Hydrol. Earth Sys. Sci. Discuss., 2, 1185–1219, 2005
www.copernicus.org/EGU/hess/hessd/2/1185/
SRef-ID: 1812-2116/hessd/2005-2-1185
European Geosciences Union

Hydrology and Earth System Sciences Discussions

Papers published in *Hydrology and Earth System Sciences Discussions* are under open-access review for the journal *Hydrology and Earth System Sciences*

A Bayesian spatial assimilation scheme for snow coverage observations in a gridded snow model

S. Kolberg[1], **H. Rue**[2], **and L. Gottschalk**[3]

[1]SINTEF Energy Research, Sem Sælands vei 11, 7465 Trondheim, Norway
[2]Department of Mathematical Sciences, NTNU, 7491 Trondheim, Norway
[3]Department of Geosciences, University of Oslo. P.O. Box 1047 Blindern, 0316 Oslo, Norway

Received: 24 June 2005 – Accepted: 6 July 2005 – Published: 25 July 2005

Correspondence to: S. Kolberg (sjur.kolberg@sintef.no)

Abstract

A spatial probability distribution of the variables in a parametric snow depletion curve (SDC) is tailored to the assimilation of satellite snow cover data into a gridded hydrological model. The assimilation is based on Bayes' theorem, in which the proposed distribution represents the a priori information about the SDC variables. From the prior gridded maps of snow storage and accumulated melt depth, the elevation gradients and the degree-day factor are separated out, creating elevation-normalised surfaces of snow storage and degree-day sum. Because the small-scale variability linked to elevation is removed, these surfaces can be described by prior distribution models with a strong spatial dependency structure. This reduction of spatial uniqueness in the prior distribution greatly increases the informational value of the remotely sensed snow coverage data.

The assimilation is evaluated in a 2400 km^2 mountainous region in central Norway (61° N, 9° E), based on two Landsat 7 ETM+ images evaluated at 1 km^2 scale. An image acquired on 11 May, a week before the peak flood, removes 78% of the variance in the remaining snow storage. Even an image from 4 May, less than a week after the melt onset, reduces this variance by 53%. Including observed discharge in the updating information improves the 4 May results, but has weak effect on 11 May. Estimated elevation gradients are shown to be sensitive to informational deficits occurring at high altitude, where snowmelt has not yet started. Caution is therefore required when using early images.

1. Introduction

The utilization of water resources in mountainous regions is often in lack of precise measurements, due to harsh climate and remote situation. This also applies to the estimation of a seasonal snow pack. The potential of remote sensing to collect spatially distributed information of the snow storage is therefore of great interest for water supply

HESSD

2, 1185–1219, 2005

or hydropower production planning, as well as flood forecasting. Satellite based measurements of snow water equivalent (SWE) is limited to coarse spatial scales, and with considerable restrictions on snow conditions. At 25 km resolution, passive microwave measurements from SSM/I are able to represent dry snow SWE fairly well in flat, low-lying regions, but show large deviations in mountainous areas (Brubaker et al., 2000). Research is currently progressing on the use of interferometric techniques on active microwave (radar) data, reporting an accuracy of 100 mm for 10×10 km resolution estimates, also restricted to dry snow (Engen et al., 2004). Further progress is expected, in particular with radar sensors planned for the near future.

In contrast, the snow covered area (SCA) can be retrieved with good accuracy at a range of scales. Both optical sensors and active microwave instruments have been used to map SCA, the latter with the advantage to measure independently of cloud coverage. Obviously, SWE and SCA are related, and using this relation to extract SWE information from SCA images represents an alternative to direct SWE observation. The simplest possible relationship is based on a binary classification of the presence of snow, which can be combined with heuristic rules for how to adjust the SWE (Rodell and Houser, 2004), to validate the simulations of a hydrological model (Sheffield et al., 2003), or to support the interpolation of SWE from point measurements (Ranzi el al., 1999; Molotch et al., 2004). These types of relationship are particularly useful in large-scale operational hydrological models, or for land surface schemes used in atmospheric models.

Multispectral techniques have proved effective in mapping sub-pixel SCA. Rosenthal and Dozier (1996) report a regression tree-based estimation of SCA from Landsat TM images, calculating fractional SCA within 30-m grid cells with a general accuracy of around 7%. Salomonsson and Appel (2004) extract fractional SCA at 500 m resolution with an accuracy of 10% using the NDSI index from MODIS data. Thus in alpine regions with deep, heterogeneous snow packs, sub-pixel SCA can be observed at scales where the SWE-SCA relation is strong, enabling monitoring and updating of the snow cover mass balance during the melt season. The main difficulty for utilising this

relation is that it depends on several variables within a spatial observation unit, whereas the SCA information is a single number. Thus, the estimation of model state suffers from a large information deficit, and the adjustment required to achieve the measured SCA is not uniquely defined.

The main aim of the current paper is to overcome the information deficit by identifying processes leading to similarity among grid cells, thus reducing the effective number of independent state variables to be updated. The development builds upon the Bayesian method presented by Kolberg and Gottschalk (2005), and introduces spatially invariant and spatially connected variables in the prior distribution. The assimilation is presented in an updating context, but also addresses variables normally considered as calibration constants rather than dynamically simulated states. The analysis is not tied to any specific hydrological model, but assumes the existence of a general, grid distributed model, which in each grid cell maintains the snow cover mass balance by a snow depletion curve (SDC) relating the mass balance to the snow covered area.

2. Assimilation by Bayesian updating

Bayes' theorem expresses how measured data $\boldsymbol{D}$ change a statistical distribution of a parameter vector $\boldsymbol{\theta}$, which may include any unknown or uncertain variable related to $\boldsymbol{D}$. The prior knowledge about $\boldsymbol{\theta}$ is expressed as a joint distribution of parameters $p(\boldsymbol{\theta})$, and the relation between the parameters and the observed data is also given in terms of a statistical distribution $p(\boldsymbol{D}|\boldsymbol{\theta})$, expressing the probability of the measured data given the parameters. Considered as a function in $\boldsymbol{\theta}$ with $\boldsymbol{D}$ as parameters, this is called the likelihood function, and its multiplication with the prior $p(\boldsymbol{\theta})$ gives the posterior distribution $p(\boldsymbol{\theta}|\boldsymbol{D})$:

$$p(\boldsymbol{\theta}|\boldsymbol{D}) \propto p(\boldsymbol{\theta}) \cdot p(\boldsymbol{D}|\boldsymbol{\theta}). \quad (1)$$

The proportionality symbol indicates the absence of the normalisation constant, which must be found by integrating the right hand side of Eq. (1). Properly normalised, the

HESSD
2, 1185–1219, 2005

posterior distribution contains our knowledge of $\boldsymbol{\theta}$ conditioned on $\boldsymbol{D}$, and is the basis for estimating elements of $\boldsymbol{\theta}$, typically by expectation (minimum squared error), mode (maximum likelihood), or median (minimum absolute error).

In the current application, $\boldsymbol{D}$ contains the satellite information, and possibly some observations of runoff. The vector $\boldsymbol{\theta}$ contains the variables describing the SDC in each grid cell in the region. The prior distribution $p(\boldsymbol{\theta})$ has a joint expectation $E[\boldsymbol{\theta}]$, typically consisting of the simulated values in the hydrological model. Similarly, the expectation $E[\boldsymbol{\theta}|\boldsymbol{D}]$ of the posterior $p(\boldsymbol{\theta}|\boldsymbol{D})$ represents updated estimates with which we can re-initialize the hydrological model. The prior $Var[\boldsymbol{\theta}]$ and posterior $Var[\boldsymbol{\theta}|\boldsymbol{D}]$ variance describe the uncertainty before and after updating, respectively, and from these we can compute the fractional reduction of variance, indicating the informational value of $\boldsymbol{D}$.

