

VOLUME 13, ISSUE 1, ARTICLE 15

Chabot, D., C. Dillon, and C. M. Francis. 2018. An approach for using off-the-shelf object-based image analysis software to detect and count birds in large volumes of aerial imagery. *Avian Conservation and Ecology* 13(1):15. https://doi.org/10.5751/ACE-01205-130115 Copyright © 2018 by the author(s). Published here under license by the Resilience Alliance.

Methodology

An approach for using off-the-shelf object-based image analysis software to detect and count birds in large volumes of aerial imagery

Dominique Chabot¹, Christopher Dillon¹ and Charles M. Francis²

¹droneMetrics, ²Canadian Wildlife Service, Environment and Climate Change Canada

ABSTRACT. Computer-automated image analysis techniques can save time and resources for detecting and counting birds in aerial imagery. Sophisticated object-based image analysis (OBIA) software is now widely available and has proven effective for various challenging detection tasks, but there is a need to develop accessible and readily adaptable procedures that can be implemented in an operational context. We developed a systematic, repeatable approach using commercial off-the-shelf OBIA software, and tested its effectiveness and efficiency to detect and count Lesser Snow Geese (*Chen caerulescens caerulescens*) in large numbers of images of breeding colonies across the Canadian Arctic that present a variety of landscapes, numerous confounding features, and varying illumination conditions and exposure levels. Coarse-scale review of analysis results was necessary to remove conspicuous clusters of commission errors, thus rendering the technique semiautomated. It was effective for imagery with spatial resolutions of 4-5 cm, producing overall accurate estimates of goose numbers compared to manual counts (R2 = 0.998, regression coefficient = 0.974) in 41 test images drawn from several breeding colonies. The total automated count (19,920) across all test images exceeded the manual count (19,836) by just 0.4%. We estimate the typical time required to review images for errors to be only 5-10% of that required to count birds manually. This could reduce the person-time required to many other bird detection tasks in aerial imagery by anyone possessing at least basic skills in image analysis and geographic information systems.

Une approche pour l'utilisation d'un logiciel d'analyse d'images orientée objet grand public afin de détecter et de compter les oiseaux dans un grand volume d'images aériennes

RÉSUMÉ. Les techniques d'analyse d'images comptées automatiquement par ordinateur peuvent sauver du temps et de l'argent pour détecter et compter les oiseaux sur des photographies aériennes. Des logiciels sophistiqués d'analyse d'images orientée objet (OBIA; object-based image analysis en anglais) sont maintenant disponibles à grande échelle et leur efficacité a été démontrée pour diverses tâches de détection complexes; toutefois, le besoin de créer des procédures accessibles et facilement adaptables qui puissent être mises en place en contexte opérationnel demeure. Nous avons élaboré une approche systématique et répétable à partir d'un logiciel OBIA grand public et avons testé son efficacité et sa précision à détecter et à compter les Petites Oies des neiges (Chen caerulescens caerulescens) sur un grand nombre d'images de colonies de reproduction dans l'Arctique canadien, territoire qui présente une variété de paysages, de nombreux éléments confondants et des degrés divers de conditions lumineuses et d'exposition. Il a été nécessaire de faire une revue des résultats d'analyses à une échelle grossière afin d'enlever les agrégations évidentes d'erreurs de commission, permettant ainsi de semi-automatiser la technique. Le logiciel s'est avéré efficace avec les images ayant une résolution spatiale de 4-5 cm, produisant des estimations globales du nombre d'oies dont la précision était comparable à celle de comptes manuels (R2 = 0.998, coefficient de régression = 0, 974) réalisés sur 41 images tests tirées de plusieurs colonies. Le compte automatisé total (19 920) à partir de toutes les images tests a excédé le compte manuel (19 836) de 0,4 % seulement. Nous avons estimé que le temps généralement requis pour repasser les images à la recherche d'erreurs était de 5 à 10 % seulement de celui requis pour compter les oiseaux manuellement. L'usage de ce logiciel permettrait de réduire le temps-personne requis pour analyser les photos aériennes des principales colonies d'Oies des neiges dans l'Arctique de plusieurs mois à plusieurs jours. Notre approche pourrait être adaptée à de nombreuses autres tâches de détection d'oiseaux sur des images aériennes par toute personne possédant au moins des notions de base en analyse d'images et en systèmes d'information géographique.

Key Words: abundance estimation; aerial surveys; automation; census techniques; data analysis; image processing; population monitoring; remote sensing

INTRODUCTION

Monitoring the distribution and abundance of bird populations and how they are changing over time provides the foundation for conservation planning and management. Trends in populations are used to set conservation priorities and identify species in need of conservation action (e.g., NABCI 2016, Rosenberg et al. 2017). Population monitoring is also important for evaluating the effectiveness of management actions, including those aimed to control overabundant species such as some geese (e.g., Leafloor et al. 2012, Lefebvre et al. 2017).

The use of aerial surveys to monitor bird populations is a wellestablished and widespread practice, enabling rapid coverage of large areas and facilitating access to remote or otherwise challenging places such as undeveloped areas, wetlands, marine environments, and polar regions. They are commonly used to survey waterbirds (Kingsford and Porter 2009), as well as raptors (Good et al. 2007) and upland game birds (Butler et al. 2007). Bird identification and counts are performed either in real time by airborne observers or using on-board cameras to collect aerial imagery that is reviewed and analyzed later. Several studies have indicated that aerial image counts can be more accurate and consistent than live-observer counts (Boyd 2000, Frederick et al. 2003, Buckland et al. 2012), but a major drawback of the former is the significant time and effort required to manually analyze large volumes of imagery (Woodworth et al. 1997, Béchet et al. 2004). With burgeoning use of small low-flying unmanned aircraft, or drones, as a means of collecting very high-resolution aerial imagery of birds (Chabot and Bird 2015), as well as increasing possibilities to census birds in satellite imagery (LaRue et al. 2017), the challenge of analyzing large volumes of imagery to detect and count subjects has been receiving increased attention.

The use of computer-automated techniques to count birds in aerial imagery dates back three decades (Gilmer et al. 1988), but did not gain much traction until the 21st century when a combination of advancements in image analysis software, computer processing performance, and digital camera technology have progressively made the techniques more accessible (Chabot and Francis 2016). When the color of subjects contrasts sharply with image backgrounds, they can in some cases be automatically isolated and counted using simple spectral thresholding in general-purpose image editing software (Chabot and Bird 2012). Additional size-based filtering of features isolated by thresholding can help reduce erroneous counting of nontarget features with similar colors to the birds (Bajzak and Piatt 1990, Laliberte and Ripple 2003, Trathan 2004). Descamps et al. (2011) developed an application for counting large aggregations of birds on the ground in aerial images that automatically detects and tallies ellipse-shaped features that contrast with the background.

More versatile techniques, such as object-based image analysis (OBIA; Blaschke 2010), are required when subjects contrast weakly with the background, vary in size or shape, and/or are more sparsely distributed throughout large numbers of images with varying backgrounds and numerous confounding features. OBIA is founded on segmenting images into a patchwork of spectrally distinct, multipixel objects that tend to correspond to coherent "real-world" features. Objects can then be analyzed in elaborate ways based on a large variety of spatial, spectral, and texture attributes. OBIA has notably been used to detect birds at sea (Groom et al. 2007, 2013) and on inland waters (Groom et al. 2011, Liu et al. 2015), as well as terrestrial mammals in a nature park (Chrétien et al. 2015, 2016). In recent years, commercial offthe-shelf OBIA software has become increasingly accessible, but there is a need to develop standardized and readily adaptable procedures for using the software to detect and count birds (Chabot and Francis 2016).

We evaluated the utility and efficiency of automated image analysis to count Lesser Snow Geese (Chen caerulescens caerulescens) from aerial surveys of breeding colonies throughout the Canadian Arctic. Population monitoring by the Canadian Wildlife Service (CWS) and the United States Fish and Wildlife Service (USFWS) has been ongoing for two decades (Kerbes et al. 2006, 2014) following action aimed at reducing Snow Goose numbers to minimize damage to Arctic habitats (Batt 1997, Leafloor et al. 2012). Population estimates from aerial photographic surveys of major known breeding colonies indicated an increase from 4.3 million birds in 1995-1998 to 5.1 million in 2005-2008 (Kerbes et al. 2014). To date, estimates have been obtained by manually counting geese in thousands of photos across all colonies, constituting a very labor-intensive task that can take several months of a technician's time, even when the images are subsampled. Following a transition to digital photography in 2009, the imagery is more readily amenable to computer-automated analysis.

