A topological constrained model-based approach to correspondence problems for branched structures

Davide Floriello
University of Canterbury
Christchurch
davide.floriello@pg.canterbury.ac.nz

Tom Botterill
University of Canterbury
Christchurch

Richard Green
University of Canterbury
Christchurch

Abstract—**Finding correspondences between images representing branched or tree structures is a challenging task. Branched structures are often encountered in many disciplines: anatomy, medicine, agronomy and engineering are only few examples of such fields. The difficulties come from the self-similarity of the structure and frequent occlusions and overlappings, which increase the number of ambiguous matchings. The aim of our research is to correspond images showing vine structure for an application in robotics. In this paper we propose a model-based approach to the problem, together with a constraint on adjacent branches, which guarantees a topological correctness of the reconstructed structure. Our work shows that the method we tested can find almost all of the correct correspondences and outperforms a maximum likelihood algorithm in terms of precision and recall.**

Keywords—**Network problems, Trees, Computer vision**

Fig. 1. An illustration of the automated pruning machine.

I. INTRODUCTION

Solving correspondence problems is one of the most important tasks in computer vision. It is a fundamental step in applications such as 3D reconstruction, optical flow estimation or alignment. Many years of research in computer vision produced a large number of general algorithms to solve correspondence problems between images. Among those, we find dense and sparse stereo algorithms, feature-based methods or RANSAC. They have been used for many applications, *e.g.* road scene and building reconstruction [8], hand tracking [1], or reconstruction of anatomical structures [6]. However, there still exist classes of images for which these algorithms show poor performances, *e.g.* [5] or [15]. An example is represented by images showing many self-similar patterns, like plant structures.

Plant images occur in applications of robotics and computer vision to agriculture. Precisely in this framework lies the reason of our research: we are part of a team of researchers that is designing a robot to prune vines automatically. Our machinery, which can be seen in Figure 1, has three cameras taking pictures of the vine in three different positions. It is also endowed with structured light, that gives some depth information. For each image, the 2D structure of the vines is extracted [3]. The extracted structure is the input for our correspondence problem. Subsequently, the 3D structure of the

Fig. 2. An example of two images of the same vine, as caught by the machine. The structured light can be seen in the red spots.

vine is recovered and passed to the AI system of the robot that decides which branches are to be cut. Examples of vine images are shown in Figure 2.

The previous closest works similar to our problem are represented by the papers of Zhang and Kosecka [15] and Dey *et al.* [7]. The former is a generalized RANSAC algorithm, while the latter is an application of a featured-based method. However, neither of these approaches exploit in full

the underlying knowledge of the structures they are dealing with. Moreover, some information regarding the topology of the structure are neglected, partly because of their specific problem. For instance, considering vines, suppose that we have an unambiguous match for a branch and that we want to find a correct correspondence to a cane branching out of it. Then, we can limit our search only to candidates whose reconstruction gives a connected structure. This, not only can eliminate some ambiguities, but even infer correct matches. Moreover, a knowledge of the topology of the plant helps in dealing with interdependent probabilities.

In this paper we propose a model-based method with a topological constraint and compare it to a maximum likelihood procedure to find the largest number of correct correspondences between images representing vines. The model-based method rely on the computation of the probability distribution of the correspondences. We compute the probabilities of matchings being correct, rather than a maximum likelihood matching, through a Gibbs sampling algorithm. Thanks to knowledge of the structure, we can immediately eliminate wrong matches.

The rest of the paper is organized as follows: in Section 2 we describe previous general methods to solve correspondence problems and we consider in particular the works by Zhang and Kosecka [15] and Dey et al. [7]. Section 3 describes our model for the problem, a general framework for finding correspondences and how we estimate the probabilities involved. The experimental evaluation is discussed in Section 4. We conclude with a final discussion in Section 5.

