



**Alberta Biodiversity
Monitoring Institute**

www.abmi.ca

Animal Density from Camera Data

Dave Huggard

Document Version: 2018-03-20

1. Overview: Simple explanation and assumptions

Density is the number of objects (trees, animals, etc.) per unit area. If a 100 m² plot contains one tree, the density is 1 tree/100 m², or 10,000 trees per km² (**Figure 1**). Similarly, if a camera has a field-of-view of 100 m² and there is always one animal in the field-of-view for the whole time that the camera is operating, the density of that species is 1 animal/100 m², or 10,000 animals per km². It doesn't matter if the animal is moving around within the field-of-view, as long as it stays in the field-of-view for the whole time. On the other hand, if that camera only has an animal in the field-of-view 1/10,000 of the time that it is operating, there is 1/10,000 animal/100 m², or 1 animal per km². If the camera has two animals together for 1/10,000 of the time, this gives 2/10,000 animals/100 m², or 2 animals per km². This is how we use cameras to calculate density.

$$\text{Density} = \Sigma (\text{number of individuals} * \text{time in field-of-view}) / (\text{area of field-of-view} * \text{total camera operating time})$$

The units are animal-seconds per area-seconds = animals per area = density.




Figure 1. Illustration of density for a tree quadrat and camera surveys.

For a given density of animals, this simple measure is independent of home range sizes or movement rates. If home ranges were twice as big, they would have to overlap twice as much to maintain the same density. Therefore, an individual would be in a particular camera's field-of-view half as often (because its home range is bigger – it has more other places to be), but there would be twice as many individuals at that camera. If movement rates were twice as fast, an individual would pass by the camera twice as often, but would spend half as much time in the field-of-view (because it is moving faster). For the simple example above, there would be two visits to the camera each occupying 1/20,000 of the time the camera is operating, rather than one visit for 1/10,000 of the time. The other way of putting this is that only the total animal-time in the field-of-view matters, whether that comes from one long visit by one individual, several short visits by one individual, or several short visits each by different individuals. In all those cases, the density is the same; it is only the home range size and overlap and/or movement rates that are changing.

Two features of cameras require us to do some additional data processing to use this simple density measure: 1) Cameras do not survey fixed areas, unlike quadrats. The probability of an animal triggering the camera decreases with distance. We therefore have to estimate an effective detection distance for the cameras, as is done for unlimited-distance point counts for birds or unlimited distance transect surveys. This effective distance can vary for different species, habitat types, and times of year.

2) Cameras take a series of images at discrete intervals, rather than providing a continuous record of how long an animal is in the field-of-view. The discrete intervals need to be converted to a continuous measure to show how long the animal was in the field-of-view, accounting for the possibility that a moderately long interval between images might be from an animal present but not moving much, and therefore not triggering the camera, versus an animal that left the field-of-view and returned. The analyses dealing with these complications are outlined below.

There are a number of strong assumptions involved in using this measure to estimate the density of a species. Two big assumptions are: 1) The cameras are a random or otherwise representative sample of the area. The density estimate applies to the field-of-view of the cameras. To make inferences about a larger region, the cameras need to be surveying a random or representative (e.g., systematic, systematic-random, random stratified) sample of the larger region. In particular, if cameras are intentionally placed in areas where species are more common, such as game trails, then the density estimate only applies to those places, not to a larger region. 2) Animals are not attracted to or repelled by the cameras (or posts used to deploy the cameras, etc.). That also means that they do not spend more or less time in front of the camera because of the presence of the camera. The effect of lures or other attractants needs to be explicitly measured and accounted for. There are additional assumptions involved in the procedures to estimate effective detection distance, including an assumption that all animals within a certain distance of the camera are detected, and in converting the discrete images into time in field-of-view. These assumptions are discussed below. Because the world is complicated, assumptions are never met perfectly. The important thing is to consider – and, ideally, design auxiliary tests to measure – whether the violations are serious enough to impact the answer to whatever question(s) the cameras are being used to answer. In many cases, absolute density estimates may not be accurate, but the results can still serve as a useable index of relative density, if assumptions are violated about equally in whatever units are being compared (habitat types, experimental treatments, years for long-term trend, etc).

A final consideration is the sampling distribution of density estimates. Because individual cameras sample tiny areas compared to the home ranges of the species they survey, the resulting sampling distribution can be horrendous – most cameras never detect the species at all (density = 0), a few cameras record the species passing by once or twice for brief periods (low densities), and a very few cameras record long durations as the animals forage or rest or play in front of the camera, or revisit a favourite spot repeatedly (very high densities). Longer duration camera deployments can help smooth out some of that extreme variation, but ultimately large numbers of cameras are required for precise estimates. Appropriate questions, rigorous study designs and modest expectations are required for camera-based studies.

The following sections give details of components of this density estimate (how the ABMI collects the necessary information, results, assumptions, tests – **Sections 2 to 5**), other factors that need to be considered in some designs – **Sections 6 to 8**), and further discussion of the basic assumptions and dealing with the sampling distribution (**Sections 9 and 10**).

