

Semantics through Time: Semi-supervised Segmentation of Aerial Videos with Iterative Label Propagation

Alina Marcu^{1,2}, Vlad Licaret¹, Dragos Costea^{1,2}, and Marius Leordeanu^{1,2}

¹ University Politehnica of Bucharest, 313 Splaiul Independentei, Bucharest

² Institute of Mathematics of the Romanian Academy, 21 Calea Grivitei, Bucharest
{alina.marcu,vlad.licaret,dragos.costea,marius.leordeanu}@upb.ro

Abstract. Semantic segmentation is a crucial task for robot navigation and safety. However, current supervised methods require a large amount of pixelwise annotations to yield accurate results. Labeling is a tedious and time consuming process that has hampered progress in low altitude UAV applications. This paper makes an important step towards automatic annotation by introducing *SegProp*, a novel iterative flow-based method, with a direct connection to spectral clustering in space and time, to propagate the semantic labels to frames that lack human annotations. The labels are further used in semi-supervised learning scenarios. Motivated by the lack of a large video aerial dataset, we also introduce *Ruralscapes*, a new dataset with high resolution (4K) images and manually-annotated dense labels every 50 frames - the largest of its kind, to the best of our knowledge. Our novel SegProp automatically annotates the remaining unlabeled 98% of frames with an accuracy exceeding 90% (F-measure), significantly outperforming other state-of-the-art label propagation methods. Moreover, when integrating other methods as modules inside SegProp's iterative label propagation loop, we achieve a significant boost over the baseline labels. Finally, we test SegProp in a full semi-supervised setting: we train several state-of-the-art deep neural networks on the SegProp-automatically-labeled training frames and test them on completely novel videos. We convincingly demonstrate, every time, a significant improvement over the supervised scenario.

1 Introduction

While ground vehicles are restricted to movements in 2D, aerial robots are free to navigate in three dimensions. This allows them to capture images of objects from a wide range of scales and angles, with richer views than the ones available in datasets collected on the ground. Unfortunately, this unconstrained movement imposes significant challenges for accurate semantic segmentation, mostly due to the aforementioned variation in object scale and viewpoint. Classic semantic segmentation approaches are focused on ground scenes. More recent work tackled imagery from the limited viewpoints of specialized scenes, such as ground-views of urban environments (from vehicles) and direct overhead views (from orbital

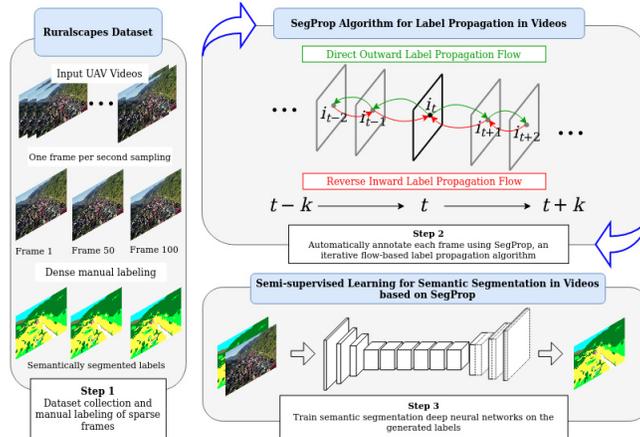


Fig. 1. SegProp: our method for automatic propagation of semantic labels in the context of semi-supervised segmentation in aerial videos. **Step 1.** First, we sample the UAV videos, at regular intervals (e.g. one or two frames per second). The resulting frames are then manually labeled. **Step 2.** We automatically propagate labels to the remaining unlabeled frames using our SegProp algorithm - based on class voting, at the pixel level, according to inward and outward label propagation flows between the current frame and an annotated frame. The propagation flows could be based on optical flow, homography transformation or other propagation method, as shown in experiments. SegProp propagates iteratively the segmentation class voting until convergence, improving performance over iterations. **Step 3.** We then mix all the generated annotations with the ground truth manual labels to train powerful deep networks for semantic segmentation and significantly improve performance in unseen videos.

satellites). Nevertheless, recent advances in aerial robotics allows us to capture previously unexplored viewpoints and diverse environments more easily. Given the current state of technology, in order to evaluate the performance of autonomous systems, the human component is considered a reference. However, the manual segmentation annotations in supervised learning is a laborious and time consuming process. In the context of video segmentation, it is impractical to manually label each frame independently, especially considering there is relatively little change from one to the next. In this context, the ability to perform automatic annotation would be extremely valuable.

SegProp - automated semantic propagation in videos. In this paper we present SegProp (Sec. 2), an iterative flow-based method to propagate, through space and time, the semantic segmentation labels to video frames that lack human annotations. SegProp propagates labels in an iterative fashion, forward and backward in time from annotated frames, by looping several times through the video and accumulating class votes at each iteration. At convergence the majority class wins. From a theoretical point of view, SegProp relates to spectral MAP labeling in graphical models and has convergence and improvement guarantees (Sec. 2.1). In practice, we demonstrate the effectiveness

of SegProp in several ways (Sec. 4). First, we show that SegProp is able to propagate labels to unlabeled frames with an accuracy that outperforms competition by a significant margin (Sec. 4.1). Second, we show that other methods for label propagation could be immediately integrated as modules inside the SegProp propagation loop, with a significant boost in performance (Sec. 4.3). And third, we demonstrate SegProp’s effectiveness in a semi-supervised learning scenario (Sec. 4.2), in which several state-of-the-art deep networks for semantic segmentation are trained on the automatically annotated frames and tested on novel videos, with an important improvement over the supervised case.

