Ambrosio-Tortorelli Segmentation of Stochastic Images

Torben Pätz[1,2] and Tobias Preusser[1,2]

[1] School of Engineering and Science, Jacobs University Bremen
[2] Fraunhofer MEVIS, Bremen, Germany

Abstract. We present an extension of the classical Ambrosio-Tortorelli approximation of the Mumford-Shah approach for the segmentation of images with uncertain gray values resulting from measurement errors and noise. Our approach yields a reliable precision estimate for the segmentation result, and it allows to quantify the robustness of edges in noisy images and under gray value uncertainty. We develop an ansatz space for such images by identifying gray values with random variables. The use of these stochastic images in the minimization of energies of Ambrosio-Tortorelli type leads to stochastic partial differential equations for the stochastic smoothed image and a stochastic phase field for the edge set. For their discretization we utilize the generalized polynomial chaos expansion and the generalized spectral decomposition (GSD) method. We demonstrate the performance of the method on artificial data as well as real medical ultrasound data.

Keywords: Image processing, segmentation, uncertainty, stochastic images, stochastic partial differential equation, polynomial chaos, generalized spectral decomposition.

1 Introduction

In many applications images are used for quantitative measurements, e.g. to determine the size or distance of objects. As image acquisition itself (e.g. by digital camera, CT, MR or Ultrasound) involves measurements of physical or chemical quantities or properties it is good scientific practice that these measurements are equipped with error estimates and that these error estimates are propagated through all analysis steps, including quantitative image processing. The goal is a reliable precision estimate for the final result. In quantitative medical imaging this for example can support the evaluation of the treatment response in chemotherapy. There the growth or shrinkage of tumors must be detected robustly on base of noisy contrast enhanced CT scans. As a matter of fact small measurement errors due to noise and uncertainty in the gray values can result in huge variations in the computed tumor volume, thus being a source for erroneous therapy-response indications.

Quantitative image processing is often related to the segmentation of an object inside an image. The main idea is to detect the shape of an object inside an image

K. Daniilidis, P. Maragos, N. Paragios (Eds.): ECCV 2010, Part V, LNCS 6315, pp. 254–267, 2010.

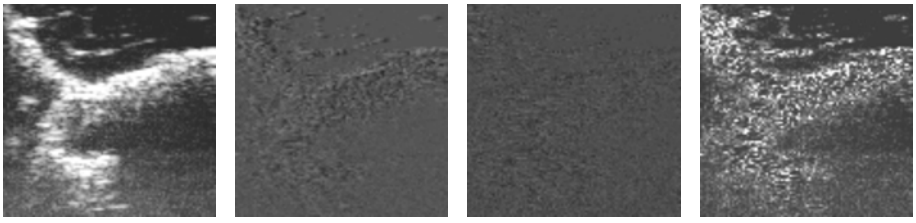

Fig. 1. From left to right: The first mode (=mean), second mode, fifth mode and the variance of a stochastic image are shown

and to separate it from the background. In the past, a multitude of methods based on PDEs have been developed. Among them are level set (sharp interface) and phase field (diffuse interface) approaches, which use implicit representations of the object boundaries that are computationally much easier to handle than explicit representations.

In [10] a speed function for the level set evolution is proposed, which depends on the image gradient. The evolution stops when the level set reaches an edge inside the image. This method was improved by Caselles et. al. [4] by introducing an additional term, which forces the level set to stay at the boundary. The idea of the Chan and Vese approach [5] is to segment homogeneous regions inside an image. An evolution equation is solved until the level set separates homogeneous regions of the image. This allows to segment objects with and without sharp edges. The well known Mumford-Shah approach [12] is based on the minimization of an energy functional, which measures the smoothness of the segmented objects as well as the length of the object boundaries. An often used regularization of the Mumford-Shah functional is the method proposed by Ambrosio-Tortorelli [1] leading to a phase field model for the description of object boundaries.

Error propagation is very difficult for classical image processing algorithms and in particular for the level set or phase field segmentation methods mentioned above. In the literature a lot of authors deal with error estimates, which have several restrictions: Weber et. al. [17] presented a method were the input data is presumed to be Gaussian distributed. Nestares et. al. [13] were able to derive bounds for the error and Bruhn et. al. [2] derived confidence measures for the error. In [15] a method is presented, which assumes values of the image pixels not to have fixed gray values but distributions of gray values. Thus, pixels are random variables (RVs), which model the errors in the image acquisition process. An image containing such RVs as pixels/voxels is then called a stochastic image. A few modes of a stochastic image are pictured in Fig. 1.

In the work presented here we extend this approach of stochastic images and combine it with Mumford-Shah segmentation in the spirit of the Ambrosio-Tortorelli phase field approximation. For an input image with uncertain gray values our approach provides a stochastic edge representation in the form of a stochastic phase field as well as a stochastic image as the representation of the smoothed input image. It allows for precise error estimates beyond the assumption of Gaussian gray value distributions. In fact, the evaluation of the stochastic

modes of the phase field allows e.g. to estimate the variance of the edge location or the confidence for the presence of edges at certain locations for arbitrary noise models and distributions of gray values.

The use of the notion of stochastic images in variational image processing leads to stochastic partial differential equations (SPDE). The numerical solution of SPDEs is a challenging problem, because intrusive methods like the stochastic finite element method (SFEM) [8] lead to high dimensional systems of equations, which are difficult to treat on contemporary hardware. We utilize the recently developed generalized spectral decomposition (GSD) [14], which allows to break down the systems of equations into a series of smaller systems by choosing optimal small subspaces in the stochastic dimension. This results in an enormous speedup of the computation, a saving of memory, and in an algorithm, which is much faster than classical sampling techniques like Monte Carlo.