Previous studies successfully used basic image thresholding and size-filtering techniques to count white-phase Snow Geese (Gilmer et al. 1988, Bajzak and Piatt 1990, Laliberte and Ripple 2003, Chabot and Bird 2012), but these studies mostly involved dense concentrations of staging or wintering geese, few total images, processing of images one at a time to adjust threshold levels to varying illumination/exposure conditions, and/or manually cropping images to include only areas containing birds. In contrast, we were interested in processing large numbers of images capturing a wide variety of landscapes, numerous confounding features, and bird concentrations ranging from very dense in a minority of images to comparatively sparse in the majority of images, and completely absent in others. Although Snow Geese are sexually monomorphic, there are two color morphs: white-phase geese are largely white viewed from above, while blue-phase geese are largely blue-gray with a white head, and contrast much less against the background. With surveys having been flown over the course of several years at different times of day and under a variety of sky conditions, there is also significant variation in illumination and exposure across the image sets, resulting in varying spectral-radiometric characteristics of geese.

We therefore investigated an OBIA approach for detecting the geese, aiming for a solution that enabled fully automated processing of large batches of images. To this end, we developed a systematic approach to establish optimal image segmentation and classification parameters, which could be readily applied to processing aerial images of other birds or wildlife targets. We used a commercial remote sensing image analysis program, ENVI 5.3 (Exelis Visual Information Solutions, Boulder, CO, USA), bundled with the OBIA-capable Feature Extraction module and the IDL (Interactive Data Language) programming application to execute batch-processing, although our general approach could also be implemented in other similar software packages. To maximize its accessibility, we largely limited ourselves to the use of basic functions and avoided overly complex or customized operations. In this paper, we present and discuss (1) the steps involved in our approach as they were carried out in our trial with imagery of Snow Geese, providing background on certain foundational image analysis concepts for the benefit of others who may wish to adopt the approach; (2) the results and performance of the developed analysis routine on a large batch of imagery.

METHODS

Description of the imagery

We used existing digital aerial imagery collected and postprocessed by the USFWS with the assistance of the CWS. The imagery was collected from 2009 to 2014 over Lesser Snow Goose breeding colonies throughout the Canadian Arctic, divided into four main regions (Fig. 1): the Western Arctic (surveyed in 2009 and 2013), Baffin Island (2011), Southampton Island (2014), and West Hudson Bay (2014). Three-band truecolor (RGB) photos were acquired with a DSS 439 (Applanix, Richmond Hill, ON, Canada) 39-megapixel aerial camera and direct georeferencing system through a 40- or 60-mm focal-length lens at spatial resolutions ranging from 2 to 39 cm. Photos were postprocessed with POSPac MMS (Applanix) to convert them to 8-bit TIFF format and correct for lens vignetting. Inpho OrthoMaster (Trimble, Sunnyvale, CA, USA) was then used to correct for lens distortion and orthorectify images to the ASTER GDEM (global digital elevation map) or GDEM2. In some cases, images were combined into orthomosaics containing up to a dozen photos using Inpho OrthoVista; in other cases, image files consisted of single, nonoverlapping orthophotos. Colonies covering smaller areas were generally imaged in their entirety, whereas those covering larger areas were sampled with strip transects. Traditional manual counts of geese in the imagery have been performed in a subset of images from each colony, with the results then extrapolated to the measured or estimated total area of the colony to produce a population estimate.

Fig. 1. Canadian Arctic regions where digital aerial imagery of Lesser Snow Goose (*Chen caerulescens caerulescens*) breeding colonies was collected from 2009 to 2014: (A) Western Arctic, (B) Baffin Island, (C) Southampton Island, and (D) West Hudson Bay.



Upon preliminary review of the imagery, we judged that spatial resolutions coarser than 6 cm were unlikely to be amenable to computer-automated detection of geese because of the decreasing confidence with which birds could be manually identified. We therefore focused our trial on resolutions of 3 cm (1736 total files, mostly single photos, from Southampton Island and West Hudson Bay), 4 cm (52 multiphoto mosaic files from the Western Arctic), 5 cm (238 files of both types from Baffin Island and the Western Arctic), and 6 cm (612 mosaic files from Southampton Island and West Hudson Bay), enabling us to evaluate how the capacity to automatically detect geese is affected by resolution.

Depending on posture, a goose on the ground covers an area of \sim 500–1000 cm² when viewed from above. Thus, a 3-cm spatial resolution results in \sim 55–110 pixels per bird, while a 6-cm resolution results in \sim 14–28 pixels per bird.

Selecting training images and manually identifying birds

We reviewed all available imagery at the selected resolutions by opening all image files in OGIS 2.18 (OGIS Development Team). one colony at a time, and panning through them at low zoom levels (1:2000-1:8000) on a 1920 x 1080-pixel monitor. Throughout this process, we noted features that could potentially be misidentified as Snow Geese based on spectral and spatial similarities. Predominant confounding features included patches of snow and ice, bright rocks and boulders, clumps of froth or ice along the edges of water, and glint reflected from water surfaces. We then selected a limited number of "training images" to be used to develop and perform initial testing of OBIA segmentation and classification parameters, deliberately selecting seemingly challenging images that contained both geese and confounding features. We selected a single training image for colonies with \leq 50 total image files, and additional images for colonies containing larger numbers of files. Overall, we selected 21 training images at 3-cm resolution, one image at 4-cm resolution, seven images at 5-cm resolution, and 13 images at 6cm resolution.

We manually identified all geese in each training image by creating a point shapefile in QGIS and marking each goose with a point. This was done by panning through the entire area of each image file at zoom levels of 1:250–1:500, occasionally zooming in further if necessary to confidently identify geese in more challenging landscapes. Although geese could be confidently identified throughout most of the training images, some omissions as well as false identifications may have occurred, especially in areas with dense bright rocks interspersed with geese and in the coarser 6cm imagery. We identified a total of 900, 1350, 772 and 5138 geese in the 3-, 4-, 5- and 6-cm training images, respectively.

Establishing image segmentation parameters

Effective segmentation of target image features is crucial to their successful subsequent classification based on spatial, spectral, and texture attributes. Improper delineation of feature boundaries, oversegmentation (target features are divided into multiple objects), and omission (target features are lumped into larger objects containing other adjacent features) can all hinder classification performance.

ENVI's Feature Extraction module uses a "watershed" segmentation algorithm (Roerdink and Meijster 2001), which first transforms the input image into a contrast gradient map in which higher values are assigned to pixels along the boundaries of spectrally contrasting features, e.g., the contour of a Snow Goose against a darker ground, while lower values are assigned to pixels in more spectrally homogeneous areas. This results in a patchwork of "basins," i.e., objects, delineating contrasting features, where the brightness of pixels separating the basins can be likened to the height of dams (Fig. 2). The contrast gradient map is then "flooded" from the bottom up, with the adjustable "scale level" parameter determining to what height the image is flooded along a normalized scale of 0–100. Increasing the scale level results in

Fig. 2. Simplified illustration of how a watershed segmentation is performed on a portion of an image containing a bright object and a dark object against a uniform background. The image is first transformed into a contrast gradient map in which dark "basins" represent spectrally homogeneous parts of the image, separated by "dams" of varying height, i.e., brightness, proportionate to the magnitude of spectral contrast between the adjacent image features. Segmentation of the image is achieved by "flooding" the basins to a specified level. In this example, progressive flooding from the bottom up will cause the dark object to be amalgamated with adjacent background objects before the bright object, because the former's magnitude of contrast with the background is smaller than the latter's.



the amalgamation of basins separated by shorter dams, preserving only those surrounded by increasingly tall dams, which represent image features with increasingly distinct edges (Fig. 2). Although increasing the scale level results in fewer objects of larger average size, there is no direct relationship between the scale level and the real-world size of the objects. Next, a "merge" step can be used to further combine spectrally similar adjacent objects. This tends to be useful for reducing clutter and oversegmentation so as to generate objects that properly capture target features. Increasing the "merge level" parameter between 0 and 100 results in preserving only objects that are increasingly spectrally dissimilar from adjacent or encapsulating objects. Finally, once the image has been segmented, a series of spatial, spectral, and texture attribute values are calculated for each object, which can be used for subsequent object classification.

To achieve our aim of batch-processing images, we needed to select a set of standard segmentation parameters that was suitably effective across many different images. As a result of variation in scene composition, illumination and exposure among images, optimal segmentation parameters to capture a given type of feature vary from image to image. For example, a given set of parameters that yields optimal segmentation of Snow Geese in one image might tend to oversegment or omit geese in others (Fig. 3). To find a suitable compromise, we tested each training image under varying segmentation parameters, detailed below. Segmentation is computationally intensive, often requiring ≥ 30 min to fully process a multigigabyte orthomosaic file on our highend desktop computer. To accelerate the process, we used the segmentation previewing function in the Feature Extraction workflow, which displays the objects generated by the currently defined parameters in a small window that can be panned over the original image and is updated in real time when parameters are modified. Given the large numbers of geese in many of the images, we focused parameter testing on a sample of birds in each image, particularly ones that were observed to have a tendency to be oversegmented or omitted.