Label propagation methods. Our method is not limited to single-object mask propagation [1]. Methods that are most related to ours perform total scene, multi-class label propagation [2]. One such method, also propagates labels between two frames, in the context of ground navigation and low resolution images (320 x 240) [3]. They employ an occlusion-aware algorithm coupled with an uncertainty estimation method. Their approach is less useful in our case, where we have very high resolution images (4k) at a high frame rate (50fps) and dense optical flow can be accurately computed. Earlier works, exploring the idea of propagating ground truth labels using an optical flow based method [4], have shown that it could be useful to treat pseudo-labels differently than the ground truth ones, during training. That idea builds upon other work that addresses occlusion errors [5]. More recent methods for automatic label propagation use a single human annotated frame and extend the label to nearby frames, such as it is the work of Zhu et al. [6], with state-of-the-art results on Cityscapes and KITTI [7]. The main limitation of [6] is that the performance drops drastically when increasing the number of propagated frames, as we confirm in our tests (Fig. 5). Budvytis et al. [8] use semi-supervised learning to improve the intermediate labels and Reza et al. [9] integrate depth and camera pose and formulate the problem as energy minimization in Conditional Random Fields.

Ruralscapes Dataset for Semantic Segmentation in UAV Videos. In this paper we also introduce *Ruralscapes*, the largest high resolution video dataset (20 high quality 4K videos) for aerial semantic segmentation, taken in flight over rural areas in Europe (Sec. 3). We manually annotate a relatively small subset (2%) of frames in a video and use SegProp, our novel iterative label propagation algorithm, to automatically annotate the whole sequence. Given a start and an end frame of a video sequence, SegProp finds pixelwise correspondences between labeled and unlabeled frames, to assign a class for each pixel in the video based on an iterative class voting procedure. In this way we generate large amounts of labeled data (over 50k densely segmented frames) to use in semi-supervised training deep neural networks and show that training on the automatically generated labels, boosts the performance at test time significantly. Our pipeline can be divided into three steps (see Figure 1). The first and most important is the data labeling step. We leverage the advantages of high quality 4K aerial videos, such as small frame-to-frame changes (50 frames per second) and manually annotate a relatively small fraction of frames, sampled at 1 frame per second. We then automatically generate labels for each intermediate frame,

between two labeled ones using SegProp, our proposed algorithm (Section 2), such that the whole video is labeled. In our last step, we use the manually and automatically annotated frames together for semi-supervised training.

Datasets for semantic segmentation in video. Since most work is focused on ground navigation, the largest datasets with real-world scenarios are ground-based. Earlier image-based segmentation datasets, such as Microsoft’s COCO [10], contained rough labels, but the large number of images (123k) and classes (80), made it a very popular choice. Cityscapes [11] was among the first large-scale dataset for ground-level semantic and instance segmentation. Year after year, the datasets increased in volume and task complexity, culminating with Apolloscape [12], which is, to the best of our knowledge, the largest real ground-level dataset. Compared to its predecessors, it also includes longer video shots, not just snippets. It comprises of 74,555 annotated video frames. To help reduce the labeling effort, a depth and flow-based annotation tool is used. Aerscapes [13] is a UAV dataset that contains real-world videos and semantic annotations for each frame and it is closer to what we aim to achieve. Unfortunately, the size of the dataset is rather small, with video snippets ranging from 2 to 125 frames. The most similar dataset to ours is UAVid [14]. The dataset contains set of 4K UAV videos, that captures urban street scenes, with 300 images manually-labeled with 8 classes, compared to our dataset that has 60% more frames, manually-labeled with 12 classes. Since labeling real-world data (especially video) is difficult, a common practice is to use synthetic videos from a simulated environment. Such examples are Playing for Benchmarks [15], for ground-level navigation and the recently released Mid-air [16], for low-altitude navigation. Mid-air has more than 420k training video frames. The diversity of the flight scenarios and classes is reduced - mostly mountain areas with roads - but the availability of multiple seasons and weather conditions is a plus.

Main contributions: **1)** We present SegProp, an iterative semantic label propagation method in video, which outperforms the current state-of-the-art (Sec. 2). **2)** We introduce Ruralscapes, the largest high-res (4K) video dataset for aerial semantic segmentation with 50,835 fully annotated frames and 12 semantic classes (Sec. 3). **3)** SegProp can be easily integrated with other label propagation methods and further improve their initial segmentation results (Sec. 4.3). **4)** We test SegProp in semi-supervised learning scenarios and compare with state-of-the-art deep neural nets for semantic segmentation (Sec. 4.2).

