of a KNN classifier and demonstrate how this budgeting mechanism preserve the classification power while reducing the number of samples. This budgeting scheme serves as a forgetting mechanism for the KNN classifier as well. If a distractor is very similar to the target in the feature space, the samples obtained from it will be labeled positive and added again to the KNN classifier, reinforcing the classifier false belief about the label of such samples. Such cases act as local minima for feature space in KNN classifier and can only be resolved if the oracle investigates and disprove them. To rescue the KNN classifier from its local minima, it requires a forgetting mechanism such as the proposed budgeting scheme. This concept is illustrated in Figure 4.

## D. Uncertainty Sampling

When using a probabilistic model for binary classification, target/non-target in the case of tracking-by-detection, uncertainty sampling simply queries the instance whose posterior probability of being is halfway between positive values and negative values [25]. Uncertainty sampling strategies may also be employed with non-probabilistic algorithms like memory-based classifiers [29]. Inspired by these studies, we calculated the score of a sample as the KNN classifier confidence score, i.e., we allow the k nearest neighbors to vote on the class label of  $\mathbf{p}_t^i$ , and the sum of these votes representing the score.

As mentioned earlier, UST queries the m most uncertain samples (having the closest scores to 0) from the oracle. To handle a large number of uncertain samples, we decided to query all the samples in the range of  $(\tau_l, \tau_p)$  from the oracle. This backup mechanism along with forgetting mechanism realized by the budgeting helps the KNN detector to escape from its local minima induced by feature-space similar objects and partial occlusions (Figure 4).

# IV. EVALUATION

This section reports on a set of quantitative experiments comparing the UST with relevant algorithms. To evaluate the performance of the proposed tracker, the experiments is conducted on 100 challenging video sequences from [1]. These sequences include many of the visual tracking challenges such as scale variation, fast motion and motion blur, illumination variations, in-plane and out-of-plane rotations, low resolution and shear problem, background clutter and various types of occlusion. The performance of the tracker is measured by the area under the surface of its success plot. A tracker in time t succeed to track the object if its response  $\hat{p}_t$  overlaps with the ground truth  $p_t^*$  more than a threshold  $\tau_o$ . Success plot, graphs the success of the tracker against different values of the threshold  $\tau_o$  and its AUC is

$$AUC = \frac{1}{T} \int_0^1 \sum_{t=1}^T \mathbb{1} \left[ \frac{|\hat{\mathbf{p}}_t \cap \mathbf{p}_t^*|}{|\hat{\mathbf{p}}_t \cup \mathbf{p}_t^*|} > \tau_o \right] d_{\tau_o}$$
(7)

where  $\tau$  is the length of sequence, |.| denotes the area of the region,  $\cap$  and  $\cup$  stands for intersection and union of the regions respectively. Since UST has non-deterministic sampling parts, we run it 5 times and the average of the results is reported.

#### A. Comparison with Baseline

This experiment strives to demonstrate the advantages of the proposed tracker in comparison to the trackers that consists of either of its detectors in isolation. To this end, we construct several trackers from the components of this tracker to serve as the baselines for this experiment. In all of these trackers, the ROI detection and input sampling are similar to the UST tracker. KNN(10) and KNN(25) trackers utilize only the feature-based nearest neighbor detector for the tracking with the neighborhood size of k = 10 and k = 25 respectively. KNN+(10) and KNN+(25) trackers additionally incorporate the proposed budgeting mechanism into their detector. SVM tracker employs the oracle to track the target whereas SVM+ include the classifier update in its framework. Figure 5 compares the performance of these baseline tracker with that of UST, and demonstrates that the uncertainty sampling data exchange, effectively connect the classifiers to construct a robust and efficient tracker.



Figure 5: Quantitative comparison of UST with its baselines



Figure 6: Quantitative comparison of the proposed tracker, UST, with the state-of-the-art trackers using success plot

#### B. Comparison with state-of-the-art

To establish a fair comparison with the state-of-theart, we select some of the most popular discriminative trackers based on [1] and perform a benchmark on the whole videos of the dataset, along with partial subsets of the dataset with a distinguishing attribute to evaluate the tracker performance under different situations. These trackers are BSBT [6], CSK [12], CT [3], CXT [10], DFT [34], FOT [35], FRAG [2], LOT [4], LSHT [36], LSK [37], MIL [7], SBT [8], STRUCK [9], TLD [11], and VR [18].