The regarding of hydrological model parameters as stochastic variables, with distributions subject to updating by conditioning on measurements, was introduced by Binley and Beven (1991). Their GLUE technique (Generalized Likelihood Uncertainty Estimation), follows a Bayesian approach, but relaxes the formal rigour of likelihood construction. A principal discussion of such relaxations appears in Beven and Young (2003) and Gupta et al. (2003). Engeland and Gottschalk (2002) apply Bayes' theorem to the determination of parameters in a regional model, based on streamflow observed in several catchments within the region. Recently, methods for combined parameter and state estimation by data assimilation have been presented by Moradkhani et al. (2005) and Vrugt et al. (2005).

By their spatial nature, satellite data are particularly valuable for calibration and validation of distributed models. Blöschl et al. (1991) used observed snow cover patterns to evaluate model performance, by comparing simulated and observed maps. By relating the patterns to process-governing terrain features, they were able to link qualitatively the observation errors to specific process representations in the model. Kirnbauer et al. (1994) concluded that spatial observations of snow cover patterns were clearly able to reveal poor model assumptions, to which the measured streamflow time

series was insensitive. The current approach can be used to extend such inference to quantitative estimates of errors and uncertainty, and to express the severity of ill-posed assumptions in terms of likelihood of observations.

3. Approach and method

The snow depletion curve (SDC) concept is extensively discussed by Liston (1999, 2004), Luce et al. (1999), and Luce and Tarboton (2004). Applied to a spatial unit with heterogeneous snow storage, the SDC describes how the snow covered area (SCA) changes with the progress of melt. Assuming melt rate to be sub-grid homogeneous, the SDC is also equivalent to the sub-grid cumulative spatial distribution of point snow storage at the start of the melt season (Liston, 1999; Luce et al., 1999). In the SDC formulation, most of the variables are static descriptors of the end-of-winter snow pack, with the accumulated melt depth (termed λ) as the only necessary dynamic variable. As argument to the SDC, λ defines the bare ground fraction y, the accumulated melt runoff Q, and the remaining snow storage SWE (Fig. 1).

Luce and Tarboton (2004) analyse nine years' empirical SDC for a wind-swept, low-vegetation mountainous area, and conclude that normalised SDCs are remarkably stable between years, and that the choice of a parametric model is less important than the estimation of a coefficient of variation. Both of these results are important for applying a SDC model to each grid cell in a large region, because even an extensive survey will only cover a small fraction of the cells. It is thus necessary to use a parsimonious statistical model for the SDC. Several parametric models have been proposed, for instance 2-parameter Lognormal (Essery et al., 1999; Liston, 2004), 3-parameter Lognormal (Donald et al., 1995), 3-parameter Beta (Brubaker and Menoes, 2001), or 3–5-parameter weighted combinations of Normal and/or Lognormal distributions (Bruland et al., 2001; Marchand and Killingtveit, 2004).

In this study, a 3-parameter mixed distribution is selected to describe the SDC in each grid cell. The parameter y_0=prob(x=0) gives the fractional area of bare ground

in the cell at the time of melt onset. In the remaining, snow covered part of the cell, the distribution of SWE is described by a 2-parameter Gamma model. This Gamma model is parameterized by the mean value m and coefficient of variation cv, both characterizing the snow pack at the end of the accumulation season. This choice makes it easier to specify prior distributions, both because m and cv more directly reflect the typically available prior information, and because the spatial distributions of the traditional shape and scale parameters are shown to be strongly dependent (Kolberg, 2001). The snow depletion curve gives the bare-ground fraction y as a function of the accumulated melt depth λ:

$$y = SDC\left(\lambda | m, cv, y_0\right) = y_0 + (1 - y_0)\, y' \tag{5}$$

$$y' = \int_0^{\lambda} f\left(x; m, cv\right) dx = \gamma\left(\frac{1}{cv^2}, \frac{\lambda}{m \cdot cv^2}\right) \tag{6}$$

Here, $\gamma(\cdot,\cdot)$ is the incomplete gamma function, expressing the cumulative Gamma distribution with shape cv^{-2} and scale $m^{-1}cv^{-2}$, evaluated in λ. It is worth noting that the main mass balance governing variables m and λ appear only as a ratio, and thus that an observation of the bare-ground fraction y only provides information on the relative magnitude of the two.

4. The spatial prior distribution

Kolberg and Gottschalk (2005) presented a Bayesian method for updating the state of a snow cover model formulated as above, applied separately to each grid cell in a region. The results in terms of variance reduction were moderate, and in particular the mass balance components variables λ and m did not achieved substantial increase of precision. Introducing spatial dependency was identified as one possible route to increase the value of the observations (Kolberg and Gottschalk, 2005). The current

paper presents a spatial model for the mass balance variables, transforming the accumulated melt depth λ and the end-of-winter snow storage m into new variables for which the spatial surfaces are smoother.

4.1. Identification of spatially constant variables

Removal of trends, periodicities and other deterministic features from a surface prior to fitting a spatial dependency structure is common in geostatistics. For the average snow storage m we identify an elevation gradient g_m, which is considered global, that is, having no spatial variance. Expressed in relative increase per unit elevation, this gradient may be approximated by the similar gradient in precipitation, neglecting pre-spring melt events and different accumulation periods. Removing the effect of g_m from m results in an elevation-normalized expected snow storage m^*. The transformation from m to $m^*$is:

$$m = m^* (1 + g_m)^{(h-h^*)} \quad ; \quad g_m = \frac{1}{m}\frac{\partial m}{\partial h} \tag{7}$$

Here, h is elevation, and the asterisk denotes a reference altitude; typically the elevation which is best represented by measurements. As is common for the precipitation lapse rate, the elevation gradient in m is expressed in relative terms, and thus the elevation normalization must be multiplicative to ensure consistency and independence of the reference altitude.

Regarding the accumulated melt depth, λ is first decomposed into a degree-day factor Cx and a sum of positive temperatures z. The degree-day sum z is further decomposed into an elevation gradient g_z and an elevation-normalized degree-day sum z^*. The chosen decomposition does not require the hydrological model to be based on these concepts, although this will frequently be the case. Formally, the transformations from λ to z^* are:

$$\lambda = Cx \cdot z \quad ; \quad z = MAX\,(0, z^* + g_z\,(h - h^*)) \quad ; \quad g_z = \frac{\partial z}{\partial h} \tag{8}$$

Again, h is elevation, the asterisk indicates a reference altitude, while Cx is the degree-day factor and z the local degree-day sum. The elevation gradient g_z applies to the aggregated sum of positive degree-days, not to momentary temperature values, and consequently it is a dynamic variable. Its formulation as a linear trend is an approximation, not a result of its connection to the temperature lapse rate. The calculation of z may include commonly applied nonlinearities like a zero-melt threshold temperature, a refreezing efficiency factor, or a maximum cold content in the snow pack.

Over a region of n grid cells, the prior distribution now consists of $4n+3$ variables; namely the four gridded maps of m^*, z^*, cv, and y_0, and the three global variables g_m, g_z, and Cx. The global variables are constant in space, but they are still stochastic variables, specified by distributions and subject to Bayesian updating. A priori, we assign Normal distributions to each of the two elevation gradients, and a Gamma distribution to the temperature index. The sub-grid coefficient of variation cv and the initial bare ground fraction y_0 are neither transformed nor given any spatial connectivity in their distributions. A Gamma prior distribution is applied to to cv in each grid cell, and y_0 is similarly given local Lognormal priors. It remains to specify the spatially connected prior distributions for m^* and z^*.

4.2. Constructing spatial distributions of m^* and z^*

With the elevation gradients removed, the spatial surfaces of z^* and m^* naturally contain less variance than the corresponding surfaces of λ and m. More important, the variance components removed with the global effects mainly follow the spatial scale of the terrain, which is small compared to the spatial scale of the reference-altitude surfaces. Thus, the transformed surfaces not only have reduced variance, they are also smoother than the original. Spatial data sets with this property are ideal for modelling by spatially connected distributions.