2 SegProp: Semantic Propagation through Time

We now present SegProp, our iterative, voting-based label propagation method (Alg. 2), which takes advantage of the temporal coherence and symmetry present in videos. Before presenting the full method, we first show how labels are propagated between two labeled frames to the intermediate initially unlabeled ones (in one iteration). Let P_k be an intermediate (initially unlabeled) video frame between two (manually) labeled key frames P_i and P_j . We first extract optical flow both forward $F_{i \rightarrow j}$ and backward $F_{j \rightarrow i}$ through time, between subsequent

frames. Then, we use the dense pixel motion trajectories formed by the optical flow and map pixels from the annotated frames towards the unlabeled P_k . Since the optical flow mapping is not bijective (mapping from P_k to P_j could differ from mapping from P_j to P_k), we take both forward and backward mappings into account. Thus, for each pixel in P_k we have 4 correspondence maps that will vote for a certain class: two votes are collected based on the direct outward maps from P_k to its nearby key labeled frames ($P_{k \rightarrow i}, P_{k \rightarrow j}$) and two are based on the reversed inward maps ($P_{i \rightarrow k}, P_{j \rightarrow k}$). Since motion errors are expected to increase with the length of the temporal distance between frames, we weigh these votes with exponential decay, decreasing with the distance (Alg. 1).

Notation. P_k is a 3-dimensional segmentation map, of the same two dimensions as the frame, but with a third dimension corresponding to the class label. Thus, votes for a given class are accumulated on the channel corresponding to that specific class. With a slight notation abuse, by $P_{k \rightarrow i}$ we denote either the flow propagation directed from frame k to frame i as well as the class vote cast by the labeled frame (in this case, i) onto the unlabeled frame k at the corresponding locations in k , according to the flow propagation map $P_{k \rightarrow i}$.

Algorithm 1 Label propagation between labeled frames in one iteration

Input: Two labeled frames P_i and P_j , optical flow maps $F_{i \rightarrow j}$ and $F_{j \rightarrow i}$.

Output: P_k , an intermediate, automatically labeled frame ($i < k < j$).

1) Compute 4 segmentations P_{kn} ($n = \overline{1, 4}$) by casting class votes ($P_{k \rightarrow i}, P_{k \rightarrow j}$) and ($P_{i \rightarrow k}, P_{j \rightarrow k}$) according to direct outward flow and reverse inward flow, respectively.

2) Gather weighted votes from all four P_{kn} for each pixel (x, y) :

$$p_k(x, y, \cdot) = \sum_{n=1}^4 w_n p_{kn}(x, y, \cdot),$$

where $w_n \propto e^{-\lambda * \text{dist}_n(k, q_n)}$ and $\text{dist}_n(k, q_n)$ is the distance between frame k and corresponding labeled frame $q_n \in i, j$. Weights w_n are normalized to sum to 1.

3) Compute final P_k by class majority voting for each pixel.

Iterative SegProp Algorithm. A similar label propagation procedure (as in Alg. 1) could be repeated for several iterations (as in Alg. 2) by considering all frames labelled from previous iterations and cast votes between nearby ones. The intuition is that after the initial voting, we can establish better temporal coherence among neighbouring frames and improve consistency by iteratively propagating class votes between each other. The iterative SegProp (Alg. 2) results in better local consensus, with smoother and more accurate labels. In Sec. 2.1 we also show that SegProp has interesting theoretical properties such as convergence to an improved segmentation objective score.

Integrating other propagation methods into SegProp. We can use SegProp as a meta-procedure on top of other label propagation solutions (such as [6] or homography-based propagation), resulting in further improvement of the initial results, as shown in our experiments (Sec. 4). Segprop could in principle start from any initial solution (soft or hard) and then, at each iteration, replace or augment the class voting with votes from any other label propagation module.

Algorithm 2 Iterative SegProp Algorithm for Label Propagation

- 1) For a given frame k , perform **Step 1** from **Algorithm 1** considering of $2f$ neighbouring frames, symmetrically spaced around k . *Optional: replace or augment the optical flow votes by another propagation procedure.*
 - 2) Gather all available votes for each class, including the ones from frame k 's previous iteration:

$$p_k(x, y, :) = \sum_{n=1}^{4f+1} w_n p_{kn}(x, y, :)$$
 - 3) Return to Step 1 and repeat several iterations, for all frames, until convergence.
 - 4) Return final segmentation by class majority voting: $class_k(x, y) = \max(p_k(x, y, :))$
-

Final segmentation space-time 3D filtering. As final post-processing step we smooth out the segmentation noise as follows: we propagate P_k along optical flow vectors for a number of steps, forward and backward through time, and concatenate the results into a local 3D spatiotemporal voting volume, one per class. We then apply a 3D (2D + time) Gaussian filter kernel to the 3D volume and obtain an average of the votes, one per each class, independently. Then we finally set hard per-pixel classes by class majority voting.

2.1 Mathematical interpretation and properties of SegProp

From a more theoretical point of view, SegProp can be seen as a Maximum A Posteriori (MAP) label inference method for graphical models [17], strongly related to other, more classical iterative optimization techniques for labeling problems with pairwise terms, such as relaxation labeling [18], deterministic and self annealing [19], spectral MAP inference [20] and iterative conditional modes [21]. Conceptually, we could think of the video as a graph of pixels in space and time, with a node for each pixel. Then, each node in the graph can get one class label out of several. A multi-class segmentation solution at the entire video level, could be represented with a single vector \mathbf{p} , with $N \times C$ elements (N - total number of nodes, C - number of possible classes per node). Thus, for a unique pixel i in the video and potential label a , we get a unique index ia . A final hard segmentation could then be expressed as an indicator vector \mathbf{p} , such that $p_{ia} = 1$ if pixel i has class a and $p_{ia} = 0$, otherwise.

