

Automatic labeling of EEG electrodes using combinatorial optimization

Mickaël Péchaud*, Renaud Keriven*, Théo Papadopoulo*, Jean-Michel Badier†

*ENPC, ENS, INRIA, Odyssee Project Team, Paris, France

mickael.pechaud@ens.fr, keriven@certis.enpc.fr, papadopoulo@sophia.inria.fr

†INSERM, U751, Marseille, France

Abstract—An important issue in electroencephalography (EEG) experiments is to measure accurately the three dimensional (3D) positions of electrodes. We propose a system where these positions are automatically estimated from several images using computer vision techniques. Yet, only a set of undifferentiated points are recovered this way and remains the problem of labeling them, i.e. of finding which electrode corresponds to each point. This paper proposes a fast and robust solution to this latter problem based on combinatorial optimization. We design a specific energy that we minimize with a modified version of the Loopy Belief Propagation algorithm. Experiments on real data show that, with our method, a manual labeling of two or three electrodes only is sufficient to get the complete labeling of a 64 electrodes cap in less than 10 seconds. However, it is shown to be robust to missing electrodes in the reconstructed data.

I. INTRODUCTION

Electroencephalography (EEG) is a widely used method for both clinical and research purposes. Conventionally, EEG readings were directly used to investigate brain activity from the evolution of the topographies on the scalp. Nowadays, it is also possible to reconstruct the brain sources that gave rise to such measurements, solving a so-called inverse problem. To this purpose, it is necessary to find the electrode positions and to relate them to the head geometry recovered from an anatomic MRI. Current techniques to do so are slow, tedious, error prone (they require to acquire each of the electrodes in a given order with a device providing 3D coordinates[1]) and/or quite expensive (a specialized system of cameras is used to track and label the electrodes[2]). Our goal is to provide a cheap and easy system for electrode localization based on computer vision techniques.

In modern EEG systems, the electrodes (64, 128 or even 256) are organized on a cap that is placed on the head. Our system takes as inputs multiple pictures of the head wearing the cap from various positions. As a preliminary step, electrodes are localized and their 3D positions are computed from the images by self-calibration (a technique that recovers the cameras' positions from the image information [3]) and triangulation. These are standard techniques that can provide 3D point coordinates with a quite good accuracy. There remains the problem of electrode identification which labels each 3D position with the name of the corresponding electrode. Finding a solution to this last problem is the focus of this paper. Note, that a good labeling software can also improve current systems by removing acquisition constraints (such as the recording of the electrodes in a given order) and

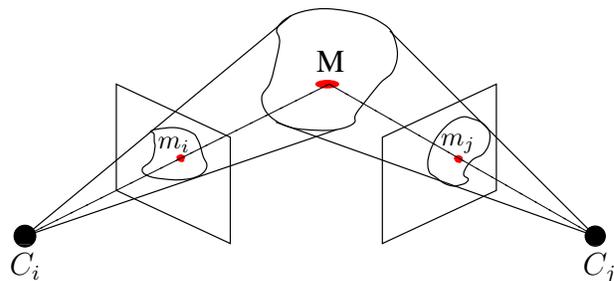


Fig. 1. Projection of an object onto two cameras' planes. The 3D position of point M can be recovered from its two projections m_i and m_j

by providing better user interfaces.

We propose a method that recovers this labeling from just a few (two or three) manually annotated electrodes. The only prior is a reference, subject independent, 3D model of the cap. Our framework is based on combinatorial optimization (namely on an extension of the Loopy Belief Propagation algorithm[4]) and is robust to soft deformations of the cap caused both by sliding effects and by the variability in subjects' head geometry, as well as to missing electrodes in the 3D reconstruction.

II. RECONSTRUCTION OF THE ELECTRODES POSITIONS

While the problem of retrieving the position of the electrodes is not tackled by this article, we give a brief insight of methods that we use to perform this task.

It is impossible to find the 3D position of a point M from one image m_i of it, taken by a single camera C_i . Yet, if two images m_i and m_j of M are known, as well as the positions from where the pictures were taken, it is possible to recover M (figure 1). Starting from this *stereovision* principle, the computer vision community had to face two major problems: (i) recovering the camera positions (the *calibration task*) and (ii) associating the corresponding 2D points (the *matching problem*). State of the art methods use a *self-calibration* procedure, and are able, from enough overlapping images, to solve these two problems all-together [3]. We are currently finalizing such a system and plan to make freely available to the EEG community. This system recovers the 3D positions of the electrodes from a reduced number of digital pictures.

III. PROBLEM DEFINITION

The inputs of our method consist of:

- a template EEG cap model providing labeled electrodes, along with their 3D positions (in fact, as we will explain further, an important feature of our method is that only the distances between close electrodes are used). \mathcal{L} will denote the set of labels (e.g. $\mathcal{L} = \{Fpz, Oz, \dots\}$), and $C = \{C_l, l \in \mathcal{L}\}$ will be their corresponding 3D positions. C_l could be for example the average position of electrode l among a variety of prior measures. However, in our experiments, it was just estimated on one reference acquisition.
- the measured 3D positions of the electrodes to label, obtained by 3D reconstruction from images. We will denote by $M = \{M_i, i \in [1..n]\}$ these n 3D points.