Preprocessing

Following initial testing, we applied two preprocessing operations to the training images in ENVI to improve segmentation results.

First, we excluded the black rectangular backdrops encasing the imagery in each file—composed of pixels with a value of 0 in all image bands—by setting a "no data" value of 0. The backdrops skewed segmentation results by introducing an artificial sharp contrast around the edges of the imagery. Removing them better normalized the response of the images to varying segmentation scale and merge levels. Second, we applied a low-pass smoothing filter that adjusted the value of every pixel to the mean of a 9-pixel square box centered on each pixel (Laliberte and Ripple 2003). This had the effect of reducing spectral contrasts within individual geese that tended to cause oversegmentation, while preserving their strong contrast with the background.

Segmentation bands

We assessed segmentation effectiveness when performed on varying combinations of the three image bands. The small image silhouettes of geese were subject to fine-scale spatial misalignment of the bands, with the red and blue bands often appearing slightly shifted to either side of the green band. In addition, whereas the contrast of geese with the background appeared generally consistent across images in the green and blue bands, it was noticeably more variable in the red band, appearing much sharper in images captured under sunny/clear sky conditions than under cloudy/overcast conditions. Consequently, we found that entirely excluding the red band from segmentation and merging resulted in improved delineation of geese across the training images. We determined that performing segmentation on a combination of the green and blue bands produced good delineation of geese while minimizing omission of blue geese, and restricting merging to only the green band minimized oversegmentation.

Merge algorithm

The default "full lambda schedule" merge algorithm (Redding et al. 1999) is effective at creating a relatively clean patchwork of objects representing predominant image features, e.g., broad patches of vegetation or other land cover, tending to merge small contrasting objects that might be regarded as noise into larger encapsulating objects, and in our case increasing omission of geese. Instead, we used the "fast lambda" algorithm (Exelis Visual Information Solutions), which merges adjacent objects based on their degree of spectral similarity as well as on the length of their **Fig. 3**. Example of an oversegmented bird (A) and an omitted bird (B) among adequately segmented birds and other segmented features in aerial imagery of Canadian Arctic-breeding Lesser Snow Geese (*Chen caerulescens caerulescens*).



shared border, such that larger contrasting features are more likely to be merged than smaller ones. Thus, small bright Snow Goose objects tended to be preserved at high merge levels while larger features tended to be merged into a single common background object, with the added benefit of reducing the total number of objects generated, and consequently processing time.

Scale and merge levels

With the above operations and settings applied, we found that optimal results across the 4- and 5-cm training images were achieved at a segmentation scale level of 90 and merge level of 90. For the 6-cm images, a lower scale level of 85 and merge level of 85 worked better, reflecting the fact that geese were noticeably darker and less contrasting at this coarser resolution. We did not succeed in achieving satisfactory segmentation of geese across the 3-cm training images using a common set of parameters because the increased level of detail revealed at this resolution was markedly more prone to causing oversegmentation in certain images. As a result, we ceased analyses of 3-cm imagery at this stage, although we consider potential approaches for analyzing these data in the discussion.

Texture kernel

Finally, we adjusted the size of the texture kernel (Warner 2011), a square box of pixels serving as a sampling grid that is applied to every pixel in each object to compute the object's texture attributes. The value of each attribute is calculated within the kernel for each pixel, then averaged across all pixels in the object to arrive at a single value for the object. For pixels along the edges of the object, the kernel crosses the boundary of the object and includes outside pixels in the texture calculations (Fig. 4). We exploited this "spillover" to quantify the magnitude of contrast between geese and their immediate surroundings via the texture range and variance attributes (Warner 2011), reasoning that while

the absolute brightness of geese varied across the imagery because of illumination/exposure variation, their relative contrast with the background should remain more consistent and consequently of value for their classification. To capture an ample amount of background surrounding each goose object, we increased the texture kernel size from the default 3 x 3 pixels to 5 x 5 pixels for the 5- and 6-cm imagery, and to 7 x 7 pixels for the 4-cm imagery. This ensured that the texture kernel extended outside the object when applied to virtually every pixel inside it, resulting in the birdbackground contrast being strongly reflected in the texture range and variance attributes (Fig. 4). In ENVI, the kernel size must be set prior to segmenting the image, and is then applied to texture attribute calculations once objects have been formed. It should be noted that a larger texture kernel size requires more processing time for attribute calculations.

Compiling attribute data of segmented bird objects

Once segmentation parameters were established, we executed the full segmentation process on all training images, outputting for each image a polygon shapefile delineating all segmented objects and an associated attribute table containing the spatial, spectral, and texture attribute values for each object. We then overlaid the segmentation polygon file, the point file of manually identified Snow Geese, and the image raster of each training image in QGIS. Using the point layer's attribute table, we successively zoomed to every manually identified bird and cumulatively selected each corresponding object in the segmentation polygon layer. We then copied and pasted the selected rows of the polygon layer's attribute table, i.e., the attribute values for all objects representing geese, into a spreadsheet. In this manner, we compiled a dataset of attribute values for all manually identified geese throughout the training images, with the exception of those that were omitted by the segmentation. We also excluded oversegmented geese when **Fig. 4.** Typical pixel dimensions of a Lesser Snow Goose (*Chen caerulescens caerulescens*) object following segmentation of 5-cm (left) and 4-cm (right) resolution aerial imagery, showing the surrounding area included in object texture calculations when using a kernel size of 5 x 5 pixels and 7 x 7 pixels, respectively. The values assigned to the pixels represent the cumulative number of times each pixel is sampled by the kernel when applied to every pixel inside the object in the process of calculating texture attributes. With the total number of times outside pixels are sampled approaching the number of times inside pixels are sampled, the magnitude of contrast between birds and background is strongly reflected in the texture range and variance attributes.



there was no clearly dominant object that covered most of the bird, as well as occasional flying geese because of their highly variable shapes and sizes. We grouped the attribute data by image resolution, and for each attribute at each resolution we generated a frequency histogram and calculated the mean, standard deviation, range, minimum, maximum, and a series of percentiles ranging from the 1st to the 99th. These statistics were then used to guide the development of classification rule sets to distinguish geese from other image objects.

Developing bird object classification rule sets

We generated an initial classification rule set template using the ENVI Feature Extraction workflow, containing a single "Snow Goose" class that consisted of a single rule incorporating all object attributes for which geese in the training images showed a clustered and well-defined distribution of values, totalling 11 spatial, 12 spectral, and 12 texture attributes (see Appendix 1, Table A1.1 for a listing of attributes and their descriptions). We then saved the rule set as a text file, created a dedicated file for each image resolution, and performed all further modifications in a text editor. We adjusted the "class threshold" so that objects would only be classified as geese if all 35 attribute criteria were satisfied, and set the "operation" for each attribute to "between" to express each attribute criterion as a minimum-to-maximum range of values.

We determined appropriate attribute value ranges based on the descriptive statistics for the attributes of the manually identified Snow Geese in the training images. We assumed that the distributions of the spatial attributes describing their size and shape remain consistent among images at a given resolution, and that the values recorded among geese in the training images were therefore representative of the broader population. With a view to minimize omission errors, we set the lower and upper thresholds for each spatial attribute near the minimum and maximum values recorded in the training images, conspicuous outliers excluded. Examples of the distribution of values for four spatial attributes among geese in the 4-cm resolution imagery are shown in Fig. 5.

Fig. 5. Frequency histograms showing the distribution of values for four spatial attributes among 1335 manually identified Lesser Snow Geese (*Chen caerulescens caerulescens*) in 4-cm resolution aerial imagery.



Unlike the spatial attributes, we expected that the values of certain spectral attributes of Snow Geese, namely the mean, minimum, and maximum, could vary beyond the ranges recorded in the training images because the broader image sets potentially contained darker/lower exposure as well as brighter/higher exposure imagery than captured among the training images. We therefore set looser ranges for these attributes to accommodate additional variation, with the lower and upper thresholds set to

one to three standard deviations below the 1st percentile and above the 99th percentile, respectively, of values recorded in the training images. Because of the overall greater spectral consistency of geese in the green band than in the red and blue bands (Fig. 6), this resulted in generally narrower ranges for green band attributes.

Fig. 6. Frequency histograms showing the distribution of values for the "spectral mean" attribute in each of the three image bands among 744 manually identified Lesser Snow Geese (*Chen caerulescens caerulescens*) in 5-cm resolution aerial imagery. Greater spectral consistency of geese in the green band is evident compared to the red and blue bands.