We consider the space-time graph edge structure as given by the flow based links between neighbouring pixels as presented in Algorithms 1 and 2. Thus, any two pixels (i, j) connected through an optical flow link that vote for the same class (where pixel i is from one frame and pixel j from another) establish an undirected edge between them. These edges define the structure of the graph, with adjacency matrix \mathbf{M} .

Then, the weighted class voting can be expressed by correctly defining \mathbf{M} , such that $\mathbf{M}_{ia,jb} = w_{ij} = e^{-\lambda * dist(i,j)}$ if and only if (i, j) are connected and class a is the same as class b (class a from frame of pixel i can only vote for the same class in frame of pixel j). One can then show that the iterative voting procedure (Alg. 2) can be written in simplified matrix form as:

$$\mathbf{p}_{ia}^{(t+1)} \leftarrow \sum_{j \in N_i} \mathbf{M}_{ia,ja} \mathbf{p}_{ja}^{(t)} \quad (1)$$

Note that the sum above is exactly the accumulation of votes coming from the neighbouring frames (labeled at the previous iteration). Also note that as the number of votes per node over all classes is constant (equal to the total number of votes), the vector \mathbf{p} remains $L1$ -normalized from one iteration to the next, both at the local level of nodes ($\sum_a \mathbf{p}_{ia} = N_{votes}$) and overall $\sum_{ia} \mathbf{p}_{ia} = N_{pixels} * N_{votes}$. One could immediately observe that Equation 1 above is the power iteration method for computing the principal eigenvector of matrix \mathbf{M} , which must have positive elements since matrix M has positive elements, according to Perron-Frobenius theorem. It also means that the final solution of SegProp should, in principle, always converge to the same solution regardless of initialization as it depends only on \mathbf{M} , which is defined by the propagation flow of labels (e.g. optical flow) and the initial manually labeled frames. Note that those frames should never change their solution and never accumulate votes, a condition that can be easily enforced through the way we set up \mathbf{M} .

It is also well known that the principal eigenvector of \mathbf{M} maximizes a segmentation score $S_L(\mathbf{p}) = \mathbf{p}^T \mathbf{M} \mathbf{p}$, under $L2$ -norm constraints on \mathbf{p} . In other words, SegProp should converge to $\mathbf{p}^* = \operatorname{argmax}_{\mathbf{p}} S(\mathbf{p}) = \sum_{ia} \sum_{jb} \mathbf{M}_{ia,jb} \mathbf{p}_{ia} \mathbf{p}_{jb}$. Now, if we use the definition of our pairwise terms in \mathbf{M} we can also show that: $S_L(\mathbf{p}) = \sum_{j \in N_i} \mathbf{M}_{ia(i),ja(i)} \mathbf{p}_{ia(i)} \mathbf{p}_{ja(i)} = \sum_i N_i(a(i))$, where $N_i(a(i))$ are the number of neighbours of node i , which have the same label $a(i)$ as i . Thus, maximizing the segmentation score has a natural and intuitive meaning: we will find the segmentation \mathbf{p} that encourages connected nodes to have the same label.

SegProp theoretical properties. In summary, we expect SegProp to converge and maximize the quadratic soft-segmentation score with pairwise links $S_L(\mathbf{p}) = \mathbf{p}^T \mathbf{M} \mathbf{p}$, under $L2$ -norm constraints on \mathbf{p} . It should do so regardless of the initialization as it only depends on \mathbf{M} , defined by the propagation flow and the initial manually labeled frames. Initialization should, however, affect the speed of convergence, as also observed in experiments (Sec. 4.1). Since the segmentation \mathbf{p} has constant $L1$ norm, we also expect it to converge to the stationary distribution of the random walk, as defined by the transition adjacency matrix \mathbf{M} . Thus, the solution, which is the principal eigenvector of \mathbf{M} , strongly relates SegProp to spectral clustering [22] and spectral MAP inference [20], a fact that could help us better understand its behaviour in practice.

3 The Ruralscapes Dataset

Manual annotation tool. In order to manually annotate the sampled frames, we designed a user-friendly tool that facilitates drawing the contour of objects (in the form of polygons). For each selected polygon we can assign one of the 12 available classes, which include background regions: forest, land, hill, sky, residential, road or water, and also foreground, countable objects: person, church, haystack, fence and car. In the context of total scene segmentation, we assume

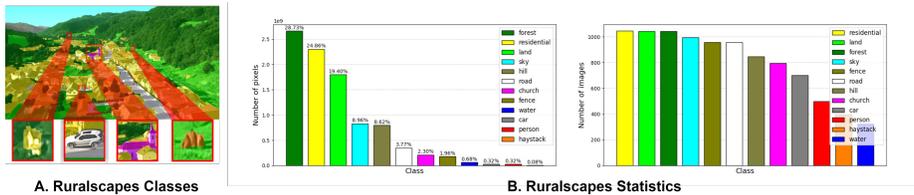


Fig. 2. A. Ruralscapes classes. Labels overlaid over RGB image with detail magnification, offering a good level of detail. Ruralscapes also offers large variation in object scale. **B. Ruralscapes statistics.** (Left) Class pixels’ distribution. Being a rural landscape, the dominant classes are buildings, land and forest. Due to high altitude, smaller classes such as haystack, car and person hold a very small percentage. (Right) Number of labeled images in which each class is present.

that the image needs to be fully segmented (e.g., no ‘other’ class). While there are other annotation tools available ([23], [24]), ours has several novel convenient features for rapid annotation that go beyond simple polygonal annotation.