2. Tagging images and checking accuracy

ABMI cameras have generated about 10 million images so far. Some of these are images taken at fixed time intervals to document the site conditions and verify that the camera was operating. The majority of the images are motion-triggered, and need to have their content identified. We use several automated processes to reduce the number of images that human taggers need to look at:

- 1) Neural network “auto-taggers” are trained to recognize images with no animals (“None” images) and images with domestic cows.
- 2) A two-stage subsampling procedure is used with series of 6 or more images separated by less than a two-minute cut-off: 5 randomly selected images are tagged by humans and used to fill in the rest of the series, unless a native mammal is detected in any of the 5 images, in which case all images in the series are tagged by people.
- 3) Additional auto-filling rules are used when a substantial number of images in a series are tagged by an auto-tagger, and for images that occur at the start or end of a deployment that are staff setting up cameras.

Cut-off scores for the auto-taggers, the subsampling approach and the other rules were all developed using data from millions of images and were set so that they are expected to miss <0.1% of images of native animals. With these procedures, humans end up tagging about 15% of images.

After all images have been tagged, we do a second species-by-species check of all identified native mammals to look for errors, including in difficult to separate species, such as marten and fisher, or white-tailed and mule deer (and we include categories for unknown deer or other groups).

Finally, we do another test of 5,000 randomly selected images from each of the different processes – auto-tagged None, auto-tagged Cows, filled-in subsampled series, human-tagged None, human-tagged cows, human-tagged native mammals – to assess final accuracy. The automated processes all miss < 0.1% of native mammals. Human taggers have a somewhat higher error rate, generally for ambiguous images. These error rates are about 100+ times better than we find for surveys of birds or plants. Details of the tagging processes and the accuracy tests are in separate documents.

3. Effective detection distances of cameras

One option to define the area surveyed by cameras is to define a fixed maximum distance, using a pole or other marker at 5m or 10m or whatever is appropriate. Animals beyond that distance are not counted. The assumption, which should be tested, is that all target species are detected if they are within that distance. The downside of this simple approach is that it excludes data from animals detected in the potentially long tail of greater distances where they are partially detectable.

ABMI uses all images (unlimited distance), with a procedure to estimate the effective detection distance of cameras. “Effective detection distance” is the fixed distance that would give the same number of detections as observed if all animals up to that distance were perfectly detectable and none were detectable further away. This approach is used for any point counts or transects with unlimited distances or with distance bands beyond the distance of perfect detectability. In the ABMI protocol, we place a prominently coloured pole 5m from the camera. All native mammals are recorded as being closer than the pole or farther than the pole, with additional categories for animals that are uncertain (near 5 m but not directly in line with the pole), investigating the pole or investigating the camera. Simple geometry gives the effective detection distance from the proportion of locations that are <5m away versus > 5 m (excluding the uncertain and investigating images): $EDD (m) = 5 / \sqrt{1-p_{>5m}}$, where $p_{>5m}$ is the proportion of images with the species > 5 m away. The area surveyed by a camera is $\pi \cdot EDD^2 \cdot \text{angle}^\circ / 360$, where angle is the angle of the camera’s field-of-view in degrees (42° for the Reconyx cameras that ABMI uses).

Detection distances are expected to differ for different species, by habitat types and possibly by time of year (e.g., on snowpacks versus in summer shrubs). We therefore used the results to develop detection-distance models for 11 species groups (see individual results in Appendix 1) and eight broad habitat types: deciduous forest, upland conifer forest, upland grass, shrub, lowland forest, wet grass,

water and human footprint. BIC-based model selection examined 7 models with those habitat types grouped into broader categories, and 7 more that added a smoothing spline for Julian date.

Four species groups showed little difference in detection distances among habitat types, while there were small to moderate differences for the other seven groups. Grass, wet grass, water, human footprint and sometime shrubs tended to have longer detection distances than forested habitats for these groups. The differences in detection distance were not large in most cases, but the area surveyed is proportional to the square of the distance, so the bigger differences represent about a two-fold or greater difference in the area surveyed in each habitat type. Julian date had little effect on detection distance for any species, although we had only a few winter samples for this analysis. The example results for coyote + foxes (**Figure 2**) show more extreme differences than most species, with almost double the detection distance (and hence almost quadruple the survey area) in open habitats. Seasonal changes are negligible.

