Hydrology and
Earth System
Sciences

Using crowdsourced web content for informing water systems operations in snow-dominated catchments

Matteo Giuliani[1], **Andrea Castelletti**[1,2], **Roman Fedorov**[1], **and Piero Fraternali**[1]

[1]Department of Electronics, Information and Bioengineering, Politecnico di Milano, Piazza L. da Vinci, 32, 20133 Milano, Italy
[2]Institute of Environmental Engineering, ETH, Wolfgang-Pauli-Str. 15, 8093 Zurich, Switzerland

Correspondence to: Matteo Giuliani (matteo.giuliani@polimi.it)

Abstract. Snow is a key component of the hydrologic cycle in many regions of the world. Despite recent advances in environmental monitoring that are making a wide range of data available, continuous snow monitoring systems that can collect data at high spatial and temporal resolution are not well established yet, especially in inaccessible high-latitude or mountainous regions. The unprecedented availability of user-generated data on the web is opening new opportunities for enhancing real-time monitoring and modeling of environmental systems based on data that are public, low-cost, and spatiotemporally dense. In this paper, we contribute a novel crowdsourcing procedure for extracting snow-related information from public web images, either produced by users or generated by touristic webcams. A fully automated process fetches mountain images from multiple sources, identifies the peaks present therein, and estimates virtual snow indexes representing a proxy of the snow-covered area. Our procedure has the potential for complementing traditional snow-related information, minimizing costs and efforts for obtaining the virtual snow indexes and, at the same time, maximizing the portability of the procedure to several locations where such public images are available. The operational value of the obtained virtual snow indexes is assessed for a real-world water-management problem, the regulation of Lake Como, where we use these indexes for informing the daily operations of the lake. Numerical results show that such information is effective in extending the anticipation capacity of the lake operations, ultimately improving the system performance.

1 Introduction

Snow accumulation and melting are fundamental components of the hydrological cycle in many watersheds across the world (e.g., Mote et al., 2005; Holko et al., 2011). Approximately 40–50 % of the Northern Hemisphere is covered by snow (Pepe et al., 2005) and snow plays a key role in mountain areas, which, in Europe, account for 40 % of the total surface (Schuler et al., 2004).

In such contexts, an accurate characterization of snow availability and its evolution in time can be extremely valuable for a variety of operational purposes, from avalanche prediction (e.g., Perona et al., 2012; Schweizer et al., 2009), water-systems operations through medium to long-term streamflow forecast (e.g., Wood and Lettenmaier, 2006; Anghileri et al., 2016), or drought risk management (e.g., Staudinger et al., 2014). The projected temperature increase induced by climate change, with consequent reductions of large volumes of snowpacks and acceleration of the water cycle in many mountainous areas, will further amplify the importance of better understanding snow dynamics (Barnett et al., 2005; Kunkel et al., 2016).

Snow processes are generally monitored through both ground monitoring networks (e.g., Brown and Braaten, 1998; López-Moreno and Nogués-Bravo, 2006) and remote sensing (for a review, see König et al., 2001; Dietz et al., 2012, and references therein). Yet both sources have serious limitations in alpine contexts, mainly related to the high spatial (e.g., Newald and Lehning, 2011) and temporal variability of snow-related processes (Blöschl, 1999; Egli, 2008; Gleason et al., 2016). Ground stations are generally very un-

evenly distributed distributed. Satellite products provide data on a denser grid but are diversely constrained depending on the sensors installed (Muñoz et al., 2013). High spatial- and temporal-resolution imagery (i.e., daily maps with spatial resolution of about 500 m) can be derived from Moderate Resolution Imaging Spectroradiometer (MODIS) products, which are, however, strongly affected by the weather because optical sensors cannot see the earth surface when clouds are present (Parajka and Blöschl, 2008). Space-board passive microwave radiometers (e.g., AMSR-E) penetrate clouds but have coarse spatial resolution (25 km). Finally, the use of active microwave systems (e.g., RADARSAT) is so far limited to the detection of liquid water content.

The last few years have seen a rising interest in complementing traditional observations by using cameras and short-range visual content analysis techniques (Bradley and Clarke, 2011), which allow improvement of the temporal and spatial resolutions for specific applications. Many case studies showed that the use of one or several time-lapse cameras allows mapping of both the spatial and temporal patterns of a variety of snow characteristics, including glacier velocity, snow-cover changes, or detailed monitoring of snowfall interception (see Parajka et al., 2012, and references therein). However, most of these systems generally rely on cameras designed and positioned ad hoc (e.g., Hinkler et al., 2002), possibly including in the camera view some specific objects, such as flags or sticks, which simplifies the calibration of geometry and colors (e.g., Floyd and Weiler, 2008; Laffly et al., 2012; Garvelmann et al., 2013). In addition, the use of these cameras is generally very expensive and often requires intensive manual efforts in the image-processing phase. This latter includes a variety of crucial, time-consuming operations, such as the selection of photographs with good meteorological and visibility conditions, the photo-to-terrain alignment and orientation, and the labeling of snow-covered pixels for estimating the total snow cover (e.g., DeBeer and Pomeroy, 2009; Farinotti et al., 2010).

The availability on the web of large volumes of public, low-cost, and spatiotemporally dense data raises the question of whether it is possible to use such data as a supplement, or at least as a complement, to traditional monitoring systems in operational contexts. The main advantage of such public data, albeit collected for completely different purposes and with much lower quality standards, is that they can significantly increase the spatial and temporal coverage at little-to-no cost (Jacobs et al., 2009; Graham et al., 2010). This idea is part of a growing application of so-called "citizen science" approaches to water resources systems operation (Buytaert et al., 2014) and, more generally, to diverse environmental problems (Fraternali et al., 2012). Crowdsourced observations may act as low-cost virtual sensors in a variety of environmental contexts (Lowry and Fienen, 2013), for example, contributing to monitoring the dynamics of forests (e.g., Daume et al., 2014), storms (e.g., Good et al., 2014), or streamflow (e.g., Michelsen et al., 2016), with potential ben-

efit in terms of the prediction of flood events and of the timely delivery of alarms (e.g., Smith et al., 2015; Mazzoleni et al., 2015a, b; Fohringer et al., 2015; Le Boursicaud et al., 2016). However, despite this interest in environmental public web- and user-generated data (Vitolo et al., 2015), most works focus on data collection and analysis, with limited assessment of the practical value of such crowdsourced information.

In this paper, we explore the potential for web- and crowdsourced data to retrieve relevant information on snow availability and dynamics in a river basin, and assess the utility of such information in informing a real-world decision-making problem. More precisely, we contribute a novel crowdsourcing procedure for extracting snow-related information from public web images, either produced by users or generated by touristic webcams, and we quantify the operational value of this information compared to other more traditional snow information, such as ground observations and a hybrid mix of satellite-retrieved information, ground data, and model outputs. Our procedure employs an articulated architecture (Fedorov et al., 2016), which automatically crawls content from multiple web data sources with a content-acquisition pipeline integrating public webcams and user-generated photographs posted on Flickr. Next, the procedure retains only geo-tagged images containing a mountain skyline with high probability and identifies the visible mountain peaks in each image, using a digital elevation model (DEM). Then, a supervised learning classifier extracts a snow mask from each image, which distinguishes the image pixels as snow or not-snow. Finally, the resulting snow masks are post-processed to derive time series of virtual snow indexes (VSIs) representing a proxy of the snow-covered area.

The extracted VSIs are used to inform water system operations. The evaluation is performed in the snow-dominated catchment of Lake Como, a regulated lake in northern Italy, where snow melt is the most important contribution to the seasonal storage. The VSI operational value is quantified by comparing, via simulation, the performance of the lake operating policies designed using crowdsourced and traditional snow information, with the performance of the baseline policy obtained by regulating the lake without snow information (Giuliani et al., 2015). This form of assessment provides an indirect validation of the utility of web- and crowdsourced information as the VSIs extracted from general-purpose mountain images and the traditional observational data collected with dedicated tools are not directly comparable due to the difference in their physical interpretation and spatiotemporal resolution (e.g., geo-located photos allow estimating the presence of snow, but not the physical measures usually employed in snow process models, such as the snow water equivalent).