Our software is suited for high resolution images. Furthermore, it offers support for hybrid contour and point segmentation - the user can alternate between point-based and contour-based segmentation during a single polygon. The most time-saving feature is a ‘send to back’ functionality to copy the border from the already segmented class to the new one being drawn. This tool is mostly useful in cases when smaller objects are on top of bigger ones (such as cars on the road). Instead of delineating the area surrounding the car twice, one can firstly contour the road and on top that polygon, segment the car.

Dataset details and statistics. We have collected 20 high quality 4K videos portraying rural areas. Ruralscapes comprises of various landscapes, different flying scenarios at multiple altitudes and objects across a wide span of scales. The video sequence length varies from 11 seconds up to 2 minutes and 45 seconds. The dataset consists of 17 minutes of drone flight, resulting in a total of 50,835 fully annotated frames with 12 classes. Of those, 1,047 were manually annotated, once every second. To the best of our knowledge, it is the largest dataset for semantic segmentation from real UAV videos. The distribution of classes in terms of occupied area is shown in Figure 2 (B). Background classes such as forest, land and residential are dominant, while smaller ones such as person and haystack are at the opposite spectrum.

Details regarding the annotation process. Labels offer a good level of detail, but, due to the reduced spatial resolution of the small objects, accurate segmentation is difficult, as seen in the sample label from Figure 2 (A). Some classes, such as haystack, are very small by the nature of the dataset, others such as person, also feature close-ups. Manual labeling is a time consuming process. Based on the feedback received from the 21 volunteers that segmented the dataset, it took them on average 45 minutes to label an entire frame. This translates into 846 human hours needed to segment the manually labeled 1,047 frames.

4 Experimental Analysis

We evaluate the performance of our proposed method and compare the results to the current state-of-the-art for label propagation [6]. We also train three widely adopted segmentation networks on the automatically generated segmentation labels and report the results compared to the baseline supervised training.

Dataset split. The whole 20 densely labeled video sequences are divided into 13 training and 7 testing video subsets. We divided the dataset in such a way to be representative enough for the variability of different flying scenarios. We selected the videos in such a manner that the scenes were equally distributed between the training and testing sets. We never have the same exact scene in both train and test. When two videos have similar scenes (w.r.t classes, viewpoint and altitude), we select the longer video for train and the other for test (happened only once). The 7 test videos ($\approx 29.61\%$ of the total frames from the dataset) have 311 manually-labeled frames (used for evaluation metrics) out of a total of 15,051 frames. The 13 training videos have 736 manually-labeled frames out of a total of 35,784 frames that we automatically annotate, starting from the initial manually labeled ones, using SegProp for the semi-supervised learning tests.

For experimental purposes, we sample manually labeled frames every 2 seconds (every 100th frame, starting with the first) from the training set, and term this set *TrainEven*. The remaining manually-labeled frames, the ones at odd seconds marks, form the *TrainOdd* set, which are used, as explained later, to test the performance of label propagation on the training set itself, before semi-supervised deep learning and evaluation on the unseen test videos. We conducted a more detailed analysis of the influence of having larger temporal gaps between labeled frames over the segmentation performance, shown in Figure 5 (B).

4.1 Comparisons to other label propagation methods

We use every pair of consecutive ground truth labels, from *TrainEven*, to populate with segmentation labels the remaining 99 frames in between and evaluate on the center frame (from *TrainOdd*), the one that is maximally distant from both manual labels and for which we have ground truth.

We compare our label propagation results with a state-of-the-art method recently proposed by Zhu et al. [6]. We use their method to similarly propagate ground truth labels. Since Zhu et al.’s method works with only one temporal direction at test time, we extract label estimations first forward and then backwards in time, up to a maximum of 50 frames, thus populating all 99 frames in between while keeping the propagation distance minimal. For a fair comparison, we used the same optical flow as [6], namely FlowNet2 [25]. We test our SegProp (Sec. 2) method against [6] and provide results in Table 1. In Figure 3 we also present some visual comparative results. We also show the effect of initialization on SegProp, when we start it with the solution from [6] vs. initializing with Alg 1. SegProp improves in both cases and converges towards the same solution, but at different speeds. Finally, we apply our 3D filtering as a final step to further remove noisy labels and observe another final jump in F-measure.

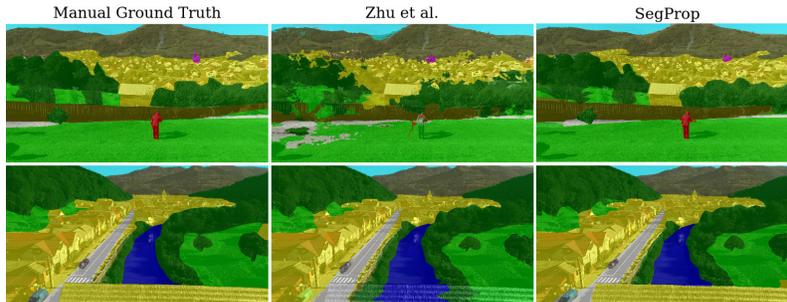


Fig. 3. Qualitative results of our label propagation method. Our iterative SegProp method provides labels that are less noisy and more consistent for larger propagation distances. Also, by looking both forward and backward in time we can better handle occlusion: this is easily visible on the second row in the bottom of the image where forward camera movement obscures a bridge.