Similar to the spatial attributes, we set the lower thresholds for the texture range and variance attributes, used as proxies for the magnitude of contrast between geese and background, near the minimum values recorded among geese in the training images, outliers excluded. As was observed while establishing image segmentation parameters, the contrast of geese was more consistent in the green and blue bands than in the red band (Fig. 7). By setting loose lower thresholds for spectral attributes while maintaining stricter lower thresholds for these texture attributes, the rule set allowed geese in darker imagery to be correctly classified as long as their relative contrast with the background was sufficient. It also prevented misclassification of nongoose objects in brighter imagery that met the spectral criteria but did not contrast sufficiently with the background to meet the texture range and variance criteria. We set looser thresholds for the other texture attributes using the same approach as with the spectral attributes.

We performed initial tests of the classification rule sets developed for each image resolution on the training images, prompting us to make a series of minor adjustments to attribute value ranges in an effort to improve the balance between omission and commission error rates. At this stage, we observed that classification performance was strikingly poor in the 6-cm resolution imagery, with automated counts exceeding manual counts by an average of 3460% per image (range = 30-13,154%, median = 783%), even after several adjustments to the rule set. We attributed this poor performance to the significantly darker and less distinct appearance of Snow Geese in the 6-cm imagery than at higher resolutions (Fig. 8), which was already evident in the need to use lower segmentation scale and merge levels to capture them as objects. Consequently, we ceased further testing of 6-cm imagery and focused on the 4- and 5-cm imagery.

Batch-processing full image sets

We used IDL to write a script to sequentially process all image files contained in a given directory in accordance with our established **Fig. 7**. Frequency histograms showing the distribution of values for the "texture range" attribute, used as a proxy for magnitude of contrast between object and background, in each of the three spectral bands among 744 manually identified Lesser Snow Geese (*Chen caerulescens caerulescens*) in 5-cm resolution aerial imagery. Contrast is sharpest and most consistent in the green band, and least consistent in the red band.



Fig. 8. Frequency histograms showing the distribution of values for the "spectral mean" attribute in the green band among manually identified Lesser Snow Geese (*Chen caerulescens caerulescens*) in 4-cm (n = 1335), 5-cm (n = 744), and 6-cm (n = 5024) resolution aerial imagery. Geese appear considerably darker in the 6-cm imagery compared to the 4- and 5-cm imagery.



preprocessing operations, segmentation parameters, and classification rule sets. We used the script to perform initial testing of our classification rule sets on the training images and, following refinement of the rule sets, to process all available imagery at 4- and 5-cm resolutions (totalling 290 files and 234 GB; see sample images of the different colonies in Appendices 2-9) for further evaluation of our routine's performance. An example of the full script code is shown and detailed in Appendix 10, Fig. A10.1. For each image, the script outputs a polygon shapefile delineating all objects classified as Snow Geese. When opened in a geographic information system (GIS), the total bird count can be obtained via the layer's attribute table or summary statistics, and the shapefiles for multiple images, e.g., all image files comprising a colony, can be merged into a single file to streamline tallying of counts. The time required to batchprocess images depends on the number of files to be processed, the size of each file, the number of objects generated by the segmentation of each image, the kernel size for texture attribute calculations, and the computer hardware's performance. In some cases, several days of continuous processing may be required, although once initiated,

Fig. 9. Examples of typical commission errors resulting from automated object-based analysis of large volumes of aerial imagery to detect Lesser Snow Geese (*Chen caerulescens caerulescens*) throughout Canadian Arctic breeding colonies: small patches of snow around the periphery of large patches (A), clumps of ice or froth along the edges of water (B), speckles of glint on the surface of water (C), and large bright rocks (D). Objects classified as geese are delineated in red, all of them commission errors except those indicated by green arrows.



the process is fully automated and can be left unattended until all files have been processed.

RESULTS

The automated analysis classified a total of 85,267 objects as Snow Geese across all available 4- and 5-cm resolution image files. As noted above, we found that geese had an excessive propensity for oversegmentation in 3-cm imagery and excessively poor classification performance in 6-cm imagery, so we did not process the full image sets at those resolutions. To evaluate the performance of the automated analysis, we selected at least two images per colony (more for colonies with larger numbers of image files) in addition to the training images, for a total of 41 test images spanning a wide range of automated counts (0-5285), in which we performed manual counts for comparison (Table 1).

Across the test images, the automated analysis identified 22,307 total objects as potential Snow Geese, compared to 19,836 manually counted geese, indicating significant numbers of commission errors. The most common sources of error were small patches of snow or ice, clumps of froth or ice along the edges of water, speckles of glint on water surfaces, and large bright rocks (Fig. 9). In many cases, errors of the former three types were conspicuously clustered; for example, large patches of snow with

Table 1. Comparison of computer-automated Lesser Snow Goose (*Chen caerulescens caerulescens*) counts, adjusted counts in which conspicuous clusters of commissions errors have been removed, and fully manual counts in 41 aerial images of breeding colonies in the Canadian Arctic. Image filenames beginning with "or" are single orthophotos whereas others are multiphoto orthomosaics.

Colony (year)	Spatial resolution	Image filename	Auto count	Adj. count	Manual count	Error (%)
Banks Island- Egg River	4 cm	CanadaUTM10_004GSD_06222009_112 CanadaUTM10_004GSD_06222009_52	281 619	264 576	290 561	-9.0 +2.7
(2009)		CanadaUTM10_004GSD_06222009_92	919	770	681	+13.1
		CanadaUTM10_004GSD_06222009_132	1010	991	961	+3.1
		CanadaUTM10_004GSD_06222009_111	1615	1454	1387	+4.8
		CanadaUTM10_004GSD_06222009_131 [†]	1919	1432	1350	+6.1
		CanadaUTM10_004GSD_06222009_91	4641	3593	3826	-6.1
		Colony total	11,004	9080	9056	+0.3
Cape Dominion	5 cm	or17668560	0	0	0	0.0
+ Bowman Bay (2011)		or17668704	3	3	2	+50.0
		or17668680	7	7	7	0.0
		or17667858	10	10	9	+11.1
		or17663301	12	12	14	-14.3
		or17668384	15	9	9	0.0
		or17669001	17	17	14	+21.4
		or17668158	22	22	22	0.0
		or17667702	25	25	23	+8.7
		or17668977	27	27	27	0.0
		or17668985	32	32	33	-3.0
		or17668288	40	40	38	+5.3
		or17668264	41	41	43	-4.7
		or17668170	43	31	29	+6.9
		or17668360	48	35	18	+94.4
		or17662957	66	63	63	0.0
		FWS_SGS44_005GSD_06272011_M06	67	67	83	-19.3
		or17668376	70	27	19	+42.1
		or17668240	72	72	74	-2.7
		or17668200	91	91	88	+3.4
		or17668232	104	102	96	+6.3
		or17668248	107	107	107	0.0
		or17668208	189	189	194	-2.6
		or17668224	347	347	357	-2.8
		Colony total	1455	1376	1369	+0.5
Cory Bay (2011)	5 cm	FWS_SGS44_005GSD_06192011_06	192	192	153	+25.5
		FWS_SGS44_005GSD_06192011_05	293	265	238	+11.3
		FWS_SGS44_005GSD_06192011_03	608	331	193	+71.5
		Colony total	1093	788	584	+34.9
Taverner Bay (2011)	5 cm	FWS_SGS44_005GSD_06262011_12	51	51	38	+34.2
		$FWS_SGS44_005GSD_06262011_02^{+}$	88	88	79	+11.4
		FWS_SGS44_005GSD_06262011_07	104	104	101	+3.0
		FWS_SGS44_005GSD_06262011_01	127	127	119	+6.7
		Colony total	370	370	337	+9.8
Banks Island-	5 cm	$FWS_SG46_005GSD_06222023_BANKS31^{\dagger}$	408	339	328	+3.4
Egg River (2013)		FWS_SG46_005GSD_06222023_BANKS25	2692	2684	2843	-5.6
		FWS_SG46_005GSD_06222023_BANKS35	5285	5283	5319	-0.7
		Colony total	8385	8306	8490	-2.2

[†]Denotes images used to initially train and test automated analysis parameters.

multiple small patches misclassified as geese scattered around the periphery, or numerous speckles of glint over the surface of a single water basin. These clusters of errors were usually evident at coarse zoom levels and could be rapidly deleted from the classification polygon file (Fig. 10). For each test image, we consequently calculated an "adjusted count" equal to the automated count minus manually removed clusters of commission errors, and compared manual counts to these adjusted counts in the final evaluation (Table 1).

