Layer-finding in Radar Echograms using Probabilistic Graphical Models

David J. Crandall Geoffrey C. Fox
School of Informatics and Computing
Indiana University
{*djcran,gcf*}*@indiana.edu*

John D. Paden
Center for Remote Sensing of Ice Sheets
The University of Kansas
paden@ku.edu

Abstract

Ground-penetrating radar systems are useful for a variety scientific studies, including monitoring changes to the polar ice sheets that may give clues to climate change. A key step in analyzing radar echograms is to identify boundaries between layers of material (such as air, ice, rock, etc.). In this paper, we propose an automated technique for identifying these boundaries, posing this as an inference problem on a probabilistic graphical model. We show how to learn model parameters from labeled training data and how to perform inference efficiently, as well as how additional sources of evidence, such as feedback from a human operator, can be naturally incorporated. We evaluate the approach on 827 echograms of the Antarctic ice sheets, measuring error with respect to hand-labeled ground truth.

1. Introduction

In recent years, scientists in a wide range of disciplines have used ground-penetrating radar systems to study subterranean structures, including searching for landmines and other buried objects [8], characterizing the internal structure of planetary bodies [3], and monitoring changes in polar ice sheets on Earth [1]. These systems are often mounted on airplanes or spacecraft, producing radar echograms that give a cross-section of underground structures along the vehicle's path.

Figure 1 shows a sample echogram produced by the Multichannel Radar Depth Sounder instrument onboard an aircraft flown over part of Antarctica [1]. An echogram is an image in which the horizontal axis is distance along the flight path, the vertical distance represents depth, and pixel intensity represents degree of radar signal return. The dark line near the top of the echogram is the boundary between the air and the ice sheet, while the more irregular lower line is the boundary between the bottom of the ice and the terrain underneath. This basal terrain is quite varied, and may be smooth or mountainous, wet or dry, and made of till, sand, soil and/or rock. This large variability together with the uncertainty inherent of the radar sensing process (e.g. unpredictable signal scatter and reflection patterns) causes echograms to be noisy and difficult to interpret. Human experts thus analyze these echograms, marking the positions of layer boundaries by hand, and then subsequent processing produces derivative data useful for scientists. Unfortunately, manual layer-finding is tedious and labor-intensive, especially as ever greater quantities of echograms are produced.

In this paper we present an automatic technique for finding layer boundaries in noisy echogram images. Our approach poses layer-finding as an inference problem on a probabilistic graphical model. The model incorporates several types of evidence and constraints, including that layer boundaries should lie along areas of high image contrast and that layer boundaries should be continuous and not intersect. An advantage of this probabilistic approach is robustness to noise, because it explictly models uncertainty and combines evidence from multiple sources before making decisions. We learn the model parameters from labeled training data and show how to perform inference efficiently for new images. We test our approach on a large echogram dataset, demonstrating an improvement over several baselines. We also show how human feedback can be added to the model to further improve its accuracy.

2. Related work

Work on subsurface imaging (including both hardware systems and data processing) is too vast to review here; see [13] for an extensive treatment. We are not aware of work studying the specific problem of layer-finding in echogram images, but two papers are particularly related to our work at least in spirit. Frigui *et al* [8] use graphical models on echograms as we do, but their task is to detect landmines and their models are thus quite different. Ferro & Bruzzone [7] identify subterranean features in Martian radar data using an itera-

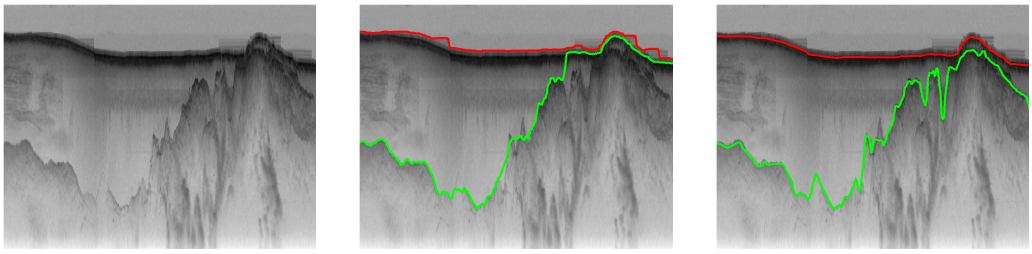

Figure 1. Sample radar echogram of part of the Antarctic ice sheet: (left) raw echogram, (center) layer boundaries found with our approach, and (right) ground truth. The upper (red) boundary is between air and ice layers, and the lower (green) boundary is between ice and terrain.

tive region-growing approach and pre-processing steps. Our approach is distinct in that we avoid intermediate thresholding and local search by solving a single (albeit approximate) inference problem on a graphical model.