2 Stochastic Images

It is popular in PDE based image processing to model an image $f : D \to \mathbb{R}$ on a domain $D \subset \mathbb{R}^d$, $d = 2, 3$ using a finite element space and a representation

$$f(x) = \sum\nolimits_{i \in \mathcal{I}} f^i P_i(x) \ , \tag{1}$$

where $f^i \in \mathbb{R}$ is the value of the i-th pixel from the pixel set $\mathcal{I}$ and P_i the shape function (e.g. tent-function) of the i-th pixel. In a stochastic image a single pixel has no longer a fixed value. Instead it depends on a vector of RVs $\boldsymbol{\xi}(\omega) = (\xi_1(\omega), \ldots, \xi_n(\omega))$ and on a random event $\omega \in \Omega$. Here Ω denotes an event space, $\mathcal{A} \subset 2^\Omega$ a σ-algebra and Π a probability measure. Note that the concept of stochastic images can also be combined with other spacial discretizations, e.g. finite difference schemes.

2.1 Polynomial Chaos Expansion

Based on the fundamental work of Wiener [18], Xiu and Karniadakis [20] developed the generalized polynomial chaos (gPC) expansion for the representation of a RV with finite second-order moments by a polynomial basis.

Following Cameron/Martin [3], every RV $X(\omega) \in L^2(\Omega, \mathcal{A}, \Pi)$ can be represented by

$$X(\omega) = \sum\nolimits_{\alpha=1}^{\infty} a_\alpha \Psi^\alpha(\hat{\xi}(\omega)) \ , \tag{2}$$

where $\hat{\xi} = (\xi_1, \xi_2, \ldots)$ is a sequence of RVs with known probability density function ρ_j and Ψ^α are polynomials in $\hat{\xi}$ forming a basis of $L^2(\Omega, \mathcal{A}, \Pi)$. For the numerical treatment an approximation with prescribed polynomial degree p and a fixed number of RVs $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_n)$ is chosen, thus

$$X(\omega) \approx \sum\nolimits_{\alpha=1}^{N} a_\alpha \Psi^\alpha(\boldsymbol{\xi}(\omega)) \ . \tag{3}$$

This representation involves a mapping $\boldsymbol{\xi} : \Omega \to \Gamma$ of events $\omega \in \Omega$ to $\boldsymbol{\xi}(\omega) \in \Gamma$, where $\Gamma = \bigtimes_{j=1}^{n} \xi_j(\Omega)$ is finite dimensional.

The number N of basis functions depends on the number n of RVs and the maximal polynomial degree p of the approximation. As usual for a polynomial basis the number of basis functions is given by $N = \binom{n+p}{p}$. Thus, the number of basis functions grows rapidly with the number of RVs n and the polynomial degree p of the approximation.

It is most convenient to choose polynomials Ψ^α which are pairwise orthogonal with respect to the corresponding probability measure of the ξ_j. Thus, in the case of Gaussian RVs ξ_j the Ψ^α are products of one-dimensional Hermite polynomials. In the case of uniformly distributed RVs the Ψ^α are products of Legendre polynomials. In our work presented here we use uniformly distributed RVs ξ_j involving Legendre polynomials. For arithmetic operations needed for the use of the gPC expansion in numerical schemes we use the methods from [6].

The expectation, variance and analogously higher stochastic moments of the approximated RV $X(\omega)$ are evaluated as

$$\mathbb{E}(X) \approx \int_{\Gamma} \sum_{\alpha=1}^{N} a_\alpha \Psi^\alpha(\boldsymbol{\xi})\, d\Pi_{\boldsymbol{\xi}}\,, \ \mathrm{Var}(X) \approx \int_{\Gamma} \left(\sum_{\alpha=1}^{N} a_\alpha \Psi^\alpha(\boldsymbol{\xi}) - \mathbb{E}(X) \right)^2 d\Pi_{\boldsymbol{\xi}} \ , \quad (4)$$

where $d\Pi_{\boldsymbol{\xi}} = \prod_{j=1}^{n} \rho_j(\xi_j)\, d\xi_j$ is the transformed probability measure.

2.2 Polynomial Chaos for Stochastic Images

Following [15], the representation of an image whose pixels values are RVs is obtained from (1) by replacing the fixed f^i by RVs $f^i(\boldsymbol{\xi})$, thus

$$f(x, \boldsymbol{\xi}) = \sum\nolimits_{i \in \mathcal{I}} f^i(\boldsymbol{\xi}) P_i(x) \ . \quad (5)$$

Note that here and in the following we omit denoting the dependence of $\boldsymbol{\xi}$ on ω for reasons of simplicity of the presentation. The gPC expansion (3) allows to approximate any second order RV $f^i(\boldsymbol{\xi})$ by a weighted sum of orthogonal multidimensional polynomials. This leads to

$$f(x, \boldsymbol{\xi}) = \sum\nolimits_{i \in \mathcal{I}} \sum\nolimits_{\alpha=1}^{N} f^i_\alpha \Psi^\alpha(\boldsymbol{\xi}) P_i(x) \quad (6)$$

as the representation of a stochastic image, i.e. an image whose gray values are RVs. For fixed α we call the coefficient f^i_α a *stochastic mode* of the pixel i. The set $\{f^i_\alpha\}_{i \in \mathcal{I}}$ collects the stochastic modes of all pixels for fixed α. This set can be visualized as a classical image, which is done in Fig. 1 where three modes of a sample image are shown.