Figure 6 depicts the performance of all of the investigated trackers. As it is evident from this plot, UST outperforms the other trackers by having the highest AUC. Interestingly, UST also outperform these trackers when facing illumination, scale and shape changes that show the resilience of this double appearance model used by the two detectors (Table I). Rotations, shear and fast motions are well addressed by the proposed tracker and only STRUCK handled motion blur as good as UST. However, background clutter and low resolution targets challenged the UST. Not equipped with the means to handle LR, the results seems acceptable (Figure 7). In the case of background clutter, the use of context information in CXT and local sparsity in LSK outperforms our proposed tracker, shedding some light on the future direction of research.

Finally, it is prudent to note that UST achieved an average speed of 28.3 fps on a Pentium 4 Core i7 @ 3.2 GHz by Matlab/C++ implementation. This experiment demonstrated that with a adequate information exchange in co-tracking, it is possible to balance a good trade-off between speed and accuracy, while the tracker is capable of properly under various tracking challenges.

Table I: Quantitative evaluation of five best trackers under different visual tracking challenges using AUC of success plot. The best performance for each attribute is **bold**.

Attribute	CSK	TLD	LSK	STRK	CXT	UST
Illumination Variation	0.40	0.49	0.50	0.46	0.52	0.52
Deformation	0.36	0.32	0.38	0.41	0.32	0.47
Occlusion	0.36	0.42	0.44	0.44	0.40	0.53
Scale Variation	0.34	0.44	0.46	0.43	0.45	0.51
In-plane Rotation	0.43	0.50	0.46	0.51	0.53	0.59
Out-of-plane Rotation	0.39	0.43	0.45	0.48	0.45	0.56
Out-of-View	0.32	0.45	0.39	0.44	0.38	0.52
Low Resolution	0.29	0.37	0.39	0.39	0.38	0.38
Background Clutter	0.42	0.40	0.45	0.39	0.45	0.41
Fast Motion	0.39	0.45	0.42	0.52	0.44	0.53
Motion Blur	0.32	0.42	0.37	0.48	0.38	0.46
ALL	0.41	0.46	0.46	0.48	0.48	0.52



Figure 7: Sample tracking results of evaluated algorithms on several challenging video sequences. (from top) *FaceOcc2* with severe occlusions, *Deer* with fast motion, *Ironman* with in-plane and out-of-plane rotations, and *Board* with background clutter. In these sequences the orange box depicts the UTS, the yellow dashed line indicates the ground truth, and the rest depict the result of other trackers.

### V. CONCLUSION

The key component is a co-tracking framework consisting of a frequently updated (KNN-based) classifier and a more conservative (part-based) detector. Built upon uncertainty sampling foundation, samples deemed uncertain by the KNN classification are labeled by the part-based detector (the oracle). A memory budgeting mechanism keep classifier updates tractable. This accounting method, along with an optical-flow based ROI detection ensures that the proposed tracker, UST, meet th ereal-time criteria for tracking. Experimental results on challenging video sequences demonstrated that the UTS tracker achieve comparable accuracy to the state-of-the-art trackers, while outperforming them in terms of efficiency and robustness.

## ACKNOWLEDGMENT

This article is based on results obtained from a project commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

#### REFERENCES

[1] Y. Wu, J. Lim, and M.-H. Yang, "Object tracking benchmark," *PAMI*, vol. 37, no. 9, pp. 1834–1848, 2015.