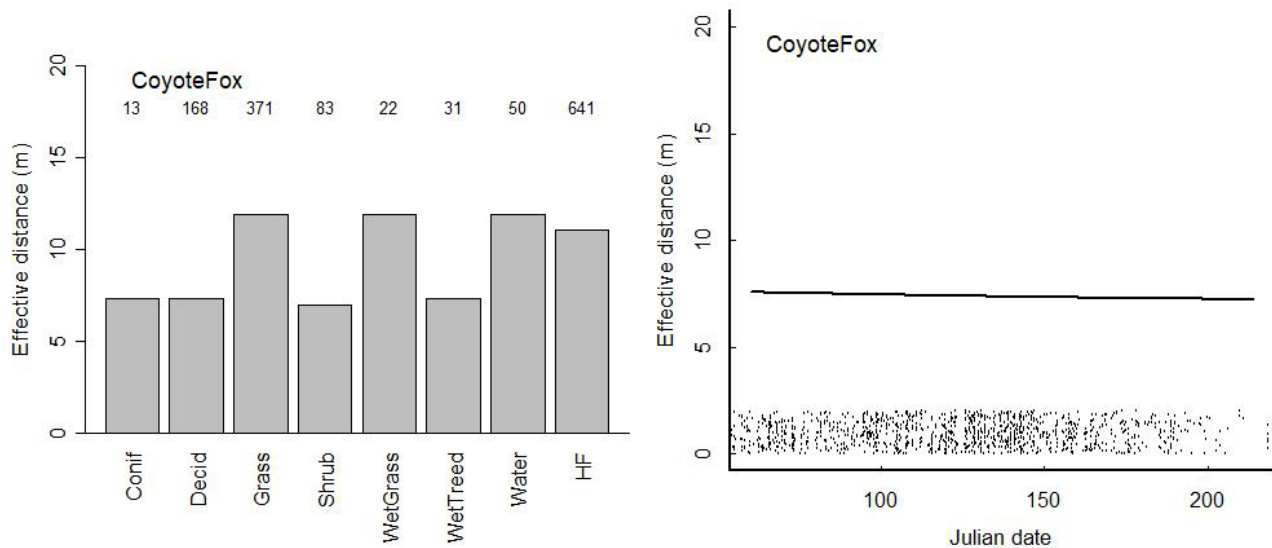


Figure 2. Model of effective detection distance for coyotes + foxes in different broad habitat types (left; HF = human footprint; numbers are sample size) and over time during the year (right; dots at bottom show seasonal distribution of data).

Overall, the effective detection distances are surprisingly low for some of the big mammals (e.g., **Figure 3** for big forest mammals, primarily moose but also some bison). The data come only from camera deployments without lure, but even then, several species spend time investigating the camera and the pole. Moose and bears are notable in that regard. Even though we exclude the “investigating” images from the distance calculation, that behaviour would produce more images near the camera as the animal approaches or leaves, or walks back and forth between the camera and the pole. In winter, several carnivore species clearly follow the field crews’ snowshoe trails, which are also mainly between the camera and the pole. These behaviours violate an assumption of the analysis, probably resulting in an underestimate of the effective detection distance, and therefore an overestimate of absolute density. As an estimate of relative density, we need to assume that this bias is the same in each habitat type (for habitat modeling), and that it will remain the same in future years (for trend). We could possibly use more detailed examination of animal positions and behaviours in the images to assess how much the assumptions are violated in different habitats, or eventually in different years.

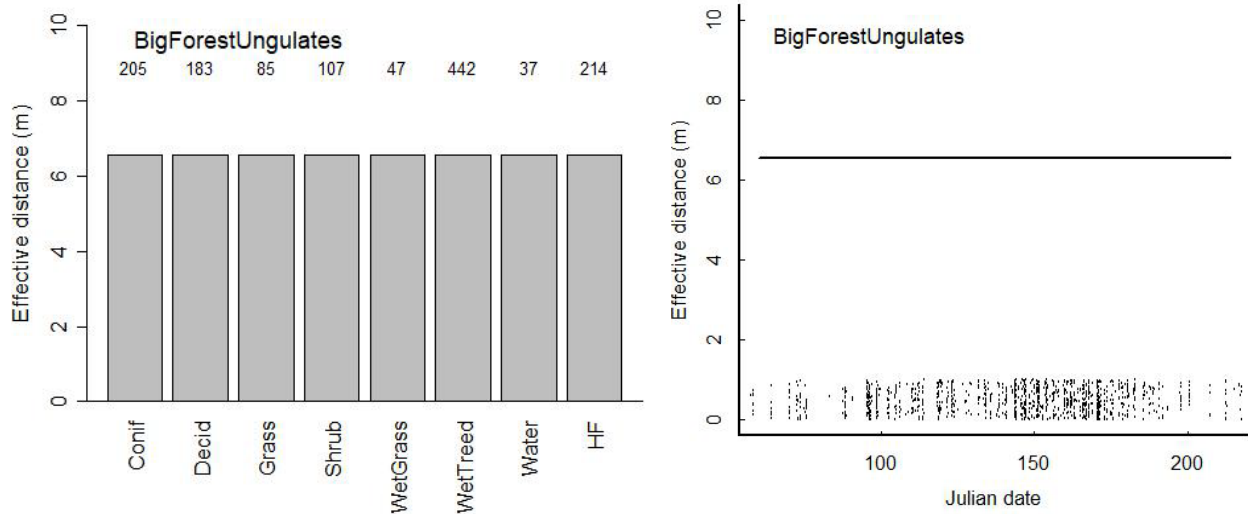


Figure 3. Model of effective detection distance for moose (and bison) in different broad habitat types (left; HF = human footprint) and over time during the year (right).

Another assumption of the distance analysis is that all animals < 5 m from the camera are detected. We ran a small test using simultaneous deployments of our motion-activated cameras and cameras taking images every 3 seconds. These showed nearly complete detection of animals larger than coyotes within 5 m, but the motion detectors clearly missed some smaller animals. For the smaller animals, we are certainly underestimating density. We are also concerned that those underestimates of small species may differ by habitat types – in particular, missing more animals in habitats with complicated ground cover. Seasonal data (**Section 6**) also suggests that those animals are more detectable on snow packs, when they are closer to the 1 m height of the camera and not obscured by vegetation. We are uncertain what effects this may have on long-term trend estimation for these small species, as vegetation and snowpacks can be changed by land use, climate change, etc.

We have fit these detection-distance models using data on position relative to the 5 m pole from all unlured sites in 2015, with supplemental data for rarer species in subsequent years. These models are used to estimate the detection distance for each species at each camera, based on the species, the camera’s habitat type and the dates the camera was operating. We use the models – rather than trying to use direct empirical data on pole position at each camera – because sample sizes are far too low for reliable estimates at each camera. Also, this calculation needs to be done for all species at all cameras, whether or not the species was detected there, to determine the total area being surveyed by each camera for each species. Applying the models only requires knowing the habitat type at each deployment and the dates the camera was operating (if there is a seasonal component to the detection distance model for the species). The habitat types for these detection distance models are broad enough that they can be determined from the camera images if detailed GIS information is not available.