From the gPC expansion (6) it is straight forward to compute stochastic moments of the images. With the use of our orthogonal set of basis functions, the Legendre Polynomials, we have $\mathbb{E}(\Psi^1) = 1$ and $\mathbb{E}(\Psi^\alpha \Psi^\beta) = 0$ if $\alpha \neq \beta$. Thus, the mean and the variance of a stochastic image are computed as

$$\mathbb{E}\left(f(x_i, \cdot)\right) = f^i_1\,, \qquad \mathrm{Var}\left(f(x_i, \cdot)\right) = \sum\nolimits_{\alpha=2}^{N} \left(f^i_\alpha\right)^2 \mathbb{E}\left(\left(\Psi^\alpha\right)^2\right) \ . \quad (7)$$

Other stochastic moments are obtained in a similar way. In Fig. 1 the mean and the variance of a stochastic image are shown.

Note that the representation of stochastic images presented here differs from the one discussed in [15]. There, an image space is used in which every pixel depends on one RV only. However, for many image acquisition processes and image processing methods the assumption that the noise is independent for every pixel is not true, thus we let every pixel depend on a vector of RVs $\boldsymbol{\xi}$.

2.3 From Samples to Input Distributions

To use the notion of stochastic images developed in the previous sections for image processing, we need to obtain the coefficients of the representation (6) for our image undergoing the analysis. Let $u^{(1)}, \ldots, u^{(M)}$, with $u^{(k)} \in \mathbb{R}^r$, $r = |\mathcal{I}|$, denote sample images, e.g. resulting from repeated acquisition. The goal is to identify these image samples as the samples of some vector of independent RVs $\boldsymbol{X}$. To this end the empirical Karhunen-Loeve decomposition [9] yields

$$u^{(k)} = \bar{u} + \sum\nolimits_{j=1}^{r} \sqrt{s_j} U_j X_j^{(k)} \quad , \tag{8}$$

where $\bar{u}$ is the mean of the input samples. The pairs (s_j, U_j) for $j = 1, \ldots, r$ are the eigenpairs sorted in descending order of the $r \times r$ covariance matrix

$$C := \frac{1}{M-1} \sum\nolimits_{k=1}^{M} (u^{(k)} - \bar{u})^T (u^{(k)} - \bar{u}) \ . \tag{9}$$

Moreover, the

$$X_j^{(k)} = (s_j)^{-1/2} U_j^T (u^{(k)} - \bar{u}) \tag{10}$$

are samples of the desired vector of RVs $\boldsymbol{X} = (X_1, \ldots, X_n)$, where $n < r$.

The estimation of the coefficients of the gPC expansion (3) of the random vector $\boldsymbol{X}$ from these samples can be achieved by inverting the discrete empirical cumulative distribution function (CDF) F_{X_j}, which is based on the samples $X_j^{(k)}$. This leads to a staircase-like approximation of the RV X_j. Following [16] we get $X_{j,\alpha}$ from the projection on Ψ^α via

$$X_{j,\alpha} = \mathbb{E}(X_j \Psi^\alpha) = \int_\Gamma F_{X_j}^{-1}\left(F_\xi(y)\right) \Psi^\alpha(\xi(y)) d\Pi_\xi(y) \ . \tag{11}$$

Note that the assumption of independence allows us to work with those basis functions, which depend on one RV only, i.e. $\Psi^\alpha(\boldsymbol{\xi}) = \Psi^\alpha(\xi)$. The empirical CDF and the empirical inverse of the CDF are obtained by

$$\begin{aligned} F_{X_j}(x) &= \frac{1}{M} \sum_{k=1}^{M} I\left(X_j^{(k)} \leq x\right) \quad , \\ F_{X_j}^{-1}(y) &= \min\left\{ x \in \left\{ X_j^{(k)} \right\}_{k=1}^{M} \mid F_{X_j}(x) \geq y \right\} \quad , \end{aligned} \tag{12}$$

where I is the indicator function attaining value 1 for true arguments and 0 else. Note that the RVs X_j are related to the eigenpairs (s_j, U_j) of the Karhunen-Loeve decomposition via (10). Using the expression for the inverse $F_{X_j}^{-1}$ together with a numerical quadrature associated with the measure Π_ξ allows to compute the gPC expansion coefficients X_j^α independently from each other.

We emphasize that the assumption of independence of the RVs X_j is very strong and in general not true. However, following [16] in particular for the case of few input samples this assumption is reasonable.

Also note that the estimation of the typically dense covariance matrix is only feasible for images of low dimension r. For large image dimensions we must model the gray value distribution using characteristics of the acquisition process and not via the analysis of samples. This noise modeling is part of ongoing work.

3 A Phase Field Model for Segmentation on Stochastic Images

We now focus on the combination of the notion of stochastic images with the segmentation approach in the spirit of Ambrosio and Tortorelli [1].

3.1 Classical Mumford-Shah and Ambrosio-Tortorelli Segmentation

For a given initial image u_0 on the domain D Mumford and Shah [12] proposed to obtain an edge set $K \subset D$ and a smooth representation u of u_0 as the minimizers of the energy

$$E_{\mathrm{MS}}(u, K) := \int_{D\setminus K} (u - u_0)^2 dx + \mu \int_{D\setminus K} |\nabla u|^2 dx + \nu \mathcal{H}^{d-1}(K) \ , \tag{13}$$

where μ and ν are positive weights, and $\mathcal{H}^{d-1}(K)$ the $d-1$-dimensional Hausdorff measure. Roughly speaking, the minimizer u must be an image, which is close to the initial u_0 away from the edges (then $\int_{D\setminus K}(u-u_0)^2 dx$ is small) and smooth away from the edges (then $\int_{D\setminus K} |\nabla u|^2 dx$ is small). Moreover the length of edges K must be small (then $\mathcal{H}^{d-1}(K)$, measuring the length of the edge set, is small).