General-purpose image segmentation has also been studied extensively. Active contour models [9] are related to our work in that they combine image features with continuity constraints, but require good initialization. Hough transforms [4] cannot be applied to echograms because layer boundaries do not obey a parametric model. Our approach is related to Markov Random Field-based image segmentation [5], and in particular to the tiered labeling approach of [6]. We develop an alternative model that captures unique assumptions of layer-finding in echogram images, explained below.

3. A graphical model for layer-finding

Given a radar echogram I of size $m \times n$ pixels, we wish to find the location of each of K layer boundaries in each column of I. More formally, let l_k^i denote a random variable corresponding to the row index of the k-th layer boundary in the i-th column of the image. Since the boundary could occur at any pixel within the column, there are m possible discrete values that could be assigned to each l_i^k. Let $L_i = (l_i^1, ..., l_i^K)$ denote all of the partition labels for the i-th image column, and $L = (L_1, ..., L_n)$ denote a labeling of the entire image.

We pose layer-finding as an inference problem on a statistical graphical model. Our goal is to find a choice of L (denoted L^*) that maximizes the probability of the labels given the image data, $P(L|I)$. From Bayes' Law,

$$L^* = \arg\max_L P(L|I) = \arg\max_L P(I|L)P(L). \quad (1)$$

This Bayesian factorization has a very natural interpretation in our problem: the likelihood $P(I|L)$ measures how well a given labeling agrees with the observed image data in I (capturing, for example, how well the layer boundaries fall along edges in the image) while $P(L)$ captures prior constraints on the labeling, such as that boundaries should be continuous and should not cross.

The space of possible labelings in the maximization in equation (1) is enormous ($O(n^{m^K})$), so maximizing it is not possible in general. Fortunately the prior and likelihood terms can be factored using reasonable assumptions about echograms. We assume that $P(I|L)$ can be factored over the columns of the image,

$$P(I|L) = \prod_{i \in [1,n]} P(I_i|L_i),$$

where I_i refers to the i-th image column. We define a template model T_k of fixed size $1 \times s$ (we use $s = 11$ pixels) for the vertical profile of each type of layer boundary k. For each pixel p in the template, we estimate a mean μ_p^k and variance σ_p^k on greyscale intensity assuming that the template is centered at the location of the boundary. We also estimate a background model for non-boundary pixels, with parameters μ_0 and σ_0. Then $P(I_i|L_i)$ is given by,

$$\prod_{p \in I_i} P(I(p)|\mu_0, \sigma_0) \prod_{k \in [1,K]} \prod_{p \in T_k} \frac{P(I(p + l_i^k)|\mu_p^k, \sigma_p^k)}{P(I(p + l_i^k)|\mu_0, \sigma_0)},$$

which can be computed efficiently using convolutions. This factorization assumes that image data across columns is independent given layer locations, and that layer boundary templates do not intersect; neither of these are strictly true but we have found that the model still yields reasonable results in practice.

We design the prior $P(L)$ to encourage continuous boundaries that do not intersect. We make a Markovian assumption that the prior can be factored as a product over both neighboring columns within the same layer and over neighboring layers within the same column,

$$P(L) = \prod_{i \in [2,n]} \prod_{k \in [1,K]} P(l_i^k|l_{i-1}^k) P(l_i^k|l_i^{k-1}).$$

The two probabilities in this factorization serve different purposes. $P(l_i^k|l_{i-1}^k)$ is the pairwise conditional probability distribution on neighboring pixels within the

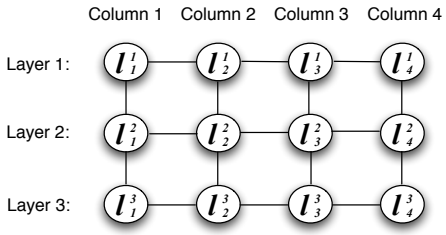

Figure 2. Graphical model for finding 3 layers in a 4-column image. Each l_i^k is a hidden variable, horizontal edges are $P(l_i^k|l_{i-1}^k)$, and vertical edges are $P(l_i^k|l_i^{k-1})$. Not shown is the observed image data connected to each hidden node.

same layer boundary, which we model as a Gaussian with zero mean and variance σ_k. This distribution encourages neighboring pixels to have similar vertical labels, thus encouraging layer boundaries to be continuous lines. The other term, $P(l_i^k|l_i^{k-1})$, is the pairwise conditional probability distribution on neighboring layers of the same column. To prevent label boundaries from crossing, we use a repulsive probability distribution here that is 0 if $l_i^k < l_i^{k-1}$ and is uniform otherwise.