Effective detection distance calculations would ideally use more detailed distance information. Five distance bands are often recommended. ABMI uses only the two bands defined by the 5 m pole because of field-logistical constraints. We tried an alternative approach measuring the height of animals on the screen in each image compared to their height when they were at the 5 m pole. This gives detailed distance distributions, but is difficult when animals are in different postures, and was too time-consuming for available budgets.

4. Time cameras are operating

For many cameras, the total operating time is simply the time from just after the initial set-up to just before the final collection. However, some cameras fail before they are picked up, most often because they run out of memory or battery power after taking many pictures of cows, vegetation or occasionally perching birds. Other cameras fail because they fall over or are knocked over by cows. We use specific criteria for determining when a camera is too displaced to use – if the 5 m pole is no longer in the image field, or the camera is tilted > 30° from horizontal. Some of those may be re-established either during a mid-season visit or simply when the camera tilts back (e.g. if it is on a post that cows lean on), leading to a gap in operating time for the camera.

Times each ABMI camera was operating are illustrated in **Figure 4**. Many off-grid cameras, which target specific habitat types to improve habitat models or address other specific questions, are set out later than the on-grid ones. 2016 and 2017 included some cameras that were deployed in the fall and left out through the winter and following summer. Variation in start and end dates has increased over the three years, due to operational logistics (crew availability and site access). Some cameras in each year were picked up in October or November, versus July for most cameras. Cameras in the south-central part of the province were more likely to have short operating times because of cows. Besides creating different total operating times that need to be included in the density estimates, the varying deployment periods interact with seasonal variation in detectability of different species, which we need to address in the analysis (**Section 6**).




Figure 4 Cont'd next page.

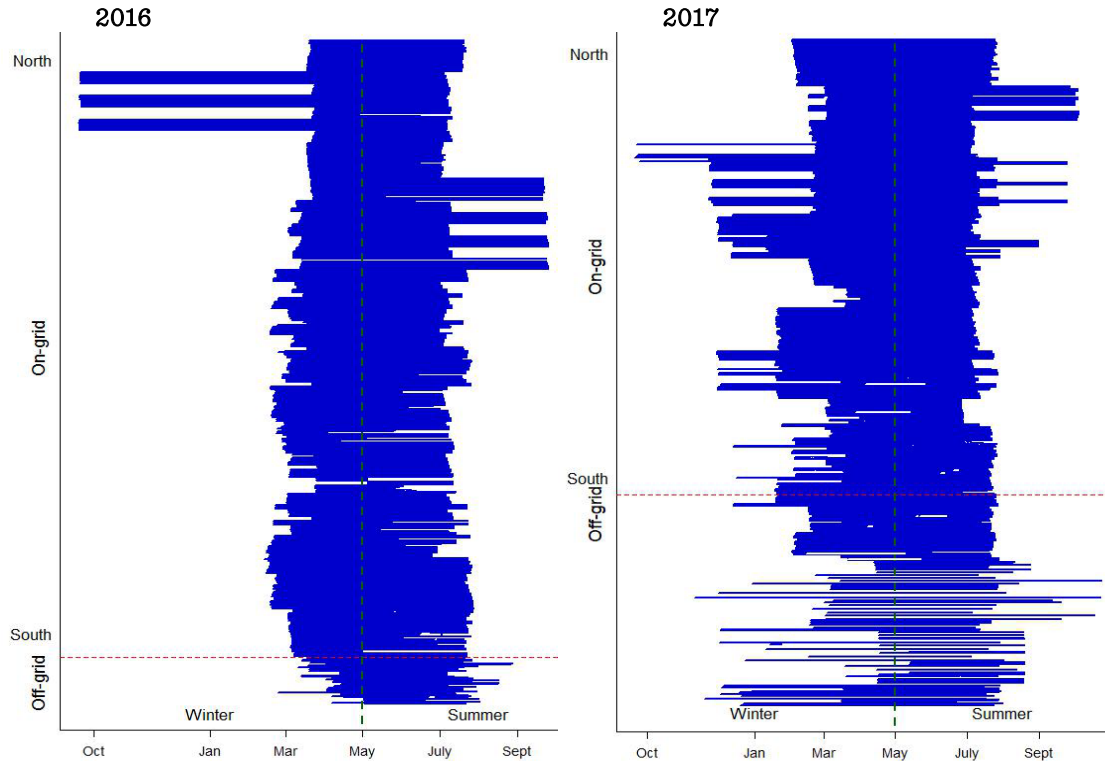


Figure 4. Times each camera deployment was operating in 2015, 2016 and 2017. Each individual horizontal blue line represents a deployment. Deployments are arranged from north (top) to south (lower) for on-grid ABMI sites, then for off-grid sites below those. Note that the x-axis scale is different for the 2015 figure – deployments in that year covered a narrower range of dates.

5. Total time a species is in the field-of-view

To derive this measure of animal use of each camera (“animal-seconds in the field-of-view”), we need to know how many animals were in each image, and we need to convert the individual discrete images into series representing periods of continuous occupancy of the field-of-view.