A convenient class of spatial models is the Gaussian Markov Random Field (GMRF). Contrary to a full multi-normal distribution, which in a region of n grid cells would have a $n \times n$ covariance matrix, the Markov property states that all spatial dependency is

described by a limited neighbourhood (Rue and Tjelmeland, 2002). Hence, spatially connected prior distributions can be specified in terms of conditional moments, each cell depending only on its neighbours. In this case the neighbourhood is set to contain the four directly adjacent cells, and for z^*, the conditional moments are given as:

$$E\left[z_i^*|z_{\{j\}}^*\right] = E\left[z_i^*\right] + \frac{\sigma_i}{n_j}\sum_{j=1}^{n_j}\left(\frac{z_j^* - E\left[z_j^*\right]}{\sigma_j}\right) \quad ; \quad Var\left[z_i^*|z_{\{j\}}^*\right] = \sigma_i^2 \Big/ n_j \tag{9}$$

Here, index i denotes the local cell, $\{j\}$ denotes its n_j neighbours, and σ_i^2 denotes the marginal variance in grid cell i. Formally, this is a first order Intrinsic GMRF applied to the standardized variable, standardisation being necessary because both marginal moments vary over space. The full stochastic model for λ_i is then given by:

$$\lambda_i = Cx \cdot z_i \tag{10}$$

$$Cx \sim Gamma\left(\frac{E^2[Cx]}{VAR[Cx]}, \frac{E[Cx]}{VAR[Cx]}\right), \text{ (by method of moments)} \tag{11}$$

$$z_i = MAX\left\{0\ ,\ z_i^* + g_z(h_i - h^*)\right\} \tag{12}$$

$$\frac{\partial z}{\partial h} \sim N\left(E[g_z], VAR[g_z]\right) \tag{13}$$

$$z_i^*|z_{\{j\}}^* \sim N\left(E\left[z_i^*\right] + \frac{\sigma_i}{n_j}\sum_j\left(\frac{z_j^* - E\left[z_j^*\right]}{\sigma_j}\right)\ ,\ \sigma_i^2 \Big/ n_j\right) \tag{14}$$

The conditional moments of m^* are defined identically as for z^*, and the full stochastic model for m is developed similarly to that for λ, except that the Cx term is excluded,

HESSD
2, 1185–1219, 2005

the lapse rate is in relative terms (and thus multiplicative), and the MAX function is obsolete.

The introduction of spatial connectivity not only reduces the effective number of variables to be estimated, but also aids in moving information between cells. This is important because the information content of an observation varies with both the prior and the observed conditions (Kolberg and Gottschalk, 2005). A grid cell with prior $E[\lambda]$ close to 0 receives little or no mass balance information if the observed y is close to y_0. But if grid cells within the range of spatial influence have larger λ due to a lower elevation, or a higher cv giving higher y for the same λ, these cells do receive such information, affecting the first cell by the spatial dependence. A special effect of this is that updating may take place also in grid cells with missing observations. Optical satellite images are notoriously sensitive to cloud coverage, and the value of an updating routine greatly increases with the usefulness of partially cloudy scenes.

4.3. Estimating the prior distributions

Estimating the prior moments of the transformed variables in an actual situation, the prior $E[z]$ is provided either from degree-day model simulations of λ with a known Cx, or from interpolating time-aggregated air temperature data. The prior $E[g_z]$ is then estimated from the map of $E[z]$, using only grid cells with substantially positive z. Similarly, prior $E[m]$ may be taken either from the hydrological mass balance simulations or from a geostatistical analysis of point precipitation and/or snow survey data; and $E[g_m]$ is estimated from the prior $E[m]$ map. Estimates of prior marginal variance of m^* and z^* are provided by geostatistical analysis, or can be assessed subjectively if these variables are taken from the hydrological simulations. The conditional variances are then the marginal scaled down by a factor n_j (Eq. 9).

Prior moments for cv are assessed from published and unpublished data from historical snow surveys in Norwegian mountains (Kolberg, 2001). For y_0, the assessment is largely subjective (Kolberg and Gottschalk, 2005). Provided that these parameters exhibit some regularity in their inter-annual variation, information accumulated from a

HESSD
2, 1185–1219, 2005

large number of images would otherwise provide a good basis for estimating the local prior distributions of cv and y_0. For the three global variables, the prior variance is also assessed and set subjectively in the current analysis, but could be supported by multi-altitude measurements, or Monte Carlo-based calibration by methods presented by Binley and Beven (1991) or Vrugt et al. (2003).

5. The likelihood function and the observations

5.1. SCA-based likelihood

Having specified the prior, the second term on the right hand side of Bayes' theorem is the likelihood of the parameters, given the observations. For a single observed bare-ground fraction y^{obs} in grid cell i and the image acquired on date t, a Beta distribution is selected to represent the likelihood:

$$P\left(y_{it}^{obs}|m_i^*, g_m, z_{it}^*, g_{zt}, Cx, cv_i, y_{0i}\right) = \frac{\Gamma\left(\varphi_{it}+\psi_{it}\right)}{\Gamma\left(\varphi_{it}\right)\Gamma\left(\psi_{it}\right)}\left(y_{it}^{obs}\right)^{\varphi_{it}-1}\left(1-y_{it}^{obs}\right)^{\psi_{it}-1} \quad (15)$$

Its parameters are determined by the method of moments:

$$\varphi_{it}=E\left[y_{it}\right]\left(\frac{E\left[y_{it}\right]\left(1-E\left[y_{it}\right]\right)}{Var\left[y_{it}^{obs}\right]}-1\right) \quad ; \quad \psi_{it}=\left(1-E\left[y_{it}\right]\right)\left(\frac{E\left[y_{it}\right]\left(1-E\left[y_{it}\right]\right)}{Var\left[y_{it}^{obs}\right]}-1\right) \quad (16)$$

The likelihood expectation $E[y_{it}]$ is the simulated y value given all SDC parameters:

$$E\left[y_{it}|m_i^*, g_m, z_{it}^*, g_{zt}, Cx, cv_i, y_{0i}\right] =$$

$$y_{0i}+\left(1-y_{0i}\right)\cdot\gamma\left(\frac{1}{cv^2}, \frac{Cx\cdot MAX\left\{0, z_{it}^*+g_{zt}\Delta h\right\}}{cv^2\cdot m_i^*\left(1+g_m\right)^{\Delta h}}\right) \quad (17)$$

The variance $Var[y_{it}]$ is the estimated observation variance from the satellite image analysis.

Since all the y values for a single cell over n observations are considered conditionally independent, the multi-observation likelihood for a single cell is simply given by the product of the single-event likelihoods. Correspondingly, since the observations in different cells are also considered conditionally independent, the likelihood for the total set of observations is the product of all the single-cell likelihoods.

5.2. Discharge-based likelihood

The transition of the assimilation task from a cell-by-cell independent routine to a spatial problem with both global and spatially connected variables, increases the relevance of any runoff measurements in the region. Engeland and Gottschalk (2002) have used measured discharge time series in Bayesian inference, and developed likelihood models including autoregressive terms. In this paper, however, only the accumulated runoff Q_{acc} to date is used. The difference between catchment runoff and outlet discharge, corresponding to the change in the internal water storage S_t, is modelled by a simple linear tank. The tank content is assumed to be zero at the melt onset, and estimated from the measured discharge on the actual day.