We use labeled training data to learn the parameters of the MRF. For the priors, we model $P(l_i^k|P_{i-1}^k)$ for each k as a zero-mean Gaussian distribution, learning its variance σ_k by maximum-likelihood estimation from the layer boundaries in the training data. We similarly use maximum-likelihood estimation for the parameters in the templates of the likelihood functions.

3.1. Inference

By plugging these likelihood and prior factorizations into equation (1), we see that we need to maximize

$$\prod_{i\in[1,n]} P(I_i|l_i) \prod_{i\in[2,n]} \prod_{k\in[1,K]} P(l_i^k|l_{i-1}^k)P(l_i^k|l_i^{k-1}). \tag{2}$$

This is a first-order Markov Random Field (MRF) model [10], with observed variables consisting of image data, hidden variables consisting of the unknown row labels for each column and layer (l_i^k), and pairwise probability functions between neighboring hidden variables. Figure 2 presents an illustration of this MRF. Unfortunately, exact inference on an MRF is NP hard in general. Approximate algorithms yield good results for some MRFs [2, 11], but are relatively slow and some (like loopy BP) are not even guaranteed to converge.

Here we take an alternative approach that breaks the MRF into a set of non-loopy graphs, and then greedily performs exact inference on each of these graphs in sequence. This is sensible because the only constraints between rows of the MRF in Figure 2 are the repulsion factors that prevent boundaries from crossing. We thus perform inference on each row of the MRF in sequence, fixing the labels found for the row above it and then solving only for the hidden variables in that row. The graph in each of these subproblems is a chain, or a Hidden Markov Model (HMM).

We perform inference on each of these HMMs using the Viterbi algorithm [12]. Viterbi uses dynamic programming to compute a max-marginal distribution for each variable based on the distribution from the previous variable in the chain, somewhat similar to Dijkstra's algorithm. The usual implementation of Viterbi takes time quadratic in the number of possible labels per variable, which in our case is m (the number of rows in the image). To speed this up, we use the linear-time generalized distance transform of [5] to compute the max-marginal distributions efficiently; this is possible because of our assumption that the prior $P(l_i^k|l_{i-1}^k)$ is a Gaussian. The running time for each HMM is thus $O(mn)$, yielding an overall running time of $O(Kmn)$.

Table 1. Evaluation of our technique and baselines, in terms of columnwise mean and mean-squared error.

(a) Comparing our approach with baselines:

	Air-ice boundary		Ice-terrain boundary	
	Mean Err	Mean SE	Mean Err	Mean SE
Fixed	32.1	5326.7	65.0	17327.6
AppearOnly	20.9	2463.1	47.8	10580.0
Our approach	18.1	1879.6	32.3	6044.6

(b) Results with user-provided constraint points:

Ours, 1 pt	14.3	993.1	15.0	1215.2
Ours, 2 pts	13.8	972.8	12.6	871.0
Ours, 3 pts	13.5	965.3	10.1	535.4

4. Experimental results

We tested our layer-finder on radar echograms from the NASA Operation Ice Bridge program in Fall 2009. This data was collected with the airbourne Multichannel Coherent Radar Depth Sounder system [1]. We used a total of 827 echograms, each with a resolution of 700×900 pixels (which corresponds to about 30km of data on the x-axis, and ice thickness of about 0 to 4 km on the y-axis). For these images we also have human-labeled layer boundaries, which we use as ground truth. This data is publicly available online at `https://www.cresis.ku.edu/data`, and we plan to also release our source code upon paper acceptance.