5.1. Individuals in each image

The number of individuals of each species visible in each image are simply counted. Juveniles – often in company of their mothers – are counted and therefore included in the density estimates. We note age class of each animal so that we could exclude non-adults if we wanted for some analysis (assuming we are confident in that designation in all images). Distant animals are sometimes captured in images triggered by an individual closer to the camera. We count all individuals, because it is not clear how we would know an individual was too distant to trigger the camera. To be consistent, when we are collecting data on animals in front of or behind the 5 m pole, we also include all individuals, even distant individuals when a closer animal clearly triggered the image. Distant animals therefore increase the effective detection distance, although only to the same extent as any animal beyond the 5 m pole.

Series of images showing continuous presence of a species in front of the camera (defined in the next subsection) can have different numbers of individuals. Because the density measure only requires total animal-seconds – i.e., it doesn’t matter which individual(s) is/are in front of the camera at any one time – we simply take the average number of individuals in each photograph in the series. A series with

1, 1, 2 and 1 individual in its four photos would have 1.25 individuals present on average. This may produce a slight upward bias in the density index, because it is possible that the camera would be triggered more often when there are more individuals in front of the camera (because one individual is more likely to be moving, and there is more likely to be an individual closer to the camera's motion detectors). More individuals may therefore be associated with shorter intervals between images than the average for the series. A more precise approach might be to do a time-weighted average number of individuals in the series, where the time for each image is the interval from the mid-point between it and the previous image, and the mid-point with the subsequent image. This more complicated time-weighted average has not yet been implemented at ABMI, because it probably makes little difference to the estimate, even when variable numbers of individuals are present in a series.

5.2. Time in field-of-view

There are two issues in calculating how long animals are in the camera's field-of-view: accounting for animals that leave the field-of-view and then return, and adding time before the first image in a series and after the last image. The latter is particularly important for single-image series, which would otherwise represent 0 seconds in front of the camera.

5.2.1. Animals leaving the field-of-view and returning, and "probabilistic gaps"

From a pilot study, we determined that if there is a gap of < 20 seconds between images of the same species at a camera, the animal is almost always still in the view (no evidence of it walking out and returning). Missing the odd time when it leaves the view for less than 20 seconds has little effect on estimates of the total time it is in the field-of-view. At the other end, if there is a gap of > 120 seconds between images of the same species, this almost always represented animals leaving and then returning (i.e., the animal is seen walking out of the field-of-view, then walking back in). Gaps of 20–120 seconds are uncertain. These relatively long periods when the animal could be in the field-of-view or not are important when estimating the total durations animals are in the field-of-view.

For the 2015 ABMI images, we checked each 20–120-second gap in series of native mammals for evidence of the animal leaving and returning. For 2016 and 2017 ABMI images, we checked 20–120-second gaps only for less common species where we had low sample size from 2015. We looked at several images on either side of gaps of 20–120 seconds. In each sequence, the animal was designated as having left the field-of-view during the 20–120-second gap if there was clear evidence of it walking out of the field-of-view and then returning (or a different individual entering the field-of-view). If the animal stayed in one location within the field-of-view, or sequential images showed the animal in disconnected places (as often happens with smaller animals), the animal was assumed to have stayed.

We used the data from 2015 to 2017 to develop models of the probability of a species leaving the field-of-view during a 20–120-second gap as a function of the gap duration. Smoothing splines were fit to the probability of leaving as a function of gap length, using a logit-linked binomial model. This was done separately for each species with enough examined gaps (**Figure 1**). Examining the resulting curves and trying to keep similar species together, species were grouped into 6 "gap groups": all ungulates except moose, moose, bears, cougars + canids, all other carnivores (mostly mustelids), all other mammals (small species).

Across all species, there is a 27% chance that the species left when the gap was only 20 seconds long (**Figure 5**). That implies that they may also have left in slightly shorter gaps, where we assume they all stayed. However, it makes little difference to the total duration to miss a few shorter gaps where the animals left (i.e., to add that extra little bit of time in), especially because we add time to the start and end of a series (next section). At the other end, there was evidence that animals left in about 70% of the 120-second gaps. Snowshoe hares, porcupines and black bears lowered that average. The first two are small

mammals that trigger the camera erratically, making it hard to figure out which way they are moving, so that they generally don't present definitive evidence of leaving the field-of-view and returning. Bears, and wolves to a lesser extent, often lie/loll around, mainly at lured deployments, and could go for relatively long periods without triggering the camera. Missing long periods when bears are present but not moving would lead to underestimates of their densities. For trend estimates and habitat modeling, we would have to assume that the prevalence of those misses does not change over time or by habitat type. The other species had higher leaving rates for long gaps, so that we are probably not too inaccurate to assume that they left during any gap > 120 seconds.

