

Learning to Segment and Track in RGBD

Alex Teichman and Sebastian Thrun

Abstract We consider the problem of segmenting and tracking deformable objects in color video with depth (RGBD) data available from commodity sensors such as the Kinect. We frame this problem with very few assumptions - no prior object model, no stationary sensor, no prior 3D map - thus making a solution potentially useful for a large number of applications, including semi-supervised learning, 3D model capture, and object recognition. Our approach makes use of a rich feature set, including local image appearance, depth discontinuities, optical flow, and surface normals to inform the segmentation decision in a conditional random field model. In contrast to previous work, the proposed method *learns* how to best make use of these features from ground-truth segmented sequences. We provide qualitative and quantitative analyses which demonstrate substantial improvement over the state of the art.

1 Introduction

The availability of commodity depth sensors such as the Kinect opens the door for a number of new approaches to important problems in robot perception. In this paper, we consider the task of propagating an object segmentation mask through time. We assume a single initial segmentation is given and do not allow additional segmentation hints. In this work the initial segmentation is provided by human labeling, but this input could easily be provided by some automatic method, depending on the application. We do not assume the presence of a pre-trained object model (*i.e.* as the Kinect models human joint angles), as that would preclude the system from segmenting and tracking arbitrary objects of interest. As such, this task falls into the category of *model-free* segmentation and tracking, *i.e.* no prior class model is assumed. Similarly, we do not assume the sensor is stationary or that a pre-built static

environment map is available. There are several reasons a complete solution to this task would be useful in robotics, computer vision, and graphics.

First, model-free segmentation and tracking opens the door for a simple and effective method of semi-supervised learning in which a large number of unlabeled tracks of objects are used in conjunction with a small number of hand-labeled tracks to learn an accurate classifier. This method, known as tracking-based semi-supervised learning, was demonstrated in the autonomous driving context using laser range finders to learn to accurately recognize pedestrians, bicyclists, and cars versus other distractor objects in the environment using very few hand-labeled training examples [23]. However, model-free segmentation and tracking was more or less freely available because objects on streets generally avoid collision and thus remain depth-segmentable using simple algorithms. To extend tracking-based semi-supervised learning to the more general case in which objects are not easily depth-segmentable, more advanced model-free segmentation and tracking methods such as that of this paper are necessary.




Fig. 1 Given an initial seed labeling (first frame; cat on a coffee table), the goal is to produce a detailed segmentation of the object over time (subsequent frames) without assuming a rigid object or a pre-trained object model. Foreground points are shown in bold.

Second, a new approach to object recognition with compelling benefits would be made possible. Currently, most object recognition methods fall into semantic segmentation or sliding window categories. In this new approach, objects are segmented and tracked over time and the track as a whole is classified online as new segmented frames stream in. This enables the classifier to use multiple views over time and to use track descriptors such as average object speed or maximum angular velocity. Some examples of this track-based approach exist, but only when special assumptions about the environment can be made, such as in [22].

Third, training data for object recognition tasks could easily be acquired with such a system. A few sparse labels provided by a user can be used to segment the remainder of the object track, providing new training examples with different object poses, view angles, lighting conditions, and so on, with little additional human effort. This same method could be used for 3D model-capture in unstructured environments, with all of the segmentation masks of a sequence being synthesized into one coherent object seen from all sides.

2 Previous work

While there is much previous work on tracking in general, our problem's online nature, output of detailed segmentation masks rather than bounding boxes, and lack of simplifying assumptions restrict the directly-related previous work substantially. The most similar and recent work to ours is HoughTrack [7], which uses Hough Forests and GrabCut to segment consecutive frames. We provide a quantitative comparison to HoughTrack in Section 4.1. Older previous work on this problem include [3] and their multi-object extension [4], which use a generative model and level sets; [11], which tracks local object patches using Basin Hopping Monte Carlo sampling; and [15], which uses a conditional random field model and loopy belief propagation. All these works limit themselves to only a small number of features and do not use learning methods to find the best general-purpose tracker. Additionally, none of these works consider depth information.

We now briefly review previous work on related problems.

Model-based tracking - For some tasks, it is appropriate to model a specific object that is to be tracked, either with an explicit 3D model as in [14] or with pre-trained statistical models that are specific to the particular object class. Examples of the latter approach include the human pose tracking of the Kinect [18] and model-based car tracking in laser data for autonomous driving [13].

Interactive segmentation - Human-assisted object segmentation has been the subject of much work including interactive graph cuts [5], GrabCut [16], and Video SnapCut [2]. These methods are largely intended for graphics or assisted-labeling applications as they require a human in the loop.