The paper is organized as follows: in the next section, we introduce our methodology for the computation of VSIs based on public web content and the assessment of their operational value. Section 3 describes the Lake Como study site, followed by the discussion of the numerical results. The last
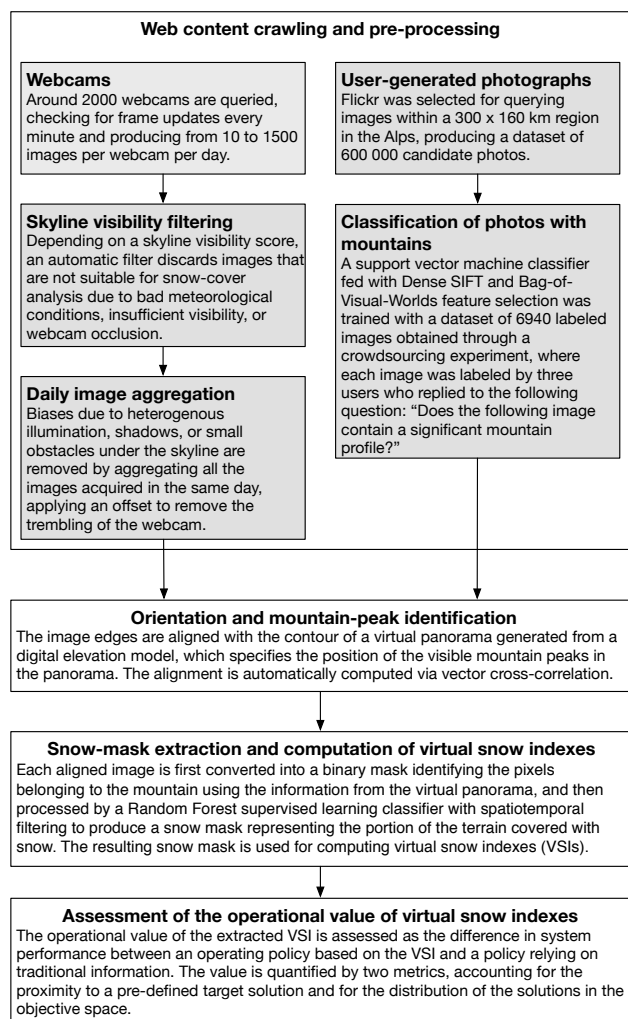
Web content crawling and pre-processing

Webcams
Around 2000 webcams are queried, checking for frame updates every minute and producing from 10 to 1500 images per webcam per day.

User-generated photographs
Flickr was selected for querying images within a 300 x 160 km region in the Alps, producing a dataset of 600 000 candidate photos.

Skyline visibility filtering
Depending on a skyline visibility score, an automatic filter discards images that are not suitable for snow-cover analysis due to bad meteorological conditions, insufficient visibility, or webcam occlusion.

Classification of photos with mountains
A support vector machine classifier fed with Dense SIFT and Bag-of-Visual-Worlds feature selection was trained with a dataset of 6940 labeled images obtained through a crowdsourcing experiment, where each image was labeled by three users who replied to the following question: "Does the following image contain a significant mountain profile?"

Daily image aggregation
Biases due to heterogenous illumination, shadows, or small obstacles under the skyline are removed by aggregating all the images acquired in the same day, applying an offset to remove the trembling of the webcam.

Orientation and mountain-peak identification
The image edges are aligned with the contour of a virtual panorama generated from a digital elevation model, which specifies the position of the visible mountain peaks in the panorama. The alignment is automatically computed via vector cross-correlation.

Snow-mask extraction and computation of virtual snow indexes
Each aligned image is first converted into a binary mask identifying the pixels belonging to the mountain using the information from the virtual panorama, and then processed by a Random Forest supervised learning classifier with spatiotemporal filtering to produce a snow mask representing the portion of the terrain covered with snow. The resulting snow mask is used for computing virtual snow indexes (VSIs).

Assessment of the operational value of virtual snow indexes
The operational value of the extracted VSI is assessed as the difference in system performance between an operating policy based on the VSI and a policy relying on traditional information. The value is quantified by two metrics, accounting for the proximity to a pre-defined target solution and for the distribution of the solutions in the objective space.

Figure 1. Flowchart of the methodology adopted in this study.

section concludes with final remarks and directions for further research.

2 Methods and tools

This section describes the methodology adopted in this work, which is illustrated in Fig. 1. Details about each phase of the procedure are provided in the following sub-sections. A detailed technical analysis of the outputs of the image processing architecture is reported in Fedorov et al. (2016).

2.1 Web content crawling and pre-processing

Two types of public web content are considered, namely touristic webcams and mountain photographs from Flickr. In particular, webcams produce a temporally dense series of images of the same view, while crowdsourced photos have better spatial distributions but lower time coverage.

2.1.1 Public webcams

A webcam is a standalone camera positioned at a fixed known location, usually with a fixed orientation, which captures frames with a certain frequency and exposes them via a web service. In contrast to surveillance webcams, which can provide real time updates (several frames per second, resulting basically in a video stream), public webcams deployed for touristic, meteorological, and publicity reasons update the current frame with lower frequency, typically from one minute to one hour. The public webcam processing phase consists of three main steps:

1. **Webcam image crawling**: public webcams most often expose a single fixed URL for the current frame, and change the image itself over time. This method simplifies the crawling, which amounts to checking the URL of the webcam periodically and downloading the image, when it changes with respect to the last acquisition. We collected the address of more than 3500 webcams in the European Alps and manually inspected them, discarding those that do not frame a significant mountain profile, retaining nearly 2000 webcams. Since December 2014, a crawler acquires all the images of these webcams, checking for frame updates every minute, thus obtaining from 10 to 1500 images per webcam per day, depending on the update frequency.

2. **Skyline Visibility Filtering**: webcams are crawled independently of the weather conditions. As a consequence, although the temporal density of webcam images guarantees a high number of input frames, filtering must be applied to discard unsuitable images that may bias the VSI computation (e.g., an image of a mountain covered by fog can be considered as completely covered by snow in the next steps). A random sampling of 1000 images from four webcams in our data set revealed that 67 % of the images were not suitable for snow-cover analysis due to adverse weather conditions (e.g., fog, heavy snowfall, or rain), insufficient visibility, or the presence of mobile obstacles such as cars or persons. Therefore, the implemented filter automatically discards unsuitable images, identified by checking for occlusions of the mountain skyline. In practice, for each webcam, the pixels that belong to the skyline $\mathcal{L}$ are first identified manually on a sample frame through a crowdsourcing experiment. Then, the binary skyline neighborhood mask L, which identifies pixels $p = (x, y)$ close to the skyline, is determined as follows:

$$L(p) = \begin{cases} 1 & \text{if } \exists p' \in \mathcal{L} : \|p - p'\| \leq \tau \\ 0 & \text{otherwise} \end{cases}, \tag{1}$$

where $\|\cdot\|$ is the Euclidean norm and τ is a distance threshold. In other words, L is a binary mask of the same dimension as the webcam image containing a dilated skyline profile.

Then, for each webcam image, its binary edge map E is computed by the Compass algorithm (Ruzon and Tomasi, 2001), where a pixel is marked as an edge when it corresponds to an abrupt color variation. The binary matrix $E \odot L$, where $\odot$ denotes the pixel-wise product between two images of the same size, represents the edges of the image that belong to the skyline. To check for occlusions, we compute a skyline visibility score v defined as

$$v = f(E \odot L)/f(L), \tag{2}$$

where $f(\cdot)$ is a function that, given an image, returns the number of columns containing at least one non-zero entry. The value of v ranges between 0 and 1, and can be intuitively seen as the percentage of the skyline which is visible in the given image. After setup trials, we discard images with $v < \bar{v}$, where $\bar{v}$ is a fixed threshold equal to 0.75. The experimental validation of the filtering algorithm on 1000 manually annotated images (i.e., frames manually classified as having "good visibility" or "bad visibility") showed that the algorithm achieves a true positive rate (TPR) equal to 87.4 %, while having a false positive rate (FPR) equal to 3.5 %.

3. **Daily Image Aggregation**: the images selected by the skyline visibility filter can still present several undesirable features due to shadows, solar glare, or temporary obstacles below the skyline (e.g., people standing in front of the camera). To attenuate such biases, assuming that the snow cover does not vary during a day significantly, we produce a single image for each webcam per day by aggregating all the images acquired in the same day. Such a daily median image (DMI) is obtained as the median of every pixel across all the daily images accepted by the filter. Given a daily sample of N images $I_1, \ldots, I_N$, the DMI is formally defined as

$$\mathrm{DMI}(x, y) = \mathrm{med}\{I_1(x, y), I_2(x, y), \ldots, I_N(x, y)\}, \tag{3}$$

where $\mathrm{med}\{\cdot\}$ is the median operator applied to the image pixel values. Figure 2 shows a DMI obtained from 11 daily images: it attenuates transient light conditions and removes the people standing in front of the webcam.

A second challenging aspect of DMI creation is the presence of webcam trembling (Latecki et al., 2005). The webcam orientation is not perfectly constant in time but may change slightly, especially in windy regions and when webcams are fixed to poles. To overcome this problem, we extract edge maps of all daily images and calculate their cross-correlation to quantify not only the similarity between the two edge maps, but the entire set of similarities at every possible position. The value of cross-correlation is then used to derive the best offset of every image with respect to the reference represented by the first image of the day. We consider a maximum

offset of 10 pixels to reduce the computation time and avoid possible correction errors. Finally, the DMI is determined from images normalized with such an offset. Intuitively, this procedure can be seen as applying a small displacement to each image in order to obtain the best possible overlap between its edges and the edges of the first image of the day.

2.1.2 User-generated photographs

The second source of mountain images are the photographs generated by common people and publicly available on social networks and photo sharing platforms. Although the volume of user-generated photographs can not obviously reach the number of webcam images, user-generated photographs present higher spatial density. The webcams are located in a few fixed locations, whereas the user photographs can be potentially acquired in any place.