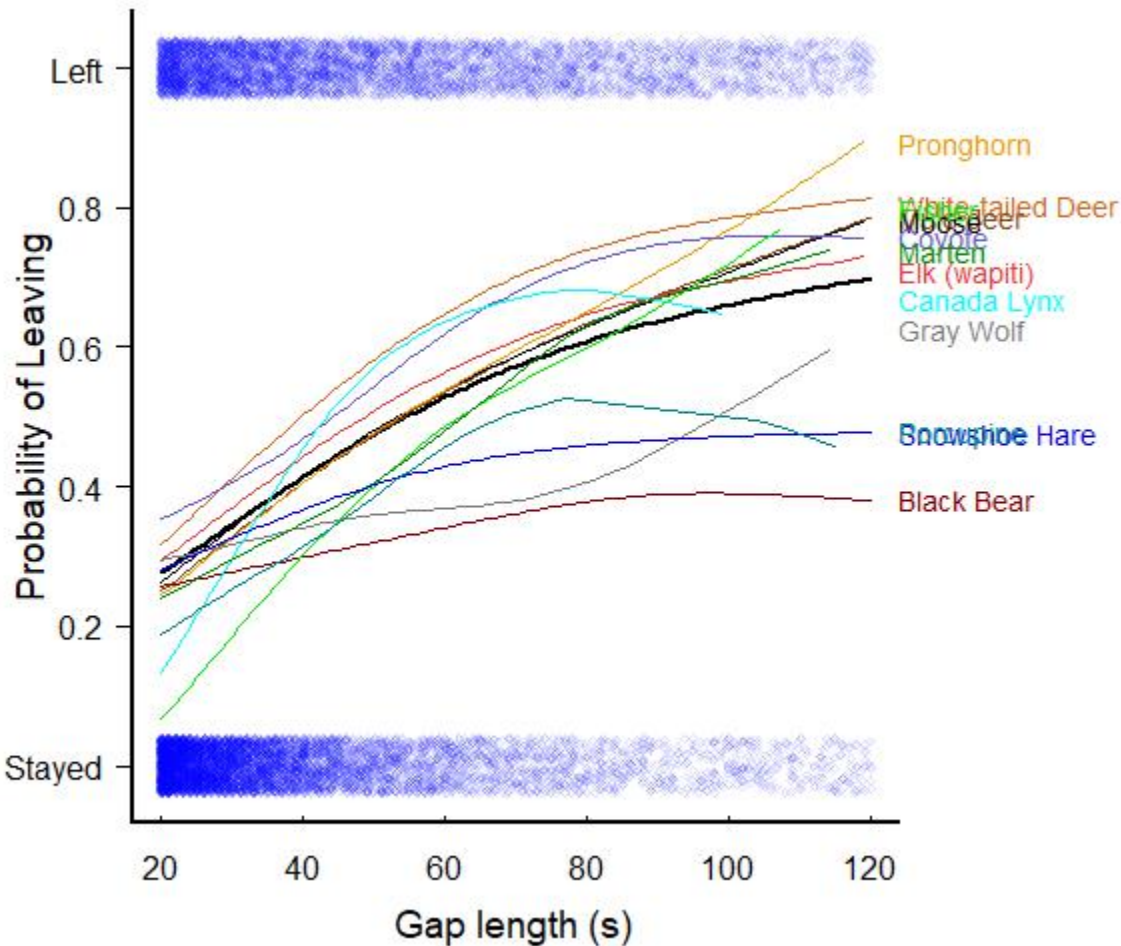


Figure 5. Probability that an animal leaves the field-of-view, based on clear evidence from the images, for gaps between images of 20–120 seconds. Blue points are data for all species, black line is pooled curve for all species, and coloured lines are for the individual species.

Ideally, each gap between images would be checked to see if the animal left or not. However, the process of checking each gap is time-consuming. When we have direct information on gap-leaving, we define a series as any set of images separated by < 120 seconds, unless the animal was observed to leave the field-of-view (in gaps of 20–120 seconds that were checked) (illustrated in **Figure 6**). If it did leave in a gap, then the series ends at the preceding image, and a new series starts when the animal (or a different individual) returns in the subsequent image. The species is in the field-of-view from the first to the last image in the field-of-view (plus end buffers; next sub-section).

When we haven't examined 20–120-second gaps, a series is all images separated by < 120 seconds. However, we then use the above models of gap-leaving probabilities to prorate the 20–120-second gaps for the probability that the species left for a gap of that length. Instead of the full 20–120-second gap length, we only add the duration of the gap * (1 – probability of leaving) to the total series length (illustrated in **Figure 6**, bottom). For example, if there were 4 images separated by 10 seconds, 5 seconds, 60 seconds and 10 seconds, and the model for that species showed a 40% chance that it left in a gap of 60 seconds, then the total time in the field-of-view for that series is $10 + 5 + 60 \cdot (1 - 0.4) + 10 = 61$ seconds. (And, if there were an average of 1.25 individuals in those 4 images, the total animal-seconds would be $1.25 \cdot 61 = 76.25$ animal-seconds). The next subsection completes the calculation of time in field-of-view for a series by accounting for time the animal was present before the first image in the series and after the final image (striped sections in **Figure 6**).