II. Related work

Matching features between different images is one of the most investigated problem in computer vision. Therefore, a large number of methods have been already proposed to solve it, *e.g.* stereo algorithms and feature-based methods. We will focus on specialised methods by Zhang and Kosecka [15] and Dey *et al.* [7] that are the most closely related to our problem. In stereo matching algorithms a 3D model of the scene is built by first finding corresponding pixels in the images and then recovering the 3D depth by triangulation techniques. Stereo methods are generally solved by seeking for an estimation of a disparity map, which describes the relation between corresponding pixels. The process of finding a disparity map is done through the minimization of a functional usually composed by a *data term*, representing a similarity measure between the pixels, and *spatial terms*, which impose some penalizations in order to have smooth solutions. Recent stereo algorithms have modified or additional terms that account for occlusions. In case of half-occluded pixels, modern algorithms allow for the possibility of not matching parts. In other works, such as those by Cho *et al.* [4], Berg *et al.* [2] or Torresani *et al.* [13], terms for the geometrical consistency of the matching parts are considered. Our proposed way to verify the topological correctness of the correspondences is an attempt to generalise the concept of the geometric consistency term used in these works.

Stereo methods rely on assumptions on the structure of the observed scene. Stereo matching approaches are among the most used to solve general correspondence problems, because of their versatility and good performances. However, there are still some applications where stereo algorithms are not totally reliable. Examples of such cases are when the assumptions made on the world are not correct or in presence of depth discontinuities. In the case of vine structures, stereo algorithms perform poorly. One of the problems is in the choice of the prior probability of the matching pixels. Generally, this term is used to force a smooth solution. However, with vine structures, we are in presence of thin objects with many depth discontinuities, which provide an unacceptable disparity map as a solution.

Feature-based methods find a sparse correspondence to estimate a 3D reconstruction of a scene. Some relevant features are extracted from the images and then the matching is achieved by first corresponding the sets of extracted features. Feature-based methods do not necessarily seek for an estimation of a disparity map. The matching of the features is usually done computing suitable distances or objective functions. Featured-based methods perform well in frameworks where the self-similarity of the scene is not high. Moreover, it can happen that some features are detected, but they are not useful for the correspondence. In this setting, we find the work by Dey *et al.* [7]. They use a feature-based method to reconstruct a 3D model of vines. A number of images of vine plants is collected and a dense three dimensional point cloud is generated in such a way that it captures the overall shape of the scene. The main goal, before the reconstruction of the scene, is to classify which of the points obtained belong to grapes, foliage or branches. The authors use a support vector machine to classify points. The information the classifier relies on are mostly given by shape features. Even though the reconstruction is sufficiently accurate for their aim, which is the estimation of crop yield, the structure is inaccurate, because of too ambiguous matchings. Dey *et al.* do not exploit knowledge of the structure, because the reconstruction is not their primal aim and because they work with vines with leaves, so that many parts of the vines are hidden.

Another often used method to find correspondences is RANSAC, or modifications of it. One of the most useful aspect of RANSAC is the possibility of jointly estimating the largest set of mutual compatible correspondences between two or more views. In this framework, we cite the work by Zhang and Kosecka [15]. They are aware of the drawbacks of RANSAC with images representing repetitive structures or with widely separated views. They claim that the failure of RANSAC in these cases is due to the fact that similarity measures are used to find matchings and, with repetitive structures or widely separated views, the chosen descriptors can change dramatically. As a result, the nearest neighbour association rule could no more be appropriate. For this reason, Zhang and Kosecka design a generalized RANSAC algorithm so that it keeps a set of admissible matchings, rather than keeping only the best match. They analyse the relationship between the fraction of points, σ_k, which have a correct correspondence in a set of k matches and k itself. Not surprisingly, they find that the higher is k, the higher is σ_k. Moreover, the nearest neighbour match is the one that has the highest probability of being the correct match, even if this probability may be low. Their work is focused on sampling strategies that allow for sets of matchings, rather than a single match. However, the problem of how to choose the final best set of correspondences still remains. Moreover, it is generally likely that the nearest neighbour match will be chosen if some probabilistic criteria are involved. A limitation is that Zhang and Kosecka do not exploit additional image structure.

In our problem, we adopt a model-based method which uses knowledge of vines to find the correct matchings. We believe that a model-based approach is more useful in our, or similar, framework. Indeed, by defining a suitable model, we can discard a large number of incorrect matches. Moreover, we only use knowledge of variables concerning the vines, thus avoiding assumptions on the scene. We also propose a way to take into account a topological correct reconstruction of the structure. A comparison is obtained by applying a maximum likelihood algorithm to our problem. Our results show that our approach outperforms the maximum likelihood method.