The storage term is assumed to be the major source of uncertainty in the catchment runoff as estimated from measured discharge, and is subjectively assigned a standard deviation of half its size. The relative uncertainty in the measured accumulated discharge is, again subjectively, assessed to 5%. The uncertainty contributions are rescaled to variance and added together:

$$\sigma_{qs}^2 = Var\left[Q_{acc} + S_t\right] = Var\left[Q_{acc}\right] + Var\left[S_t\right] = \left(Q_{acc} \cdot 0.05\right)^2 + \left(S_t \cdot 0.5\right)^2 \quad (18)$$

Similarly to the SCA-based likelihood, the discharge-based likelihood expectation is given from the SDC parameter values, and the variance from the observations. The discharge-based likelihood is assumed to follow a Normal model, and with observed

HESSD
2, 1185–1219, 2005

accumulated runoff Q_{acc} and storage S_t estimated from actual-day discharge Q_t, the likelihood for the sum $Q_{acc}+S_t$ is:

$$P\left(Q_{acc}+S_t|\bar{Q}_t\right)=\frac{1}{\sqrt{2\pi}\sigma_{qs}}\exp\left[-\frac{1}{\sigma_{qs}^2}\left(Q_{acc}+S_t-\bar{Q}_t\right)^2\right] \quad ; \quad \bar{Q}_t=\frac{1}{n}\sum_{i=1}^{n}Q_i(\theta_{i,t}) \qquad (19)$$

where $Q_i(\theta_{i,t})$ is accumulated runoff at cell i and to date t, given the local SDC parameter set $\theta_{i,t}$. Whether or not the discharge-based likelihood term is included, the posterior distribution is analytically intractable, and is sampled by Markov Chain Monte Carlo techniques using the Metropolis-Hastings algorithm (Chib and Greenberg, 1995).

6. Site and data

The assimilation of observed snow covered area into spatially distributed snow depletion curves is evaluated over a 60×40 km rectangular region covering the Vinstra and Sjoa catchments in Jotunheimen, central Norway, at 61.4° N, 8.6° E. The elevation ranges from 710 m a.s.l. to 2240 m a.s.l measured at grid scale. Sparse forest and lakes each cover about 10% of the area, small glaciers are present at high altitudes. The accumulation season usually starts in November, and snowmelt in late April or May, depending on elevation.

From two Landsat 7 ETM+ images acquired during the 2000 melt season, gridded maps of fractional bare ground y and its standard error are estimated at 30 m scale, using the decision tree algorithm of Rosenthal and Dozier (1996). Minor adjustments are made to adapt the Landsat 5 TM-based routine to the Landsat 7 ETM+ sensor, which has a different sensitivity in some of the bands. Cloud and lake masks were applied at the same scale, discarding the corresponding 1 km cell if occupying more than 20% of its area.

Runoff observations exist in 6 catchments, which total area covering approximately half the total region. Five of these are nested, heavily regulated subcatchments in

the Vinstra drainage area, with a total area of 744 km^2; whereas Sjoa is 473 km^2 and unregulated. Naturalized flow series from regulated catchments are known to be of reduced quality on daily and sub-catchment scale, but accumulated over a larger period and area, the uncertainty in the discharge observations from the naturalized series approaches the unregulated.

For the 2000 melt season, 9 updating experiments are run; with satellite data sets from 4 and 11 May used separately and combined; and runoff information either not included, aggregated from the whole gauged part of the region, or from the Sjoa and Vinstra catchments separately. Simulated melt depths suggest a melting front at approximately 1500 m a.s.l. on 4 May, and at approx. 1800 m a.s.l. on 11 May.

7. Results

7.1. Direct comparison of prior and observed bare-ground fraction y

Before starting to analyse the results of the Bayesian updating, Figs. 2 and 3 both compare the observations to the simulated, a priori snow state. Figure 2 shows that on both dates, the observed bare-ground fraction y^{obs} is underestimated by prior $E[y]$ both at high and low elevation, but is approximately correct around 1100 m a.s.l. Figure 3 displays the spatial histogram of prior cumulative probabilities of the observed value on 11 May. A slight over-representation of cumulative probabilities above 0.8 reflects the tendency to underestimate y^{obs}, but the striking feature of Fig. 3 is the large frequency at the right tail. The prior distribution is obviously too restrictive towards the multi-dimensional tail corresponding to high y values. The tendency at the left tail is opposite, but weaker. Bearing in mind that the prior is subjectively specified with emphasis on marginal moments, the quantitative posterior estimates should be assessed with care in grid cells where the observation is unlikely high.

7.2. Changes to the end of winter snow storage $E[m]$

In the following, all graphical illustrations are based on the updating experiment using the 11 May image, and not using any runoff-based likelihood term. Figures 4a–4d shows how the expected pre-melt snow storage $E[m]$ and its reference-elevation counterpart $E[m^*]$ are updated. Comparing the prior and posterior $E[m]$ maps (top row), it is evident that the posterior elevation gradient g_m is stronger than the prior. In addition, there are local changes to the reference-altitude $E[m^*]$ in the eastern part of the region (bottom row). Note the smoothness of the $E[m^*]$ maps compared to those of $E[m]$. The changes to $E[m]$ and $E[m^*]$ agree with the similar changes on 4 May with correlation coefficients of 0.87 and 0.81, respectively.

For the snow storage gradient g_m, the prior and posterior expectations and standard deviations are given in Table 1. The increased elevation dependency seen in Fig. 4 is confirmed by the 11 May g_m increasing from approx. 2.6%/100 m to more than 4%/100 m. Even more dramatic, the 4 May image results in a posterior $E[g_m]$ of more than three times the prior; far out on the prior distribution tail. The use of both images yields an intermediate $E[g_m]$, but closer to the moderate increase of 11 May. Including observed discharge in the likelihood produces a slight reduction in posterior $E[g_m]$, most notably for the 4 May image.

7.3. Changes to the accumulated melt depth E[λ]

Figure 5 shows the updating of $E[\lambda]$. Here, the elevation dependency is attenuated, with melt extending to a greater altitude in the posterior than in the prior case. Again, there are local changes as well, for instance in the south-eastern part of the area, but the gradient driven changes dominate. The changes to the dynamic $E[\lambda]$ and $E[z^*]$ on 11 May agree with the changes to these variables on 4 May with correlation coefficients of 0.85 and 0.89, respectively.

Table 2 shows the updating of the degree-day factor Cx and the degree-day sum gradient g_z, for the actual date. For $E[g_z]$, all experiments result in a large increase,

i.e. an attenuation of the negative gradient, the 4 May image more than on 11 May. Observe that the two updates address different g_z, since the degree-day sum z is dynamic. Using both images strengthens the 4 May $E[g_z]$, but further attenuates the 11 May $E[g_z]$, compared to the single-image posterior estimates. Adding discharge information has only marginal effect. For Cx, only moderate adjustments are made, and posterior $E[Cx]$ is well within the prior uncertainty.

Figure 6 shows the prior and posterior $E[cv]$, quantifying the sub-grid heterogeneity of snow storage and determining the shape of the Gamma model used for the SDC. The main effect of updating is an increase in the small-scale variability, evident as more scatter in the image. Also for this variable, there are instances of local neighbourhoods showing a common response. Since the cv prior has no spatial structure, this is either a direct result of the observations, or a result of posterior dependencies between cv and m^* or z^*, which do have spatial priors. For $E[cv]$, the correlation between 4 May and 11 May updates is 0.68.

The prior and posterior y_0 maps provide little valuable spatial information, and are omitted. $E[y_0]$ changes are generally small, and negative in the majority of the cells. However, the 11 May update cause an increase in $E[y_0]$ of more than 0.15 in approximately 50 grid cells, and the maximum posterior value is as high as 0.72. On 4 May, only 11 pixels have $E[y_0]$ increased by more than 0.15, with 0.46 as the maximum posterior value. The correlation coefficient between the 4 May and the 11 May updates is 0.57.

7.4. Uncertainty reduction

A comparison of prior and posterior uncertainty is useful to quantify the information content in the SCA observations, with respect to the SDC parameters and the derived snow states. Table 3 shows that the 11 May image provides more information on the mass balance governing m and λ. The earlier 4 May image provides more information on the variables cv and y_0, which mainly affect the sub-grid snow distribution. The two images used together generally reduce variance more than any of them alone,

except for y_0 from the 4 May image and λ from the 11 May image. The inclusion of likelihood terms from discharge measurements provide little or no extra information for the 11 May and the 4 and 11 May updates, but significantly improves the mass balance precision on 4 May.