After reviewing the test images to remove clusters of commission errors (totalling 2307, or 11% of all objects classified as geese) there remained 19,920 geese identified by the automated analysis, **Fig. 10**. Coarse-scale overview of a mosaic of aerial images captured along two parallel flight lines over the Banks Island-Egg River Lesser Snow Goose (*Chen caerulescens caerulescens*) breeding colony in 2009, with the automated classification layer overlaid, showing all features classified as geese in red. A dense cluster of commission errors caused by speckles of glint on the surface of water is evident near the top right-hand corner, as well as errors caused by small patches of snow around the periphery of the large patches in the bottom part of the image. Deleting these errors from the classification layer took 2 min, compared to 107 min to manually count all geese in the image.



that is, 84 more than the manual count, for an overall error rate of +0.4%. The overall error rate was +0.3% (automated = 9080, manual = 9056) in the 4-cm imagery and +0.6% (automated = 10,840, manual = 10,780) in the 5-cm imagery, and the overall error rate within individual colonies was between $\pm 10\%$ for all but one colony (Table 1). The total automated count across the eight training images had a higher overestimation rate (+12%) than that of the overall test set, which was an expected consequence of having selected them on the basis of containing numerous potential confounding features that could cause commission errors. Overall, there was a strong linear relationship between adjusted automated counts and manual counts across the test set (Fig. 11; $R^2 = 0.998$, P < 0.001, $\beta \pm SE = 0.974 \pm 0.007$).

The time required to review automated analysis results and remove clusters of commission errors in all 41 test images was 35.5 min (mean = 0.9 min/image), ranging from ~10 s for a single orthophoto containing no conspicuous clusters of errors (mean = 0.3 min/image across all single photos) to 4 min for a large orthomosaic containing multiple clusters of errors (mean = 1.6 min/image across all mosaics). By comparison, it took 690 min (~20 times longer) to manually count the geese in all test images (mean = 16.8 min/image), with count time varying considerably among images, from ~1 min for a single orthophoto containing no geese (mean = 3.6 min/image across all single photos) to 116 min for the most populous orthomosaic containing 5319 geese (mean = 33.8 min/image across all mosaics). Manual count rate in terms of birds per minute generally increased with density of geese, and decreased in more challenging landscapes, notably those cluttered with bright rocks.

DISCUSSION

We demonstrated how an automated analysis routine developed using commercial off-the-shelf OBIA software can be used to detect and count Snow Geese in large volumes of aerial imagery much more efficiently, in terms of person-time, than through manual processing. Based on a systematic analysis of manually identified geese in a small subsample of images, we established standard image segmentation and classification parameters that could be applied through an automated batch-processing script



Fig. 11. Scatter plot of semiautomated counts vs. manual counts of Lesser Snow Geese (*Chen caerulescens caerulescens*) in 41 aerial images of breeding colonies in the Canadian Arctic.

to produce accurate counts of geese in images taken in a range of habitats and lighting conditions (Appendices 2–9). This represents one of the few examples to date of successful use of automated aerial image analysis to detect animals over such a large and varied geographic extent (Hollings et al. 2018).

Owing to large numbers of commission errors in some images, it was necessary to perform a quick manual review and correction of results, so our solution may be best described as semiautomated or supervised. Although most of the errors were conspicuously clustered and could therefore be rapidly identified and eliminated at coarse zoom levels, bright goose-sized rocks were more problematic because they tended to be scattered throughout images and could not be efficiently removed. This resulted in greater overestimates of bird numbers notably in images from the Cory Bay colony (Table 1), which has a particularly rocky landscape. Images from other colonies that produced significant overestimates similarly contained scattered rather than clustered commission errors. An alternative approach to reviewing results in situations like this might be to successively zoom to and verify each object classified as a goose using the classification layer's attribute table. Depending on the commission error rate, this is likely to still be faster than manually counting geese, especially because geese are more challenging and time-consuming to identify visually in landscapes cluttered with confounding small bright objects.

There was also a variety of sources of omission errors, leading to net underestimates of bird numbers in certain images (Table 1) but no overall bias. Some geese were omitted in the segmentation process, typically particularly dark and poorly contrasting birds, oftentimes blue geese, and occasionally birds with shadows cast over them. Further omissions resulted from the object classification process. Flying geese were often omitted because their size and shape tended to fall outside the ranges defined in the rule sets. Oversegmented geese were omitted if neither of their constituent objects met the spatial attribute criteria (but were included if at least one object met all criteria; if both objects met the criteria, they were automatically merged into a single feature). In other cases, automatic merging of adjacent goose objects led to two side-by-side birds being merged into one. Occasionally, very close-together geese were segmented as a single object and subsequently omitted in the classification process on account of being too large. Finally, geese with extreme values for various object attributes were susceptible to omission, including extreme shapes and sizes, birds in very dark/low-exposure or very bright/ high-exposure portions of imagery, and blue geese that were insufficiently bright or contrasting.

Our approach was achievable using 4- and 5-cm resolution imagery, which yielded Snow Goose objects of ~20-60 pixels in size. We found that 3-cm resolution imagery was too detailed to achieve consistent segmentation of geese using the parameters we tested. Nevertheless, more detailed imagery could be valuable for manual confirmation of species identification, especially in situations where multiple species may be present. Excessively high-resolution imagery could potentially be resampled or subjected to more aggressive smoothing filters to make it more amenable to automated processing. Imagery with a 6-cm resolution, on the other hand, was too coarse to reliably distinguish geese from other bright objects, notably rocks, when applying our standard classification rule set to multiple images from different colonies. This result was not surprising because even manual identification of geese at this resolution is noticeably more challenging than at finer resolutions, at which their distinctive bright white "cores" become much more visually striking. Among all the imagery of Snow Goose colonies collected from 2009 to 2014, only 17% (by file size) was at a resolution of 5 cm or finer, with 40% at 6 cm and 43% at coarser resolutions (much of it 10 cm or larger). In landscapes with no confounding features such as boulders or patches of snow, it is possible to manually count white objects in these coarser images and assume that most of them are geese, although blue geese generally cannot be distinguished. Although it may be possible to achieve effective automated detection of geese in coarser imagery, classification parameters would likely need to be adjusted for each colony depending on background and lighting conditions, which is less efficient than batch-processing all imagery with a standard set of parameters. Collecting finer resolution imagery by flying closer to the ground results in less area coverage for the same flying time; for the same camera and lens, the strip width of 5-cm imagery is 17% narrower compared to 6 cm, and 50% narrower compared to 10 cm. However, the ability to perform wholesale automated processing of images makes it worthwhile to potentially sacrifice area coverage in favor of higher spatial resolution. Because traditional manual counts are only performed in a subset of images, much of the imaged area goes unanalyzed to begin with, so a reduction in total area coverage would not necessarily decrease the statistical confidence of extrapolated colony population estimates if it becomes possible to analyze all collected images.

Overall, our results suggest the potential for considerable time savings when analyzing large volumes of data compared to manual counting, which requires thorough systematic viewing of images at high zoom levels, even those containing few or no geese. In our sample of test images, the time required for manual review of automated results was only ~5% of that required to manually count all geese. The actual time savings in a broader survey would depend on the proportion of images with different conditions and densities of birds. Even if the ratio were closer to 10%, use of the semiautomated procedure could reduce the typical person-time required to analyze a comprehensive aerial photographic survey of Arctic Snow Goose colonies, including thousands of individual photos, from several months to several days. Other such semiautomated solutions to detect birds or their nests in aerial imagery have similarly proven to be less time-consuming than fully manual analysis (McNeill et al. 2011, Andrew and Shephard 2017). Analysis of time savings also needs to consider initial development time: the time required to review available imagery, select training images, manually identify subjects, establish segmentation parameters, compile and analyze subject attribute data, and create classification rule sets. This process could take from a few days to several weeks depending on the dataset. There is also significant computer processing time required, although when executed in a batch mode, it can run in the background without consuming person-time. Ultimately, a method like ours will be most beneficial where large volumes of imagery and/or repeated monitoring of subjects over time are involved, such that the savings exceed the set-up time.

