Rejoinder

Lasse Holmström and Leena Pasanen

Department of Mathematical Sciences
University of Oulu
Finland
e-mail: lasse.holmstrom@oulu.fi*;* leena.pasanen@oulu.fi

We heartily thank the Editors of the International Statistical Review for their decision to publish our review as a discussion paper. We are grateful to the discussants for their useful comments that point out interesting current developments in statistical scale space methods not covered by our article, suggest new applications where they might be useful, and also raise relevant questions about the comparative merits of the various approaches.

Godtliebsen, Skrøseth and Wei in their contribution suggest using scale space representation as a feature generating process for example in reinforcement learning algorithms. A potential application could be patient self-management and health care delivery for a chronic disease, such as type 1 diabetes. They also ask how scale-space methods can be utilized in a Bayesian reinforcement learning approach.

The discount rate γ that determines the weight given to future returns in the reinforcement learning paradigm has a natural interpretation as a temporal scale parameter. As the return R_t depends γ, so do all other quantities that are based on it, such as the optimal action-value function and the optimal policy obtained by maximizing it. Thus, instead of an optimal action-value function and an optimal policy, we have a family of them, each member corresponding to a particular discount rate, that is, a future time horizon of rewards. For example in health care or patient self management, instead of focusing on a fixed future time horizon, decisions can then be made taking account expected returns in different time scales.

As noted by the discussants, reinforcement learning can also be formulated in a Bayesian framework. Application domain knowledge can then be incorporated in the learning process and, instead of point estimates, (posterior) distributions of the quantities of interest, such as value functions are obtained making it possible also to judge uncertainties in the estimates. Bayesian Q-learning paradigms have been proposed (e.g. Chapter 11 of Wiering and van Otterlo (2014)). However, the Gaussian process based approach described in Engel et al. (2003, 2005, 2007) and also in Wiering and van Otterlo (2014), might offer a particularly suitable setting for a scale space treatment. The action-value function under a stationary policy is regarded as a random process $Q(s, a)$, $(s, a) \in \mathcal{S} \times \mathcal{A}$. An linear generative model similar to that used in the Bayesian scale space approaches is used where the role of observations is given to the realized rewards and the unknown quantity of interest is $Q(s, a)$,

$$\mathbf{R} = \mathbf{A}\mathbf{Q} + \boldsymbol{\varepsilon}.$$

Here the constant matrix $\mathbf{A}$ depends on the discount rate γ, $\mathbf{R}$ and $\mathbf{Q}$ include rewards and action-values corresponding to a sequence of realized states and actions, and $\boldsymbol{\varepsilon}$ represents cor-

*Research supported by Academy of Finland grant numbers 250862 and 276022.

related random errors; see the papers of Engel and his co-workers for more details. The process $Q(s,a)$ is taken to be zero-mean Gaussian *a priori* with a separable covariance function $C((s,a),(s',a')) = C_{\mathcal{S}}(s,s')C_{\mathcal{A}}(a,a')$, making the prior of $\mathbf{Q}$ also Gaussian. Under such assumptions the posterior distribution of $Q(s,a)$ is Gaussian and the value of the posterior mean function $(s,a) \mapsto \mathbb{E}(Q(s,a)|\mathbf{R})$ has a simple explicit formula. The covariance functions $C_{\mathcal{S}}$ and $C_{\mathcal{A}}$ control the prior smoothness of the action-value function and therefore offer an opportunity for scale space analysis. Taking for example

$$C_{\mathcal{S}}(s,s') = c_s \exp[-\lambda_s \mathrm{dist}_{\mathcal{S}}(s,s')], \qquad C_{\mathcal{A}}(a,a') = c_a \exp[-\lambda_a \mathrm{dist}_{\mathcal{A}}(a,a')],$$

where $c_s, c_a > 0$, are constants and $\mathrm{dist}_{\mathcal{S}}$, $\mathrm{dist}_{\mathcal{A}}$ are appropriate distance measures in $\mathcal{S}$ and $\mathcal{A}$, respectively, the prior smoothness and therefore also the posterior smoothness of $Q(s,a)$ can be controlled by the parameters $\lambda_s > 0$ and $\lambda_a > 0$. Finding a good action in a state s can be based on maximizing $a \mapsto \mathbb{E}(Q(s,a)|\mathbf{R})$ and, in scale space spirit, a robust optimizer is one which is not overly sensitive to prior smoothness assumptions and therefore works in a range of scales. In some cases, e.g. when $\mathcal{A} \subset \mathbb{R}$ or $\mathcal{A} \subset \mathbb{R}^2$, perhaps even scale space inference maps could be produced to help decision making.