The output will be a labeling of the electrodes, i.e. a mapping ϕ from $[1..n]$ to \mathcal{L} . In practice, ϕ must be an injection: two candidate electrodes should not share the same label. When $n = |\mathcal{L}|$, this constraint leads to ϕ being one-to-one, but in general this allows for missing electrodes in the measured set (i.e. $n < |\mathcal{L}|$).

IV. MOTIVATION

In this section, we discuss other possible approaches for the electrode labeling problem. As it will be detailed in section VII, we have tried some of these methods without any success. This will motivate our energy-based combinatorial approach.

A simple method could consist of a 3D registration step, followed by a nearest-neighbor labeling.

Let T be a transformation that sends M into the spatial referential of C . A straight labeling could be:

$$\phi(i) = \arg \min_{l \in \mathcal{L}} d(C_l, T(M_i))$$

where $d(A, B)$ denotes the Euclidean distance between points A and B .

Actually, we first tested two direct ways of obtaining an affine registration T :

- *moment-based affine registration*: in this case, we computed first and second order moments of the sets of points M and C and choose T as an affine transformation which superimposes these moments.
- *4 points manual registration*: here, we manually labeled 4 particular electrodes in M and took for T the affine transformation which exactly sends these 4 electrodes to the corresponding positions in C .

As explained in section VII, we observed that these two approaches give very bad average results. One could argue that this might be caused by the quality of the registration. A solution could be to use more optimal affine registration methods, like Iterative Closest Points[6], [7]. Yet, a close look at what caused bad labeling in our experiments, reveals that this would not improve the results. The main reasons are indeed that (i) the subject whose EEG has to be labeled does not have the same head measurements than the template, and moreover that (ii) the cap is a soft structure that undergoes non-affine deformations from one experiment to an another.

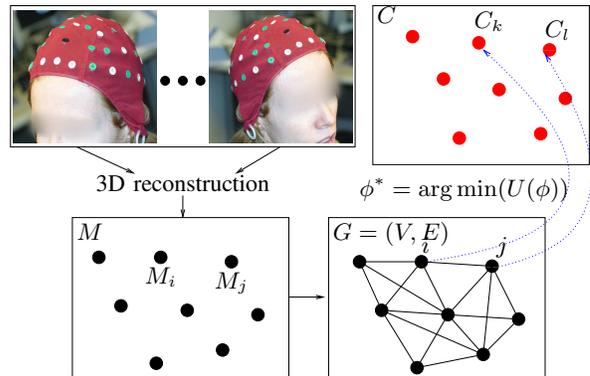


Fig. 2. Complete pipeline : we obtain 3D positions M (bottom left) by reconstruction from several (usually 10) pictures (top left). A graph G then is constructed from these positions (bottom right). Considering a template cap and associated positions C (top right), we label the measured electrodes by estimating $\phi^* = \arg \min(U(\phi))$. In this example, $\phi(i) = k$, $\phi(j) = l$.

It is clear that only a non-affine registration could send M close to C . However, modeling the problem in term of space deformation is not suitable. For instance, a Thin-Plate Spline[8], [9] based algorithm would not be suitable. Actually, a more suitable framework could be a deformable shape matching one. We could see our problem as a shape registration one, based on shape deformation and *intrinsic* shape properties[10], rather than on deforming the ambient space in order to make the shapes match. Because of the topology of the electrodes on the cap, relations between points are of crucial importance. In that sense, our problem is close to the one investigated by Coughlan et al. [11], [12], which they solve recovering both deformations and soft correspondences between two shapes. Yet, in our case, we see two main differences: (i) labeling, rather than shape matching, is the key issue, and (ii) enforcing relational constraints between points are more important than regularizing deformations. For these reasons, we propose a method based on optimal labeling for which the only (soft) constraints are the distances between nearby points, without modeling any deformation.

In the remaining of the article, we first state our model and the associated energy; we then discuss our choice for an energy minimization algorithm. Finally, we validate our method giving qualitative and quantitative results on real experiments.