Bounding box tracking - Discriminative tracking [8, 19, 10] addresses the online, model-free tracking task, but where the goal is to track arbitrary objects in a bounding box rather than provide a detailed segmentation mask.

Rigid object tracking - Rigid object tracking using depth information, such as the open source method in PCL [17], addresses a similar but simpler problem, as it is not designed to work with deformable objects.

Offline methods - The work of [6] takes as input segmentations for the first and last frames in a sequence, rather than just the first frame. While [24] takes as input only the segmentation of the first frame, they construct a CRF on the entire video sequence, with CRF labels corresponding to object flows. Here we are interested in a method that can update as new data streams in, thus making it applicable for robotics or other online vision tasks.

Background subtraction - Finally, background subtraction approaches can greatly simplify the segmentation task. For example, [1] assumes a stationary camera to produce fine-grained segmentation masks for multiple objects. With depth sensors, background subtraction methods can also operate while the sensor is in motion: Google's autonomous car project uses pre-built 3D static environment maps to subtract away all 3D points except the moving obstacles [26]. This enables simple depth segmentation to be effective as long as moving objects do not touch, but assumes a high-precision localization system as well as a pre-learned static environment map.

3 Approach

There are a large number of possible cues that could inform a segmentation and tracking algorithm. Optical flow, image appearance, 3D structure, depth discontinuities, color discontinuities, etc., all provide potentially useful information. For example:

- An optical flow vector implies that the label of the source pixel at time $t - 1$ is likely to propagate to that of the destination pixel at time t .
- Pixels with similar color and texture to previous foreground examples are likely to be foreground.
- The shape of the object in the previous frame is likely to be similar to the shape of the object in the next frame.
- Points nearby in 3D space are likely to share the same label.
- Nearby points with similar colors are likely to share the same label.

Previous work generally focuses on a few particular features; here, we advocate the use of a large number of features. The above intuitions can be readily encoded by node and edge potentials in a conditional random field model. However, this additional complexity introduces a new problem: how much importance should each feature be assigned? While using a small number of features permits their weights to be selected by hand or tested with cross validation, this quickly becomes impractical as the number of features increases.

The margin-maximizing approach of structural SVMs [21, 25], adapted to use graph cuts for vision in [20], provides a solution to learning in these scenarios. Intuitively, this method entails running MAP inference, and then adding constraints to an optimization problem which assert that the margin between the ground truth labeling and the generated (and generally incorrect) labeling should be as large as possible. The application of this approach will be discussed in detail in Section 3.2.

3.1 Conditional random fields and inference

A conditional random field is an undirected graphical model that can make use of rich feature sets and produce locally-consistent predictions. It aims to spend modeling power on only the distribution of the target variables given the observed variables. In particular, the conditional random field takes the form

$$\mathbb{P}(y|x) = \frac{1}{Z(x)} \exp(-E(y,x)), \quad (1)$$

where Z is the normalizer or partition function, $y \in \{-1, +1\}^n$ is the segmentation for an image with n pixels, and x is a set of features defined for the RGBD frame, to be detailed later. The energy function E contains the features that encode various intuitions about what labels individual pixels should take and which pixels should

share the same labels. In particular, the energy function is defined as

$$E(y, x) = \sum_{i \in \Phi_v} w_i \sum_{j \in \mathcal{V}_i} \phi_j^{(i)}(y, x) + \sum_{i \in \Phi_e} w_i \sum_{(j, k) \in \mathcal{N}_i} \phi_{jk}^{(i)}(y, x). \quad (2)$$

Here, Φ_v is the set of node potential indices (*i.e.* one for each type of node potential such as local image appearance or 3D structure alignment), \mathcal{V}_i is the set of all node indices for this potential type (normally one per pixel), and $\phi_j^{(i)}$ is the node potential of type i at pixel j . Similarly, Φ_e is the set of edge potential indices, \mathcal{N}_i is the neighborhood system for edge potential type i (normally pairs of neighboring pixels), and $\phi_{jk}^{(i)}$ is the edge potential between pixels j and k for edge potentials of type i . Thus, the weights apply at the feature-type level, *i.e.* w_i describes how important the i th feature is.

The goal of MAP inference in this model is to choose the most likely segmentation y given the features x by solving

$$\underset{y}{\text{maximize}} \quad \mathbb{P}(y|x) = \underset{y}{\text{minimize}} \quad E(y, x). \quad (3)$$

During inference, the weights w remain fixed. The solution to this problem can be efficiently computed for $y \in \{-1, +1\}^n$ and submodular energy function E using graph cuts [5].