We selected Flickr as the content source because it contains a high number of public photographs with associated geo-tags (e.g., Serdyukov et al., 2009). Furthermore, Flickr does not remove the EXIF (exchangeable image file format) information present in the original images (Tesic, 2005); the EXIF container specifies several photo-related details, in particular the GPS location, the camera model and manufacturer, and optical information, such as the focal length used during the shot. This information is fundamental for the peak-detection algorithm (see Sect. 2.2). A continuous search system was set up for querying images within a 300×160 km region in the Alps. At present, the Flickr search pipeline examined around 600 000 candidate photos. Differently from registered webcams, which produce thousands of images of the same view, the user-generated photographs are taken at unknown locations and may have irrelevant content, and thus must be classified as relevant one-by-one. To this end, a supervised content-based classifier was developed to perform mountain-image detection. The classifier was trained on a set of 6940 images randomly sampled from the very large crawled data set; the ground-truth images were classified manually through a crowdsourcing experiment. For each image, three users were asked to reply to the following question: "Does the image contain a significant mountain profile?". A web interface proposed a tutorial on how to annotate an image as positive (mountain image) or negative (non-mountain image). The experiment was conducted using an internal (non-paid) crowd, collecting a total of more than 20 000 image classification labels. The aggregated ground-truth label of each image was then derived via majority voting. Approximately 23 % of the original 6940 images were classified as positive.

The automatic classification was performed with a support vector machine (SVM) classifier fed with dense SIFT (scale invariant feature transform) and "bag-of-visual-words" (BoVW) feature selectors (Fei-Fei and Perona, 2005). This technique relies on the idea that every image
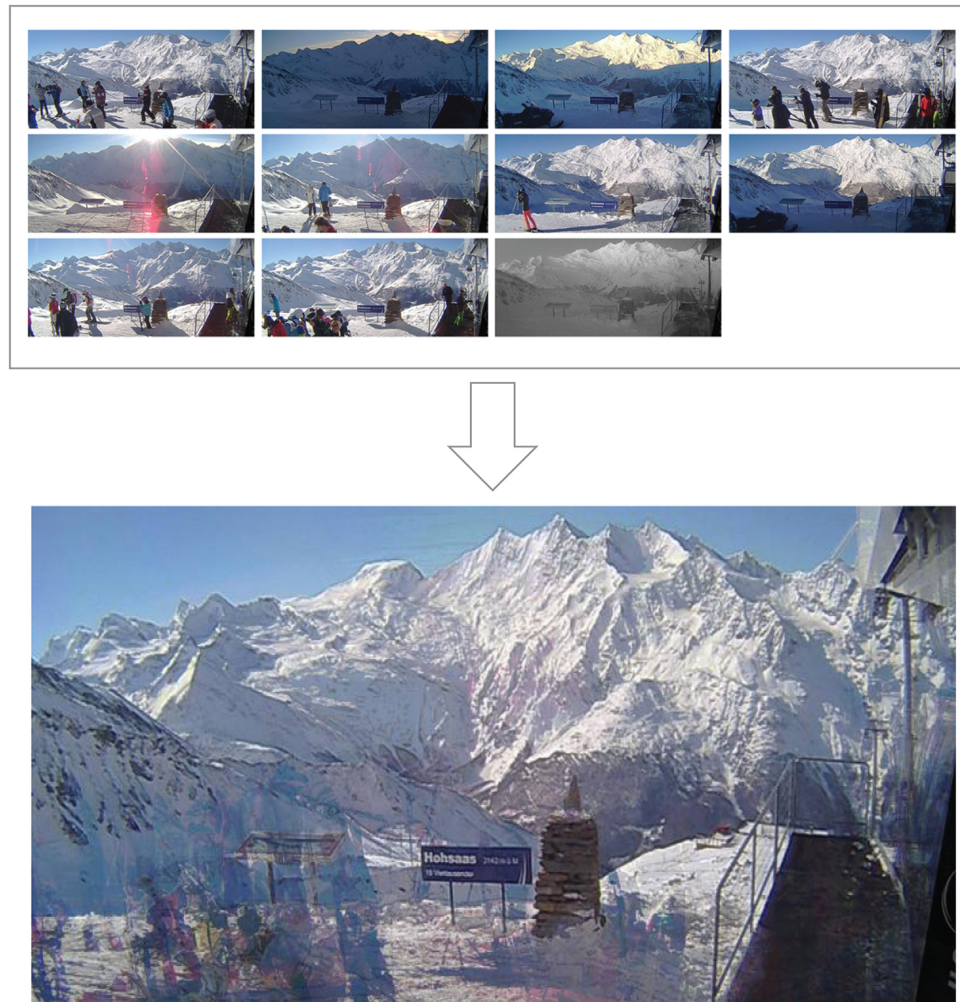
Figure 2. Example of Daily Median Image obtained from 11 images acquired in the same day.

is composed of small patches (i.e., image portions), which somehow share common features with the images in the same class (i.e., images that do contain or do not contain mountains). Since the number of possible patches to observe is very large, the patches are split into a finite number of clusters. Each patch represents a visual word, which contributes to defining the content of the image. All the visual words of the image are aggregated into a histogram, which is then used as a feature vector for the SVM classifier. To create a balanced data set, we retained all the positive samples and randomly selected the same number of negative samples. Then, we used around 70 % of these images for training and validation, and the remaining 30 % for testing. The performance attained by the classifier on the test data set is: 95.1 % accuracy, 94.0 % precision, and 96.3 % recall.

2.2 Orientation and mountain-peak identification

The orientation and mountain-peak identification procedure (see Fig. 3) is applied to the user-generated photographs classified as positive and to the median daily images of webcams. In fact, although webcams are geo-located, the information regarding the orientation of the webcam, and consequently the corresponding mountain peaks observed, is not available. In both cases, image orientation is estimated through the alignment with respect to a 360° virtual panorama generated using a digital elevation model (DEM) that specifies the position of the visible mountain peaks in the panorama.

The automatic alignment of an image to the virtual panorama requires scaling the image to achieve the same angular and/or pixel dimension. This step is performed by computing the image field of view (FOV), namely the size of the angle comprising the view. The FOV can be estimated from the image EXIF information, such as focal length, camera model, and manufacturer. The procedure extracts the edge
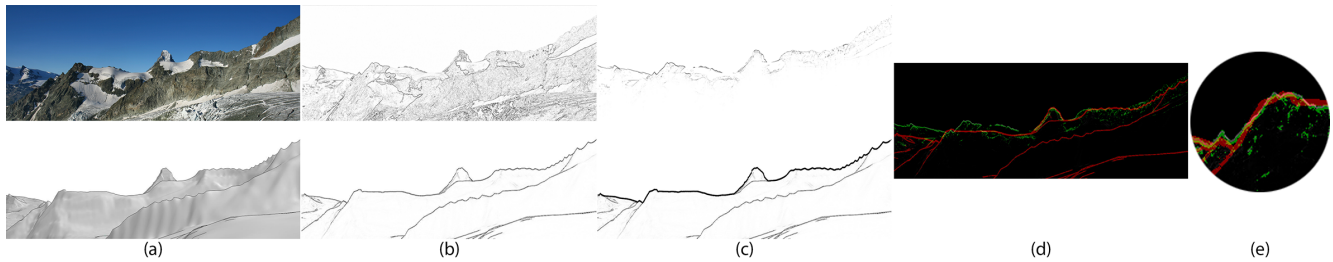
Figure 3. Example of the orientation and mountain peak identification procedure: **(a)** input image (top) and corresponding panorama (bottom); **(b)** edge maps; **(c)** skyline detection; **(d)** global alignment; **(e)** local alignment.
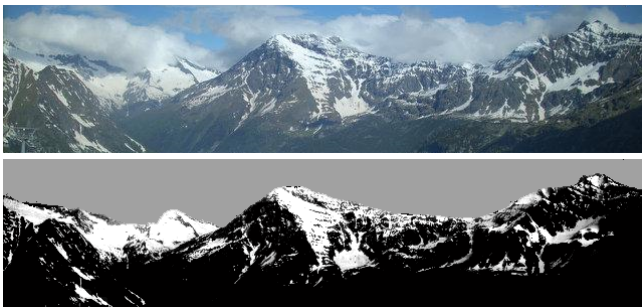

Figure 4. Example of an image (top) and the computed snow mask (bottom), where white stands for snow, black for non-snow, and gray for sky.

maps for both the scaled image and the virtual panorama. In particular, a modified version of the multi-stage graph algorithm by Lie et al. (2005) was used for extracting the skyline from the edge maps and to eliminate all the edges above the skyline (clouds and obstacles) by gradually reducing the strength of the edges below the skyline. This step was not applied to the webcams as a single annotation was sufficient for obtaining a precise skyline extraction. Then, the best overlapping position between the image and the virtual panorama is identified with a vector cross-correlation (VCC) procedure (Baboud et al., 2011). The VCC finds the best horizontal overlap position of the image with respect to the panorama by maximizing the cross-correlation score of the edges, while also considering the estimated image orientation. The identified overlap position allows projection of the peak positions from the panorama to the image, to estimate which peaks are visible, and what their coordinates are, in the image. When the image does not contain the EXIF information, the automatic orientation and mountain-peak identification procedure can not be applied, and the image requires a manual alignment with respect to the panorama. Finally, a local refinement of the alignment is obtained by repeating the VCC procedure with a maximum radius equal to 50 pixels. This is done for each mountain peak independently to adjust its position through the identification of the best match in its neighborhood region.