Ambrosio and Tortorelli [1] proposed to approximate the edge set by a phase field $\phi : D \to \mathbb{R}$, i.e. a smooth function that is zero on edges and one away from the edges. To this end they define the energy

$$\begin{aligned} E_{\mathrm{AT}}(u, \phi) &:= E_{\mathrm{fid}}(u) + E_{\mathrm{reg}}(u, \phi) + E_{\mathrm{phase}}(\phi) \\ &:= \int_D \left(u(x) - u_0(x)\right)^2 \, dx + \int_D \mu \, \left(\phi(x)^2 + k_\varepsilon\right) |\nabla u(x)|^2 \, dx \\ &\qquad + \int_D \nu \, \varepsilon \, |\nabla \phi(x)|^2 + \frac{\nu}{4\varepsilon} \left(1 - \phi(x)\right)^2 \, dx \ . \end{aligned} \tag{14}$$

The first integral ensures closeness of the smoothed image to the original u_0. The second integral measures smoothness of u apart from areas where ϕ is small,

and enforces ϕ to be small in the vicinity of edges. The parameter k_ε ensures coerciveness of the differential operator and thus existence of solutions, because ϕ^2 may vanish. The third integral drives the phase field towards one and ensures small edge sets via the term $|\nabla\phi|^2$. The parameter ε allows to control the scale of the detected edges.

For the numerical determination of a minimizing pair (u, ϕ) the Euler-Lagrange equations of (14) are solved. Thus we seek $u, \phi \in H^1(D)$ as the weak solutions of

$$-\mathrm{div}\left(\mu(\phi^2 + k_\varepsilon)\nabla u\right) + u = u_0\,, \qquad -\varepsilon\Delta\phi + \left(\frac{1}{4\varepsilon} + \frac{\mu}{2\nu}|\nabla u|^2\right)\phi = \frac{1}{4\varepsilon}\ . \quad (15)$$

In an implementation both equations can be solved alternately letting either u or ϕ vary until a fixed point as the joint solution of both equations is reached.

3.2 Ambrosio-Tortorelli Segmentation on Stochastic Images

For the segmentation of stochastic images by the phase field approach of Ambrosio and Tortorelli we replace the deterministic u and ϕ by their corresponding stochastic analogs. The stochastic energy components are then defined as the expectations of the classical components, i.e.

$$\begin{aligned}
E^s_{\mathrm{fid}}(u) &:= \mathbb{E}(E_{\mathrm{fid}}) = \int_\Gamma\int_D \left(u\,(x,\boldsymbol{\xi}) - u_0\,(x,\boldsymbol{\xi})\right)^2\ dx\,d\Pi_{\boldsymbol{\xi}} \\
E^s_{\mathrm{reg}}(u,\phi) &:= \mathbb{E}(E_{\mathrm{reg}}) = \int_\Gamma\int_D \mu\,\left(\phi\,(x,\boldsymbol{\xi})^2 + k_\varepsilon\right)|\nabla u(x,\boldsymbol{\xi})|^2\ dx\,d\Pi_{\boldsymbol{\xi}} \\
E^s_{\mathrm{phase}}(\phi) &:= \mathbb{E}(E_{\mathrm{phase}}) = \int_\Gamma\int_D \nu\,\varepsilon\,|\nabla\phi(x,\boldsymbol{\xi})|^2 + \frac{\nu}{4\varepsilon}\left(1-\phi\,(x,\boldsymbol{\xi})\right)^2\ dx\,d\Pi_{\boldsymbol{\xi}}
\end{aligned} \quad (16)$$

and we define the stochastic energy as the sum of these, i.e.

$$E^s_{AT}(u,\phi) = E^s_{\mathrm{fid}}(u) + E^s_{\mathrm{reg}}(u,\phi) + E^s_{\mathrm{phase}}(\phi)\ . \quad (17)$$

The Euler-Lagrange equations of the energy are obtained from the first variation of the above integrals. Since the stochastic energies (16) are just the expectations of the classical energies (14) the computations are straight forward and completely analog to the deterministic case. For example, we get for a test function $\theta : D \times \Gamma \to \mathbb{R}$

$$\begin{aligned}
\frac{d}{dt}E^s_{\mathrm{fid}}(u+t\,\theta)\Big|_{t=0} &= \frac{d}{dt}\int_\Gamma\int_D \Big(u(x,\boldsymbol{\xi}) + t\theta(x,\boldsymbol{\xi}) - u_0(x,\boldsymbol{\xi})\Big)^2\ dx\ d\Pi_{\boldsymbol{\xi}}\Big|_{t=0} \\
&= \int_\Gamma\int_D 2\Big(u(x,\boldsymbol{\xi}) - u_0(x,\boldsymbol{\xi})\Big)\theta(x,\boldsymbol{\xi})\ dx\ d\Pi_{\boldsymbol{\xi}}\ .
\end{aligned} \quad (18)$$

With analog computations for the remaining energy contributions we arrive at the following system of stochastic partial differential equations: We seek for $u, \phi : D \times \Gamma \to \mathbb{R}$ as the weak solutions of

$$\begin{aligned} -\operatorname{div}\left(\mu(\phi(x,\boldsymbol{\xi})^2+k_\varepsilon)\nabla u(x,\boldsymbol{\xi})\right)+u(x,\boldsymbol{\xi}) &= u_0(x,\boldsymbol{\xi}) \\ -\varepsilon\Delta\phi(x,\boldsymbol{\xi})+\left(\frac{1}{4\varepsilon}+\frac{\mu}{2\nu}|\nabla u(x,\boldsymbol{\xi})|^2\right)\phi(x,\boldsymbol{\xi}) &= \frac{1}{4\varepsilon}\ . \end{aligned} \tag{19}$$

This system is analog to the classical system (15) in which images have been replaced by stochastic images.