3.2 Learning

Given a dataset of (y_i, x_i) for $i = 1 \dots m$, the goal of CRF learning is to choose the weights w that will result in the lowest test error. While it would be desirable to learn the weights directly using the maximum likelihood approach

$$\underset{w}{\text{maximize}} \quad \prod_m \mathbb{P}(y_m | x_m), \quad (4)$$

this is not possible to do exactly because of the presence of the partition function $Z(x) = \sum_y \exp(-E(y, x))$ in the gradient. This function sums over all 2^n segmentations of an image with n pixels.

Fortunately, there is an alternative approach known as the structural support vector machine [21, 25, 20], which we now briefly review. Solving the margin maximizing optimization problem

$$\begin{aligned} \underset{w, \xi}{\text{minimize}} \quad & \frac{1}{2} \|w\|^2 + \frac{C}{M} \sum_{m=1}^M \xi_m \\ \text{subject to} \quad & \xi \geq 0 \\ & E(y, x_m) - E(y_m, x_m) \geq \Delta(y_m, y) - \xi_m \quad \forall m, \forall y \in \mathcal{Y} \end{aligned} \quad (5)$$

would result in a good solution. Here, C is a constant, Δ is a loss function, ξ_m is a slack variable for training example m , and \mathcal{Y} is the set of all possible labelings of an image. This problem, too, is intractable because it has exponentially many constraints. However, one can iteratively build a small, greedy approximation to the exponential set of constraints such that the resulting weights w are good.

Further, one can transform (5), known as the n -slack formulation, into the equivalent problem

$$\begin{aligned}
& \underset{w, \xi}{\text{minimize}} && \frac{1}{2} \|w\|^2 + C\xi \\
& \text{subject to} && \xi \geq 0 \\
& && \frac{1}{M} \sum_{m=1}^M E(\hat{y}_m, x_m) - E(y_m, x_m) \geq \frac{1}{M} \sum_{m=1}^M \Delta(y_m, \hat{y}_m) - \xi \\
& && \forall (\hat{y}_1, \dots, \hat{y}_M) \in \mathcal{Y}^M,
\end{aligned} \tag{6}$$

as described in [9]. As before, while the number of constraints is exponential, one can build a small, greedy approximation that produces good results. Known as the 1-slack formulation, this problem is equivalent and can be solved much more efficiently than (5), often by one or two orders of magnitude.

The goal is to learn the weights that will best segment an entire sequence given a single seed frame. However, the segmenter operates on a single frame at a time, so our training dataset must be in the form of (y, x) pairs. As will be discussed in Section 3.3, some of the features are stateful, *i.e.* they depend on previous segmentations. This presents a minor complication. At training time, we do not have access to a good choice of weights w , but the features depend on previous segmentations which in turn depend on the weights. To resolve this, we adopt the simple strategy of generating features assuming that previous segmentations were equal to ground truth. This dataset generation process is specified in Algorithm 1.

Algorithm 1 Training set generation

```

 $\mathbb{S} = \{S : S = ((y_0, d_0), (y_1, d_1), \dots)\}$  is a set of sequences,
  where  $y_i$  is a ground truth segmentation and  $d_i$  is an RGBD frame
 $\mathcal{D} = \emptyset$ 
for  $S \in \mathbb{S}$  do
  Initialize stateful components, e.g. the patch classifier that learns its model online
  for  $(y_i, d_i) \in S$  do
    Update stateful components using  $y_{i-1}$  as the previous segmentation
    Generate features  $x$ 
     $\mathcal{D} := \mathcal{D} \cup \{(y_i, x)\}$ 
  end for
end for
return  $\mathcal{D}$ 

```

The structural SVM solver is detailed in Algorithm 2. Because the graph cuts solver assumes a submodular energy function, non-negative edge weights w_i for

Algorithm 2 Structural SVM for learning to segment and track

\mathcal{D} is a set of training examples (y, x) , formed as described in Algorithm 1.
 C and ε are constants, chosen by cross validation.

$\mathcal{W} \leftarrow \emptyset$

repeat

Update the parameters w to maximize the margin.

$$\underset{w, \xi}{\text{minimize}} \quad \frac{1}{2} \|w\|^2 + C\xi$$

$$\text{subject to} \quad w \geq 0, \quad \xi \geq 0$$

$$\frac{1}{M} \sum_{m=1}^M E(\hat{y}_m, x_m) - E(y_m, x_m) \geq \frac{1}{M} \sum_{m=1}^M \Delta(y_m, \hat{y}_m) - \xi$$

$$\forall (\hat{y}_1, \dots, \hat{y}_M) \in \mathcal{W}$$

for $(y_m, x_m) \in \mathcal{D}$ **do**

Find the MAP assignment using graph cuts.

$$\hat{y}_m \leftarrow \operatorname{argmin}_y E(y, x_m)$$

end for

$$\mathcal{W} \leftarrow \mathcal{W} \cup \{(\hat{y}_1, \dots, \hat{y}_M)\}$$

until $\frac{1}{M} \sum_{m=1}^M \Delta(y_m, \hat{y}_m) - E(\hat{y}_m, x_m) + E(y_m, x_m) \leq \xi + \varepsilon$

all $i \in \Phi_{\mathcal{E}}$ are required. As our node potentials are generally designed to produce values with the desired sign, we constrain them in Algorithm 2 to have non-negative weights as well. The term $\Delta(y_m, y)$ is 0-1 loss, but a margin rescaling approach could easily be obtained by changing Δ to Hamming loss and making a small modification to the graph cuts solver during learning (see [20]).