V. PROPOSED FRAMEWORK

The complete pipeline of our system is depicted figure 2. As we already explained, we do not consider here the 3D reconstruction step, but only the labeling one. From the measured data M , we construct an undirected graph $G = (V, E)$, where $V = [1..n]$ is the set of vertices and E a certain set of edges which codes the relations between nearby electrodes. As it will become clear in the following, the choice of E will tune the “rigidity” of the set of points M . Practically, the symmetric k -nearest neighbors or all the neighbors closer than a certain typical distance, are two

valid choices. Given an edge $e = (i, j) \in E$ for $i \in V$ and $j \in V$, we denote by $d_{ij} = d(M_i, M_j)$ the distance between points M_i and M_j in the measured data and by $\tilde{d}_{ij} = d(C_{\phi(i)}, C_{\phi(j)})$ the reference distance between the electrodes $\phi(i)$ and $\phi(j)$ in the model. In order to preserve in a soft way the local structure of the cap, we propose to simply minimize the following energy:

$$U(\phi) = \sum_{(i,j) \in E} \rho(d_{ij}, \tilde{d}_{ij}) \quad (1)$$

where ρ is a cost-function which penalizes differences between the observed and template distances. The injection property of the mapping is not explicitly enforced in our model, but we choose the function ρ such that non injective solutions are strongly penalized.

VI. ENERGY MINIMIZATION

Following the classical framework of Markov Random Fields, this energy minimization problem can be rewritten as solving a Maximum A Posteriori (MAP) problem for a Gibbs distribution over G [13], [14], [15].

Several methods exist for solving this kind of problem (see [15], [18], [19]). We opted for a widely spread algorithm, namely *Loopy Belief Propagation* (LBP) ([4]). Briefly, it consists in propagating information through the edges of the graph: each node i sends *messages* to its neighbors k , measuring the estimated label of k from its own point of view. Messages are passed between nodes iteratively until a convergence criterion is satisfied. This algorithm is neither guaranteed to converge nor to converge to an optimal solution. However, it behaves well in a large variety of early vision problems. Empirical and theoretical convergence of this family of methods were studied for instance in [21], [22].

VII. EXPERIMENTS

We used 6 sets of 63 electrodes (one of the 64 electrodes is not considered in our experiments). Each set consists of 63 estimated three dimensional points, acquired on different subjects with the same EEG cap and manually labeled. To test our algorithm as extensively as possible, we ran the algorithm on each set, taking successively each of the other sets as a reference. We hence simulated 30 different pairs (M, C) . At least one electrode in M was manually labeled (see further).

E was chosen the following way : we first estimated a typical neighbor distance by computing the maximum of the nearest neighbor distance for all electrodes in M , and then considered as belonging to E , every pair of distinct electrodes within less than three times this distance. In order to accelerate and enforce convergence, we used the following technical tricks:

- we added a classical momentum term ([21])
- denoting by V_f the subset of V of the manually labeled electrodes, we added the set of edges $V_f \times V$ to E , allowing accurate information to propagate quickly in the graph.

Different experiments were carried out. First, the prior consisted in manually labeling electrodes Fpz , Oz , and $T8$. In that case, our method recovers all the electrodes, which was, as expected, not at all the case with an affine registration+nearest neighbor approach (see figure 3). Actually, we observed that labeling $(Oz, T8)$ seems sufficient. Yet, without any further data, we do not consider that labeling two electrodes only is reliable. Figure 4 shows a result on a case where affine registration does not work and the final 3D reconstruction with our method.

To demonstrate the robustness of our algorithm, we also tested hundreds of other conditions, in which 1, 2 or 3 randomly chosen electrodes were "manually" labeled. Non-convergence was only observed for non reasonable choices of "manually" labeled electrodes: indeed, if they are chosen on the sagittal medium line, there is an indetermination due to the left-right symmetry of the cap. This does not occur when the electrodes are set by a human operator. The classification error rates are low (see figure 3 again) but not negligible. This makes us plead for a manual labeling of two or three fixed and easy to identify electrodes, e.g. $(Fpz, Oz, T8)$.

Finally, we also successfully tested cases for which $n < |\mathcal{L}|$, i.e. when some electrodes are missing : if a few electrodes were forgotten in the 3D reconstruction process, our algorithm should still be able to label the detected ones. This should allow us to find which electrodes were forgotten, to compute their approximate 3D position from the template cap model and to use this information to detect them back in the pictures. To carry our experiments, we removed randomly from 1 to 10 electrodes in the data sets to be labeled. Labeling was performed using the $(Fpz, Oz, T8)$ prior as explained above. Results are synthesized figure 3.

	<i>NC</i>	misclassified labels
<i>Affine registration</i>		
- Moment based	-	48.7%
- 4 manual points	-	21.3%
<i>Our method</i>		
- $(Fpz, Oz, T8)$ given	0%	0%
- $(Oz, T8)$ given	0%	0%
- 3 random elect. given	0%	0.03%
- 2 random elect. given	0.3%	0.2%
- 1 random elect. given	4.2%	3.7%
<i>Our method with missing electrodes</i>		
- 1 missing elect.	0%	0%
- 2 missing elect.	0%	0%
- 3 missing elect.	0%	0.01%
- 10 missing elect.	0%	1.11%

Fig. 3. Classification errors. *NC* gives the percentage of instances of the problem for which *Loopy Belief Propagation* did not converge. Misclassified labels percentages are estimated only when convergence occurs.

VIII. DISCUSSION

Experiments show that our framework leads to fast, accurate and robust labeling on a variety of data sets. We consider