- [2] A. Adam, E. Rivlin, and I. Shimshoni, "Robust fragments-based tracking using the integral histogram," in CVPR'06, 2006.
- [3] K. Zhang, L. Zhang, and M.-H. Yang, "Real-time compressive tracking," in *ECCV'12*. Springer, 2012, pp. 864–877.
- [4] S. Oron, A. Bar-Hillel, D. Levi, and S. Avidan, "Locally orderless tracking," *IJCV*, vol. 111, no. 2, pp. 213–228, 2015.
- [5] A. Taalimi, H. Qi, and R. Khorsandi, "Online multi-modal taskdriven dictionary learning and robust joint sparse representation for visual tracking," in AVSS'15. IEEE, 2015, pp. 1–6.
- [6] S. Stalder, H. Grabner, and L. Van Gool, "Beyond semi-supervised tracking: Tracking should be as simple as detection," in *ICCVw'09*.
- [7] B. Babenko, M.-H. Yang, and S. Belongie, "Visual tracking with online multiple instance learning," in *CVPR'09*, 2009.
- [8] H. Grabner, J. Matas, L. Van Gool, and P. Cattin, "Tracking the invisible: Learning where the object might be," in CVPR'10, 2010.
- [9] S. Hare, A. Saffari, and P. H. Torr, "Struck: Structured output tracking with kernels," in *ICCV'11*. IEEE, 2011, pp. 263–270.
- [10] T. B. Dinh, N. Vo, and G. Medioni, "Context tracker: Exploring supporters and distracters in unconstrained environments," in CVPR'11.
- [11] Z. Kalal, K. Mikolajczyk, and J. Matas, "Tracking-learningdetection," PAMI, vol. 34, no. 7, pp. 1409–1422, 2012.
- [12] J. F. Henriques, R. Caseiro, and J. Batista, "Exploiting the circulant structure of tracking-by-detection with kernels," in ECCV'12, 2012.
- [13] F. Tang, S. Brennan, Q. Zhao, and H. Tao, "Co-tracking using semisupervised support vector machines," in *ICCV'07*, 2007.
- [14] H. Grabner, M. Grabner, and H. Bischof, "Real-time tracking via on-line boosting." in *BMVC'06*, vol. 1, no. 5, 2006, p. 6.
- [15] S. Avidan, "Support vector tracking," PAMI, 2004.
- [16] H. Grabner, C. Leistner, and H. Bischof, "Semi-supervised on-line boosting for robust tracking," in ECCV'08, 2008.
- [17] S. Avidan, "Ensemble tracking," PAMI, 2007.
- [18] R. T. Collins, Y. Liu, and M. Leordeanu, "Online selection of discriminative tracking features," *PAMI*, 2005.
- [19] C. Bao, Y. Wu, H. Ling, and H. Ji, "Real time robust 11 tracker using accelerated proximal gradient approach," in *CVPR'12*, 2012.
- [20] Z. Hong, Z. Chen, X. Mei, and D. Tao, "Multi-store tracker (muster): A cognitive psychology inspired approach," in CVPR'15.
- [21] K. Zhang, L. Zhang, M.-H. Yang, and Q. Hu, "Robust object tracking via active feature selection," CSVT, 2013.
- [22] C. H. Lampert and J. Peters, "Active structured learning for highspeed object detection," in *PR*. Springer, 2009, pp. 221–231.
- [23] C. Li, X. Wang, W. Dong, J. Yan, Q. Liu, and H. Zha, "Active sample learning and feature selection: A unified approach," 2015.
- [24] A. Beygelzimer, S. Dasgupta, and J. Langford, "Importance weighted active learning," in *ICML'09*, 2009.
- [25] D. D. Lewis and W. A. Gale, "A sequential algorithm for training text classifiers," in ACM SIGIR'94, 1994, pp. 3–12.
- [26] D. D. Lewis and J. Catlett, "Heterogeneous uncertainty sampling for supervised learning," in *ICML'94*, 1994, pp. 148–156.
- [27] T. Scheffer, C. Decomain, and S. Wrobel, "Active hidden markov models for information extraction," in *ISIDA'01*, 2001.
- [28] B. Settles and M. Craven, "An analysis of active learning strategies for sequence labeling tasks," in *EMNLP'08*, 2008.
- [29] M. Lindenbaum, S. Markovitch, and D. Rusakov, "Selective sampling for nearest neighbor classifiers," *JML*, 2004.
- [30] S. Tong and D. Koller, "Support vector machine active learning with applications to text classification," JMLR, vol. 2, pp. 45–66, 2002.
- [31] T. Brox, C. Bregler, and J. Malik, "Large displacement optical flow," in CVPR'09. IEEE, 2009, pp. 41–48.
- [32] P. F. Felzenszwalb, R. Girshick, and D. Ramanan, "Object detection with discriminatively trained part-based models," *PAMI*, 2010.
- [33] F. Angiulli, "Fast condensed nearest neighbor rule," in ICML'05.
- [34] L. Sevilla-Lara and E. Learned-Miller, "Distribution fields for tracking," in CVPR'12. IEEE, 2012, pp. 1910–1917.
- [35] J. Matas and T. Vojir, "Robustifying the flock of trackers," in *CVWW'11*. Citeseer, 2011, p. 91.
- [36] S. He, Q. Yang, R. Lau, J. Wang, and M.-H. Yang, "Visual tracking via locality sensitive histograms," in CVPR'13, 2013.
- [37] B. Liu, J. Huang, L. Yang, and C. Kulikowsk, "Robust tracking using local sparse appearance model and k-selection," in CVPR'11.