III. PROPOSED MODEL-BASED APPROACH

In this section we describe our proposed model to choose the correct correspondences between images. Let us first introduce some notations and conventions. We present our method for $m = 2$ cameras. The generalization to $m > 2$ is straightforward. Let L and R be the left and right images, $|L|$ and $|R|$ the number of branches seen in L and R, respectively. If $|L| \neq |R|$ there will surely be some branches that cannot be matched. Therefore, at most $n = \min\{|L|, |R|\}$ matches can take place and let $M = \max\{|L|, |R|\}$. The branches in the images are detected as described in [3] and this step is performed by another team of researchers. We parametrise the branches as curves: $l = l(t) \in C([0, 1]; \mathbb{R}^2)$ and $r = r(t) \in C([0, 1]; \mathbb{R}^2)$. Moreover, we assume we can refer to the detected curves with unique indices $i = 1, \ldots, |L|$ and $j = 1, \ldots, |R|$. We indicate with $\omega_{ij} = (l_i, r_j)$ the matching between two generic branches l_i and r_j. In case a branch is not matched, we indicate it with $\omega_{i0} = (l_i, \emptyset)$ or $\omega_{0j} = (\emptyset, r_j)$. We suppose the curves to be simple, *i.e.* they have no self-intersections. We define Ω to be the space of all the sets of matchings; that is $\Omega = \{\boldsymbol{\omega} : \boldsymbol{\omega} = (\omega_{1\sigma(1)}, \ldots, \omega_{M\sigma(M)})\}$, where $\sigma : \{1, \ldots, M\} \rightarrow \{0, 1, \ldots, n\}$ is a surjective function, with the requirement that $\sigma : \{1, \ldots, M\} \rightarrow \{1, \ldots, n\}$ is injective. With this notation, we make explicit the requirement that only a single matching for each branch is possible.

One of the most challenging task, now, is to estimate the probability of each generic matching $P((l, r))$. This is a hard problem, because the matchings can be mutually dependent and we do not know the actual dependencies. However, a simpler task can be solved; namely the estimation of each correspondence given a specific set of matchings. Indeed, computing $P((l, r)|\boldsymbol{\omega})$ is easier than calculating $P((l, r))$, because we can make explicit the dependencies between the matches of a given $\boldsymbol{\omega}$. Specifically, we adopt a Markov Random Field Model for the vines. Therefore, we are supposing that $P((l, r)|\boldsymbol{\omega}) = P((l, r)|(p_L(l), p_R(r))$, where $p_L(l)$ is the parent branch from which l is branching out and analogously for $p_R(r)$ and r. The MRF model applied to vines assures a topologically correct reconstruction: it says that two branches can be corresponded only if their parents are matched. We observe here, that we are assuming that we can exploit a good knowledge of the structure of the vines.

We compute the probability of a matching using only knowledge about the vines and the robot and the goodness of the reconstructed branch implied by a specific correspondence. Each reconstructed branch is seen as a manifold in a three-dimensional space. We compute the probability of a given match by using a vector of attributes. Specifically, we can observe and reconstruct the pointwise thickness of branches,

their length, their curvature and the depth of the resulting reconstruction. We can also exploit other information given by the lasers. We assume that the listed variables are pairwise independent and specify two distributions: one in case a correspondence (l, r) is correct (then we say that (l, r) is an "inlier") and one in case of incorrect correspondence $((l, r)$ is an "outlier"). Therefore, let us indicate with $f_k^I((l, r))$ and $f_k^O((l, r))$ the PDF's of the $k-$th variable for the inlier and outlier distributions. The distributions are learnt from training cases. It is possible that the distributions show differences among different variety of vines. We checked this hypothesis and some differences in the parameters of the distributions can be observed. However, they are small and we decided to keep only one inlier and outlier distribution for each variable. We think this decision results in a broader applicable method. Then, we can define the total PDF of a correct matching as $F_I((l, r)) = \prod_k f_k^I((l, r))$ and, analogously, the total PDF of an incorrect matching as $F_O((l, r)) = \prod_k f_k^O((l, r))$. Moreover, let us define a prior distribution $\pi((l, r))$, that evaluates the prior information of (l, r) being a correct matching. π takes into account the information about the projective transformations between the images, namely, the satisfaction of the epipolar constraint. Therefore, using Bayes' formula, we define the probability of a matching (l, r) given a specific set of matchings as:

$$P((l, r)|\boldsymbol{\omega}) := \frac{\pi((l, r))F_I((l, r)) \cdot \iota}{\pi((l, r))F_I((l, r)) + (1 - \pi((l, r)))F_O((l, r))}, \quad (1)$$

where ι is an indicator variable which equals 1 if (l, r) is compatible with $\boldsymbol{\omega}$ and 0 otherwise. The meaning of (1) is the following: it is the probability of a selected match being correct within the reconstructed structure defined by a given possible set of matchings $\boldsymbol{\omega}$. In other terms, (1) is the posterior probability of a match, given a model. We also define the probability of a set of matching $\boldsymbol{\omega}$ as

$$P(\boldsymbol{\omega}) := \prod_{(l, r) \in \boldsymbol{\omega}} P((l, r)|\boldsymbol{\omega}). \quad (2)$$

Because of the MRF hypothesis we made, the compatibility check is done by considering only the match between the parents of l and r. This could be a simplification, but it allows us to consider many matchings as inadmissible and it is a dependency easy to be checked. Finally we note, that, because of (1) and (2), Ω becomes a discrete probability space.

A. General framework to select the correspondences

In the present subsection we briefly discuss a general criterion to select the correct correspondences. We need a way to compare the sets of matchings. So, let us suppose we have a distance $d_\Omega(\boldsymbol{\omega}_1, \boldsymbol{\omega}_2)$ on Ω, defined componentwise. d_Ω counts the components that differ between $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$. Let us indicate the unknown optimal set of correspondences with $\boldsymbol{\omega}^*$. The $\boldsymbol{\omega}$'s with greater $d_\Omega(\boldsymbol{\omega}, \boldsymbol{\omega}^*)$ are to be rejected, because they imply worse reconstructions of the 3D structure. Therefore, we can treat d_Ω as a loss function for choosing a certain $\boldsymbol{\omega}$. This can be seen as a special case of the pattern recognition problem, as detailed in Vapnik [14]. Then, we seek for the set of correspondences minimizing the expected loss:

$$\hat{\boldsymbol{\omega}} = \operatorname{argmin}_\Omega \mathbb{E}(d_\Omega(\boldsymbol{\omega}, \boldsymbol{\omega}^*)). \quad (3)$$

A large number of methods can be interpreted in this framework. In general, it is not too difficult to see that this criterion suggests to select a correspondence if it maximizes the marginal probability of the matching, $P((l, r))$, as specified by the model. Moreover, this is true even for more general loss functions than distances. The solution provided by this setting is an upper bound: it is the best achievable result, by this class of methods. The goodness of a solution, therefore, relies mostly on the model used for the problem. Suitable choices of the loss functions result in more restrictive criteria; that is, a match is selected only if its probability is greater than a threshold τ. However, by implementing these criteria, we get a solution that is a subset of the result found through (3). To see that this is true, let us suppose that a threshold $\tau > \frac{1}{2}$ has been defined. Then, if we obtain $P((l, r)) \geq \tau$, we also have that it is the maximum, since $P((l, \cdot))$ is a probability measure, for every fixed l. In the following, we consider a threshold $\tau = \frac{1}{2}$, because this allows us to retain only the most reliable matchings. All the correspondences that have a probability less than or equal to $\frac{1}{2}$ will be discarded.