We have tested different connectivity between frames for the iterative SegProp (Alg. 2) and settled on the structure which connects the center frame k to frames $[k - 10, k - 5, k, k + 5, k + 10]$. As computational costs increase both with the number of frames included in the set and with the distance between them, we find this to be a good compromise between width (inter-frame distance) and depth (number of iterations we can run). Note that P_{k-10} at iteration 3 will have included propagated votes from P_{k-20} on iteration 2, so width indirectly increases with iteration depth.

4.2 Semi-supervised learning with automatically generated labels

In order to assess the gain brought by the generated labels with SegProp, we train 3 different deep convolutional networks, two of which are widely adopted semantic segmentation models, namely Unet [26] and DeepLabV3+ [27], and a model that has previously shown to yield good segmentation results on similar problems, in UAV flying scenarios [28]. Our approach, however, is agnostic of the chosen architecture and could work with any semantic segmentation method. We chose to train Unet since is the de facto standard for semantic segmentation networks and has been widely applied in many scenarios. Ultimately, our goal is to be able to deploy the model and use it on the UAV. Therefore, we trained two embeddable-hardware compatible deep convolutional networks, namely DeepLabv3+ with a MobileNetv2 [29] backbone and SafeUAV-Net Large.

Supervised baselines. We also trained the same models only on the manually-labeled frames from *TrainEven*. To compensate for the differences in terms of training volume, only for baselines, we apply data augmentation in the form of random rotations, color jittering and random flips, online, during training.

Training details. Models were trained using the same learning setup. Our deep learning framework of choice is Keras with a backend of Tensorflow. We

Table 1. Automatic label propagation comparisons. We measure mean F1-score and mean IOU over all classes. We present Zhu et al. [6] (which has only 1 iteration) vs. SegProp starting either from [6] or from Alg. 1 (our full SegProp Version). An interesting result, which confirms our theoretical expectation is that SegProp (Alg. 2) seems to converge to the same global relaxed solution, regardless of initialization ([6] to Alg. 1), albeit with different convergence rates. The final output depends only on the structure of the graph (defined by the flow links) and the manually labeled frames, as expected.

Methods \ Iterations		Iterations								
			1	2	3	4	5	6	7	+ Filt.
Zhu et al. [6]	mF1	.846	-	-	-	-	-	-	-	-
	mIOU	.747	-	-	-	-	-	-	-	-
SegProp from [6]	mF1	.846	.874	.877	.885	.888	.891	.893	.896	
	mIOU	.747	.785	.790	.801	.805	.810	.813	.818	
SegProp from Alg. 1	mF1	.884	.894	.896	.897	.897	.897	.897	.903	
	mIOU	.801	.817	.819	.821	.821	.821	.821	.829	

use RMSprop optimizer with a learning rate starting from $1e-4$ and decreasing it, no more than five times when optimization reaches a plateau. Training is done using the early stopping paradigm. We monitor the error on the validation set and suspend the training when the loss has not decayed for 10 epochs. The models were trained with RGB frames at a spatial resolution of 2048×1080 px, rescaled from the original 4K resolution (4096×2160 px).

Quantitative results on the testing set are reported in Table 2. We compare our results with reference to the ground truth (manually given labels) from the testing set. The overall score was computed as mean F-measure over the whole classes. Some of the classes were not predicted at all in the supervised baseline case and were marked with .000. It is clear that all methods trained in a semi-supervised fashion on the SegProp generated labels (marked with SP) strongly benefit from the label propagation procedure. The relative performance boost, compared to the supervised case, varies from 3.8% for DeepLabv3+ with MobileNetv2 backbone, 14.5% for Unet, and up to 16.7% for SafeUAV-Net. The results also show that small classes experience a significant boost. The “secret” behind recovering classes that are completely lost lies in the label propagation algorithm that is able to add significantly more evidence for classes that are initially not well represented in the ground truth frames: appearing rarely, being very small or often occluded. The ambiguity for the land, forest and hill classes is reflected in our results. Well represented classes in the dataset such as residential areas and land, yield the best results.

Qualitative results on our testing set are shown in Figure 4. They exhibit good spatial coherency, even though the neural networks process each frame individually. The quality of segmentation is affected by sudden scene geometry changes, cases not well represented in the training videos and and motion blur.

Table 2. Quantitative results after training the neural networks on the generated labels. We report mean F-measure over all videos from the testing set, for each individual class: (1) - land, (2) - forest, (3) - residential, (4) - haystack, (5) - road, (6) - church, (7) - car, (8) - water, (9) - sky, (10) - hill, (11) - person, (12) - fence and the average over the all classes. The best results for each class and each trained model, are bolded. Results clearly show a significant performance boost over the baseline, when training with SegProp. (*) Due to space limitations we abbreviate SegProp with SP in the table.