The orientation and peak identification algorithm was tested on a data set of 162 images randomly sampled from the web and manually aligned to the corresponding virtual panoramas to create the ground-truth data. Considering a tolerance of $3°$, 75 % of the image orientations were correctly estimated. The accuracy grows to 77.6 % for photos with no clouds and to 81.6 % for photos with no mountain slopes in the very short range (the effect of GPS errors is more relevant if mountains are close to the shooting location). The average peak positioning error was $0.78°$.

2.3 Snow-mask extraction and computation of virtual snow indexes

The third step of the procedure is the conversion of the snow information contained in the aligned image into one or more VSIs associated with the mountain viewpoint portrayed in the photo. This phase requires estimating a snow mask representing the portion of the terrain that is covered with snow. Formally, let I denote an image and M a binary mask having the same size as I, where $M(x, y) = 1$ indicates that the pixel $p(x, y)$ of the image belongs to the mountain area, or $M(x, y) = 0$ otherwise. The binary mask M is derived from the alignment of the image with the virtual panorama (see the previous section), which allows for differentiation of pixels that correspond to terrain or sky.

The pair (I, M) is processed by a pixel-level binary classifier, which extracts the snow mask S by assigning to each pixel a label denoting the presence of snow ($S(x, y) = K_1$), non-snow ($S(x, y) = K_2$), and sky ($S(x, y) = K_3$), as shown in Fig. 4. We computed snow masks using the Random Forest supervised learning classifier with spatiotemporal median smoothing of the output (Liaw and Wiener, 2002). Such a classifier discriminates the presence of snow in a pixel based on its color and on the color of the neighbor pixels. Moreover, it applies a spatiotemporal median filter to smooth the snow variation and attenuate the errors. Smoothing implements the assumption that pixels close to each other in the same image and pixels in the same position in images close in time should belong to the same class (i.e., snow or non-snow). The training and testing of the supervised classifier was performed on a data set including 59 images annotated

in snow or non-snow areas, containing more than 7 million single pixel ground-truth labels. The accuracy attained by the classifier is 93.5 %, outperforming other existing methods for pixel-level classification of snow presence (Fedorov et al., 2016).

Finally, different VSIs are computed from the snow masks S, potentially also considering the altitude associated with each pixel, which can be determined from the image to virtual panorama alignment. In this work, we report the results obtained with a virtual snow index σ representing a proxy of the snow-cover area, defined as follows:

$$\sigma = \sum_{p(x,y)\in I} \Phi(p(x,y)), \quad \text{where}$$

$$\Phi(p(x,y)) = \begin{cases} 1 & \text{if } S(x,y) = K_1 \\ 0 & \text{otherwise} \end{cases}. \tag{4}$$

2.4 Assessment of the operational value of virtual snow indexes

The operational value of the extracted VSIs is assessed as the difference in system performance between an operating policy based upon the VSIs and a policy relying on more traditional information, including water availability in the lake and day of the year. In particular, the operating policies are computed by solving a multi-objective optimal control problem (Castelletti et al., 2008) formulated as follows:

$$p^* = \arg\min_p \mathbf{J} = |J^1, \ldots, J^q|, \tag{5}$$

where the policy p is defined as a closed-loop control policy that determines the release decision $u_t = p(d_t, \mathbf{x}_t, \mathcal{I}_t)$ at each time step t as dependent on the day of the year d_t, the current state of the system is $\mathbf{x}_t$ (i.e., the level of the lake at time t), and a vector of exogenous information is $\mathcal{I}_t$ (i.e., variables that are observed but are not endogenous in the problem formulation and hence are not modeled). Note that the resolution of Problem (5) does not yield a unique optimal solution but a set of Pareto optimal solutions.

The most common technique to solve Problem (5) is dynamic programming (DP) (Bellman, 1957). However, DP is severely limited by the curse of modeling in designing operating policies conditioned on exogenous information (Tsitsiklis and Van Roy, 1996) and by the curse of multiple objectives in exploring multidimensional tradeoffs (Powell, 2007). We therefore solve Problem (5) by means of an evolutionary multi-objective direct policy search (EMODPS) (Giuliani et al., 2016a), an approximate dynamic programming approach that combines direct policy search, nonlinear approximating networks, and multi-objective evolutionary algorithms. EMODPS allows the direct use of exogenous information through a partially data-driven controller tuning approach (Formentin et al., 2013). The operating policy, defined as a nonlinear approximating network, is directly conditioned on observations of exogenous information, which

cannot be accurately modeled and would produce detrimental effects on the performance of an operating policy conditioned on approximate model's outputs (Formentin et al., 2012). The selected policy parameterization strongly influences the selection of the optimization approach, as the number of parameters necessary to obtain a good approximation for the unknown optimal control policy grows with the increasing dimension of the policy's argument (Zoppoli et al., 2002). Since the optimization of the policy parameters requires searching high dimensional spaces that map to stochastic and multimodal objective function values, global optimization methods such as evolutionary algorithms are preferred to gradient-based methods (Heidrich-Meisner and Igel, 2008).

Given the Pareto optimal solutions of Problem (5), the operational value of the estimated VSI is quantified by means of two metrics (Giuliani et al., 2015). The first metric is a measure of the proximity between a pre-defined target solution $\mathbf{J}_T$ and the closest alternative in the Pareto front of the policy under examination, i.e.,

$$D_{\min} = \min_{i=1,\ldots,N} \|\mathbf{J}_T - \mathbf{J}_i\|, \tag{6}$$

where $\|\cdot\|$ stands for the (normalized) Euclidean norm, N is the number of solutions in the Pareto front under exam, and $\mathbf{J}_i$ is the performance of the ith solution in the Pareto front. The lower the value of $D_{\min}$, the closer the performance is to the target.

A more informative assessment can be done by evaluating not only how close a given policy can get to a pre-defined target solution but, more generally, how the Pareto approximate solutions distribute in the objective space. Among the commonly used metrics adopted in the literature (see Maier et al., 2014, and references therein), we adopt the hypervolume indicator (HV), which captures both the convergence of the Pareto front under examination $\mathcal{F}$ to the optimal one $\mathcal{F}^*$ as well as the representation of the full extent of tradeoffs in the objective space. The hypervolume measures the volume of objective space dominated ($\preceq$) by the considered set of solutions. This metric allows set-to-set evaluations, where the Pareto Front with higher HV is considered to be better. HV is calculated as the hypervolume ratio between $\mathcal{F}$ and $\mathcal{F}^*$, formally defined as:

$$\text{HV}(\mathcal{F}, \mathcal{F}^*) = \frac{\int \alpha_{\mathcal{F}}(\mathbf{x})\mathrm{d}x}{\int \alpha_{\mathcal{F}^*}(\mathbf{x})\mathrm{d}x}, \quad \text{where}$$

$$\alpha_{\mathcal{F}}(\mathbf{x}) = \begin{cases} 1 & \text{if } \exists \mathbf{x}' \in \mathcal{F} \text{ such that } \mathbf{x}' \preceq \mathbf{x} \\ 0 & \text{otherwise} \end{cases}. \tag{7}$$
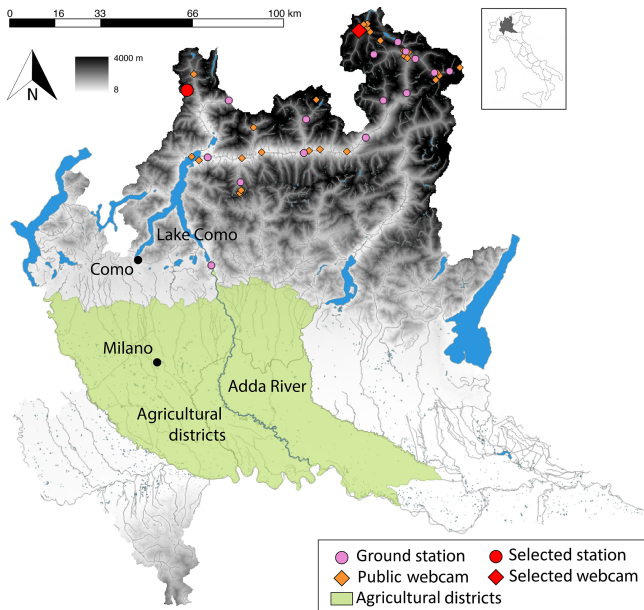
Figure 5. Adda River basin: Lake Como, Adda River, downstream agricultural districts, ground stations, and public webcams.

3 Lake Como study site

3.1 System description

Lake Como is a regulated lake in the Adda River basin, Italy (Fig. 5). The lake has an active storage capacity of $254 \, \mathrm{Mm^3}$ and is fed by a $3500 \, \mathrm{km^2}$ alpine catchment that reaches altitudes over $4000 \, \mathrm{m}$ a.s.l. Downstream from the lake, the Adda River serves a dense network of irrigation canals belonging to four agricultural districts for a total irrigated area of $1400 \, \mathrm{km^2}$ (the green area in Fig. 5). Major cultivated crops are maize and temporary grasslands, while minor crops include rice, soybean, wheat, tomato, and barley. The hydrometeorological regime in the catchment is the typical subalpine one, with scarce discharge in winter and summer, and peaks in late spring and autumn due to snowmelt and rainfall, respectively. In particular, snowmelt from May to July is the most important contribution to the formation of the seasonal storage (Fig. 6).