B. Computation of the marginal probabilities and optimal solution

From the discussion at the beginning of the current section, we need a way to compute the marginal probability of a matching. To this purpose, we use a Gibbs sampler algorithm. We refer to Robert and Casella's book [11] for a detailed description of the method. The Gibbs sampler algorithm is very effective in cases like ours, where the full conditionals, that is the quantities $P((l, r)|\omega)$, are known. We first encode a set of correspondences ω as a matrix $W(\omega) \in \mathbb{R}^{|L| \times |R|}$ of Bernoulli random variables. Each component $w_{ij}(\omega)$ is equal to 1 if the match (l_i, r_j) is in the set of correspondences ω. We impose the constraint that the sum of the elements in each row and each column is less than or equal to 1. In this way, we are assured that only at most one match for each branch is possible. Each $w_{ij}(\omega)$ has a probability of being 1 as given by (1). With the described setting, the i−th row $\mathbf{w}_i(\omega)$ of the matrix records the possible matchings to the branch l_i, whereas the j−th column $\mathbf{w}^j(\omega)$ the ones to r_j. We can verify the mutual compatibility of the matchings and discard immediately the sets of correspondences that are not admissible. The initialization of the algorithm is given by a random admissible set of correspondences. Then, at each iteration t of the Gibbs sampler, given a sequence from the step before, $\forall i = 1, \ldots, |L|, \ j = 1, \ldots, |R|$ we generate

$$(l_i, r_j)^t \sim Be(P((l_i, r_j)|\{(l_k, r_h)^t\}_{k,h}, \quad (4)$$
$$\{(l_K, r_H)^{t-1}\}_{K,H}, \text{with } k < i, h < j, K > i, H > j)).$$

In the previous expression, $\{(l_k, r_h)^t\}_{k,h}$ is the set of all matchings, updated to step t, with indices less than i and j, whereas $\{(l_K, r_H)^{t-1}\}_{K,H}$ is the set of all matchings, updated to the $(t-1)$−th step, with indices greater than i and j. In fact, because of the MRF hypothesis, the previous relations could be simplified. Then, we use the obtained sample to estimate the marginal probabilities. Indeed, thanks to the laws of large numbers, we get

$$P((l_i, r_j)) = \lim_{t \to \infty} \frac{1}{t} \sum_{\tau=1}^{t} w_{ij}(\omega^\tau). \quad (5)$$

In all our experiments 5000 samples from the Gibbs sampler were used. We found that this number of iterations of the algorithm allowed us to have a good level of precision and it is suitable for real-time applications. However, the number of iterations could depend on how complicated the vine structure is.

C. A Maximum Likelihood Method

We compared our approach, with the following Maximum Likelihood method. For each pair of stereo images, since we have a list of the detected branches, we generated N random sets ω of admissible correspondences. Then, for each possible match (l_i, r_j) of ω, we computed the likelihood of it being correct, given a model for the k−th variable:

$$L_{ij}^k = L^k(l_i, r_j) = f_k^I((l_i, r_j)|(p_L(l_i), p_R(r_j)). \quad (6)$$

Note that we used the same inlier distributions we specified for the model-based methods. Therefore, we defined the likelihood of a match as

$$L_{ij} = \prod_k L^k(l_i, r_j) \quad (7)$$

and, finally, the likelihood of a set of matchings ω as

$$L(\omega) = \prod_{ij} L_{ij}. \quad (8)$$

Then, we kept the set ω with the highest likelihood among those generated.

IV. EXPERIMENTAL RESULTS

In this section we present experimental results of our approach to the correspondence problem with branched structures and compare our method to the maximum likelihood algorithm explained above.

The comparisons were conducted on the same sets of images. The cameras of our machinery take 30 frames per second, for a total of 476 images per camera. The images collected show three different variety of vines: Müller Thurgau, Sauvignon Blanc and Riesling. The vines show different levels of complexity in their structure. The correctness of the obtained optimal set of matching is assessed by comparing them with manually labelled correspondences.

In the experiments, we used all of the three calibrated cameras of the robot, which we label as L (Left), R (Right) and T (top). The pairwise relative pose of the cameras is known. We decided to find correspondences between the right and the top cameras if the orientation of a branch is mainly horizontal. Otherwise, the matchings are sought between the right and the left cameras. This allows us to restrict the number of possible candidates for each matching and to avoid possible degenerate configurations with the epipolar lines.

We tested two slightly different versions of our method: the one we described above, with a threshold $\tau = \frac{1}{2}$ for the probability (in the following referred to as "thresholded method") and the same method, but without thresholds, so that we kept the matches maximizing the marginal probability (in the following referred to as "max-probability method").

