Extraction of arbitrarily-shaped objects using stochastic multiple birth-and-death dynamics and active contours.

No Author Given

No Institute Given

Summary (200words)

We extend the marked point process models that have been used for object extraction from images to arbitrarily shaped objects, without greatly increasing the computational complexity of sampling and estimation. From an alternative point of view, the approach can be viewed as an extension of the active contour methodology to an *a priori* unknown number of objects. Sampling and estimation are based on a stochastic birth-and-death process defined in a space of multiple, arbitrarily shaped objects, where the objects are defined by the image data and prior information. The performance of the approach is demonstrated via experimental results on synthetic and real data.

Abstract (500words)

The resolution of optical satellite and aerial images is continually increasing. At these resolutions the geometry of objects at human scale is getting clearly visible, and needs to be taken into account in analysing the images. Stochastic point process models of multiple objects and their configurations are able to include this type of geometrical information. The optimal configuration of objects is then estimated, usually using MAP estimation. In previous work, the objects involved have been represented using simple geometrical shapes, *e.g.* discs, ellipses, or rectangles.

The aim of this paper is to lift this restriction without increasing the dimension of the single-object space and thus the computational complexity of sampling and estimation. The single-object space considered remains then of small dimension, but the possible individual objects are determined not *a priori*, but by the image data and a single-object version of the model. A Gibbs probability distribution is then defined as an exponential of an energy on the configuration space of an arbitrary number of single objects. The energy includes single-object terms that relate each object to the image and control boundary smoothness, and that is used in the definition of the single-object space, and a term controling the interaction between objects. In the absence of interaction terms, the MAP estimate would thus consist of the subset of the single objects with negative energy. This would probably lead to degenerate solution configurations, however, and so an interaction term is added that controls the relation between different objects, in particular discouraging overlaps. To find the MAP estimate using the full energy, we sample from a multiple birth-and-death process embedded in an annealing scheme [1]. The use of a birth-and-death process allows the number of objects to be unknown *a priori*.

We model individual object boundaries as closed planar curves γ lying in the image domain V, and we suppose an energy functional E_d defined on a space $\mathit{\Gamma}$ of these curves. In this paper, it will take the form of a classical active contour energy. Given an initial curve $\gamma \in \mathit{\Gamma}$, we can then perform gradient descent to arrive at a local minimum of E_d, *i.e.* locally adapted to the data object denoted $\hat{\gamma} \in \mathit{\Gamma}$.

Let $\mathcal{C}$ to be a set of circles lying in V, Now, the single-object space we consider is $\mathit{\Gamma}_o = \hat{\mathcal{C}}$, and the multiple-objects space denoted $\mathit{\Omega}_{\mathit{\Gamma}_0}$ consists of all configurations $\hat{\omega}$ of zero or more objects $\hat{\gamma}$, where $\omega \in \mathit{\Omega}_{\mathcal{C}}$.

The full energy $H(\omega)$ defined in $\mathit{\Omega}_{\mathcal{C}}$ then takes the form:

$$H(\omega) = \sum_i H_1(\omega_i) + \sum_{i \neq j} H_2(\omega_i, \omega_j) \ ,$$

where

$$H_1(\omega_i) = E_d(\hat{\omega}_i) \ ,$$

and

$$H_2(\omega_i, \omega_j) = \frac{A(R(\hat{\omega}_i) \cap R(\hat{\omega}_j))}{\min(A(R(\hat{\omega}_i)), A(R(\hat{\omega}_j)))} + \delta_\epsilon(\omega_i, \omega_j) \ ,$$

with ω_i elements of ω, A the area functional and δ_ϵ a hard-core repulsion (to some tolerance ϵ).

To estimate the configuration of the objects, the sampling uses a Markov chain in $\mathit{\Omega}_{\mathcal{C}}$ consisting of a discrete-time multiple birth-and-death process. The birth step adds a number of circles lying in V, with radii in $[r_{\text{min}}, r_{\text{max}}]$ and centres at the image pixels, to the current configuration with some intensity independent of the current: temperature and configuration energy. Every circle is then evolved using gradient descent of E_d, producing configuration $\hat{\omega} \in \mathit{\Omega}_{\mathit{\Gamma}_o}$. The death step removes a number of components from the current configuration with some intensity dependent on the current (inverse) temperature and the energy difference $H(\omega_i, \omega) = H(\omega - \omega_i) - H(\omega)$. If all the components added in the birth step are removed in the following death step, then stop; if not, the temperature and time step are decreased, and the procedure is repeated.

Figure 1 shows the results of experiments on a on one band of a real colour image. Note that the initial configuration is empty, *i.e.* there are no curves. The first birth step creates a certain number of curves, but the final number is determined automatically by the convergence of the annealed birth-and-death process.

The advantage of the model is its flexibility. The marked point process models that have been used in the past for object extraction from images have incorporated only simple geometric objects. In this paper, we have proposed a method for the incorporation of arbitrarily complex shapes. At the same time, the work can be seen as an extension of the active contour methodology to an *a priori* unknown number of objects. The set of objects considered is defined using an energy that incorporates prior information as well as information coming from the data. The model could therefore be extended by including more specific prior information about object shape into the energy E_d in order to detect objects of several classes in scenes of high complexity containing overlapping objects, *e.g.* in dense forest or urban areas.

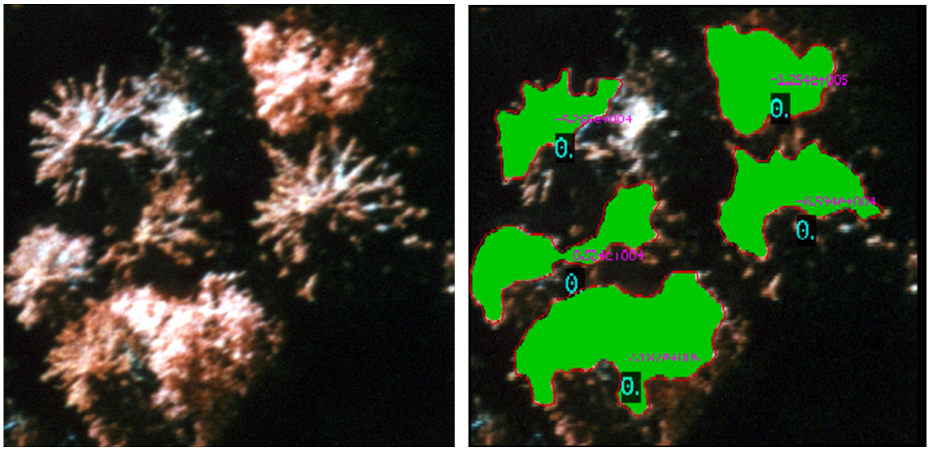

Fig. 1. Left: original real image. Right: final object configuration. The numbers in the interiors of the curves show the value of H_2 (black background), demonstrating that there is no degeneracy in the solution, and the value of H_1.

References

1. Descombes, X., Minlos, R., Zhizhina, E.: Object extraction using a stochastic birth-and-death dynamics in continuum. J. Math. Imaging Vis., (2009) **33(3)**: 347–359