We thank Dutta and Ghosh for pointing out the interesting developments in supervised and semi-supervised classification. Clearly, multi-scale methods show much promise in discrimination problems and Bayesian modeling can provide a coherent and practical framework in which such an approach can also be implemented in practice. The discussion of the novel localized version of statistical depth proposed by the authors highlights yet another context where the notion of scale has a natural role. One family of techniques not included in our review are level set tree (LST) based methods (e.g. Klemelä (2009)). Like spatial depth, such LST based methods as shape and tail trees analyze multivariate objects relative to a central location. In principle, LST methods can be used to visualize functions and data sets in spaces of arbitrary dimensions and it would be interesting to consider scale space versions of these techniques (cf. Klemelä (2009, Chapter 7)). For an example of an application of level set tree analysis to multivariate data visualization, see Karttunen et al. (2014).

Park et al. in their discussion draw attention to an emerging new data type, interval-valued observations. Such data certainly are quite commonplace and it is interesting to see how the authors introduce a scale space view into their analysis. We expect work on this topic to continue and look forward for example to a more extensive discussion of the relative merits of the proposed approach and that of Park et al. (2016).

Oliveira and her co-workers in their comments correctly point out a central obstacle in extending scale space methodologies to higher dimensions, the difficulty of devising an effective visualization of the results. This is important because the intuitive summaries provided by SiZer-like maps clearly has been an important factor behind the attractiveness of these techniques. Also, feasibility of interactive data analysis would certainly be welcome and this clearly is not possible for the kind of very large data sets suggested by the authors or even for smaller data sets when simulation based inference is used, such as in some of the Bayesian approaches. We also agree that further comparisons of various versions of scale space techniques from the point of view of quality inference and computational costs would be useful for the prospective users of these methods.

Finally, Zhang and Mei in their discussion describe a promising extension of their earlier work on scale space analysis through varying coefficient linear regression to a spatial setting. Their idea adds an interesting new tool for spatial scale space data analysis and at the same time augments geographically weighted regression (GWR) in a useful manner by addressing the issue of scale dependency.

References

Engel, Y., Mannor, S., and Meir, R. (2003). Bayes meets Bellman: The Gaussian process approach to temporal difference learning. In *Proceedings of the Twentieth International Conference on Machine Learning (ICML-2003)*, pages 154–161.

Engel, Y., Mannor, S., and Meir, R. (2005). Reinforcement learning with Gaussian processes. In *ICML '05: Proceedings of the 22nd International Conference on Machine Learning*, pages 201–208, New York, NY, USA. ACM.

Engel, Y., Mannor, S., and Meir, R. (2007). Bayesian Reinforcement Learning with Gaussian Process Temporal Difference Methods. An unpublished manuscript.

Karttunen, K., Holmström, L., and Klemelä, J. (2014). Level set trees with enhanced marginal density visualization. In *Proceedings of the 5th International Conference on Information Visualization Theory and Applications, (IVAPP 2014), Lisbon, Portugal, January 5 – 8*, pages 210–217. Available on-line at http://dx.doi.org/10.5220/0004844302100217.

Klemelä, J. (2009). *Smoothing of Multivariate Data: Density Estimation and Visualization.* Wiley Series in Probability and Statistics. Wiley, Hoboken, NJ.

Park, C., Jeon, Y., and Kang, K.-H. (2016). An exploratory data analysis in scale-space for interval-valued data. To appear in Journal of Applied Statistics.

Wiering, M. and van Otterlo, M. (2014). *Reinforcement Learning: State-of-the-Art.* Springer Publishing Company, Incorporated.