In Figure 3 an example of some correspondences obtained with our method can be seen. The picture shows the reprojection of

the computed structure on the top camera. Even though some errors are visible, almost every branch is correctly matched. Beside the intrinsic difficulty in the scenes we dealt with, we think there are some other aspects that further complicate the problem. The structure extraction step is one of those. In some cases it happens that the structure is not properly identified. For example, because of many overlappings or the presence of obstacles, some branches are not processed. In some other cases they can be joined, thus giving us a wrong structure. As an example, let us consider Figure 4. The big post in the image forces a wrong detection in the vine structure: it splits the long branch at the bottom of the image into two different branches and only one part is matched. In Figure 5, another wrong detection can be observed. The metal part on the wire is erroneously identified as a branch. Finally, another wrong detection can be observed in Figure 6. It can be seen that the terminal part of the branch is wrongly reconstructed. We think this happened in the structure detection step, because what is obtained from that step is what is used for the correspondences. In case of Figure 6, we can presume that, since the ending part of that branch is thin and not very distinguished from the background, it has not been detected.

Beside the discussed drawbacks, both the model-based methods performed well in our experiments. In the case of the thresholded method if, for a branch, the probability it is matched is never above $\frac{1}{2}$, it is not assigned. This is done because, in general, we want parts to be matched if they show enough evidence for the correspondence. Moreover, provided the structure is well detected, our method can cope well with occlusions and overlappings between branches, even in case of complex vines. Considering all the experiments, the thresholded method is better than the others in terms of precision. The max-probability method showed good performances as well, but it is more prone to false positives. The max-probability method always tries to match branches and the event "no match" is not considered. From this, we can deduce that some false positive, though maximizing the marginal probability, do not reach the threshold.

As it can be seen in Figure 7, the maximum likelihood algorithm performed worse in almost all cases. We tested it with different number of iterations. In the case of experiments with Müller Thurgau, the results were comparable to the model-based methods. We think this is due to the less complex structure shown in the images with that specific variety of vine. In the case with the number of iterations of the maximum likelihood method $N = 5000$, the two methods took about the same time. For higher N's, the maximum likelihood algorithm became slower then the model-based methods. We think that the maximum likelihood algorithm performed worse because, for some of the considered variables, the inlier and outlier distributions are very close.

We think that both the model-based approach and the maximum likelihood algorithm suffered from the structure extraction step. The relations between parents and children branches are corrupted with noise and sometimes they are not observable. In these cases, it can happen that the only available constraint for the reconstruction of the 3D vine structure is the uniqueness of matchings.
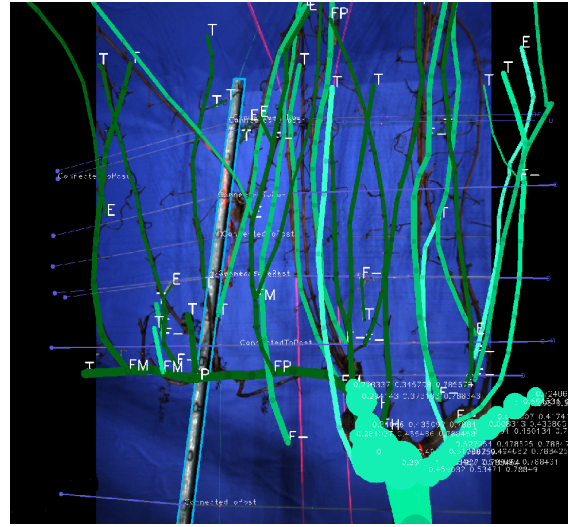