Our approach to establishing automated analysis parameters with OBIA software could be used for other surveys to detect birds or other wildlife species. With OBIA becoming increasingly accessible through a variety of off-the-shelf packages including industry-standard geospatial analysis software suites, the approach could potentially be used by anyone possessing at least basic skills in image analysis and GIS, and executed on midrange desktop computers that now possess sufficiently powerful hardware to execute OBIA operations at reasonable speed. The technique will work best with species that contrast markedly from the background in at least some spectral bands, whether brighter or darker. More weakly contrasting subjects could also be detected, but would require lower segmentation levels, resulting in greater numbers of image objects that must subsequently be filtered out at the classification stage. Subjects with relatively consistent size and shape (which is usually the case for birds on the ground or water surface) are also easier to classify, whereas birds in flight or other species with more variable sizes and shapes/ postures may be more challenging. A key advantage of our approach is that it enables identification of all object attributes for which target subjects present a well-defined distribution, thus increasing the potential to achieve good detection accuracy. In some cases, it may be possible to reduce commission error rates to the point where manual review of automated analysis results is not necessary, while in other cases, like ours, varying degrees of manual postanalysis effort may be required. Although animal contrast in thermal-infrared imagery has proven useful for automated detection of mammals (Conn et al. 2014, Chrétien et al. 2015, 2016, Seymour et al. 2017), the very coarse pixel resolution of thermal cameras compared to RGB cameras generally renders them ineffective for aerial detection of comparatively smaller birds (Chabot and Francis 2016). It should be noted that any aerial image-based survey will only allow detection of subjects that are visible from overhead and miss subjects that are, for example, concealed under canopy or diving underwater.

Further refinements to our approach could enable it to be extended to more challenging situations and/or increase its accuracy. It may be useful to try more advanced functions for building classification rule sets, such as assigning varying weights to different attributes, allowing tolerances for anomalous values for certain attributes, or creating multiple rules for a subject class that operate on an either/or basis. Multiple subject classes could be created for multispecies or polymorphic species detection, provided subjects of interest can be adequately captured using a common set of segmentation parameters. Alternatively, an image set could be batch-processed more than once with varying segmentation parameters to capture different subjects. Even if species cannot be reliably differentiated through automated classification, it may still be valuable to detect them in terms of a single or a few general classes, then manually review results one subject at a time to identify them to the species level. Groom et al. (2013) reported that this type of semiautomated solution was more efficient than fully manual analysis for surveying multiple species of birds at sea that were sparsely distributed across large numbers of images. Other OBIA software may also be better suited to overcome some of the challenges we encountered by offering different segmentation options, object attributes, and/or workflow possibilities. Options include other commercial packages such as eCognition (Trimble), ERDAS IMAGINE (Hexagon Geospatial, Madison, AL, USA), and Geomatica (PCI, Markham, ON, Canada); freeware packages such as Orfeo ToolBox (CNES, Paris, France) and InterIMAGE (Laboratório de Visão Computacional, Rio de Janeiro, Brazil); or third-party plugins for GIS software such as Feature Analyst (Textron Systems, Providence, RI, USA). Ultimately, in an operational context, an analyst must consider the time required to refine automated detection performance in relation to time savings in postanalysis review of images.

Finally, a potential variation on our approach would be to substitute the development of a classification rule set via manual examination of subject attributes with the use of a machinelearning classification algorithm (Ma et al. 2017). In this case, the manually identified target objects, i.e., birds, from the segmented training images would serve as the input training data for the algorithm. Selection and parameterization of attributes that distinguish target objects would then be determined by the algorithm. Although some wildlife detection studies have previously employed machine learning (Torney et al. 2016, Andrew and Shephard 2017, Longmore et al. 2017), a concern we had in our case was the presumed challenge of a machinelearning classifier to anticipate subject brightness variation beyond that of the training set, which we considered crucial to coping with the highly variable illumination and exposure levels throughout our image set. Nevertheless, it should be noted that advances in machine learning are proceeding at a rapid pace, with particular promise shown by flourishing "deep learning" algorithms (Ball et al. 2017), including some early work on wildlife detection in remotely sensed imagery (Maire et al. 2015, Maussang et al. 2015, Marburg and Bigham 2016).

Responses to this article can be read online at: http://www.ace-eco.org/issues/responses.php/1205

Acknowledgments:

This research was funded by the Canadian Wildlife Service, Environment and Climate Change Canada. We thank K. Dufour for helping to arrange access to the imagery, and M. Mitchell for providing details on imagery acquisition and postprocessing. The work presented in this paper was entirely conducted using preexisting aerial imagery of birds that was collected by federal wildlife agencies in accordance with appropriate licenses and protocols.

LITERATURE CITED

Andrew, M. E., and J. M. Shephard. 2017. Semi-automated detection of eagle nests: an application of very high-resolution image data and advanced image analyses to wildlife surveys. *Remote Sensing in Ecology and Conservation* 3:66-80. http://dx. doi.org/10.1002/rse2.38

Bajzak, D., and J. F. Piatt. 1990. Computer-aided procedure for counting waterfowl on aerial photographs. *Wildlife Society Bulletin* 18:125-129.

Ball, J. E., D. T. Anderson, and C. S. Chan. 2017. Comprehensive survey of deep learning in remote sensing: theories, tools, and challenges for the community. *Journal of Applied Remote Sensing* 11:042609. http://dx.doi.org/10.1117/1.JRS.11.042609

Batt, B. D. J. 1997. Arctic ecosystems in peril: report of the Arctic Habitat Working Group. Arctic Joint Venture Special Publication, U.S. Fish and Wildlife Service, Washington, D.C., USA, and Canadian Wildlife Service, Ottawa, Ontario, Canada.

Béchet, A., A. Reed, N. Plante, J.-F. Giroux, and G. Gauthier. 2004. Estimating the size of the Greater Snow Goose population.

Journal of Wildlife Management 68:639-649. http://dx.doi. org/10.2193/0022-541X(2004)068[0639:ETSOTG]2.0.CO;2

Blaschke, T. 2010. Object based image analysis for remote sensing. *ISPRS Journal of Photogrammetry and Remote Sensing* 65:2-16. http://dx.doi.org/10.1016/j.isprsjprs.2009.06.004

Boyd, W. S. 2000. A comparison of photo counts versus visual estimates for determining the size of Snow Goose flocks. *Journal of Field Ornithology* 71:686-690. http://dx.doi.org/10.1648/0273--8570-71.4.686

Buckland, S. T., M. L. Burt, E. A. Rexstad, M. Mellor, A. E. Williams, and R. Woodward. 2012. Aerial surveys of seabirds: the advent of digital methods. *Journal of Applied Ecology* 49:960-967. http://dx.doi.org/10.1111/j.1365-2664.2012.02150.x

Butler, M. J., W. B. Ballard, M. C. Wallace, S. J. Demaso, and B. K. McGee. 2007. Aerial surveys for estimating Wild Turkey abundance in the Texas Rolling Plains. *Journal of Wildlife Management* 71:1639-1645. http://dx.doi.org/10.2193/2006-254

Chabot, D., and D. M. Bird. 2012. Evaluation of an off-the-shelf unmanned aircraft system for surveying flocks of geese. *Waterbirds* 35:170-174. http://dx.doi.org/10.1675/063.035.0119

Chabot, D., and D. M. Bird. 2015. Wildlife research and management methods in the 21st century: where do unmanned aircraft fit in? *Journal of Unmanned Vehicle Systems* 3:137-155. http://dx.doi.org/10.1139/juvs-2015-0021

Chabot, D., and C. M. Francis. 2016. Computer-automated bird detection and counts in high-resolution aerial images: a review. *Journal of Field Ornithology* 87:343-359. http://dx.doi. org/10.1111/jofo.12171

Chrétien, L.-P., J. Théau, and P. Ménard. 2015. Wildlife multispecies remote sensing using visible and thermal infrared imagery acquired from an unmanned aerial vehicle (UAV). *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XL-1/W4:241-248. http://dx. doi.org/10.5194/isprsarchives-XL-1-W4-241-2015

Chrétien, L.-P., J. Théau, and P. Ménard. 2016. Visible and thermal infrared remote sensing for the detection of white-tailed deer using an unmanned aerial system. *Wildlife Society Bulletin* 40:181-191. http://dx.doi.org/10.1002/wsb.629

Conn, P. B., J. M. Ver Hoef, B. T. McClintock, E. E. Moreland, J. M. London, M. F. Cameron, S. P. Dahle, and P. L. Boveng. 2014. Estimating multispecies abundance using automated detection systems: ice-associated seals in the Bering Sea. *Methods in Ecology and Evolution* 5:1280-1293. http://dx.doi.org/10.1111/2041-210X.12127

Descamps, S., A. Béchet, X. Descombes, A. Arnaud, and J. Zerubia. 2011. An automatic counter for aerial images of aggregations of large birds. *Bird Study* 58:302-308. http://dx.doi. org/10.1080/00063657.2011.588195

Frederick, P. C., B. Hylton, J. A. Heath, and M. Ruane. 2003. Accuracy and variation in estimates of large numbers of birds by individual observers using an aerial survey simulator. *Journal of Field Ornithology* 74:281-287. http://dx.doi.org/10.1648/0273-85-70-74.3.281

Gilmer, D. S., J. A. Brass, L. L. Strong, and D. H. Card. 1988. Goose counts from aerial photographs using an optical digitizer. *Wildlife Society Bulletin* 16:204-206.