3.3 Weak Formulation and Discretization

The system (19) contains two elliptic SPDEs, which are supposed to be interpreted in the weak sense. To this end we multiply the equations by a test function $\theta : D \times \Gamma \to \mathbb{R}$, integrate over Γ with respect to the corresponding probability measure and integrate by parts over the physical domain D. For the first equation in (19) this leads us to

$$\begin{aligned} \int_\Gamma\int_D \mu\left(\phi(x,\boldsymbol{\xi})^2+k_\varepsilon\right)\nabla u(x,\boldsymbol{\xi})\cdot\nabla\theta(x,\boldsymbol{\xi})+u(x,\boldsymbol{\xi})\theta(x,\boldsymbol{\xi})\,dx\,d\Pi_{\boldsymbol{\xi}} \\ = \int_\Gamma\int_D u_0(x,\boldsymbol{\xi})\theta(x,\boldsymbol{\xi})\,dx\,d\Pi_{\boldsymbol{\xi}} \end{aligned} \tag{20}$$

and to an analog expression for the second part of (19). Here we assume Neumann (natural) boundary conditions for u and ϕ such that no boundary terms appear in the weak form. For the existence of solutions for these SPDEs, the constant k_ϵ is supposed to ensure the positivity of the diffusion coefficient $\mu(\phi^2+k_\epsilon)$. In fact, there must exist $c_{\min}, c_{\max} \in (0,\infty)$ such that

$$P\left(\omega\in\Omega \mid \mu\left(\phi(x,\boldsymbol{\xi}(\omega))^2+k_\varepsilon\right)\in[c_{\min},c_{\max}]\ \forall x\in\overline{D}\right)=1\ . \tag{21}$$

Finally, solutions u and ϕ will be random fields, i.e. RVs, which are indexed by a spatial coordinate and such that $u(\cdot,\boldsymbol{\xi}),\phi(\cdot,\boldsymbol{\xi})\in H^1(D)$ almost sure. Thus, for almost every realization (in the sense of the measure $\Pi_{\boldsymbol{\xi}}$) the stochastic images u and ϕ have weak derivatives in $L^2(D)$.

The weak system (20) is discretized with a substitution of the gPC expansion (6) of the image and the phase field. As test functions products $P_j(x)\Psi^\beta(\boldsymbol{\xi})$ of spatial basis functions and stochastic basis functions are used. Denoting the vectors of coefficients by $U^\alpha=(u^i_\alpha)_{i\in\mathcal{I}}\in\mathbb{R}^{|\mathcal{I}|}$ and similarly for the phase field ϕ and the initial image u_0 we get the fully discrete systems

$$\begin{aligned} \sum_{\alpha=1}^N\left(M^{\alpha,\beta}+L^{\alpha,\beta}\right)U^\alpha &= \sum_{\alpha=1}^N M^{\alpha,\beta}(U_0)^\alpha \qquad \forall\beta\in\{1,\dots,N\} \\ \sum_{\alpha=1}^N\left(\varepsilon S^{\alpha,\beta}+T^{\alpha,\beta}\right)\Phi^\alpha &= \sum_{\alpha=1}^N A^\alpha \qquad \forall\beta\in\{1,\dots,N\} \end{aligned} \tag{22}$$

where $M^{\alpha,\beta}, L^{\alpha,\beta}, S^{\alpha,\beta}$ and $T^{\alpha,\beta}$ are the blocks of the system matrix, defined as

$$\left(M^{\alpha,\beta}\right)_{i,j}=\mathbb{E}(\Psi^\alpha\Psi^\beta)\int_D P_i\,P_j\,dx\ ,\quad \left(S^{\alpha,\beta}\right)_{i,j}=\mathbb{E}(\Psi^\alpha\Psi^\beta)\int_D\nabla P_i\cdot\nabla P_j dx \tag{23}$$

and

$$\begin{aligned}\left(L^{\alpha,\beta}\right)_{i,j} &= \sum_k \sum_\gamma \mathbb{E}(\Psi^\alpha \Psi^\beta \Psi^\gamma)\ (\widetilde{\phi^2})^k_\gamma \int_D \nabla P_i \cdot \nabla P_j P_k \, dx\,, \\ \left(T^{\alpha,\beta}\right)_{i,j} &= \sum_k \sum_\gamma \mathbb{E}(\Psi^\alpha \Psi^\beta \Psi^\gamma)\ u^k_\gamma \int_D P_i P_j P_k \, dx \ .\end{aligned} \tag{24}$$

Here, $(\widetilde{\phi^2})^k_\gamma$ denotes the coefficients of the gPC expansion of the Galerkin projection of ϕ^2 onto the image space (cf. [6]). Finally, the right hand side vector is defined as

$$\left(A^\alpha\right)_i = \int_\Gamma \Psi^\alpha d\xi \int_D \frac{1}{4\varepsilon} P_i \, dx = \begin{cases} \displaystyle\int_D \frac{1}{4\varepsilon} P_i \, dx & \text{if } \alpha = 1, \\ 0 & \text{else} \ . \end{cases} \tag{25}$$

Note that the expectations of the products of stochastic basis functions involved above can be precomputed in advance, since these do only depend on the choice of basis functions. Analog to the classical finite element method the systems of linear equations can be treated by an iterative solver like the method of conjugate gradients.