8. Discussion

Evaluating the relative variance reductions (Table 3), substantial improvement is achieved for all variables. Even with the 4 May image, less than a week after melt onset, the variance in remaining snow storage SWE is reduced by more than 50% on average. On 11 May, the average variance reduction is more than 75%, i.e. the posterior variance is less than a quarter of the prior. The similar method used with independent priors in each grid cell (Kolberg and Gottschalk 2005) reduced the prior variance in SWE on 4 May and 11 May by only 12% and 15%, respectively. Obviously, the spatial prior largely compensates the information deficit apparent for each individual cell, at the cost of some spatial detail.

8.1. Gradient estimation difficulties

For the 4 May image, the posterior snow storage gradient $E[g_m]$ of nearly 10%/100 m (Table 1) is beyond credibility. The poor variance reduction in snow storage m on 4 May (Table 3) is due to this large value, because g_m is in relative terms and thus increases the m uncertainty at high altitude. Also the posterior degree-day sum gradient g_z (Table 2) is extreme. Noticing that the observed changes in $E[g_m]$ and $E[g_z]$ affect $E[y]$ in opposite directions, the two gradients appear to attain unrealistic values while compensating each others effect on the posterior y. Recalling that the likelihood only relates y to the mass balance through the λ/m ratio (Eq. 6), changes in λ and m easily compensate. With elevation as a major source of variability in both λ and m, this compensation could transfer to the two gradients. In the posterior distribution, g_m and

g_z are indeed positively dependent with an R^2 of 0.30 on 4 May, 0.40 on 11 May. This dependency is noticeable, but cannot alone explain the extreme values, in particular noticing the low posterior variance.

More than a problem of weak identification due to mutual dependency, the unlikely gradients on 4 May is a problem of extrapolation. Above the melting front, y is insensitive both to m and to any realistic change in λ, leaving the SCA observations almost non-informative. Figure 2 indicates that the elevation dependency of y is a priori underestimated below 1100 m a.s.l., and overestimated between 1100 m and 1400 m, while above 1400 m, y is close to y_0. With no restrictions from high-altitude information, g_m and g_z adapt freely to the situation below 1400 m a.s.l., and combine to mimic the observed gradient both below and above 1100 m. Even below 1400 m, there are plenty of cells contributing to the likelihood, ensuring the low posterior variance. In contrast, the poor two observations in the runoff likelihood only suffice to slight moderations in the posteriors which they strongly contradict.

8.2. Comparing the updates to measurements

In Table 4, the posterior end of winter snow storage $E[m]$ and runoff to date Q produced by the different updating experiments are averaged for the two measured catchments, and compared to their prior expectations as well as observed values. Average snow storage is compared to the accumulated runoff at the end of the melt season, corrected for precipitation during the snow melt period. $E[m]$ is underestimated in the prior state, and increased by all updating experiments, in particular through a higher gradient. For the 4 May image, the high gradient discussed above yields an over-correction of $E[m]$; and in Sjoa, the posterior absolute error is even larger than the prior except when two observed runoff series are used.

Concerning accumulated runoff to date Q, all the updating experiments result in increased estimates. For the Vinstra catchment, the posterior $E[Q]$ is largely confirmed by the observed values, again with the changes on 4 May slightly too large. For Sjoa,

however, observed runoff on both dates corresponded well to the prior $E[Q]$, whereas the 4 May update produces a three-fold increase in both runoff and melt depth. Less dramatic, but still too large, is the runoff increase produced by the 11 May update. These changes are mainly caused by the attenuation of the g_z gradient discussed above. The situation improves slightly by including runoff observation in the likelihood, for the 4 May case also by using both images.

A temperature station at 1600 m a.s.l. is not used in the prior estimation, and combined with one at 850 m a.s.l. it provides an independent evaluation of the posterior g_z estimates. Calculated gradients g_z are −7.73 and −13.67 degree-days/100 m on 4 May and 11 May, respectively. The prior distributions of g_z capture these observations within +/−1 standard deviation on both dates (Table 2), but the posteriors fail to do so, both by larger errors and by smaller standard deviations.

8.3. Small-scale variability and the role of y_0

The spatial restrictions on m and λ prevent these variables from being too hardly altered by unexpected observations in isolated cells. Consequently, the local variables cv and y_0 are more likely to respond to small-scale variability, even if this may be caused by local m or λ effects like a net transport loss of snow, or an increased melt due to sun exposure. This may be the reason why y_0 is poorer defined by the use of both images, than by the 4 May image only (Table 4). Figure 7 shows how y_0 is the responding posterior dimension when extremely high y values are observed. No other variable is changed this far out of normal range by extreme observations. Neither does any single variable respond similarly to low y observations, which are more easily captured by a moderate changes in all the parameters.

Recognising that y_0 does not contribute to any relationship between snow coverage and melt depth, one might consider removing this parameter or fixing its value close to 0. For predictive purposes, observations beyond reasonable prior credibility could be simply disregarded, leaving the actual grid cells to neighbourhood updating. For investigative purposes, the increase in observed y between the two dates at high alti-

tude, with no elevation dependency above 1700 m a.s.l. (Fig. 2), suggests that neither the "initial" property of y_0, nor the degree-day model consequence of a distinct melting front, are well justified.

9. Conclusions

Identifying processes that operate similarly on many grid cells, is a powerful way to increase the informational value of remotely sensed snow coverage data. In this study, the isolation of two elevation gradients and a temperature melt index as spatially constant variables enabled a strong spatial dependence to be built into the prior distribution of the mean snow storage m and the accumulated melt depth λ. The transformed, spatial prior distribution produced largely better results than a non-spatial prior model previously reported.

The information content of a snow coverage observation depends on the situation it images. Early observations contain more information about the snow cover heterogeneity, in this case the sub-grid coefficient of variation cv and the pre-melt bare ground fraction y_0. In particular, observations prior to the melt onset are close to non-informative except for y_0. However, the spatial prior enables mass balance information gained in some grid cells, to be transferred to neighbouring or similar-altitude grid cells where the SCA observation was less informative.

The inability of a local SCA observation to contain mass balance information beyond relative terms, is to some extent reflected in the spatial data as compensation between the elevation gradients. In this case an early image leads to some unrealistic results, largely because the two gradients combine to adapt to a non-linear feature within a limited elevation interval. In the common situation with meteorological measurements concentrated at low altitude, an early image lacking effective information at high-altitude is likely to produce estimates sensitive to extrapolation.

Combination of two images usually provides a greater reduction of variance than any of the images alone. The inclusion of observed runoff in the likelihood, however, gen-

HESSD
2, 1185–1219, 2005

erally supplies weak explanatory power, compared to the satellite image. This is due to the discharge being observed in only two catchments, whereas the SCA likelihood variance is greatly reduced by multiplying more than 2000 single-cell likelihood distributions. However, in the 4 May case, where the mass balance information of the image is smallest, the runoff observation is able to reduce some of the error in the gradient estimation, as well as the posterior variance in the mass balance variables.

References

Beven, K. J. and Young, P.: Comment on "Bayesian recursive parameter estimation for hydrologic models" by Thiemann, M., Trosset, M., Gupta, H., and Soorooshian S., Water Resources Research, 39, 5, doi:10.1029/2001WR001183, 2003.

Binley, A. M. and Beven, K. J.: Physically-based modelling of catchment hydrology: A likelihood approach to reducing predictive uncertainty, in: Computer modelling in the environmental sciences, edited by: Farmer, D. G. and Rycroft, M. J., Clarendon, Oxford, 1991.

Blöschl, G., Kirnbauer, R., and Gutknecht, D.: Distributed Snowmelt Simulations in an Alpine Catchment 1. Model Evaluation on the Basis of Snow Cover Patterns, Water Resources Research, 27, 12, 3171–3179. 1991.

Brubaker, K. L., Jasinski, M., Chang, A. T., and Josberger, E.: Interpolating sparse surface measurements for calibration and validation of satellite-derived snow water equivalent in Russian Siberia, in: Remote sensing and Hydrology 2000, edited by: Owe, M., Brubaker, K., Ritchie, J., and Rango, A., IAHS publ. No. 267, 93–98. 2001.