Methods	SP*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	All
Unet	✗	.681	.497	.834	.000	.000	.000	.000	.000	.967	.000	.000	.000	.248
[26]	✓	.757	.544	.838	.000	.556	.672	.000	.000	.900	.454	.000	.000	.393
DeepLab	✗	.500	.416	.745	.000	.220	.073	.000	.000	.909	.242	.000	.000	.259
v3+ [27]	✓	.570	.452	.776	.022	.369	.122	.007	.000	.926	.272	.004	.043	.297
SafeUAV	✗	.713	.475	.757	.000	.371	.640	.000	.000	.953	.260	.000	.003	.348
Net [28]	✓	.783	.488	.836	.364	.552	.748	.031	.428	.973	.176	.481	.610	.515

4.3 Ablation studies: the effect of the propagation module

Homography propagation module. Even state-of-the-art optical flow is prone to noise. In order to obtain more robust results, we test with the idea of incorporating geometric constraints to improve the class propagation. Thus, we compute two additional class voting maps coming from connected class regions in the labeled frames that are transformed with a homography and placed on the current frame of interest. The homography is estimated in a robust way, with RANSAC, using as correspondences the already computed flow maps between the labeled frames and the current one. Adding the homography based votes to the optical flow votes improves the results (Tab. 3) even from the first propagation iteration. Then, by applying our 3D filtering step on top, we further improve performance. While the homography based voting is clearly superior it is also much more computationally intensive, reason for which we did not include it in the other experiments presented in the paper. Note that voting propagation based on homography is particularly useful for edge preservation, where the CNN-based optical flow generally lacks precision (see Fig. 5 (A)).

Other vote propagation modules. As mentioned in Sec. 2 we could in principle use any label propagation method to bring in more votes. Thus, in the same way we added homography voting to the initial optical flow ones, we also added two more class votes by using the method of Zhu et al. [6] to propagate class labels from the manually labeled frames to each unlabeled one. We weighted the votes with a validated parameter ($w = 0.25$) and observed another additional performance gain (see Tab. 3), even from the first iteration. Since it would have been computationally expensive to re-apply the method in [6] for voting, we have only tested with one iteration. Optical flow voting, while not the most accurate, remains very fast (computed only once at the start) and enables SegProp to achieve a significant boost over iterations.

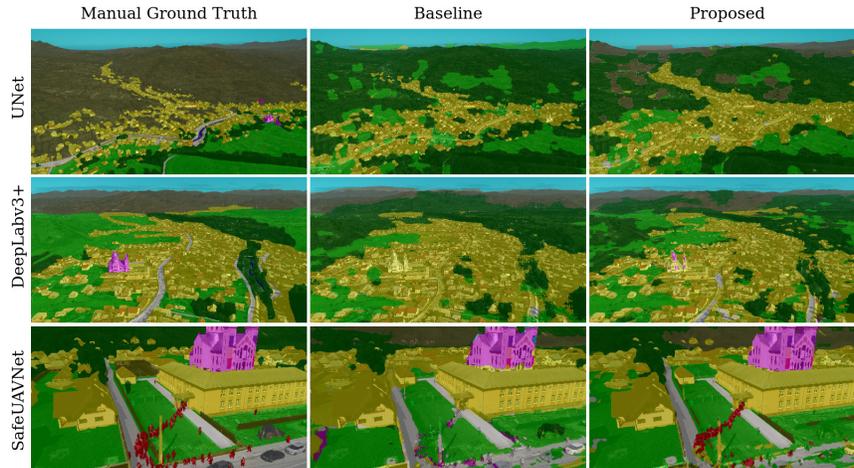


Fig. 4. Qualitative results on the testing set. The results show that our proposed method leads to significantly more accurate segmentation in the semi-supervised scenario rather than the supervised case. SegProp clearly benefits the smaller, not well represented, classes such as person (third row).

Influence of temporal propagation length. We measured the degradation in performance as the propagation temporal length increases, from 25 frames to 150 frames and also compared with Zhu et al. [6] (Fig. 5 (B)). We performed the study on a subset of clips that are annotated every 25 frames (1/3 of the dataset), such that the evaluation can be done at every 25 frames as the propagation period increases. We measure mean F-measure over all classes from the selected videos. Note that our performance degrades slower than that of [6].

Table 3. Ablation studies comparison. We run SegProp including other votes next to our optical flow based mappings, measuring mean F-measure over all classes. For the version with homography voting we also run the final filtering step. The bolded values are the best results.

Method	Iteration	Overall
Zhu et al. [6]	1	.846
SegProp	1	.884
SegProp + Zhu et al. [6]	1	.892
SegProp + Homography	1	.894
SegProp + Homography + Filtering	1	.904

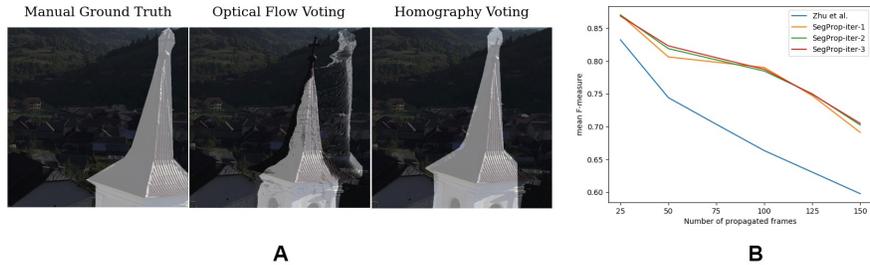


Fig. 5. A. Label propagation example showing typical optical flow voting difficulties. From left to right: RGB frame with manual white label overlaid, flow-based voting, homography-based voting. **B.** The influence of increasingly larger temporal gaps between labeled frames over the segmentation performance (mean F-measure over all classes on a subset of videos labeled with a frequency of 25 frames).