The alpine orography constrains the accurate monitoring of snow dynamics. The existing ground stations (46 over the $10\,500 \, \mathrm{km^2}$ alpine area in the Lombardy region) provide a very coarse coverage of the region and are not sufficient to reliably monitor the snow coverage and the associated water content. This is instead estimated by the Regional Agency for Environmental Protection (Agenzia Regionale per la Protezione dell'Ambiente – ARPA), which produces estimates of snow water equivalent (SWE) through a hybrid procedure combining snow height and temperature data from ground stations, measures of snow density in few specific locations, satellite retrieved data of snow cover from MODIS,

and model outputs for spatially interpolating these data. As a result of this complex procedure, ARPA elaborates a weekly estimate of SWE. Such reports are delivered only weekly due to the well known limitations of snow products derived from optical sensors associated with the frequent satellite occlusion by cloud coverage. This limitation is particularly restrictive in the alpine region, where previous studies observed an average cloud occlusion of 63 % over a five-year monitoring period (Parajka and Blöschl, 2006), with critical episodes of cloud coverage lasting for more than 25 days per month in wintertime. In contrast, webcams are less affected by cloud coverage and can provide observations during cloudy days as shown illustratively in Fig. 7. In this study, we contrast the operational value in informing the lake operation of three different snow-related data sources: (i) daily observations of snow height from coarsely distributed ground stations; (ii) weekly SWE estimate provided by ARPA; (iii) daily values of the VSI σ extracted from public web images.

The existing regulation of the lake is driven by two primary, competing objectives: water supply, mainly for irrigation, and flood control in the city of Como, which is the lowest point of the lake shoreline. In particular, the agricultural districts downstream would like to store the snowmelt volume for the summer water-demand peak, when the natural inflow is not sufficient to satisfy the irrigation requirements (see the magenta area in Fig. 6). Yet, storing such water increases the lake level and, consequently, the flood risk, which would be instead minimized by keeping the lake level as low as possible. On the basis of previous works (e.g., Castelletti et al., 2010; Giuliani and Castelletti, 2016; Culley et al., 2016; Giuliani et al., 2016b), these two objectives are formulated as follows:

- *Flood control*: the average annual number of flooding days in the evaluation horizon H, defined as days when the lake level h_t is higher than the flooding threshold ($\bar{h} = 1.24 \, \mathrm{m}$):

$$J^{\mathrm{flood}} = \frac{1}{H/365} \sum_{t=1}^{H} \Lambda(h_t), \quad \text{where}$$

$$\Lambda(h_t) = \begin{cases} 1 & \text{if } h_t > \bar{h} \\ 0 & \text{otherwise} \end{cases}. \quad (8)$$

- *Irrigation supply*: the daily average quadratic water deficit between the lake release r_{t+1} and the daily water demand w_t of the downstream system, subject to the minimum environmental flow constraint q^{MEF} to ensure adequate environmental conditions in the Adda River:

$$J^{\mathrm{irr}} = \frac{1}{H} \sum_{t=1}^{H} \max\left(w_t - \max(r_{t+1} - q^{\mathrm{MEF}}, 0), 0\right)^2. \quad (9)$$
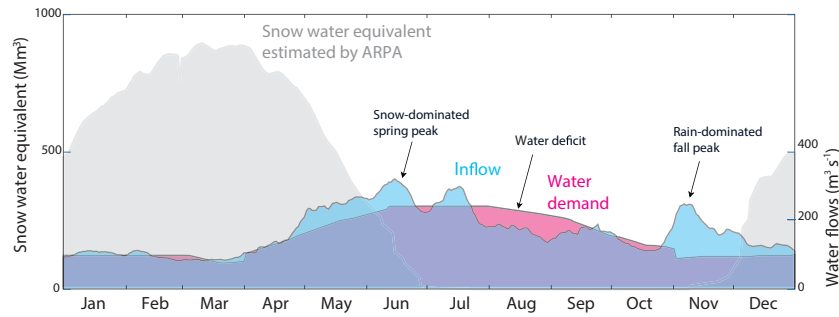
Figure 6. Hydro-meteorological regime of Lake Como.
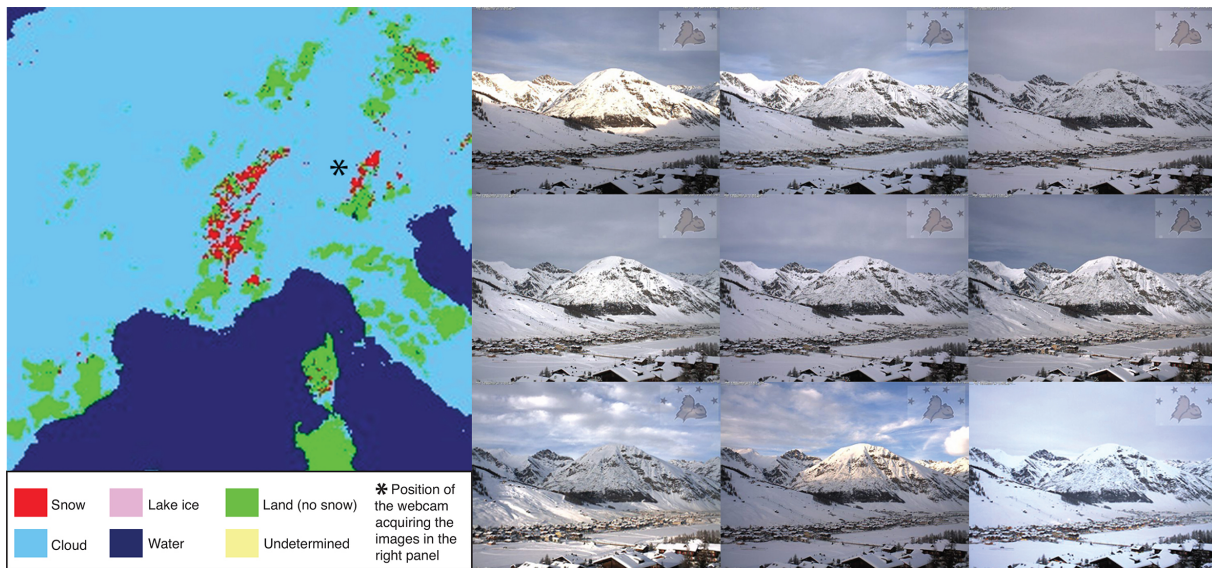

Figure 7. Comparison of MODIS daily snow-cover map (left panel) with the images acquired by a webcam (right panel) on 9 January 2014 at the location denoted by the asterisk in the map.

This quadratic formulation aims to penalize severe deficits in a single time step, while allowing for more frequent, small shortages (Hashimoto et al., 1982).

3.2 Experiment setting

Our assessment of the operational value of the VSIs relies on the comparison of the performance attained by informing the operating policies of Lake Como with alternative snow-related information: (i) policies P1 informed by snow height observations from ground stations; (ii) policies P2 informed by SWE estimates provided by ARPA; (iii) policies P3 informed by the virtual snow index σ. Performance is evaluated against an upper-bound solution, designed assuming perfect foresight of future inflows, and a baseline solution, corresponding to a traditional regulation conditioned on the day of the year and the lake level. The experimental setting is structured as follows:

– *Observational data*: we consider the time horizon 2013–2014 over which time series of snow height, SWE estimate, and VSIs are available. In particular, snow height data are measured at the Truzzo ground station, while the VSI derives from the images of a webcam in Livigno (see Fig. 5); both sources have time series covering the selected time horizon.

– *Informed solutions*: the operating policies P1, P2, and P3 are designed via EMODPS by parameterizing the policies as Gaussian radial basis functions, which have been demonstrated to be effective in solving these types of multi-objective policy design problems (Giuliani et al., 2014a, b), particularly when exogenous information is used for conditioning the operations (Giuliani et al., 2015). To perform the optimization, we use the self-adaptive Borg MOEA (Hadka and Reed, 2013), which has been shown to be highly robust in solving multi-objective optimal control problems, where it met or exceeded the performance of other state-of-the-art
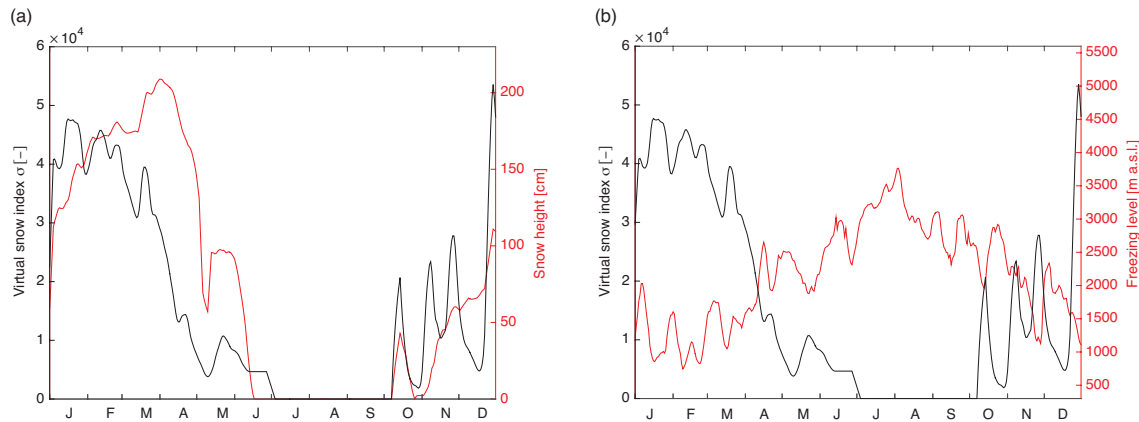
Figure 8. Comparison of the trajectories in 2013 of the Virtual Snow Index σ with the snow height measured at Oga San Colombano (left panel) and with the freezing level (right panel).