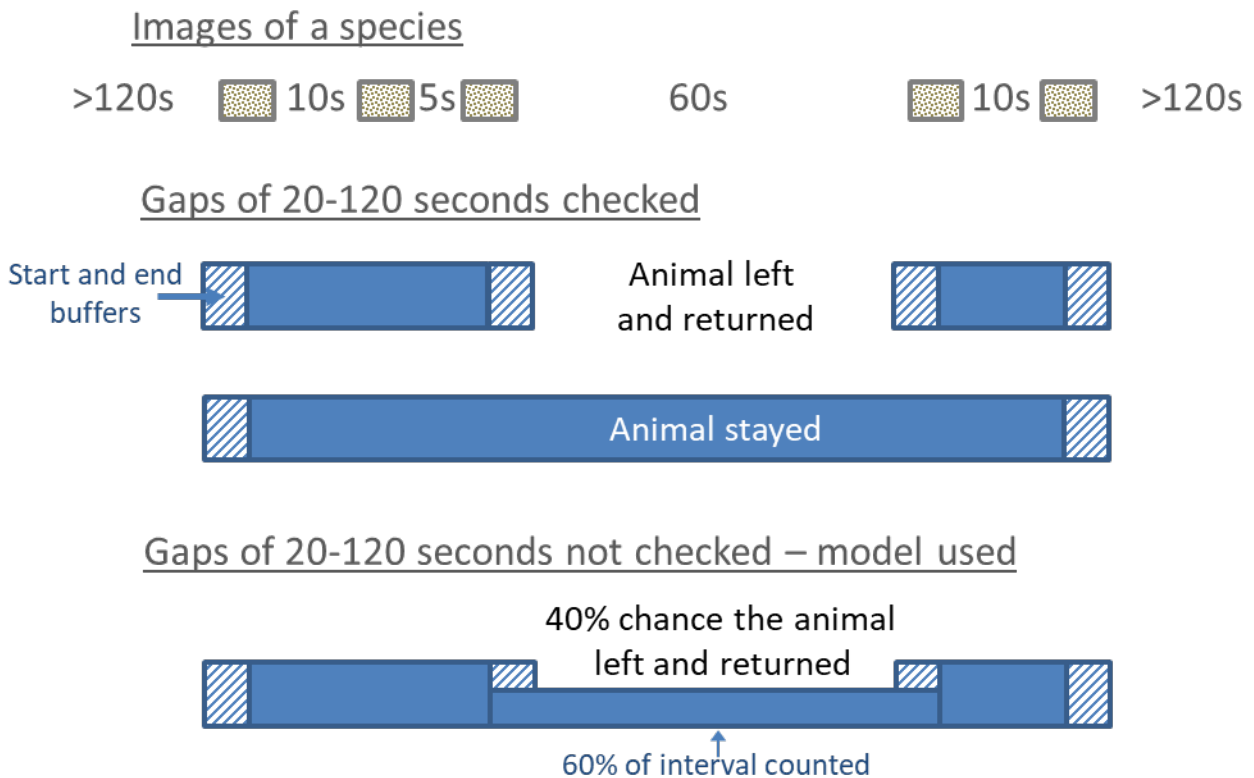


Figure 6. Illustration of how sequential individual images are converted to series, and how the species' total time in the field-of-view is calculated for each series, when we have directly examined gaps of 20–120 seconds (top) or not examined them and used the probabilistic models instead (bottom).

5.2.2. Time before first image and after last image, and time of single-image series

The cameras only take images every few seconds or longer when an animal is present. One question is therefore how long the animal is in the field-of-view before the first image in a series and after the last image. We estimate that by calculating the average time between images in all series, separately by species (**Table 1**). This is typically 4–7 seconds for larger species and somewhat longer for small species. This time is added to the duration of each series, including to single-image series, which would otherwise have a duration of 0 seconds. The assumption is that the animal is in the field-of-view for half the average inter-photo time before the first image, and after the last image. (Although animals often

appear to be further into the field-of-view when the first image is taken, and nearer the edge when the final one is taken, presumably reflecting a lag-time in the motion-detector – the assumption is that this averages to the average interval within series).

Table 1. Average time between photos within a series for each detected native mammal species.

Time between photos			Time between photos		
Species	Average	n	Species	Average	n
Badger	7.48	52	Mink	15.75	4
Beaver	1.00	2	Moose	4.52	39332
Bison	4.92	1241	Mule deer	4.72	76294
Black Bear	6.17	15694	Porcupine	4.38	797
Bobcat	4.50	2	Pronghorn	5.52	12722
Canada Lynx	5.21	672	Raccoon	10.48	21
Cougar	3.38	48	Red fox	6.62	364
Coyote	7.11	6371	Red Squirrel	7.55	96
Deer	5.33	2481	Richardson's GS	10.88	81
Elk (wapiti)	3.93	24794	Snowshoe Hare	11.46	1070
Fisher	11.45	248	Striped Skunk	8.43	56
Foxes	1.00	1	Swift fox	4.57	7
Gray Squirrel	2.00	4	Voles, Mice and Allies	1.80	5
Gray Wolf	5.11	732	Weasels and Ermine	18.00	1
Grizzly bear	3.96	190	White-tailed Deer	4.81	87260
Groundhog	9.25	4	WT Jack Rabbit	10.09	527
Least Chipmunk	4.70	10	Wolverine	7.47	81
Marten	8.23	146	Woodland Caribou	4.30	1297

Note: n here is the number of intervals between photos in the same series (i.e., excluding singletons).

When we have 20–120-second gaps that are using the probabilistic gap-leaving model, we add this extra time on the start and end of the whole series in the normal way, and we also added this time multiplied by the probability of leaving for each 20–120-second gap, to account for the cases where the animal would have left the field-of-view in those 20-120-second gaps and hence created another series (see illustration, **Figure 6** bottom).

6. Seasonal detections and seasonal effort

Within the ABMI monitoring, different cameras are deployed across different dates. Some of these seasonal differences are due to cameras failing (e.g., full memory cards), but in many cases it is just operational logistics that have cameras set out or collected at different times of year. In theory, if the assumptions of the density estimates were perfectly met, density should only change seasonally in response to true changes in the population, due to seasonal recruitment and mortality. In practice, assumptions are not met perfectly, and different species are more or less detectable in different seasons. Bears are an obvious example - they are not detected when they are hibernating in winter. Again theoretically, if we had millions of truly randomly placed cameras, a few would end up in bear dens, and (if they were triggered when the animal was present), they would produce extremely high density estimates that would average out with all the other cameras' zeroes. Clearly, we do not have that many cameras, they are not placed in bear dens, and they probably wouldn't function if they were. Other species can show more subtle seasonal changes in detection due to violations of assumptions: ungulates

likely use security and thermal cover more in hunting and winter seasons, and we do not place cameras in dense vegetation. Lynx, fishers, martens and snowshoe hares appear to be more visible on snow than in vegetation, probably reflecting a violation in snow-free seasons of the assumption of perfect detectability up to the 5m pole.