5 Conclusions

We introduced SegProp, an efficient iterative label propagation algorithm for semi-supervised semantic segmentation in aerial videos. We also introduced Ruralscapes, the largest high resolution (4K) dataset for dense semantic segmentation in aerial videos from real UAV flights - which we make publicly available alongside a fast segmentation tool and our label propagation code³, in a bid to help aerial segmentation algorithms. We have demonstrated in extensive experiments that SegProp outperforms other published labeled propagation algorithms, while also being able to work in conjunction with similar methods. Moreover, we have showed in semi-supervised learning experiments that deep neural networks for semantic image segmentation could extensively benefit (by up to 16% increase in F-measure) from the added training labels using the proposed label propagation algorithm. SegProp is fast (it only needs to compute the optical flow once) and flexible, being able to integrate other label propagation methods. Furthermore, it has provable convergence and optimality properties. We believe that our work, introducing a well-needed dataset and algorithm, with strong experimental results, could bring a solid contribution to semantic segmentation in video and UAV research.

Acknowledgements. This work was funded by UEFISCDI, under Projects EEA-RO-2018-0496 and PN-III-P1-1.2-PCCDI-2017-0734. We express our sincere thanks to Aurelian Marcu and The Center for Advanced Laser Technologies (CETAL) for providing access to their GPU computational resources.

³ We make our code, dataset and annotation tool publicly available at: <https://sites.google.com/site/aerialimageunderstanding/semantics-through-time-semi-supervised-segmentation-of-aerial-videos>

References

1. Le, H., Nguyen, V., Yu, C.P., Samaras, D.: Geodesic distance histogram feature for video segmentation. In: Asian Conference on Computer Vision, Springer (2016) 275–290
2. Galasso, F., Cipolla, R., Schiele, B.: Video segmentation with superpixels. In: Asian conference on computer vision, Springer (2012) 760–774
3. Budvytis, I., Sauer, P., Roddick, T., Breen, K., Cipolla, R.: Large scale labelled video data augmentation for semantic segmentation in driving scenarios. In: Proceedings of the IEEE International Conference on Computer Vision. (2017) 230–237
4. Mustikovela, S.K., Yang, M.Y., Rother, C.: Can ground truth label propagation from video help semantic segmentation? In: European Conference on Computer Vision, Springer (2016) 804–820
5. Chen, A.Y., Corso, J.J.: Propagating multi-class pixel labels throughout video frames. In: 2010 Western New York Image Processing Workshop, IEEE (2010) 14–17
6. Zhu, Y., Sapra, K., Reda, F.A., Shih, K.J., Newsam, S., Tao, A., Catanzaro, B.: Improving semantic segmentation via video propagation and label relaxation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2019) 8856–8865
7. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* **32** (2013) 1231–1237
8. Budvytis, I., Badrinarayanan, V., Cipolla, R.: Label propagation in complex video sequences using semi-supervised learning. In: BMVC. Volume 2257. (2010) 2258–2259
9. Reza, M.A., Zheng, H., Georgakis, G., Koščeká, J.: Label propagation in rgb-d video. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE (2017) 4917–4922
10. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision, Springer (2014) 740–755
11. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016)
12. Huang, X., Cheng, X., Geng, Q., Cao, B., Zhou, D., Wang, P., Lin, Y., Yang, R.: The apolloscape dataset for autonomous driving. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. (2018) 954–960
13. Nigam, I., Huang, C., Ramanan, D.: Ensemble knowledge transfer for semantic segmentation. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV), IEEE (2018) 1499–1508
14. Lyu, Y., Vosselman, G., Xia, G., Yilmaz, A., Yang, M.Y.: The uavid dataset for video semantic segmentation. arXiv preprint arXiv:1810.10438 (2018)
15. Richter, S.R., Hayder, Z., Koltun, V.: Playing for benchmarks. In: Proceedings of the IEEE International Conference on Computer Vision. (2017) 2213–2222
16. Fonder, M., Droogenbroeck, M.V.: Mid-air: A multi-modal dataset for extremely low altitude drone flights. In: Conference on Computer Vision and Pattern Recognition Workshop (CVPRW). (2019)

17. Li, S.Z.: Markov random field models in computer vision. In: European conference on computer vision, Springer (1994) 361–370
18. Hummel, R.A., Zucker, S.W.: On the foundations of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (1983) 267–287
19. Rangarajan, A.: Self-annealing and self-annihilation: unifying deterministic annealing and relaxation labeling. *Pattern Recognition* **33** (2000) 635–649
20. Leordeanu, M., Hebert, M.: Efficient map approximation for dense energy functions. In: Proceedings of the 23rd international conference on Machine learning. (2006) 545–552
21. Besag, J.: On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society: Series B (Methodological)* **48** (1986) 259–279
22. Meila, M., Shi, J.: A random walks view of spectral segmentation. In: AISTATS. (2001)
23. Brooks, J.: COCO Annotator. <https://github.com/jsbroks/coco-annotator/> (2019)
24. Russell, B.C., Torralba, A., Murphy, K.P., Freeman, W.T.: Labelme: a database and web-based tool for image annotation. *International journal of computer vision* **77** (2008) 157–173
25. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2017) 2462–2470
26. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention, Springer (2015) 234–241
27. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European conference on computer vision (ECCV). (2018) 801–818
28. Marcu, A., Costea, D., Licaret, V., Pîrvu, M., Slusanschi, E., Leordeanu, M.: Safeuav: Learning to estimate depth and safe landing areas for uavs from synthetic data. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 0–0
29. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2018) 4510–4520