Good, R. E., R. M. Nielson, H. Sawyer, and L. L. McDonald. 2007. A population estimate for Golden Eagles in the western United States. *Journal of Wildlife Management* 71:395-402. http://dx.doi.org/10.2193/2005-593

Groom, G., I. K. Petersen, M. D. Anderson, and A. D. Fox. 2011. Using object-based analysis of image data to count birds: mapping of Lesser Flamingos at Kamfers Dam, Northern Cape, South Africa. *International Journal of Remote Sensing* 32:4611-4639. http://dx.doi.org/10.1080/01431161.2010.489068

Groom, G., I. K. Petersen, and A. D. Fox. 2007. Sea bird distribution data with object based mapping of high spatial resolution image data. Pages 252-257 *in Proceedings of the Remote Sensing and Photogrammetry Society Annual Conference*, 11-14 September 2007, Nottingham, UK.

Groom, G., M. Stjernholm, R. D. Nielsen, A. Fleetwood, and I. K. Petersen. 2013. Remote sensing image data and automated analysis to describe marine bird distributions and abundances. *Ecological Informatics* 14:2-8. http://dx.doi.org/10.1016/j. ecoinf.2012.12.001

Hollings, T., M. Burgman, M. van Andel, M. Gilbert, T. Robinson, and A. Robinson. 2018. How do you find the green sheep? A critical review of the use of remotely sensed imagery to detect and count animals. *Methods in Ecology and Evolution* 9:881-892. http://dx.doi.org/10.1111/2041-210X.12973

Kerbes, R. H., K. M. Meeres, and R. T. Alisauskas. 2014. *Surveys of nesting Lesser Snow Geese and Ross's Geese in Arctic Canada, 2002-2009.* Arctic Goose Joint Venture Special Publication, U.S. Fish and Wildlife Service, Washington, D.C., USA, and Canadian Wildlife Service, Ottawa, Ontario, Canada.

Kerbes, R. H., K. M. Meeres, R. T. Alisauskas, F. D. Caswell, K. F. Abraham, and R. K. Ross. 2006. *Surveys of nesting mid-continent Lesser Snow Geese and Ross's Geese in eastern and central Arctic Canada, 1997-1998.* Technical Report Series no. 447, Canadian Wildlife Service, Saskatoon, Saskatchewan, Canada.

Kingsford, R. T., and J. L. Porter. 2009. Monitoring waterbird populations with aerial surveys-what have we learnt? *Wildlife Research* 36:29-40. http://dx.doi.org/10.1071/WR08034

Laliberte, A. S., and W. J. Ripple. 2003. Automated wildlife counts from remotely sensed imagery. *Wildlife Society Bulletin* 31:362-371.

LaRue, M. A., S. Stapleton, and M. Anderson. 2017. Feasibility of using high-resolution satellite imagery to assess vertebrate wildlife populations. *Conservation Biology* 31:213-220. http://dx. doi.org/10.1111/cobi.12809

Leafloor, J. O., T. J. Moser, and B. D. J. Batt. 2012. *Evaluation of* specific management measures for midcontinent Lesser Snow Geese and Ross's Geese. Arctic Goose Joint Venture Special Publication, U.S. Fish and Wildlife Service, Washington, D.C., USA, and Canadian Wildlife Service, Ottawa, Ontario, Canada. Lefebvre, J., G. Gauthier, J.-F. Giroux, A. Reed, E. T. Reed, and L. Bélanger. 2017. The Greater Snow Goose *Anser caerulescens atlanticus*: managing an overabundant population. *Ambio* 46 (Suppl. 2):262-274. http://dx.doi.org/10.1007/s13280-016-0887-1

Liu, C.-C., Y.-H. Chen, and H.-L. Wen. 2015. Supporting the annual international Black-faced Spoonbill census with a low-cost unmanned aerial vehicle. *Ecological Informatics* 30:170-178. http://dx.doi.org/10.1016/j.ecoinf.2015.10.008

Longmore, S. N., R. P. Collins, S. Pfeifer, S. E. Fox, M. Mulero-Pázmány, F. Bezombes, A. Goodwin, M. De Juan Ovelar, J. H. Knapen, and S. A. Wich. 2017. Adapting astronomical source detection software to help detect animals in thermal images obtained by unmanned aerial systems. *International Journal of Remote Sensing* 38:2623-2638. http://dx.doi.org/10.1080/014311-61.2017.1280639

Ma, L., M. Li, X. Ma, L. Cheng, P. Du, and Y. Liu. 2017. A review of supervised object-based land-cover image classification. *ISPRS Journal of Photogrammetry and Remote Sensing* 130:277-293. http://dx.doi.org/10.1016/j.isprsjprs.2017.06.001

Maire, F., L. M. Alvarez, and A. Hodgson. 2015. Automating marine mammal detection in aerial images captured during wildlife surveys: a deep learning approach. Pages 379-385 *in* B. Pfahringer and J. Renz, editors. *AI 2015: advances in artificial intelligence*. Springer International, Cham, Switzerland. http://dx.doi.org/10.1007/978-3-319-26350-2_33

Marburg, A., and K. Bigham. 2016. Deep learning for benthic fauna identification. *In* Proceedings of the MTS/IEEE Oceans Conference, 19-23 September 2016, Monterey, California, USA. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, USA. http://dx.doi.org/10.1109/OCEANS.2016.7761146

Maussang, F., L. Guelton, R. Garello, and A. Chevallier. 2015. Marine life observations using classification algorithms on ocean surface photographs. *In* Proceedings of the MTS/IEEE Oceans Conference, 18-21 May 2015, Genova, Italy. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, USA. http:// dx.doi.org/10.1109/OCEANS-Genova.2015.7271678

McNeill, S., K. Barton, P. Lyver, and D. Pairman. 2011. Semiautomated penguin counting from digital aerial photographs. *In* Proceedings of the IEEE International Geoscience and Remote Sensing Symposium, 24-29 July 2011, Vancouver, British Columbia, Canada. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, USA. http://dx.doi. org/10.1109/IGARSS.2011.6050185

North American Bird Conservation Initiative (NABCI). 2016. *The state of North America's Birds 2016*. NABCI. [online] URL: <u>http://www.stateofthebirds.org/2016/</u>

Redding, N. J., D. J. Crisp, D. Tang, and G. N. Newsam. 1999. An efficient algorithm for mumford-shah segmentation and its application to SAR imagery. *In* Proceedings of the Conference on Digital Image Computing: Techniques and Applications, 7-8 December 1999, Perth, Western Australia, Australia.

Roerdink, J. B. T. M., and A. Meijster. 2001. The watershed transform: definitions, algorithms and parallelization strategies. *Fundamenta Informaticae* 41:187-222.

Rosenberg, K. V., P. J. Blancher, J. C. Stanton, and A. O. Panjabi. 2017. Use of North American Breeding Bird Survey data in avian conservation assessments. *Condor* 119:594-606. http://dx.doi. org/10.1650/CONDOR-17-57.1

Seymour, A. C., J. Dale, M. Hammill, P. N. Halpin, and D. W. Johnston. 2017. Automated detection and enumeration of marine wildlife using unmanned aircraft systems (UAS) and thermal imagery. *Scientific Reports* 7:45127. http://dx.doi.org/10.1038/ srep45127

Torney, C. J., A. P. Dobson, F. Borner, D. J. Lloyd-Jones, D. Moyer, H. T. Maliti, M. Mwita, H. Fredrick, M. Borner, and J. G. C. Hopcraft. 2016. Assessing rotation-invariant feature classification for automated wildebeest population counts. *PLoS ONE* 11: e0156342. http://dx.doi.org/10.1371/journal.pone.0156342

Trathan, P. N. 2004. Image analysis of color aerial photography to estimate penguin population size. *Wildlife Society Bulletin* 32:332-343. http://dx.doi.org/10.2193/0091-7648(2004)32[332:IAOCAP] 2.0.CO;2

Warner, T. 2011. Kernel-based texture in remote sensing image classification. *Geography Compass* 5:781-798. http://dx.doi. org/10.1111/j.1749-8198.2011.00451.x

Woodworth, B. L., B. P. Farm, C. Mufungo, M. Borner, and J. O. Kuwai. 1997. A photographic census of flamingos in the Rift Valley lakes of Tanzania. *African Journal of Ecology* 35:326-334. http://dx.doi.org/10.1111/j.1365-2028.1997.098-89098.x



Sponsored by the Society of Canadian Ornithologists and Bird Studies Canada Parrainée par la Société des ornithologistes du Canada et Études d'oiseaux Canada



BIRD STUDIES CANADA

Appendix 1. Attributes Used to Classify Lesser Snow Geese (Chen caerulescens caerulescens).