3.4 Generalized Spectral Decomposition

A significant speedup of the solution process and an enormous reduction of the memory requirements are achieved by selecting suitable sub-spaces and a special basis which captures the dominant stochastic effects. In the GSD [14] the solution u (analogously ϕ) is approximated by

$$u(x, \boldsymbol{\xi}) \approx \sum\nolimits_{j=1}^{K} \lambda_j(\boldsymbol{\xi}) V_j(x) \ , \tag{26}$$

where V_j is a deterministic function, λ_j a stochastic function and K the number of modes of the decomposition. Thus, the GSD allows to compute a solution where the deterministic and the stochastic basis functions are not fixed a priori. The flexible basis functions allow to find a solution, which has significant less modes, i.e. $K \ll N$, but has nearly the same approximation quality.

In [14] it is shown how to achieve the modes of an optimal approximation in the energy norm $\|\cdot\|_A$ of the problem, i.e. such that

$$\left\| u - \sum\nolimits_{j=1}^{K} \lambda_j V_j \right\|_A^2 = \min_{\lambda, U} \left\| u - \sum\nolimits_{j=1}^{K} \lambda_j V_j \right\|_A \ . \tag{27}$$

Details about the GSD method, proofs for the optimality of the approximation, implementation details and numerical tests can be found in [14] and in the supplementary material of this contribution. In our implementation the power-type GSD presented in [14] is used.

4 Results

In the following we demonstrate the performance of our stochastic segmentation approach. Our first input image data set consists of $M = 5$ samples from the artificial "street sequence" [11], the second dataset consists of $M = 45$ sample images from ultrasound (US) imaging of a structure in the forearm, acquired within 2 seconds. Note that we do not consider the street sequence as an image sequence here, instead we use 5 consecutive frames as samples of the noisy and uncertain acquisition of the same object. From the samples we compute the gPC representation using $n = 10$ (US), respectively $n = 4$ (street scene) RVs with the method described in Section 2.3. Our images have a resolution of 100×100 pixels. We use a maximal polynomial degree of $p = 3$ leading to a gPC dimension $N = 286$ (US) and $N = 35$ (street scene), respectively. For the reduction of the complexity by the GSD we set $K = 6$. Furthermore, we use $\nu = 0.00075$ and $k_\varepsilon = 2.0h$ in all computations, where h is the grid spacing. To show the influence of the RVs, we have also used the US data using the mean value only ($n = 0$).

4.1 Street Image Data Set

Between the samples of the street sequence the camera position and the position of the car differs, thus the edge detection using (17) should show a high variance at edges close to the camera (thus moving much) and around the moving car. The results depicted in Fig. 2 match with these expectations. Indeed, in the

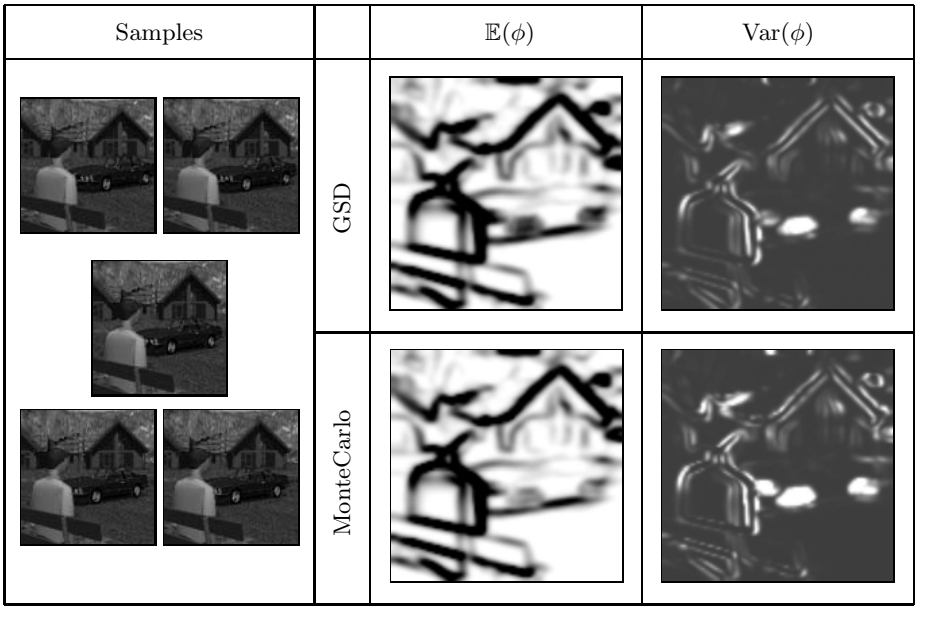

Fig. 2. Results of the segmentation of the street scene

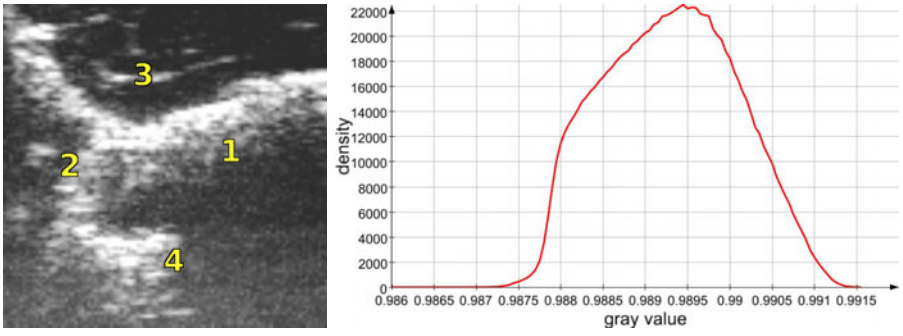

Fig. 3. Left: Denotation of image regions. The image shows a structure in the forearm. Right: Probability density function of a single pixel from the resulting phase field.