Fig. 3. In this image an example of selected correspondences, using our model-based method is shown. All the shown reprojections are onto the top camera. The vine is almost all reconstructed, even if some errors are still present.
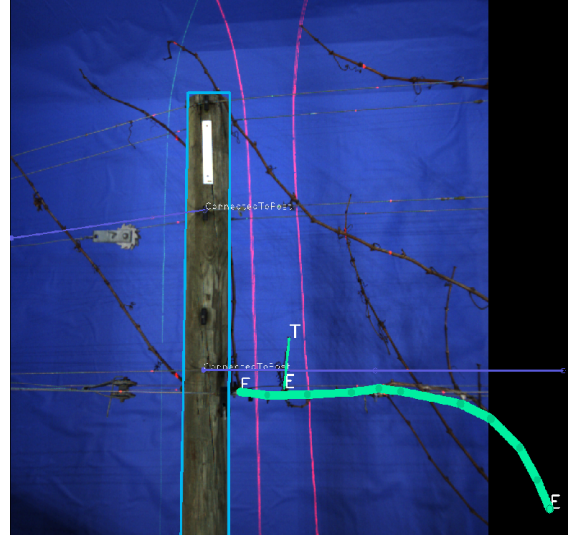

Fig. 4. In this image it can be seen how the overlappings between obstacles and vines can influence the structure detection step. The post forces the branch at the bottom to be split into two separate branches.

V. CONCLUSIONS AND FUTURE DEVELOPMENTS

We applied a model-based approach, with a topological constraint to solve the correspondence problem in the case of images showing branched structures. Our method outperformed another approach based on likelihood maximization. We believe that the model-based approach is one of the most suitable for this kind of problems and, especially, in our application. It allows us to incorporate knowledge of the structure, without any other assumption on the scene, that could not hold in general. A well-designed model can be a better way to eliminate ambiguities in matchings, rather than making additional assumptions on the correspondences.

As a future development, we aim at giving a distribution to the topology of the branches. We think that this direction of
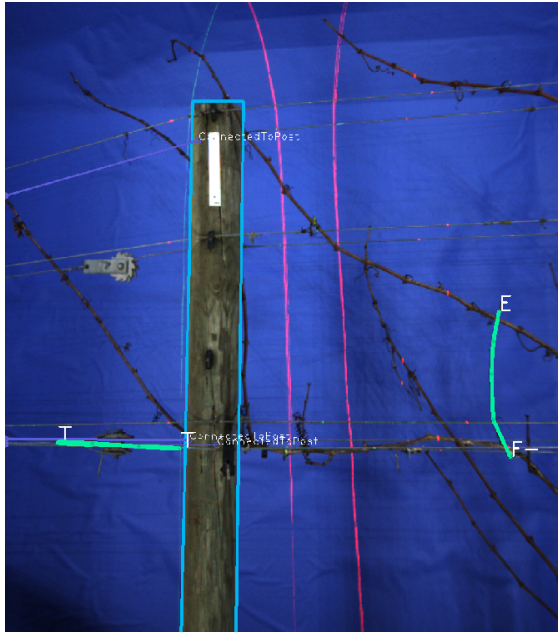
Fig. 5. A false positive correspondence. In this image, a wire with a steel part is wrongly classified as branch and matched.
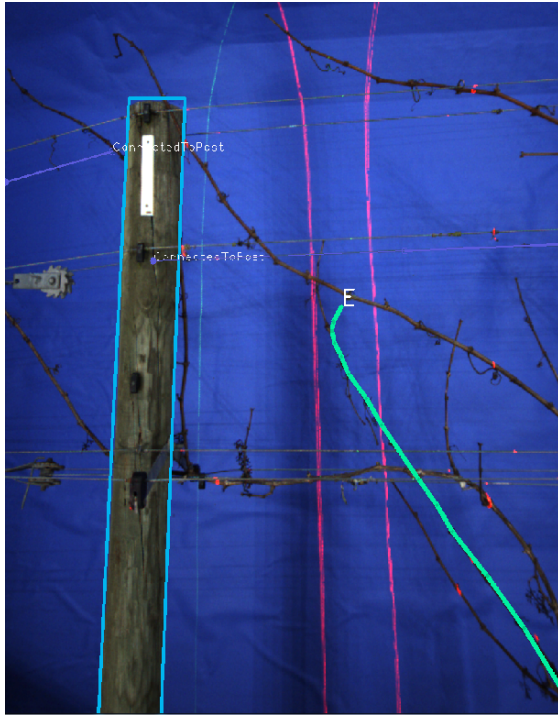

Fig. 6. Another case of misdetection in the 2D structure of the vine. The terminal part of the branch, is wrongly detected and this affects the reconstruction and therefore the reprojection of the branch.

research will bring a relevant contribution to a proper reconstruction of the vines. Indeed, we suppose that the correctly matched branches imply a smoother and less bumpy reconstruction, contrarily to the incorrect correspondences. We will investigate this claim and exploit the obtained informations.