MOEAs (Zatarain-Salazar et al., 2016). Each optimization was run for 2 million function evaluations. To improve solution diversity and avoid dependence on randomness, the solution set from each formulation is the result of 30 random optimization trials. The final set of Pareto optimal policies for each experiment is defined as the set of non-dominated solutions from the results of all the optimization trials.

– *Upper-bound solution*: this ideal set of operating policies, which assume perfect foresight of future inflows, were designed via deterministic dynamic programming over 2 years (2013–2014). The weighting method is used to aggregate the two operating objectives (i.e., flood control and irrigation) into a single objective, via convex combination.

– *Baseline solution*: the traditional regulation of the lake is represented in terms of a set of operating policies conditioned on the day of the year d_t and on the lake level h_t. Also these policies were designed via EMODPS.

4 Results and discussion

A first qualitative analysis of the virtual snow index σ defined in Eq. (4) can be performed by comparatively analyzing the trajectory of this VSI with respect to the snow height observations in the closest ground station (i.e., Oga San Colombano, located around 15 km from the webcam) or with respect to some physical variables closely related to the snow dynamics. Figure 8 contrasts the historical trajectory of σ in 2013 with the trajectories of snow height observations at Oga San Colombano station (left panel) and of the freezing level (right panel). Despite some differences due to the different locations of the webcam and the ground station, the first comparison shows similar temporal patterns: most of the snowmelt occurs between April and first half of May, fol-

lowed by a late snowfall at the end of May; no snow is present in the summertime from late June, with the first snowfall of the next winter observed in early October. The comparison between σ and the freezing level shows a negative correlation between low values of freezing level from January to March as well as in November and December, which are associated with high values of σ. On the contrary, the freezing level increases in summertime in correspondence to low and zero values of σ. Moreover, it is worth noting the consistency in the oscillations of the two trajectories especially in wintertime, when the snow accumulation is captured by increasing values of σ associated with decreasing freezing levels and, vice versa, the snow melting corresponds to decreasing values of σ and increasing freezing levels.

To further demonstrate the value of σ, we then quantified its operational value for informing the Lake Como operations (see Sect. 2.4). The performance of this set of informed operating policies (P3) is contrasted with the baseline solution, namely the traditional lake regulation conditioned on the day of the year and the lake level, and the upper-bound solution, namely an ideal set of policies designed under the assumption of perfect foresight of future inflows. The same experiment is repeated using either ground observations of snow height (P1) or SWE data provided by the ARPA (P2) in order to validate the value of the VSI information with respect to traditional data sources.

Figure 9 illustrates the performance of the different set of solutions in terms of flood control (J^{flood}) and irrigation supply (J^{irr}), evaluated over the horizon 2013–2014. The arrows indicate the direction of increasing preference, with the best solution located in the bottom-left corner of the figure. Visual comparison of the baseline (blue circles) and upper-bound solutions (black squares) shows the potential room for improvement generated by the ideal perfect information of the future inflow trajectories. A quantitative measure of this space is provided by the values of the two metrics D_{min} and HV introduced in Sect. 2.4. Table 1 shows that the normal-
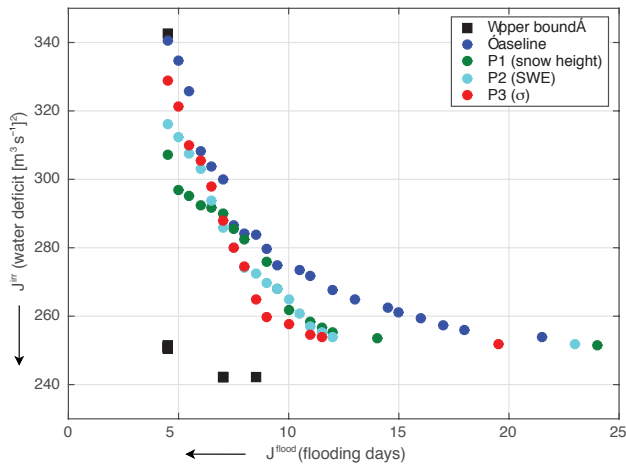
Figure 9. Performance obtained by different Lake Como operating policies informed with ground observations (P1 – green circles), SWE estimated by ARPA (P2 – cyan circles), or virtual snow indexes (P3 – red circles). The performance of these solutions is contrasted with the upper bound of the system performance (black squares) and the baseline operating policies (blue circles).

ized distance between the closest baseline solution to the target upper-bound solution is 0.342, with this gap also confirmed for the entire set of solutions by the 0.292 difference in terms of hypervolume indicator. Valuable snow-related information is hence expected to fill the gap between the baseline and upper-bound solutions. It is interesting to observe that, besides improving the performance of the operating policies with respect to both the objectives, the use of perfect information reduces the conflict between flood control and water supply, and discovers a number of solutions close to the independent optima of the two objectives, including the selected target solution $\mathbf{J}_T = (4.5; 250.6)$.

Given the references provided by the baseline and upperbound solutions, we can assess the operational value of different snow-related information by looking at the performance of informed operating policies, represented by the green, cyan, and red circles in Fig. 9. Not surprisingly, numerical results show that enlarging the information used in the lake operations by accounting for the snow dynamics in the upstream catchment is producing an improvement of the system performance. In fact, the baseline solutions are completely dominated by the sets P1, P2, and P3. These informed operating policies successfully exploit the available snow data to implicitly obtain a medium- to long-term forecast of the future water availability due to snow melt, which supports the daily operations of the lake, balancing flood protection on the short term and water supply on the long one. Overall, the three sets of Pareto optimal solutions, obtained using different snow information, attain similar performance, thus suggesting that the VSI can be considered equivalent to the other two physically based indexes. Figure 9 also shows that policies P1 are the best for very low values of J^{flood}

Table 1. Operational value of the VSIs quantified by the two metrics introduced in Sect. 2.4.

Policy	$D_{\min}$	$\Delta D_{\min}$	HV	ΔHV
baseline	0.342	–	0.708	–
P1 (snow height)	0.291	15.1 %	0.788	11.3 %
P2 (SWE)	0.290	15.2 %	0.785	10.9 %
P3 (σ)	0.238	30.4 %	0.790	11.6 %
upper bound	0.0	–	1.0	–

but high values of J^{irr}, while policies P3 result to be the best in the compromise region of the objectives space (i.e., $J^{\mathrm{flood}} < 10$ days and $J^{\mathrm{irr}} < 275\ (\mathrm{m}^3/\mathrm{s})^2$), which is likely including the most interesting solutions for the lake operator as they successfully balance the system tradeoffs.

Finally, the values of the metrics reported in Table 1 confirm this visual evaluation. The three sets P1, P2, and P3 attain similar values of hypervolume indicator, which assesses the quality of the entire set of solutions. Interestingly, the policies P3 relying on the VSI outperform the other informed solutions both in terms of proximity to the target solution (i.e., lowest value of $D_{\min}$) as well as quality of the entire Pareto front (i.e., highest value of HV). Although the differences in terms of hypervolume are limited, the operational value of σ in terms of $D_{\min}$ is relevant and improves the performance of the baseline solutions by 30 %, doubling the improvement achievable by using either snow height or SWE data.

5 Conclusions

In this paper, we present a web content processing architecture for extracting snow-related information from public web images, either produced by users or generated by touristic webcams. The images, crawled from multiple web data sources, are automatically processed to derive time series of virtual snow indexes representing a proxy of the snow-covered area. We then quantify the operational value of such data for informing the operations of Lake Como.

Numerical analysis shows that the time series of the virtual snow index extracted from a representative webcam is positively correlated with the snow height observations from ground stations and negatively correlated with the freezing level's dynamics. Moreover, our results demonstrate that the operational value of the virtual snow index meets or exceeds the one of traditional snow information. While the use of any snow information allows attainment of a 10 % increase in the hypervolume indicator with respect to the baseline system operations, the operating policies that use the virtual snow index are the closest to the target solution, selected as a good compromise between flood control and irrigation supply.