3.3 Energy function terms

We now review the particular choice of energy function (2) that we use in our implementation. The particular choice of features can be modified or added to without substantially changing the underlying framework. See Figure 2 for visualizations. All node potentials are designed to be constrained to $[-1, +1]$, and all edge potentials to $[-1, 0]$. While not strictly necessary, this aids in interpreting the weights learned by the structural SVM.

3.3.1 Node potentials

Node potentials capture several aspects of shape, appearance, and motion. All take the form




Fig. 2 Visualizations of selected edge and node potentials in the CRF. Edge potentials are zoomed in to show fine structure. Strong edge potentials are darker, while weaker edge potentials fade to original image color. Node potentials expressing a preference for foreground are shown in red, for background in green, and for neither in black. Top row: original image, canny edge potentials, color distance, depth edge potentials. Bottom row: ICP, frame alignment bilateral filter, optical flow, and patch classifier node potentials. Best viewed in color.

$$\phi_j(y,x) = \begin{cases} a_j & \text{if } y_j = 1 \\ b_j & \text{if } y_j = -1, \end{cases} \quad (7)$$

where a_j and b_j are functions of the data x , and $y_j = 1$ indicates that pixel j has been assigned to the foreground.

Seed labels - The first frame of every sequence provides a foreground/background segmentation. Seed labels can also be used for algorithm-assisted human labeling, used here only in construction of the ground truth dataset. Each pixel in the image can be labeled by the user as foreground or background, or left unlabeled. Concretely, let $s_j \in \{-1, 0, +1\}$ be the seed labels for each pixel. Then, the potential is defined by

$$\phi_j(y,x) = \begin{cases} -1 & \text{if } y_j = s_j \\ 0 & \text{otherwise.} \end{cases}$$

Optical flow - Pixels in the current frame with optical flow vectors from the previous frame are likely to have the label of the originating pixel. If no optical flow vector terminates at pixel j in the current frame, let $f_j = 0$. Otherwise, let f_j be equal to the label of the originating pixel from the previous frame. Then, the potential is defined by

$$\phi_j(y,x) = \begin{cases} -1 & \text{if } y_j = f_j \\ 0 & \text{otherwise.} \end{cases}$$

Frame alignment bilateral filter - Optical flow provides correspondences from one frame to the next; after aligning these frames, a bilateral filter given 3D position and RGB values is used to blur the labels from the previous frame into the current frame, respecting color and depth boundaries.

Specifically, alignment between the two frames is found using a RANSAC method, with 3D correspondences given by optical flow vectors. Flow vectors that originate or terminate on a depth edge are rejected, as their 3D location is often unstable. After alignment, we find the set of neighboring points N from the previous frame within a given radius of the query point, then compute the value

$$z_j = \sum_{k \in N} y'_k \exp \left(-\frac{\|c_j - c_k\|_2}{\sigma_c} - \frac{\|p_j - p_k\|_2}{\sigma_d} \right),$$

where y'_k is the $\{-1, +1\}$ label of point k in the previous segmentation, c_j is the RGB value of pixel j , and p_j is the 3D location of pixel j . The two σ s are bandwidth parameters, chosen by hand. The value z_j is essentially that computed by a bilateral filter, but without the normalization. This lack of normalization prevents a point with just a few distant neighbors from being assigned an energy just as strong as a point with many close neighbors. Referring to (7), the final energy assignment is made by setting $b_j = 0$ and $a_j = 1 - 2/(1 + e^{-z_j})$.

Patch classifiers - Two different parameterizations of a random fern patch classifier similar to [12] are trained online and used to make pixel-wise predictions based on local color and intensity information.

A random fern classifier is a semi-naive Bayes method in which random sets of bit features are chosen (“ferns”), and each possible assignment to the bit features in the fern defines a bin. Each of these bins maintains a probability estimate by simple counting. The final probability estimate is arrived at by probabilistically combining the outputs of many different randomly generated ferns. See [12] for details. We use 50 randomly generated ferns, each of which have 12 randomly generated features.