To evaluate our system, we split the data into training and test image sets (413 training images, 414 test images) and learned the model parameters from the training images. There are two layers of interest in

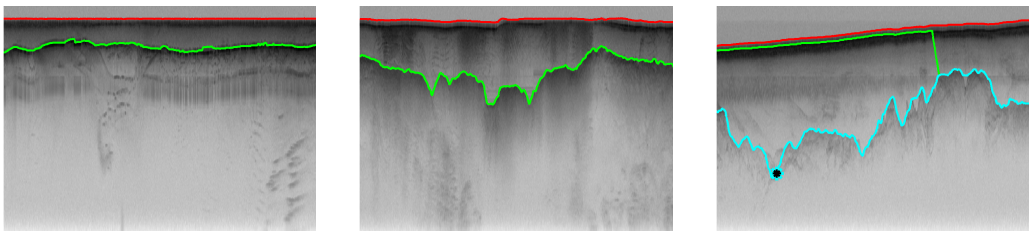

Figure 3. The automatic algorithm found correct layers in the left and center images, but in the right image found an incorrect terrain-ice boundary (green). A user provided a single constraint point (black/cyan asterisk) and layer-finding was re-run, giving a correct result (cyan line).

this dataset (the ice-terrain and ice-air boundaries), so K=2. We then ran inference on each of the test images and calculated the error with respect to human-labeled ground truth, as shown in Table 1(a). The median error is 18.1 pixels per column for the air-ice boundary, and 32.3 pixels for the ice-terrain boundary. To give some context, we compare our results to two simple baselines: Fixed always chooses a straight horizontal line for each layer at the position of its mean location in the training images, while AppearOnly uses a uniform prior in the MRF, thus choosing the layer boundaries by maximizing $P(I_i|L_i)$ for each column i independently. Our approach performs significantly better than either baseline. Our algorithm is very fast, taking an average of about 0.23 seconds per image on a 3.0GHz machine. The left two panes of Fig. 3 show some sample results.

Additional evidence can be easily incorporated into our probabilistic formulation. For example, actual ground-truth data (e.g. from ice bores) may be available for some locations, or a human operator might provide feedback by marking true layer boundaries for a sparse set of points. Either source of evidence can be incorporated by setting $P(I_i|L_i)$ for the corresponding column i such that the probability is 1 for the known label and 0 for all others. Inference is then performed as before.

We tested the effectiveness of this technique by simulating feedback from a human operator. For each image, we computed initial layer boundaries using the MRF as before. Then we compared those boundaries to ground truth and found the column at which the error was highest. We simulated user feedback by constraining the boundary to pass through the correct row at that column, by changing the likelihood function as explained above. We repeated inference to find a new solution given this constraint, and re-computed the error with respect to ground truth. Table 1(b) presents the results, showing that adding just a single correct constraint decreases mean error significantly (from 32.3 to 15.0 pixels), while adding additional constraints further decreases the error. This experiment suggests that we could increase accuracy significantly with a minimal amount of human interaction, in which a user clicks a single point on each layer boundary. Figure 3 shows a sample output that was corrected by human feedback.

5. Conclusion

We presented a technique for automatic layer finding in radar echograms using statistical graphical models. These models can robustly handle noisy echograms because they explicitly reason about uncertainty and avoid making hard decisions (like thresholding) until after all evidence has been combined together. In future work we plan to explore alternative inference strategies for this MRF, and to apply our work to other sources of radar echograms with larger numbers of layers.

References

[1] C. Allen et al. Antarctic ice depth sounding radar instrumentation for the NASA DC-8. *IEEE Aero. and Elec. Sys. Magazine*, To appear.

[2] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *PAMI*, 2001.

[3] G. Bruzzone et al. Subsurface radar sounding of the Jovian moon Ganymede. *Proc. IEEE*, 99:835–857, 2011.

[4] R. Duda and P. Hart. Use of the Hough transform to detect lines and curves in pictures. *Comm. ACM*, 1972.

[5] P. Felzenszwalb and D. Huttenlocher. Efficient belief propagation for early vision. *IJCV*, 70(1):41–54, 2006.

[6] P. Felzenszwalb and O. Veksler. Tiered scene labeling with dynamic programming. In *CVPR*, 2010.

[7] A. Ferro and L. Bruzzone. A novel approach to the automatic detection of subsurface features in planetary radar sounder signals. In *IEEE Geo & Rem Sens. Symp.*, 2011.

[8] H. Frigui, K. Ho, and P. Gader. Real-time landmine detection with ground-penetrating radar using discriminative hidden markov models. *J. Appl. Sig. Proc.*, 2005.

[9] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *IJCV*, 1(4):321–331, 1988.

[10] R. Kindermann and J. Snell. *Markov Random Fields and their Applications*. AMS, 1980.

[11] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.

[12] L. Rabiner. A tutorial on HMMs and selected applications in speech recognition. In *Proc. IEEE*, 1989.

[13] A. Turk, K. Hocaoglu, and A. Vertily. *Subsurface Sensing*. Wiley, 2011.