It is worth noting that our results require a large computing effort for crawling and processing webcams and user-

generated photos for the selected study site. For example, the generation of a $1500 \times 12\,000$ pixel panoramic view requires approximately 1000 ms with a GeForce GTX 850M graphic card. The alignment of an image to the virtual panorama requires approximately 30 000 ms on an OpenStack virtual instance with 4 2.5 GHz VCPUs and 8 Gb of RAM. Conversely, the requirements in terms of human intervention are very low. Human intervention is indeed required only for the skyline annotation and for setting up the experiment on the Lake Como basin (e.g., select the webcam to use, ensuring it has enough information). Finally, the availability of public information, either in terms of webcams or photos, represents a key point for implementing our approach. In our case, we split the 300×160 km region of the Italian and Switzerland Alps using a 5×5 km grid. We analyzed all the images acquired in the specified region over a 6 months period (from 1 December 2014 to 31 May 2015) obtaining a spatial coverage (i.e., the fraction of grid cells containing at least one image over the considered observation period) of 38 and 10 % for photographs and webcams, respectively.

Future research efforts will focus on consolidating this approach by extending the evaluation horizon and by using, at the same time, multiple webcams and photographs to better understand the system dynamics in terms of snow accumulation and melting as well as of the informed lake operations. In parallel, the amount of web content is expected to increase, potentially improving the spatial and temporal resolution of the generated snow-related information and its operational value. We have indeed developed a gamified web portal (http://snowwatch.polimi.it/) where users can cooperatively access and enrich the data set of alpine mountain images, possibly allowing a direct comparison between the information extracted from a webcam and from a user-generated photo in the same location, which is currently unfeasible because we do not have overlapping data. The gamified portal is also expected to facilitate the users' engagement, fostering a more active participation in our image collection effort. Furthermore, our intent is to transform the web platform into a unique mountain-related media repository for testing novel methods and tools. Finally, we are going to release our algorithms as an open-source implementation in order to maximize the portability of our architecture in other snow-dominated catchments where public webcams and user-generated photos are available. This will also allow for exploration of its potential in different environmental problems that may benefit from using public web content sources as low-cost virtual sensors, including sediment monitoring in river beds or vegetation monitoring in remote mountain regions.

6 Data availability

The hydrologic data used in this study were provided by Consorzio dell'Adda (http://www.addaconsorzio.it/) and Agen-

zia Regionale per la Protezione dell'Ambiente (http://ita. arpalombardia.it). The crowdsourced data are available on the SNOW WATCH portal (http://snowwatch.polimi.it).

References

Anghileri, D., Voisin, N., Castelletti, A., Pianosi, F., Nijssen, B., and Lettenmaier, D.: Value of long-term streamflow forecasts to reservoir operations for water supply in snow-dominated river catchments, Water Resour. Res., 52, 4209–4225, doi:10.1002/2015WR017864, 2016.

Baboud, L., Cadik, M., Eisemann, E., and Seidel, H.-P.: Automatic photo-to-terrain alignment for the annotation of mountain pictures, in: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, 41–48, 2011.

Barnett, T. P., Adam, J. C., and Lettenmaier, D. P.: Potential impacts of a warming climate on water availability in snow-dominated regions, Nature, 438, 303–309, 2005.

Bellman, R.: Dynamic programming, Princeton University Press, Princeton, 1957.

Blöschl, G.: Scaling issues in snow hydrology, Hydrol. Process., 13, 2149–2175, 1999.

Bradley, E. S. and Clarke, K. C.: Outdoor webcams as geospatial sensor networks: Challenges, issues and opportunities, Cartogr. Geogr. Inform., 38, 3–19, 2011.

Brown, R. D. and Braaten, R. O.: Spatial and temporal variability of Canadian monthly snow depths, 1946–1995, Atmos. Ocean, 36, 37–54, 1998.

Buytaert, W., Zulkafli, Z., Grainger, S., Acosta, L., Bastiaensen, J., De Bievre, B., Bhusal, J., Chanie, T., Clark, J., Dewulf, A., Foggin, M., Hannah, D. M., Hergarten, C., Isaeva, A., Karpouzoglou, T., Pandey, B., Paudel, D., Sharma, K., Steenhuis, T., Tilahun, S., Van Hecken, G., and Zhumanova, M.: Citizen science in hydrology and water resources: opportunities for knowledge generation, ecosystem service management, and sustainable development, Front. Earth Sci., 2, 1–21, doi:10.3389/feart.2014.00026, 2014.

Castelletti, A., Pianosi, F., and Soncini-Sessa, R.: Water reservoir control under economic, social and environmental constraints, Automatica, 44, 1595–1607, 2008.

Castelletti, A., Galelli, S., Restelli, M., and Soncini-Sessa, R.: Tree-based reinforcement learning for optimal water reservoir operation, Water Resour. Res., 46, W09507, doi:10.1029/2009WR008898, 2010.

Culley, S., Noble, S., Yates, A., Timbs, M., Westra, S., Maier, H., Giuliani, M., and Castelletti, A.: A bottom-up approach to identifying the maximum operational adaptive capacity of water resource systems to a changing climate, Water Resour. Res., 52, 6751–6768, doi:10.1002/2015WR018253, 2016.

Daume, S., Albert, M., and von Gadow, K.: Forest monitoring and social media–Complementary data sources for ecosystem surveillance?, Forest Ecol. Manag., 316, 9–20, 2014.

DeBeer, C. M. and Pomeroy, J. W.: Modelling snow melt and snow-cover depletion in a small alpine cirque, Canadian Rocky Mountains, Hydrol. Process., 23, 2584–2599, 2009.

Dietz, A. J., Kuenzer, C., Gessner, U., and Dech, S.: Remote sensing of snow–a review of available methods, Int. J. Remote Sens., 33, 4094–4134, 2012.

Egli, L.: Spatial variability of new snow amounts derived from a dense network of Alpine automatic stations, Ann. Glaciol., 49, 51–55, 2008.

Farinotti, D., Magnusson, J., Huss, M., and Bauder, A.: Snow accumulation distribution inferred from time-lapse photography and simple modelling, Hydrol. Process., 24, 2087–2097, 2010.

Fedorov, R., Camerada, A., Fraternali, P., and Tagliasacchi, M.: Estimating snow cover from publicly available images, IEEE T. Multimed., 18, 1187–1200, 2016.

Fei-Fei, L. and Perona, P.: A bayesian hierarchical model for learning natural scene categories, in: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), 2, 524–531, 2005.

Floyd, W. and Weiler, M.: Measuring snow accumulation and ablation dynamics during rain-on-snow events: innovative measurement techniques, Hydrol. Process., 22, 4805–4812, 2008.

Fohringer, J., Dransch, D., Kreibich, H., and Schröter, K.: Social media as an information source for rapid flood inundation mapping, Nat. Hazards Earth Syst. Sci., 15, 2725–2738, doi:10.5194/nhess-15-2725-2015, 2015.

Formentin, S., Heusden, K., and Karimi, A.: A comparison of model-based and data-driven controller tuning, Int. J. Adapt. Control, 28, 882–897, 2012.

Formentin, S., Karimi, A., and Savaresi, S. M.: Optimal input design for direct data-driven tuning of model-reference controllers, Automatica, 49, 1874–1882, 2013.

Fraternali, P., Castelletti, A., Soncini-Sessa, R., Vaca Ruiz, C., and Rizzoli, A.: Putting humans in the loop: Social computing for Water Resources Management, Environ. Modell. Softw., 37, 68–77, 2012.

Garvelmann, J., Pohl, S., and Weiler, M.: From observation to the quantification of snow processes with a time-lapse camera network, Hydrol. Earth Syst. Sci., 17, 1415–1429, doi:10.5194/hess-17-1415-2013, 2013.

Giuliani, M. and Castelletti, A.: Is robustness really robust? How different definitions of robustness impact decision-making under climate change, Climatic Change, 135, 409–424, 2016.

Giuliani, M., Herman, J., Castelletti, A., and Reed, P.: Many-objective reservoir policy identification and refinement to reduce policy inertia and myopia in water management, Water Resour. Res., 50, 3355–3377, 2014a.

Giuliani, M., Mason, E., Castelletti, A., Pianosi, F., and Soncini-Sessa, R.: Universal approximators for direct policy search in multi-purpose water reservoir management: A comparative anal-

ysis, in: Proceedings of the 19th IFAC World Congress, Cape Town (South Africa), 2014b.

Giuliani, M., Pianosi, F., and Castelletti, A.: Making the most of data: an information selection and assessment framework to improve water systems operations, Water Resour. Res., 51, 9073–9093, 2015.

Giuliani, M., Castelletti, A., Pianosi, F., Mason, E., and Reed, P.: Curses, tradeoffs, and scalable management: advancing evolutionary multi-objective direct policy search to improve water reservoir operations, J. Water Res. Pl.-ASCE, 142, 04015050, doi:10.1061/(ASCE)WR.1943-5452.0000570, 2016a.

Giuliani, M., Li, Y., Castelletti, A., and Gandolfi, C.: A coupled human-natural systems analysis of irrigated agriculture under changing climate, Water Resour. Res., 52, 6928–6947, doi:10.1002/2016WR019363, 2016b.

Gleason, K. E., Nolin, A. W., and Roth, T. R.: Developing a representative snow monitoring network in a forested mountain watershed, Hydrol. Earth Syst. Sci. Discuss., doi:10.5194/hess-2016-317, in review, 2016.