Brubaker, K. L. and Menoes, M.: A technique to estimate snow depletion curves from time-series data using the Beta distribution. Proceedings, 58th Eastern Snow Conference, Ottawa, Ontario, Canada, 2001.

Bruland, O., Sand, K., and Killingtveit, Å.: Snow distribution at a high arctic site at Svalbard, Nordic Hydrology, 32, 1, 1–12. 2001.

Chib, S. and Greenberg, E.: Understanding the Metropolis-Hastings algorithm, The American Statistician, 49, 4, 327–335. 1995.

Donald, J. R., Soulis, E. D., Kouwen, N., and Pietroniro, A.: A land cover-based snow cover representation for distributed hydrologic models, Water Resources Research, 31, 4, 995–1009. 1995.

Engeland, K. and Gottschalk, L.: Bayesian Estimation of Parameters in a regional Hydrological Model, Hydrol. Earth Sys. Sci., 6, 5, 883–898, 2002, SRef-ID: 1607-7938/hess/2002-6-883.

Engen; G., Guneriussen, T., and Overrein, Ø.: Delta-K interferometric SAR technique for snow water equivalent (SWE) retrieval, IEEE Geoscience and Remote Sensing Letters, 1, 2, 57–61. 2004.

Essery, R., Li, L., and Pomeroy, J.: A distributed model of blowing snow over complex terrain, Hydrological Processes, 13, 2423–2438, 1999.

Gupta, H. V., Thiemann, M., Trosset, M., and Sorooshian, S.: Reply to comment by K. Beven and P. Young on "Bayesian recursive parameter estimation for hydrologic models), Water Resources Research, 39, 5, doi:10:1029/2002WR001405, 2003.

Kirnbauer, R., Blöschl, G., and Gutknecht, D.: Entering the era of distributed snow models, Nordic Hydrology, 25, 1–24, 1994.

Kolberg, S. A.: A snow routine for Bayesian updating. Proceedings, 13th Northern Research Basins international symposium and workshop, Saarisälke, Finland, 19–24 August 2001, 121–130. 2001.

Kolberg, S. A. and Gottschalk, L.: Updating of snow depletion curve with remote sensing data, Hydrological Processes, in press, 2005.

Liston, G. E.: Interrelationships among Snow Distribution, Snowmelt, and Snow Cover Depletion: Implications for Atmospheric, Hydrologic, and Ecologic modelling, Journal of Applied Meteorology, 38, 1474–1487. 1999.

Liston, G. E.: Representing Subgrid Snow Cover Heterogeneities in Regional and Global Models, Journal of Climate, 17, 6, 1381–1397. 2004.

Luce, C. H., Tarboton, D. G., and Cooley, K. R.: Sub-grid parameterization of snow distribution for an energy and mass balance snow cover model, Hydrological Processes, 13, 1921–1933, 1999.

Luce, C. H. and Tarboton, D. G.: The application of depletion curves for parameterization of subgrid variability of snow, Hydrological Processes, 18, 1409–1422. 2004.

Marchand, W.-D. and Killingtveit, Å.: Statistical properties of spatial snowcover in mountainous catchments in Norway, Nordic Hydrology, 35-2, 101–117, 2004.

Molotch, N. P., Fassnacht, S. R., Bales, R. C., and Helfrich, S. R.: Estimating the distribution of snow water equivalent and snow extent beneath cloud cover in the Salt-Verde River Basin, Arizona, Hydrological Processes 18, 1595–1611. 2004.

Moradkhani, H., Sorooshian, S., Gupta, H. V., and Houser, P. R.: Dual state-parameter estimation of hydrological models using ensemble Kalman filter, Advances in Water Resources, 28, 135–147, 2005.

Ranzi, R., Grossi, G., and Bacchi, B.: Ten years of monitoring areal snowpack in the Southern Alps using NOAA-AVHRR imagery, ground measurements and hydrological data, Hydrological Processes, 13, 2079–2095. 1999.

Rodell, M. and Houser, P. R.: Updating a land surface model with MODIS-derived snow cover, Journal of Hydrometeorology, 5, 1064–1075, 2004.

Rosenthal, W. and Dozier, J.: Automated mapping of montane snow cover at subpixel resolution from the Landsat Thematic Mapper, Water Resources Research, 32, 1, 115–130. 1996.

Rue, H. and Tjelmeland, H.: Fitting Gaussian Markov Random Fields to Gaussian Fields, Scandinavian Journal of Statistics, 29, 31–49. 2002.

Salomonsson, V. V. and Appel, I.: Estimating fractional snow cover from MODIS using the normalized difference snow index, Remote sensing of environment, 89, 351–360, 2004.

Sheffield, J., Pan, M., Wood, E. F., Mitchell, K. E., Houser, P. R., Schaake, J. C., Robock, A., Lohmann, D., Cosgrove, B., Duan, Q., Luo, L., Higgins, R. W., Pinker, R. T., Tarpley, J. D., and Ramsay, B. H.: Snow process modelling in the North American Land Data Assimilation System (NLDAS): 1. Evaluation of model-simulated snow cover extent, J. Geophys. Res., 108, D22, 8849, doi:10.1029/2002JD003274. 2003.

Vrugt, J. A., Gupta, H. V., Bouten, W., and Sorooshian, S.: A shuffled Complex Evolution Metropolis algorithm for optimisation and uncertainty assessment of hydrologic model parameters, Water Resources Research, 39, 8, doi:10.1029/2002WR001642, 2003.

Vrugt, J. A., Diks, C. G. H., Gupta, H. V., Bouten, W., and Verstraten, J. M.: Improved treatment of uncertainty in hydrologic modeling: Combining the strengths of global optimization and data assimilation, Water Resources Research, 41, doi:10.1029/2004WR003059, 2005.

Table 1. Changes in expectation and standard deviation for the $E[SWE]$ elevation gradient, resulting from different combinations of images and discharge data in the likelihood. The gradient is in % per 100 m.

	A priori	4 May posterior			11 May posterior			4+11 May posterior		
# discharge series:		0	1	2	0	1	2	0	1	2
$E[g_m]$	2.58	9.93	9.39	8.58	4.39	4.29	4.03	5.41	5.12	4.53
$sd[g_m]$	1.5	0.73	0.71	0.66	0.64	0.61	0.60	0.51	0.57	0.55

Table 2. Changes in expectation and standard deviation for Cx and g_z, resulting from different combinations of images and discharge data. Two gauged subcatchments in the region are used separately or in sum. The gradient is in °C×d/100 m, Cx is the degree-day melt index. Observed g_z between 850 and 1600 m a.s.l. is −7.7 degree-days/100 m on 4 May and −13.7 degree-days/100 m on 11 May, respectively.

	A priori	4 May posterior			11 May posterior			4+11 May posterior		
# discharge series:		0	1	2	0	1	2	0	1	2
$E[Cx]$	3.00	3.51	3.49	3.46	2.88	2.87	2.86	3.03	3.02	2.98
$sd[Cx]$	0.75	0.06	0.06	0.06	0.04	0.05	0.04	0.04	0.04	0.04
4 May $E[g_z]$	−9.32	−3.84	−3.90	−4.02				−4.10	−4.20	−4.23
4 May $sd[g_z]$	2.15	0.2	0.2	0.1				0.13	0.13	0.12
11 May $E[g_z]$	−12.12				−7.11	−7.15	−7.25	−6.11	−6.26	−6.40
11 May $sd[g_z]$	2.80				0.22	0.20	0.21	0.20	0.22	0.19

Table 3. Fractional reduction of variance from prior to posterior state; map averages.