Table 1. Comparison of the computation times of different methods for the discretization of SPDEs (measured for the street scene)

Numerical method	Computation time	Number of samples
MonteCarlo	about 35 hours	10000
Stochastic Collocation	about 7 hours	about 2000
GSD	about 2 hours	n/a

region around the wheels of the car and around the right shoulder of the person the edge detection is most influenced by the moving camera, respectively the varying gray values between the samples at the edges. Also around the edges in the background the variance is increased due to the moving camera. A comparison of the results of our GSD implementation with a simple Monte Carlo method with 10000 sample computations shown in Fig. 2 reveals that both approaches lead to similar results. In Table 1 we report the execution times on a typical desktop PC for the Monte Carlo sampling, a more sophisticated sampling using stochastic collocation [19] and the GSD method discussed here. We see that the GSD implementation is about 20 times faster as the classical sampling, however the intrusive GSD needs more implementational effort than the non-intrusive sampling techniques, which can reuse existing deterministic code, because sampling techniques solve the classical Ambrosio-Tortorelli model for every sample and compute stochastic quantities like the variance afterwards from the results on the samples.

4.2 Ultrasound Samples

The conversion of the input samples into the gPC expansion as described in Section 2.3 leads to the representation of the stochastic ultrasound image in a 286-dimensional space. Thus, the only meaningful way of visualizing this stochastic image is via stochastic moments like mean and variance. Fig. 4 shows the mean

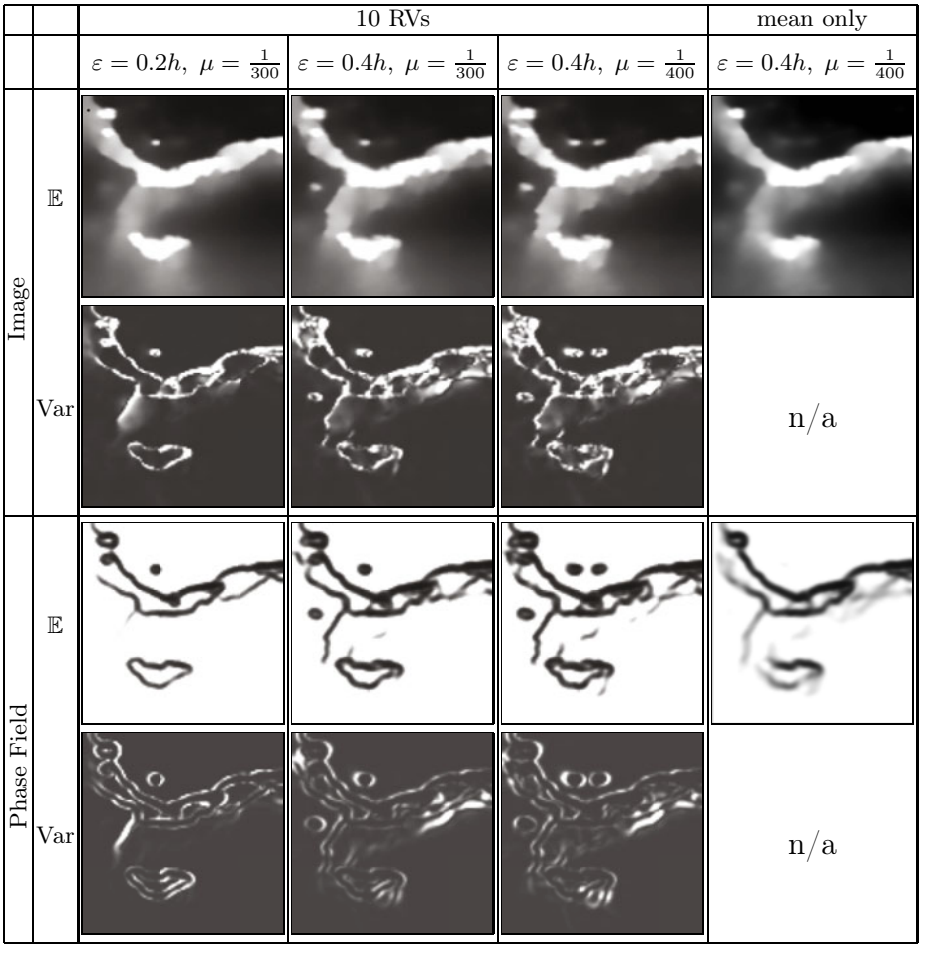

Fig. 4. The mean and variance of the resulting image and phase field for different parameter settings and numbers of RVs using the ultrasound data

and the variance of the phase field ϕ and the smoothed image u for different settings of the smoothing coefficient μ and the phase field width ε. The algorithm needs about 100 iterations, i.e. alternate solutions of (19) for u and ϕ. However, in the first steps the convergence is very fast and already after about 10 iterations no visible difference in u and ϕ can be seen (cf. movie in the supplementary material that shows the solution iterations).

From the variance image of the phase field the identification of regions where the input distribution has a strong influence on the segmentation result (areas with high variance) is straight forward. A benefit of our new stochastic edge detection via the phase field ϕ is that it allows for an identification of edges in a way that is robust with regard to parameter changes. Indeed, in particular within the four regions marked in the left picture of Fig. 3 the expectation of

the phase field is highly influenced by the choice of μ and ν as can be seen in Fig. 4. The blurred edge at position 1 is seen in the expectation of the phase field only when a narrow phase field is used. In region 2 we have a different situation in which the edge can be identified only using a widish phase field. Also the edges at positions 3 and 4 can be identified using adjusted parameters. However, note that in case one of these edges is not seen in the expectation of ϕ because of a particular choice of parameters, a high variance of ϕ indicates the possible existence of an edge. This is in particular obvious for the regions 1 and 2.