Table A1.1. List of spatial, spectral and texture attributes used to classify Snow Geese in aerial imagery of Canadian Arctic breeding colonies. Each spectral and texture attribute is calculated for each of the image's spectral bands (in this case, red, green and blue). Attribute descriptions are from ENVI 5.3 (Exelis Visual Information Solutions, Boulder, CO, USA) help files.

Attribute	Description
Spatial attributes	
Area	Total area of the object, minus the area of the holes.
Length	The combined length of all boundaries of the object, including the boundaries of the holes.
Compactness	A shape measure that indicates the compactness of the object. A circle is the most
	compact shape with a value of 1 / pi.
	Compactness = Sqrt (4 * Area / pi) / outer contour length
Convexity	A measure of object convexity. The convexity value for a convex object with no holes is
	1.0, while the value for a concave object is less than 1.0.
Calidita.	Convexity = length of convex hull / Length Δ above the area of a convex hull
Solidity	A shape measure that compares the area of the object to the area of a convex hull
	the value for a concave object is less than 1.0
	Solidity = $Area / area of convex hull$
Roundness	A shape measure that compares the area of the object to the square of the maximum
recurrencess	diameter of the object. The "maximum diameter" is the length of the major axis of an
	oriented bounding box enclosing the object. The roundness value for a circle is 1, and
	the value for a square is 4 / pi.
	Roundness = $4 * (Area) / (pi * Major_Length^2)$
Form Factor	A shape measure that compares the area of the object to the square of the total perimeter.
	The form factor value of a circle is 1, and the value of a square is $pi / 4$.
	Form_Factor = $4 * pi * (Area) / (total perimeter)^2$
Elongation	A shape measure that indicates the ratio of the major axis of the object to the minor axis of
	the object. The major and minor axes are derived from an oriented bounding box
	containing the object. The elongation value for a square is 1.0, and the value for a
	rectangle is greater than 1.0.
Pactongular Fit	Elongation – Majoi_Lengtii / Minoi_Lengtii A shape measure that indicates how well the shape is described by a rectangle. This
Rectangular Fit	A shape measure that mulcates now wen the shape is described by a rectangle. This attribute compares the area of the object to the area of the oriented bounding box
	enclosing the object. The rectangular fit value for a rectangle is 1.0 and the value for a
	non-rectangular shape is less than 1.0
	Rectangular Fit = Area / (Major Length * Minor Length)
Major Length	The length of the major axis of an oriented bounding box enclosing the object.
Minor Length	The length of the minor axis of an oriented bounding box enclosing the object.
Spectral attributes	
Spectral Mean	Mean value of the pixels comprising the object in band x.
Spectral Max	Maximum value of the pixels comprising the object in band x.
Spectral Min	Minimum value of the pixels comprising the object in band x.
Spectral STD	Standard deviation value of the pixels comprising the object in band x.
Texture attributes	
Texture Range	Average data range of the pixels comprising the region inside the texture kernel in band x.
Texture Mean	Average value of the pixels comprising the region inside the texture kernel in band x.
Texture Variance	Average variance of the pixels comprising the region inside the texture kernel in band x.
Texture Entropy	Average entropy value of the pixels comprising the region inside the kernel in band x.

Appendix 2. Sample orthomosaic from an aerial survey of the Egg River Lesser Snow Goose (*Chen caerulescens caerulescens*) breeding colony on Banks Island in the Canadian Arctic, 22 June 2009. Note that the spatial resolution of the image has been resampled from 4 cm to 20 cm to reduce file size.



Appendix 3. Sample orthophoto from an aerial survey of the Cape Dominion and Bowman Bay Lesser Snow Goose (*Chen caerulescens*) breeding colonies on Baffin Island in the Canadian Arctic, 25 June 2011. Note that the spatial resolution of the image has been resampled from 5 cm to 10 cm to reduce file size.



Appendix 4. Sample orthophoto from an aerial survey of the Cape Dominion and Bowman Bay Lesser Snow Goose (*Chen caerulescens*) breeding colonies on Baffin Island in the Canadian Arctic, 25 June 2011. Note that the spatial resolution of the image has been resampled from 5 cm to 10 cm to reduce file size.



Appendix 5. Sample orthophoto from an aerial survey of the Cape Dominion and Bowman Bay Lesser Snow Goose (*Chen caerulescens*) breeding colonies on Baffin Island in the Canadian Arctic, 25 June 2011. Note that the spatial resolution of the image has been resampled from 5 cm to 10 cm to reduce file size.



Appendix 6. Sample orthophoto from an aerial survey of the Cape Dominion and Bowman Bay Lesser Snow Goose (*Chen caerulescens*) breeding colonies on Baffin Island in the Canadian Arctic, 25 June 2011. Note that the spatial resolution of the image has been resampled from 5 cm to 10 cm to reduce file size.



Appendix 7. Sample orthomosaic from an aerial survey of the Cory Bay Lesser Snow Goose (*Chen caerulescens caerulescens*) breeding colony on Baffin Island in the Canadian Arctic, 19 June 2011. Note that the spatial resolution of the image has been resampled from 5 cm to 25 cm to reduce file size.



Appendix 8. Sample orthomosaic from an aerial survey of the Taverner Bay Lesser Snow Goose (*Chen caerulescens caerulescens*) breeding colony on Baffin Island in the Canadian Arctic, 26 June 2011. Note that the spatial resolution of the image has been resampled from 5 cm to 15 cm to reduce file size.



Appendix 9. Sample orthomosaic from an aerial survey of the Egg River Lesser Snow Goose (*Chen caerulescens caerulescens*) breeding colony on Banks Island in the Canadian Arctic, 22 June 2013. Note that the spatial resolution of the image has been resampled from 5 cm to 25 cm to reduce file size.



Appendix 10. Code for Batch-processing Images.

Fig. A10.1. Example IDL (Exelis Visual Information Solutions, Boulder, CO, USA) code to batch-process multiple aerial image files for automated detection of Lesser Snow Geese (Chen caerulescens caerulescens). The script instructs ENVI (Exelis) to sequentially open all TIFFformat raster image files in a specified directory, including in subdirectories. For each file, ENVI is instructed to ignore pixels with a data value of 0, and query the file's metadata to retrieve the image's filename, spatial dimensions and list of spectral bands. A low-pass filter is then applied to the image, and the resulting smoothed image is saved to a specified output directory with the identifier "smoothraster" appended to the end of the original filename. The smoothed image is then passed along to the 'Rule-Based Feature Extraction' procedure, which is instructed to process the image's full spatial extent and all spectral bands. The established image segmentation parameters are then applied (note that IDL identifies spectral bands 1, 2 and 3 as 0, 1 and 2, respectively) and classification of the resulting image objects is executed in accordance with the specified rule set file. Finally, a segmentation raster file as well as a polygon vector file (with the identifier "classifiedsnowgeese" appended to the end of the original filename) delineating all classified geese are saved to the output directory. Since the segmentation raster is four times the file size of the original image and its only purpose is to serve as the input for classification following segmentation, to save disk space it was not assigned a unique filename for each processed image file. This way, the segmentation raster generated by each successive image simply overwrites the file generated by the previous image.

```
PRO batch_process_5cm_originals
COMPILE OPT IDL2
e = envi (/HEADLESS)
files = File_Search('F:\Snow_goose_imagery\5cm_GSD\', '*.tif')
FOR i = 0, N_Elements(files)-1 DO BEGIN
 raster = e.OpenRaster(files[i], DATA IGNORE VALUE = 0)
 fid = ENVIRasterToFID(raster)
 envi file query, fid, sname = sname, dims = dims, nb = nb
  task = ENVITask('LowPassFilter')
 Task.Input Raster = raster
 Task.Output_Raster_URI = 'F:\Snow_goose_imagery\5cm_GSD\Classification_output\'+sname+'_smoothraster.dat'
 Task.Execute
  smooth_raster = Task.Output_Raster
  s fid = ENVIRasterToFID(smooth raster)
  envi_doit, 'envi_fx_rulebased_doit', $
    fid = s_fid, dims = dims, pos = Lindgen(nb), $
   segment bands = [1, 2], scale level = 90.0, $
   merge_algorithm = 'fast lambda', merge_bands = [1], merge_level = 90.0, $
    kernel size = 5, $
    rule_filename = 'F:\Snow_goose_imagery\5cm_GSD\5cm_ruleset.rul', $
    segmentation raster filename = 'F:\Snow goose imagery\Scm GSD\Classification output\segmentationraster.dat', $
    vector_filename = 'F:\Snow_goose_imagery\5cm_GSD\Classification_output\'+sname+'_classifiedsnowgeese.shp'
ENDFOR
```

END