Fractional variance reduction (%) (Map averages)	4 May update # discharge series			11 May update # discharge series			4+11 May update # discharge series		
	0	1	2	0	1	2	0	1	2
Cv	55.0	54.8	54.2	45.7	45.8	45.7	63.4	63.4	63.3
M	10.4	16.9	26.7	45.1	48.0	48.4	48.2	49.8	50.8
Y_0	67.7	67.2	66.2	33.1	33.3	32.1	63.1	63.3	62.6
4 May λ	47.1	50.0	56.1				65.3	66.3	68.6
4 May y	73.4	72.8	72.9				75.5	75.5	76.0
4 May swe	52.9	56.4	61.7				75.5	75.7	76.9
4 May Q	61.7	63.1	65.8				76.3	76.8	78.0
11 May λ				94.7	95.0	95.1	94.5	94.6	94.9
11 May y				82.6	83.0	83.0	88.0	87.8	88.0
11 May swe				77.9	78.5	79.1	82.6	82.7	83.5
11 May Q				86.4	86.6	86.5	89.5	89.8	90.1

HESSD
2, 1185–1219, 2005

Table 4. Prior and posterior mass balance of the Vinstra and Sjoa catchments, with the observed values estimated from the discharge series. Posterior columns identify the nine experiments defined by images used (4 May, 11 May, both) and number of discharge series used in the likelihood.

Mass balance (mm)			4 May posterior # discharge series			11 May posterior # discharge series			4+11 May posterior # discharge series		
	Prior	Obs	0	1	2	0	1	2	0	1	2
Vinstra											
$E[m]$	378	474	514	503	483	421	420	415	441	436	423
4 May $E[Q]$	55	95	106	104	100				99	98	95
11 May $E[Q]$	125	157				150	150	149	165	163	159
Sjoa											
$E[m]$	443	550	683	663	623	484	481	473	514	504	481
4 May $E[Q]$	30	35	99	97	89				85	85	79
11 May $E[Q]$	83	83				127	126	124	140	138	130

HESSD
2, 1185–1219, 2005

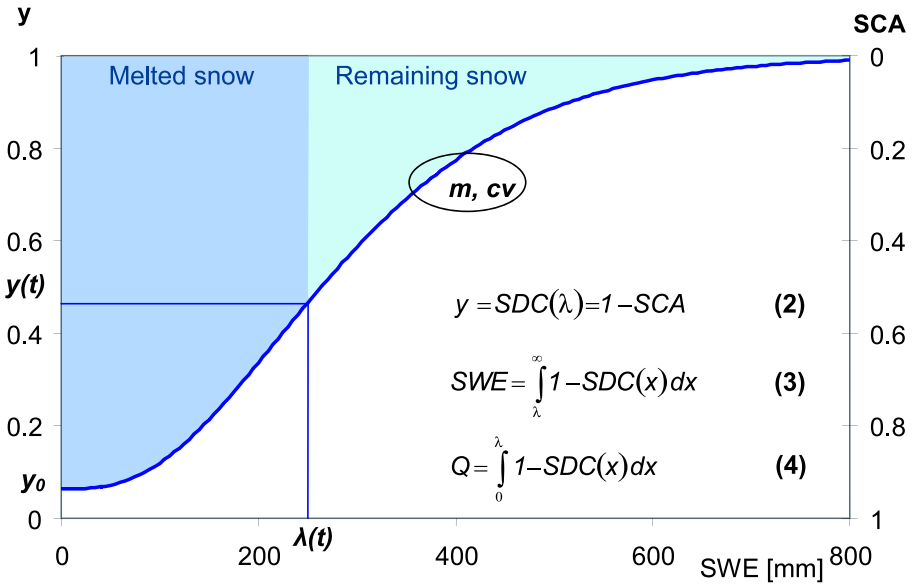

Fig. 1. Principle of the *SDC* concept. Three static parameters define the *SDC* characterizing the distribution of point snow storage x at the melt season onset; the average storage m, the sub-grid spatial coefficient of variation cv, and the initial bare round fraction y_0. The accumulated melt depth l is the only dynamically simulated variable, defining the bare-ground area fraction y, and the mass balance components Q and SWE.

HESSD
2, 1185–1219, 2005

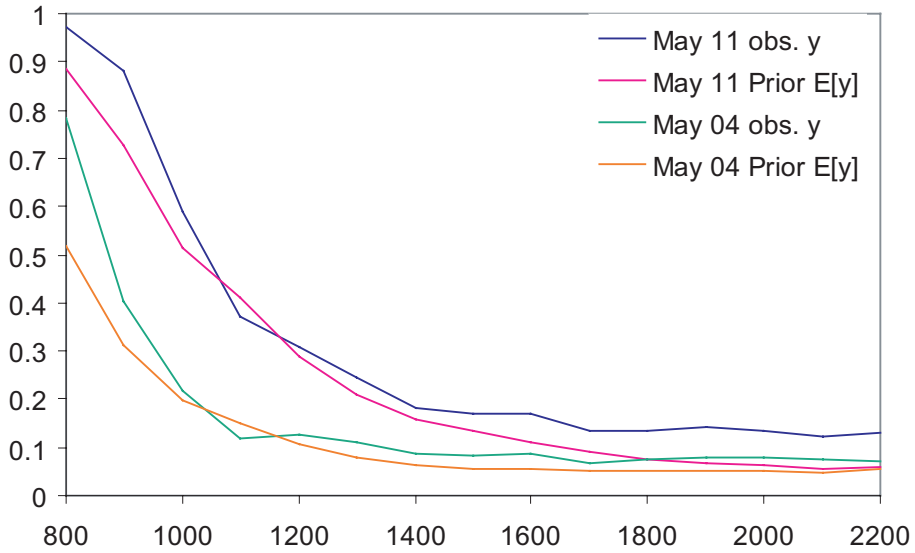

Fig. 2. Observed and a priori expected bare ground area y on 4 May and 11 May, averaged over elevation intervals of 100 m.

HESSD
2, 1185–1219, 2005

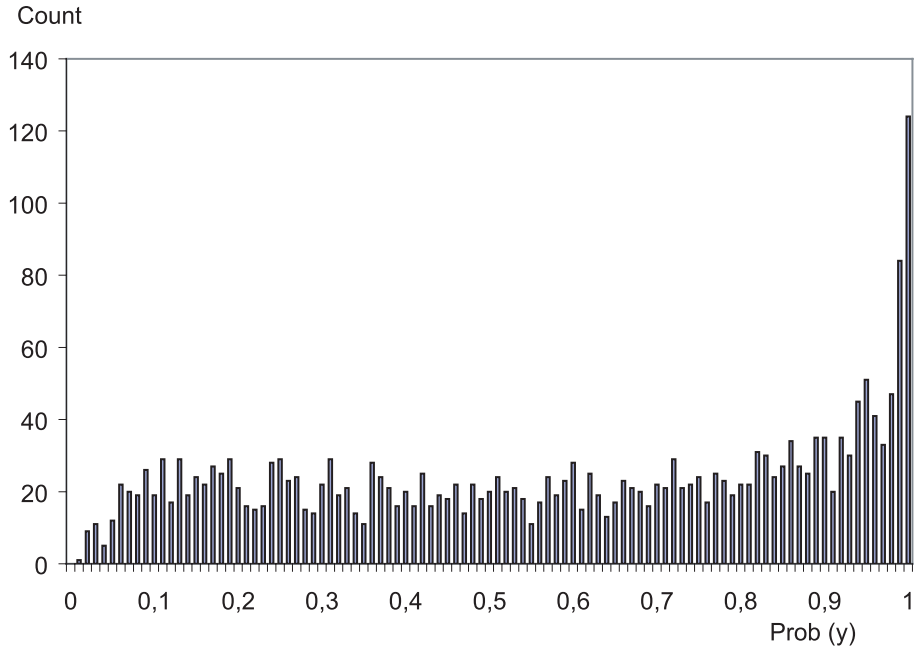

Fig. 3. Spatial histogram of a priori cumulative probability of the observation. A uniform histogram indicates that the prior distribution is well estimated, according to the observations.