Each bit feature in a fern is generated by very simple functions of the data contained in the image patch. For example, one of the bit features we use is defined by

$$z = \begin{cases} 1 & \text{if } I_{p_0} < I_{p_1} \\ 0 & \text{otherwise,} \end{cases}$$

where p_0 and p_1 are two pre-determined points in the image patch, and I_p is the intensity of the camera image at point p . Other bit features include color comparisons and Haar wavelet comparisons, efficiently computed using the integral image.

At each new frame, the classifier is updated using the graph cuts segmentation output from the previous frame as training data, then is used to compute a probability of foreground estimate for every pixel in the current frame. The value a_j is set to the probability of background, and b_j to probability of foreground.

Distance from previous foreground - Points very far away from the object’s position in the previous frame are unlikely to be foreground.

After alignment of the current frame to the previous frame using the same optical-flow-driven method as the bilateral term, we compute the distance d from the point j in the current frame to the closest foreground point in the previous frame. The final potentials are set to $a_j = 0$ and $b_j = \exp(-d/\sigma) - 1$.

ICP - The iterative closest point algorithm is used to fit the foreground object's 3D points (and nearby background points) from the previous frame to the current frame. If the alignment has a sufficient number of inliers, points in the current frame are given node potentials similar to those of the bilateral node potential, but where the ICP-fit points are used instead of the entire previous frame.

Prior term - To express a prior on background, we add a term for which $a_j = 0$ and $b_j = -1$.

3.3.2 Edge potentials




Fig. 3 While image appearance and depth information may sometimes provide little information about an object boundary, a change in surface normals can be informative. Surface normals are shown as small dark lines, seen from two angles to aid in depth perception.

Edge potentials capture the intuition that the foreground/background boundary is likely to be at an image gradient, color change, depth discontinuity, or surface normal change. All take the form

$$\phi_{jk}(y, x) = \begin{cases} a_{jk} & \text{if } y_j = y_k \\ 0 & \text{otherwise,} \end{cases}$$

where a_{jk} is a function of the data x . All edges are between neighboring points in the image.

Canny edges - All neighboring pixels are connected by an edge potential except those cut by canny edges. Concretely,

$$a_{jk} = \begin{cases} -1 & \text{if neither pixel lies on a Canny edge} \\ 0 & \text{otherwise.} \end{cases}$$

Color distance - Edge weights are assigned based on Euclidean distance between neighboring RGB values; $a_{jk} = -\exp(-\|c_j - c_k\|/\sigma)$, where c_j and c_k are the RGB values of the two pixels and σ is a bandwidth parameter, chosen by hand.

3D distance - Points nearby in 3D are likely to share the same label, especially if they lie in the same plane. Out-of-plane distance changes are penalized more heavily than in-plane distance changes. Specifically,

$$a_{jk} = -\exp\left(-\frac{|(p_j - p_k)^T n_k|}{\sigma_n} - \frac{\|p_j - p_k\|_2}{\sigma_d}\right),$$

where p_j and p_k are the neighboring 3D points, n_k is the surface normal at point p_k , and the σ s are bandwidth parameters chosen by hand.

Surface normals - Neighboring points that share the same surface normal are likely to have the same label. See Figure 3. We use $a_{jk} = -\exp(-\theta/\sigma)$, where θ is the angle between the two normals and σ is a bandwidth parameter chosen by hand.

Edge potential products - Several combinations of the above are also provided, taking the form $a_{jk} = -|\prod_i a_{jk}^{(i)}|$, where $a_{jk}^{(i)}$ is the value from one of the above edge potentials. Intuitively, this encodes an edge potential that is strong (*i.e.* favors label equality) if all of the component edge potentials are strong.

4 Experiments

We provide a quantitative analysis to demonstrate improvement over the state of the art and a qualitative discussion the strengths and weaknesses of the current implementation. All experiments use 160x120 RGBD data. The structural SVM of Algorithm 2 was implemented with standard interior point solver techniques. The graph cuts solver of [5] was used for inference.

4.1 Quantitative analysis

As there is no work we are aware of which does segmentation and tracking of non-rigid objects using learning or depth data, we compare against HoughTrack, discussed in Section 2. We used the implementation that accompanies [7]. To ensure a fair comparison, we modified the HoughTrack implementation to be initialized with a segmentation mask rather than a bounding box.

No RGBD segmentation and tracking dataset currently exists, so we generated one containing 28 fully-labeled sequences with a total of about 4000 frames. Objects include, for example, a sheet of paper, cat, jacket, mug, laptop, and hat. In general, the dataset includes rigid and non-rigid objects, and textured and non-textured objects. Ground truth was generated by hand labeling, assisted with an interactive version of the segmenter, similar to the interactive graph cuts work of [5].