Moreover, our algorithm can estimate the reliability of detected edges: A low mean phase field value and a low variance indicate, that the edge is robust and not influenced by the noise and uncertainty of the acquisition process. This is for example true for the edges on the top of the structure shown here. In contrast to that a high variance in regions with a high or low mean phase field value (e.g. the labeled regions 1-4) indicates regions, where the detected edge is highly sensitive to the noise and uncertain acquisition process.

Also, we can easily extract the distribution of the gray values for any pixel location inside the image and the phase field from the gPC expansion obtained via GSD. In Fig. 3, right, we show the probability density function of a pixel from the phase field computed via the GSD.

5 Conclusions

We have presented an extension of the well known Ambrosio-Tortorelli phase field approximation of the Mumford-Shah functional to stochastic images. Our approach allows us to propagate information about the distribution of the gray values in the input image, which result from noise or erroneous measurements, through the segmentation process, leading to a segmentation result that contains information about the reliability of the segmentation. The resulting SPDEs are discretized by the generalized polynomial chaos approach and a generalized spectral decomposition method. We have shown the application of the segmentation to artificial sample images as well as to noisy ultrasound image samples. In an ongoing work we investigate the use of our algorithm on the basis of noise models instead of multiple input image samples.

In particular for medical applications of quantitative image processing we envisage that our approach can be a basis for superior results, since it allows to measure the size of lesions including reliability estimates. But also other applications, e.g. material science, quality control, geography etc. can benefit from the reliability estimates. In the future we plan to investigate a stochastic extension of edge linking methods [7] for the Mumford-Shah functional. Also, we will study stochastic extensions of sharp interface segmentation methods like level set based approaches.

Acknowledgements. We acknowledge R.M. Kirby from the University of Utah, USA for fruitful discussions and D. Ojdanic from Fraunhofer MEVIS, Bremen, Germany for providing the ultrasound data set.

References

1. Ambrosio, L., Tortorelli, M.: Approximation of functionals depending on jumps by elliptic functionals via gamma-convergence. Comm. Pure Appl. Math. 43(8), 999–1036 (1990)
2. Bruhn, A., Weickert, J., Schnörr, C.: Lucas/Kanade meets Horn/Schunck: Combining local and global optic flow methods. Int. J. of Computer Vision 61(3), 211–231 (2005)
3. Cameron, R.H., Martin, W.T.: The orthogonal development of non-linear functionals in series of Fourier-Hermite functionals. The Annals of Mathematics 48(2), 385–392 (1947)
4. Caselles, V., Kimmel, R., Sapiro, G.: Geodesic active contours. Int. J. of Computer Vision 22(1), 61–79 (1997)
5. Chan, T., Vese, L.: Active contours without edges. IEEE T. Image. Process. 10(2), 266–277 (2001)
6. Debusschere, B.J., Najm, H.N., Pébay, P.P., Knio, O.M., Ghanem, R.G., Le Maître, O.P.: Numerical challenges in the use of polynomial chaos representations for stochastic processes. SIAM J. Sci. Comput. 26(2), 698–719 (2005)
7. Erdem, E., Sancar Yilmaz, A., Tari, S.: Mumford-Shah regularizer with spatial coherence. In: Sgallari, F., Murli, A., Paragios, N. (eds.) SSVM 2007. LNCS, vol. 4485, pp. 545–555. Springer, Heidelberg (2007)
8. Ghanem, R.G., Spanos, P.D.: Stochastic finite elements: A spectral approach. Springer-Verlag, New York (1991)
9. Loeve, M.: Probability theory, 4th edn. Springer, New York (1977)
10. Malladi, R., Sethian, J.A., Vemuri, B.C.: Evolutionary fronts for topology-independent shape modeling and recovery. In: Eklundh, J.-O. (ed.) ECCV 1994, Part I. LNCS, vol. 801, pp. 3–13. Springer, Heidelberg (1994)
11. McCane, B., Novins, K., Crannitch, D., Galvin, B.: On benchmarking optical flow. Comput. Vis. Image Underst. 84(1), 126–143 (2001)
12. Mumford, D., Shah, J.: Optimal approximations by piecewise smooth functions and associated variational problems. Comm. Pure Appl. Math. 42(5), 577–685 (1989)
13. Nestares, O., Fleet, D.J., Heeger, D.J.: Likelihood functions and confidence bounds for total-least-squares problems. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 1, pp. 523–530 (2000)
14. Nouy, A.: A generalized spectral decomposition technique to solve a class of linear stochastic partial differential equations. Computer Methods in Applied Mechanics and Engineering 196(45-48), 4521–4537 (2007)
15. Preusser, T., Scharr, H., Krajsek, K., Kirby, R.: Building blocks for computer vision with stochastic partial differential equations. Int. J. of Computer Vision 80(3), 375–405 (2008)
16. Stefanou, G., Nouy, A., Clement, A.: Identification of random shapes from images through polynomial chaos expansion of random level-set functions. Int. J. for Numerical Methods in Engineering 79(2), 127–155 (2009)
17. Weber, J., Malik, J.: Robust computation of optical flow in a multi-scale differential framework. Int. J. of Computer Vision 14(1), 67–81 (1995)
18. Wiener, N.: The homogeneous chaos. Am. J. of Mathematics 60(4), 897–936 (1938)
19. Xiu, D., Hesthaven, J.S.: High-order collocation methods for differential equations with random inputs. SIAM J. Sci. Comput. 27(3), 1118–1139 (2005)
20. Xiu, D., Karniadakis, G.E.: The Wiener–Askey polynomial chaos for stochastic differential equations. SIAM J. Sci. Comput. 24(2), 619–644 (2002)