REFERENCES

[1] L. Ballan, A. Taneja, J. Gall, L. Gool, M. Pollefeys, (2012), *Motion Capture of Hands in Action Using Discriminative Salient Points*, in Proceedings ECCV 2012, 640-653.

[2] A.C. Berg, T.L. Berg, J. Malik, (2005), *Shape matching and object recognition using low distortion correspondences*, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005.

[3] T. Botterill, R. Green, S. Mills, (2013), *Finding a vine's structure by bottom-up parsing of cane edges*, IVCNZ 2013, 112-117.

[4] M. Cho, K. Alahari, J. Ponce, (2013), *Learning Graphs to Match*, Proceedings of the IEEE Interational Conference on Computer Vision.

[5] O. Chum, J. Matas, (2005), *Matching with PROSAC - Progressive Sample Consensus*, Conference on Computer Vision and Pattern Recognition, Vol. 1, 220-226.

[6] D. Delaere, C. Smets, P. Suetens, G. Marchal, (1991) *Knowledge-based system for the three dimensional reconstruction of the blood vessels from two angiographic projections*, Journal of medical and biological engineering and computing, Vol. 29, No. 6.

[7] D. Dey, L. Mummert, R. Sukthankar, (2012), *Classification of plant structures from uncalibrated image sequences*, IEEE Workshop on the Applications of Computer Vision.

[8] D. Gallup, J.-M. Frahm, M.Pollefeys, (2010), *Piecewise planar and non-planar stereo for urban scene reconstruction*, in Proceedings CVPR 2010, 1418-1425.

[9] R. Hartley, A. Zisserman, (2004), Multiple View Geometry in Computer Vision, Cambridge University Press, Cambridge.

[10] A. N. J. Hast, A. Marchetti, (2013), *Optimal RANSAC - Towards a Repeatable Algorithm for Finding the Optimal Set*, Journal of WSCG, Vol. 21, No. 1, 21 - 30.

[11] C. P. Robert, G. Casella, (2005), *Monte Carlo Statistical Methods*, Springer-Verlag New York, Inc., Secaucus, NJ, USA.

[12] P. H. S. Torr, A. Zisserman, (2000), *MLESAC: A new robust estimator with application to estimating image geometry*, Computer Vision and Image Understanding, Vol. 78, No. 1, 138 - 156.

[13] L. Torresani, V. Kolmogorov, R. Carsten, (2008), *Feature Correspondence Via Graph Matching: Models and Global Optimization*, Proceedings of the 10th European Conference on Computer Vision, 596-609.

[14] V. N. Vapnik, (1998), *Statistical Learning Theory*, John Wiley & Sons Inc., Hoboken, NJ, USA.

[15] W. Zhang, J. Kosecka, (2006), *Generalized RANSAC Framework for Relaxed Correspondence Problems*, Third International Symposium on 3D Data Processing, Visualization, and Transmission, 854-860.
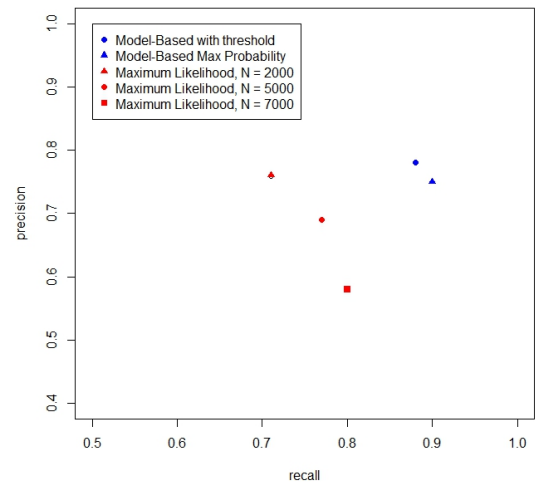
Fig. 7. Effects of changing N in the Maximum Likelihood method.