The dataset was split into a training set of about 500 frames over 10 sequences and a testing set of about 3500 frames over 18 sequences. Training of our method was run once on the training set and the resulting segmenter was used to produce results for all testing sequences. HoughTrack has no training stage, and thus did not use the training set. Testing results were produced for both by providing an initial segmentation for the first frame of each sequence, and all further frames were segmented without additional input.

Individual frames are evaluated with two different metrics. Hamming loss, or total number of pixels wrong, is simple and direct but ignores the size of the object; a Hamming loss of twenty could be negligible or a significant fraction of the object in question. To correct for this, we also report *normalized accuracy*, which is 1 if all pixels in the frame are correctly labeled and 0 if the number of incorrectly labeled pixels equals or exceeds the number of pixels in the foreground object. More precisely, normalized accuracy is $1 - \min(1, \text{num_wrong}/\text{num_fg})$. Sequence results are reported as the average of these values over all frames.

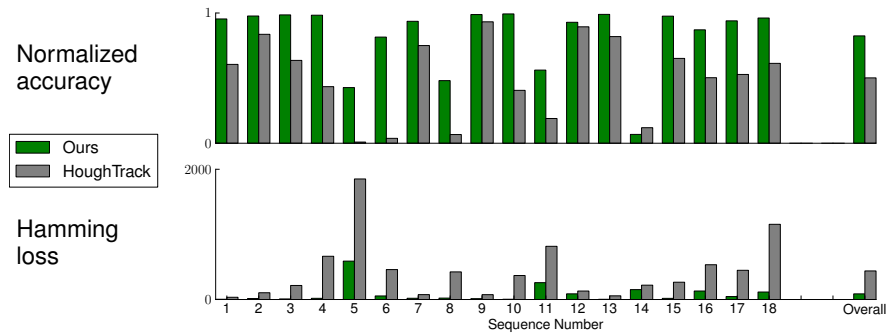


Fig. 4 Comparison of our method with the most similar state-of-the-art work.

Overall, our method demonstrates a normalized error reduction of about 65% compared to the state-of-the-art. Detailed results can be seen in Figure 4. Training of our method takes a total of about 10 minutes, and as this need only be done once, this is acceptable. At runtime, HoughTrack’s implementation takes about 200ms per frame, whereas our method’s implementation takes about 1200ms per frame.¹ While this is currently too slow for real-time usage, almost all of this time is spent computing features and only about a millisecond is spent on the graph cuts solver. This means real-time performance can be achieved with speedups of feature computation alone, most immediately by parallelization and by computing features only at the boundaries of the object where they are actually needed.

¹ Both implementations were compiled with optimizations turned on, run on the same machine, and given one thread.

4.2 Qualitative analysis

4.2.1 Strengths




Fig. 5 Visualization of results in which image appearance alone would lead to a very difficult segmentation problem, but which becomes relatively easy when reasoning about depth as well.

Our results demonstrate that this method can be effective even in sequences which include significant non-rigid object transformations, occlusion, and a lack of visually distinguishing appearance. As seen in the illustration in Figure 2, depth provides a very powerful cue as to what neighboring points are likely to share the same label. In addition to depth discontinuities, surface normals (Figure 3) can provide useful information about object boundaries, enabling the segmentation and tracking of objects for which using visual appearance alone would be extremely challenging, such as that of Figure 5. While this sequence would be relatively easy for off-the-shelf 3D rigid object trackers such as that in PCL [17], the tracking of deformable objects such as cats and humans would not be possible.

Our approach is general enough to handle both objects with few visually distinguishing characteristics and objects which are non-rigid; see Figure 6 for examples. In particular, the interface between the cat and the girl in sequences (A) and (B), and the interface between the cat and the bag in sequence (C) are maintained correctly. The method also works for objects without significant depth, such as the paper on the desk in sequence (G), and can recover from heavy occlusion as in sequence (F). The hat in sequence (H) undergoes deformation and heavy RGB blurring due to darkness of the environment, yet is tracked correctly. Sequence (D) shows the tracking of a piece of bread with Nutella while it is being eaten.

4.2.2 Weaknesses and future work

As is typical in tracking tasks, stability versus permissiveness is a tradeoff. In some cases, it is not well defined whether a recently-unoccluded, disconnected set of points near the foreground object should be part of the foreground or part of the

background. An example of this occurs in sequence (E), when the cat’s head becomes self-occluded, then unoccluded, and the system cannot determine that these points should be assigned to foreground. The error is then propagated. Similarly, the girl’s far arm in sequence (A) is lost after being occluded and then unoccluded. Rapid motion appears to exacerbate this problem. While partial occlusion can be handled, as in sequence (F), the current implementation has no facility for reacquisition if the target is completely lost. Finally, thin parts of objects are often problematic, as it seems the edge potentials are not well connected enough in these areas.