HESSD

2, 1185–1219, 2005

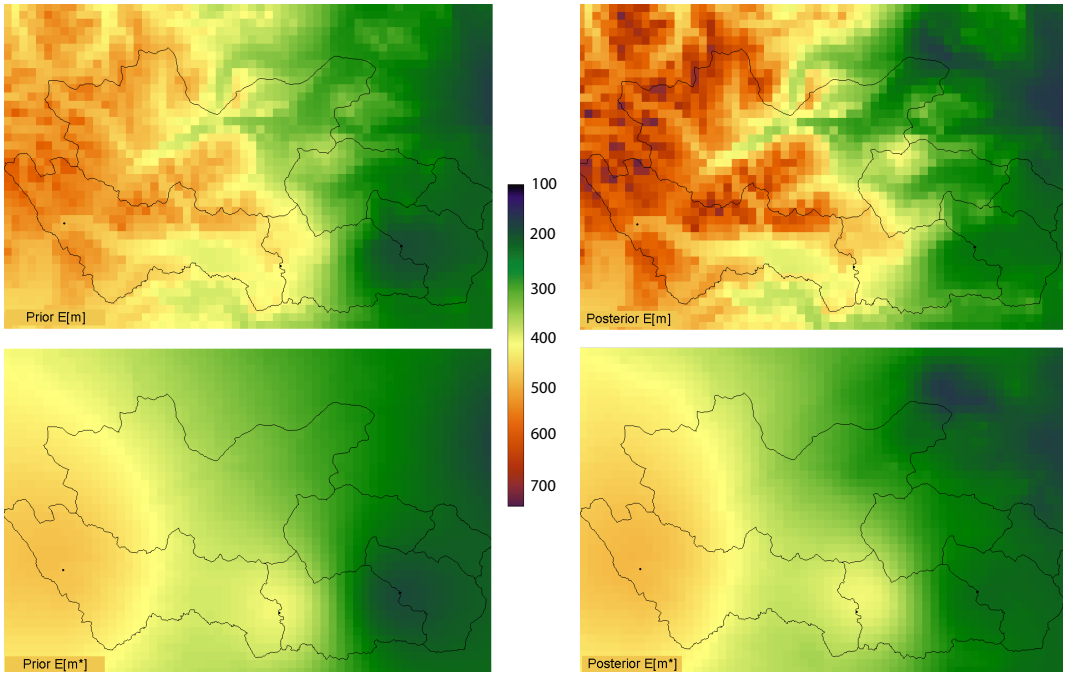

Fig. 4. Expected end of winter snow storage m and its reference-elevation component m^*, through transformation and updating. The elevation gradient is increased, and there are some changes to m^* in the eastern part of the region. Notice the smoothness of the m^* maps compared to the similar maps of m. Subcatchment boundaries are shown, with Sjoa (473 km^2) in the North, and the location of three precipitation gauges are indicated.

HESSD
2, 1185–1219, 2005

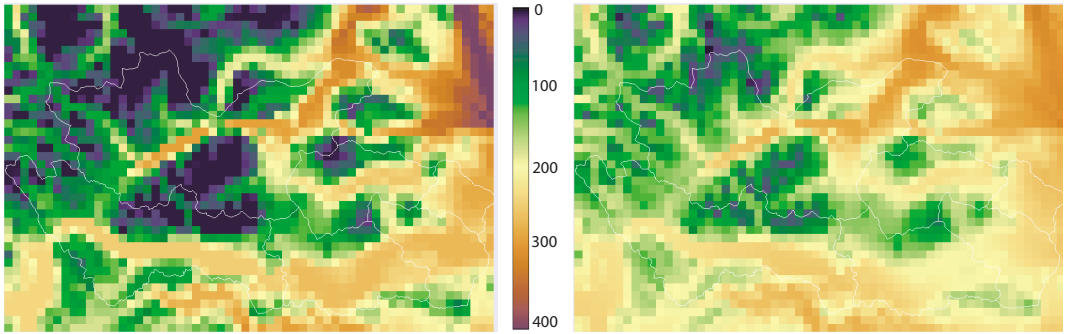

Fig. 5. Expected accumulated melt depth λ; prior (left) and posterior (right). The elevation gradient is considerably less steep, with melt extending to greater elevation, and lower melt depths in the valley to the Northeast.

HESSD

2, 1185–1219, 2005

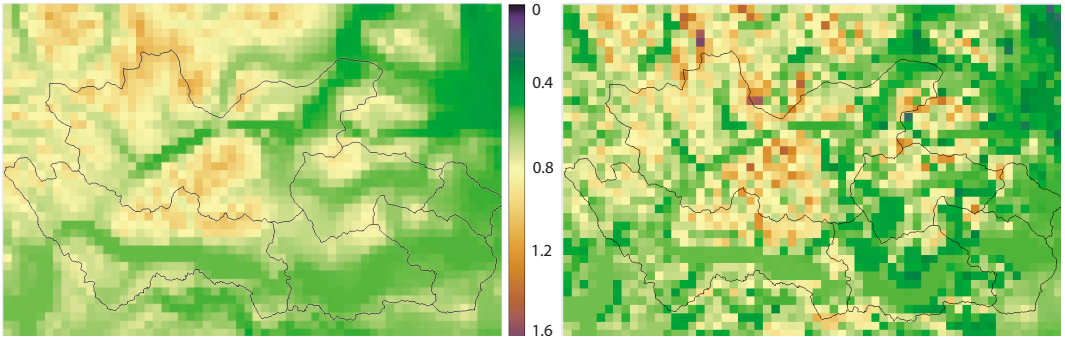

Fig. 6. Expected sub-grid SWE coefficient of variation *cv*; a priori (left) and a posteriori (right). The prior expectancy is linked to elevation, assuming that stronger winds at high altitude produce more re-distribution of snow. The posterior expectancy show more small-scale variability, but have approximately similar elevation dependency.

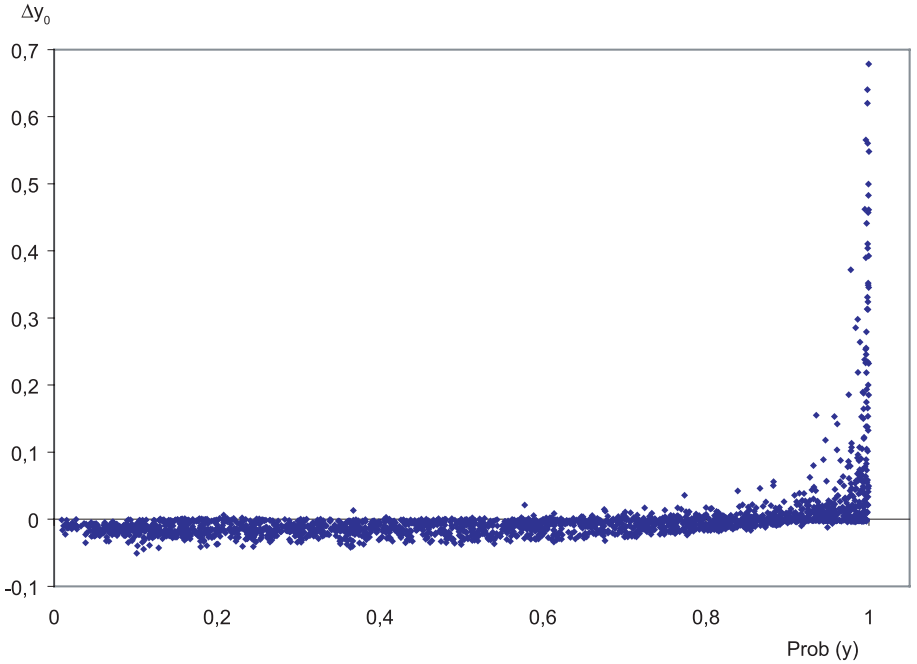

Fig. 7. Expectation changes for y_0 versus the prior cumulative probability of the observed y value. Large increases in y_0 are invariably linked to extreme y observations, but the general connection is weak over most of the probability interval.

HESSD
2, 1185–1219, 2005