Good, S. P., Mallia, D. V., Lin, J. C., and Bowen, G. J.: Stable isotope analysis of precipitation samples obtained via crowdsourcing reveals the spatiotemporal evolution of superstorm sandy, PloS one, 9, e91117, doi:10.1371/journal.pone.0091117, 2014.

Graham, E. A., Riordan, E. C., Yuen, E. M., Estrin, D., and Rundel, P. W.: Public Internet-connected cameras used as a cross-continental ground-based plant phenology monitoring system, Global Change Biol., 16, 3014–3023, 2010.

Hadka, D. and Reed, P.: Borg: An Auto–Adaptive Many–Objective Evolutionary Computing Framework, Evol. Comput., 21, 231–259, 2013.

Hashimoto, T., Stedinger, J., and Loucks, D.: Reliability, resilience, and vulnerability criteria for water resource system performance evaluation, Water Resour. Res., 18, 14–20, 1982.

Heidrich-Meisner, V. and Igel, C.: Similarities and differences between policy gradient methods and evolution strategies, in: Proceedings of the 16th European Symposium on Artificial Neural Networks, Citeseer, 149–154, 2008.

Hinkler, J., Pedersen, S. B., Rasch, M., and Hansen, B. U.: Automatic snow cover monitoring at high temporal and spatial resolution, using images taken by a standard digital camera, Int. J. Remote Sens., 23, 4669–4682, 2002.

Holko, L., Gorbachova, L., and Kostka, Z.: Snow Hydrology in Central Europe, Geography Compass, 5, 200–218, 2011.

Jacobs, N., Burgin, W., Fridrich, N., Abrams, A., Miskell, K., Braswell, B. H., Richardson, A. D., and Pless, R.: The global network of outdoor webcams: properties and applications, in: Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM, 111–120, 2009.

König, M., Winther, J.-G., and Isaksson, E.: Measuring snow and glacier ice properties from satellite, Rev. Geophys., 39, 1–27, 2001.

Kunkel, K., Robinson, D., Champion, S., Yin, X., Estilow, T., and Frankson, R.: Trends and Extremes in Northern Hemisphere Snow Characteristics, Current Climate Change Reports, 2, 65–73, 2016.

Laffly, D., Bernard, E., Griselin, M., Tolle, F., Friedt, J.-M., Martin, G., and Marlin, C.: High temporal resolution monitoring of snow

cover using oblique view ground-based pictures, Polar Rec., 48, 11–16, 2012.

Latecki, L. J., Miezianko, R., and Pokrajac, D.: Instantaneous reliability assessment of motion features in surveillance videos, in: Electronic Imaging 2005, 76–86, International Society for Optics and Photonics, 2005.

Le Boursicaud, R., Pénard, L., Hauet, A., Thollet, F., and Le Coz, J.: Gauging extreme floods on YouTube: application of LSPIV to home movies for the post-event determination of stream discharges, Hydrol. Process., 30, 90–105, 2016.

Liaw, A. and Wiener, M.: Classification and regression by random Forest, R news, 2, 18–22, 2002.

Lie, W.-N., Lin, T. C.-I., Lin, T.-C., and Hung, K.-S.: A robust dynamic programming algorithm to extract skyline in images for navigation, Pattern Recogn. Lett., 26, 221–230, 2005.

López-Moreno, J. and Nogués-Bravo, D.: Interpolating local snow depth data: an evaluation of methods, Hydrol. Process., 20, 2217–2232, 2006.

Lowry, C. S. and Fienen, M. N.: CrowdHydrology: crowdsourcing hydrologic data and engaging citizen scientists, GroundWater, 51, 151–156, 2013.

Maier, H., Kapelan, Z., Kasprzyk, J., Kollat, J., Matott, L. S., Cunha, M. C., Dandy, G. C., Gibbs, M. S., Keedwell, E., Marchi, A., Ostfeld, A., Savic, D., Solomatine, D. P., Vrugt, J. A., Zecchin, A. C., Minsker, B. S., Barbour, E. J., Kuczera, G., Pasha, F., Castelletti, A., Giuliani, M., and Reed, P. M.: Evolutionary algorithms and other metaheuristics in water resources: Current status, research challenges and future directions, Environmental Modell. Softw., 62, 271–299, 2014.

Mazzoleni, M., Alfonso, L., Chacon-Hurtado, J., and Solomatine, D.: Assimilating uncertain, dynamic and intermittent streamflow observations in hydrological models, Adv. Water Res., 83, 323–339, 2015a.

Mazzoleni, M., Verlaan, M., Alfonso, L., Monego, M., Norbiato, D., Ferri, M., and Solomatine, D. P.: Can assimilation of crowdsourced streamflow observations in hydrological modelling improve flood prediction?, Hydrol. Earth Syst. Sci. Discuss., 12, 11371–11419, doi:10.5194/hessd-12-11371-2015, 2015b.

Michelsen, N., Dirks, H., Schulz, S., Kempe, S., Al-Saud, M., and Schüth, C.: YouTube as a crowd-generated water level archive, Sci. Total Environ., 568, 189–195, 2016.

Mote, P. W., Hamlet, A. F., Clark, M. P., and Lettenmaier, D. P.: Declining mountain snowpack in western North America, B. Am. Meteorol. Soc., 86, 39–49, 2005.

Muñoz, J., Infante, J., Lakhankar, T., Khanbilvardi, R., Romanov, P., Krakauer, N., and Powell, A.: Synergistic Use of Remote Sensing for Snow Cover and Snow Water Equivalent Estimation, British Journal of Environment & Climate Change, 3, 612–627, 2013.

Newald, T. G. and Lehning, M.: Altitudinal dependency of snow amounts in two small alpine catchments: can catchment-wide snow amounts be estimated via single snow or precipitation stations?, Ann. Glaciol., 52, 153–158, 2011.

Parajka, J. and Blöschl, G.: Validation of MODIS snow cover images over Austria, Hydrol. Earth Syst. Sci., 10, 679–689, doi:10.5194/hess-10-679-2006, 2006.

Parajka, J. and Blöschl, G.: Spatio-temporal combination of MODIS images–potential for snow cover mapping, Water Resour. Res., 44, W03406, doi:10.1029/2007WR006204, 2008.

Parajka, J., Haas, P., Kirnbauer, R., Jansa, J., and Blöschl, G.: Potential of time-lapse photography of snow for hydrological purposes at the small catchment scale, Hydrol. Process., 26, 3327–3337, 2012.

Pepe, M., Brivio, P., Rampini, A., Nodari, F., and Boschetti, M.: Snow cover monitoring in Alpine regions using ENVISAT optical data, Int. J. Remote Sens., 26, 4661–4667, 2005.

Perona, P., Daly, E., Crouzy, B., and Porporato, A.: Stochastic dynamics of snow avalanche occurrence by superposition of Poisson processes, P. R. Soc. A, 468, 4193–4208, 2012.

Powell, W.: Approximate Dynamic Programming: Solving the curses of dimensionality, Wiley, NJ, 2007.

Ruzon, M. A. and Tomasi, C.: Edge, junction, and corner detection using color distributions, IEEE T. Pattern Anal., 23, 1281–1295, 2001.

Schuler, M., Stucki, E., Roque, O., and Perlik, M.: Mountain Areas in Europe: Analysis of mountain areas in EU member states, acceding and other countrie, Tech. Rep. 2002.CE.16.0.AT.136, Nordic Centre for Spatial Development, 2004.

Schweizer, J., Mitterer, C., and Stoffel, L.: On forecasting large and infrequent snow avalanches, Cold Reg. Sci. Technol., 59, 234–241, 2009.

Serdyukov, P., Murdock, V., and Van Zwol, R.: Placing flickr photos on a map, in: Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval, ACM, 484–491, 2009.

Smith, L., Liang, Q., James, P., and Lin, W.: Assessing the utility of social media as a data source for flood risk management using a real-time modelling framework, Journal of Flood Risk Management, doi:10.1111/jfr3.12154, in press, 2015.

Staudinger, M., Stahl, K., and Seibert, J.: A drought index accounting for snow, Water Resour. Res., 50, 7861–7872, 2014.

Tesic, J.: Metadata practices for consumer photos, IEEE MultiMedia, 12, 86–92, 2005.

Tsitsiklis, J. and Van Roy, B.: Feature-Based Methods for Large Scale Dynamic Programming, Mach. Learn., 22, 59–94, 1996.

Vitolo, C., Elkhatib, Y., Reusser, D., Macleod, C., and Buytaert, W.: Web technologies for environmental Big Data, Environ. Modell. Softw., 63, 185–198, 2015.

Wood, A. and Lettenmaier, D.: A test bed for new seasonal hydrologic forecasting approaches in the western United States, B. Am. Meteorol. Soc., 87, 1699–1712, doi:10.1175/BAMS-87-12-1699, 2006.

Zatarain-Salazar, J., Reed, P., Herman, J., Giuliani, M., and Castelletti, A.: A diagnostic assessment of evolutionary algorithms for multi-objective surface water reservoir control, Adv. Water Resour., 92, 172–185, 2016.

Zoppoli, R., Sanguineti, M., and Parisini, T.: Approximating networks and extended Ritz method for the solution of functional optimization problems, J. Optimiz. Theory App., 112, 403–440, 2002.