It is likely that improvements to node and edge potentials could resolve these limitations. Using longer range edges (rather than a simple 2D grid connecting neighbors only) could improve performance on thin objects and quickly moving objects, and improvements in the image-appearance-based patch classifier could result in better segmentations of objects that self-occlude parts frequently. A node potential based on LINE-MOD or a discriminative tracker such as [10] could solve the reacquisition problem, and could be run on object parts rather than the object as a whole to enable usage on non-rigid objects.

The feature-rich representation we use presents a challenge and an opportunity. There are a large number of parameters in the computation pipeline which cannot be learned via the structural SVM (*e.g.* the σ ’s discussed in Section 3.3). There is also substantial freedom in the structure of the computation pipeline. Choosing the structure and the parameters by hand (as we have done here) is possible, but onerous; there is an opportunity to learn these in a way that maximizes accuracy while respecting timing constraints.

5 Conclusions

We have presented a novel method of segmenting and tracking deformable objects in RGBD, making minimal assumptions about the input data. We have shown this method makes a significant quantitative improvement over the most similar state-of-the-art work in segmentation and tracking of non-rigid objects. A solution to this task would have far-reaching ramifications in robotics, computer vision, and graphics, opening the door for easier-to-train and more reliable object recognition, model capture, and tracking-based semi-supervised learning. While there remains more work to be done before a completely robust and real-time solution is available, we believe this approach to be a promising step in the right direction.

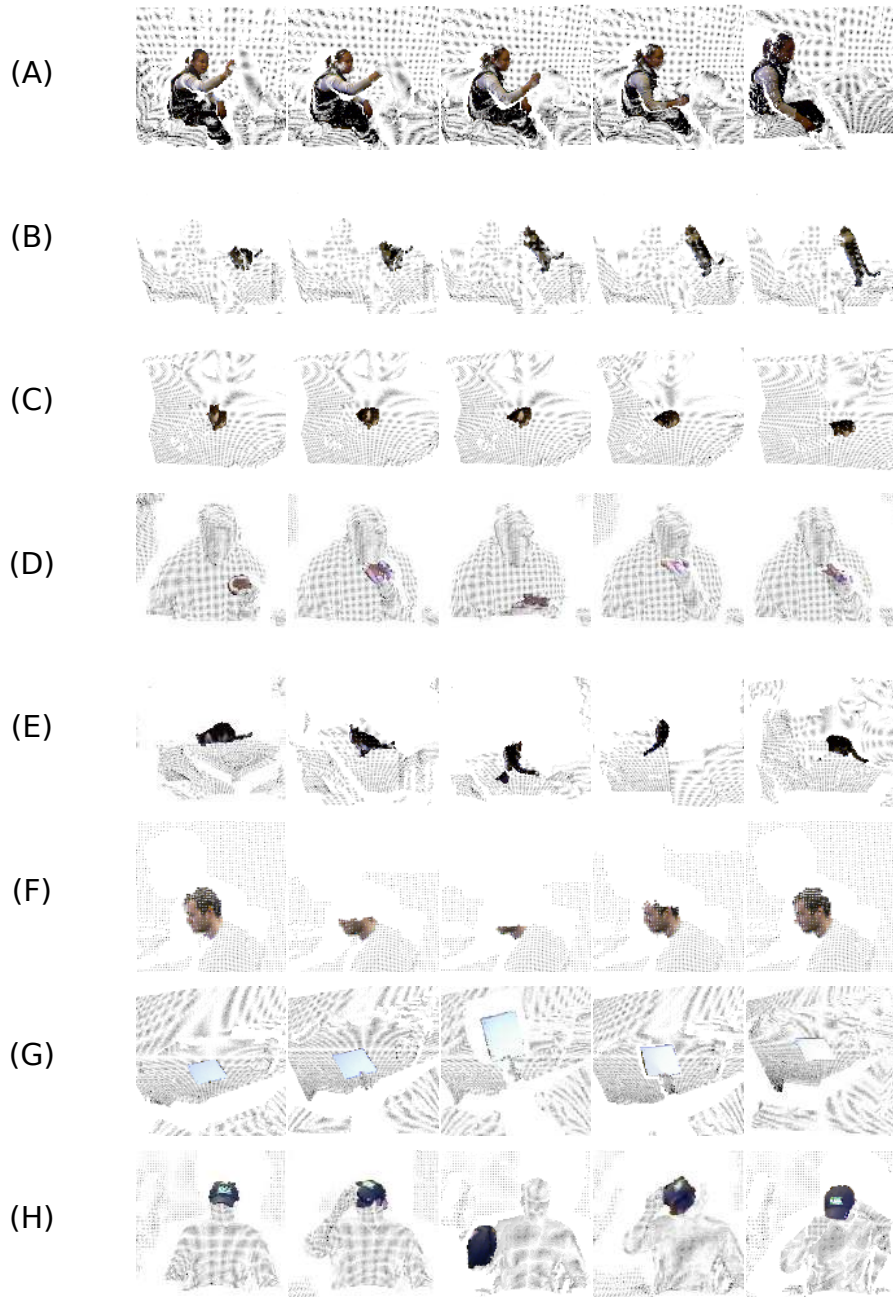


Fig. 6 Visualization of results. The first frame in each sequence (far left) is the seed frame. Foreground points are shown in bold and color while background points are shown small and gray. Best viewed on-screen.

References

1. C. Aeschliman, J. Park, and A. Kak. A probabilistic framework for joint segmentation and tracking. In *CVPR*, 2010.
2. X. Bai, J. Wang, D. Simons, and G. Sapiro. Video snapcut: robust video object cutout using localized classifiers. In *SIGGRAPH*, 2009.
3. C. Bibby and I. Reid. Robust real-time visual tracking using pixel-wise posteriors. In *ECCV*, 2008.
4. C. Bibby and I. Reid. Real-time tracking of multiple occluding objects using level sets. In *CVPR*, 2010.
5. Y. Boykov and V. Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26:359–374, 2001.
6. I. Budvytis, V. Badrinarayanan, and R. Cipolla. Semi-supervised video segmentation using tree structured graphical models. In *CVPR*, 2011.
7. M. Godec, P. Roth, and H. Bischof. Hough-based tracking of non-rigid objects. In *ICCV*, 2011.
8. H. Grabner, M. Grabner, and H. Bischof. Real-Time Tracking via On-line Boosting. In *British Machine Vision Conference*, 2006.
9. T. Joachims, T. Finley, and C.-N. Yu. Cutting-plane training of structural svms. *Machine Learning*, 77(1):27–59, 2009.
10. Z. Kalal, J. Matas, and K. Mikolajczyk. P-n learning: Bootstrapping binary classifiers by structural constraints. In *CVPR*, 2010.
11. J. Kwon and K. M. Lee. Tracking of a non-rigid object via patch-based dynamic appearance modeling and adaptive basin hopping monte carlo sampling. In *CVPR*, 2009.
12. M. Ozuysal, P. Fua, and V. Lepetit. Fast keypoint recognition in ten lines of code. In *CVPR*, 2007.
13. A. Petrovskaya and S. Thrun. Model based vehicle detection and tracking for autonomous urban driving. In *Autonomous Robots*, 2009.
14. V. A. Prisacariu and I. D. Reid. Pwp3d: Real-time segmentation and tracking of 3d objects. In *BMVC*, 2009.
15. X. Ren and J. Malik. Tracking as repeated figure/ground segmentation. In *CVPR*, 2007.
16. C. Rother, V. Kolmogorov, and A. Blake. "grabcut": interactive foreground extraction using iterated graph cuts. In *ACM SIGGRAPH 2004 Papers*, SIGGRAPH '04, pages 309–314, New York, NY, USA, 2004. ACM.
17. R. B. Rusu and S. Cousins. 3D is here: Point Cloud Library (PCL). In *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011.
18. J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake. Real-time human pose recognition in parts from single depth images. In *CVPR*, 2011.
19. S. Stalder, H. Grabner, and L. V. Gool. Beyond semi-supervised tracking: Tracking should be as simple as detection, but not simpler than recognition. In *International Conference on Computer Vision Workshop on On-line Learning for Computer Vision*, Sept. 2009.
20. M. Szummer, P. Kohli, and D. Hoiem. Learning crfs using graph cuts. In *ECCV*, 2008.
21. B. Taskar, V. Chatalbashev, D. Koller, and C. Guestrin. Learning structured prediction models: A large margin approach. In *ICML*, 2005.
22. A. Teichman, J. Levinson, and S. Thrun. Towards 3D object recognition via classification of arbitrary object tracks. In *International Conference on Robotics and Automation*, 2011.
23. A. Teichman and S. Thrun. Tracking-based semi-supervised learning. In *Robotics: Science and Systems*, 2011.
24. D. Tsai, M. Flagg, and J. Rehg. Motion coherent tracking with multi-label mrf optimization. In *BMVC*, 2010.
25. I. Tsochantaridis, T. Joachims, T. Hofmann, and Y. Altun. Large margin methods for structured and interdependent output variables. 2005.
26. C. Urmson. The Google Self-Driving Car Project, 2011.