As a result, there are clear seasonal differences in catch per unit effort (images/time) for many species (examples in **Figure 7**, but see design caveats below). Black bears show the expected absence in their hibernation season. Caribou show a spring peak, then declining abundance through the summer into the fall (and absence in winter, but we may not have had winter-active deployments in caribou ranges). Several smaller species are more commonly detected on the snowpack, including marten, fisher, lynx, coyote and snowshoe hare. Additionally, lynx, wolverine and maybe other species clearly follow the field crew's tracks in winter deployments, increasing their abundance in the many sites with those late-winter deployments. The more common ungulates tend to be roughly equally common through the year, except for changes that might be expected from the geographic changes in deployments that represent different months (e.g., very few white-tailed deer in October and November, simply because the cameras active in those months were in areas north of the species' main range).

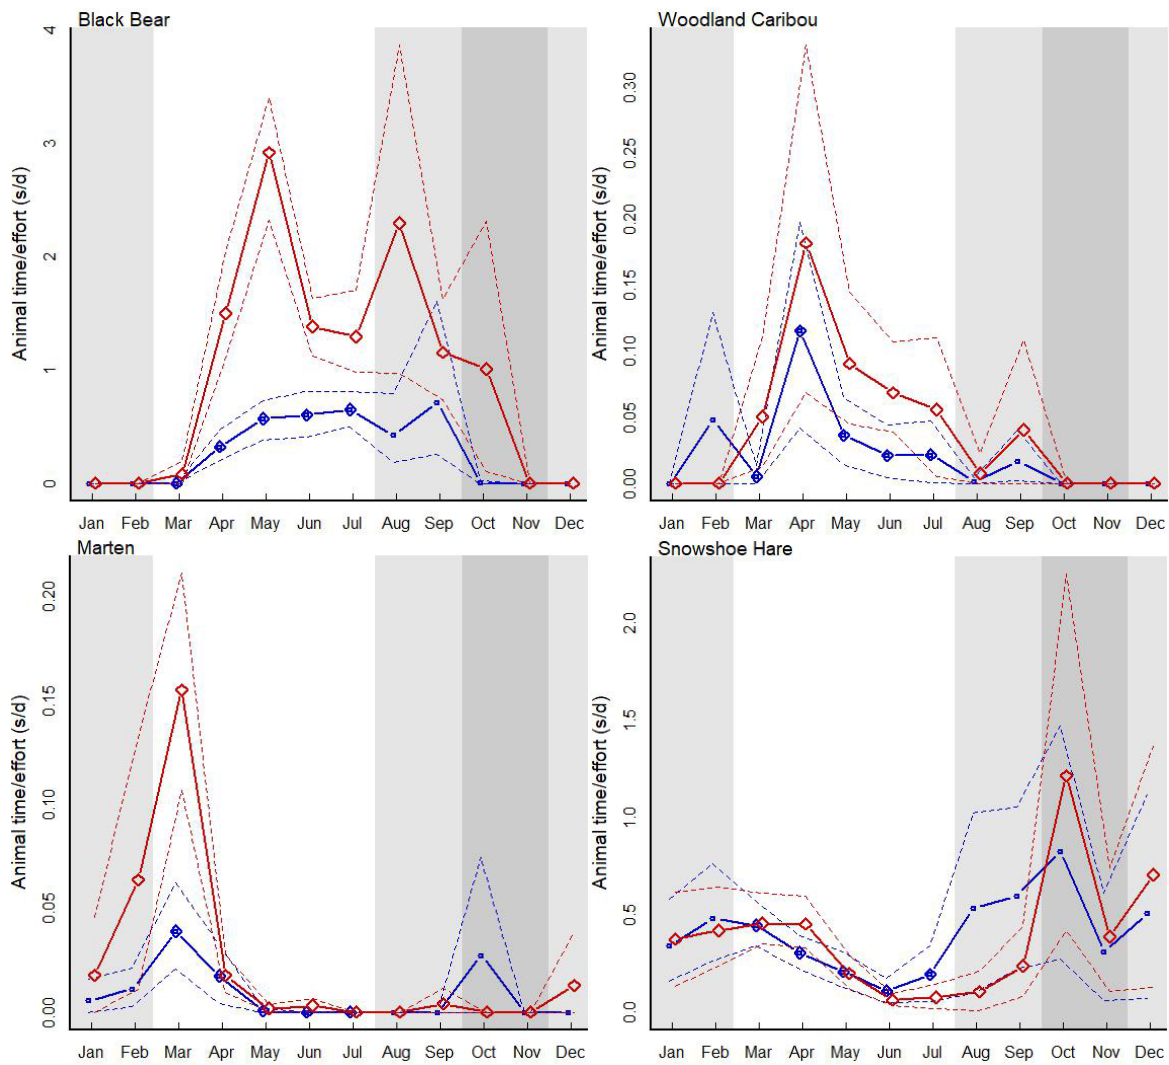


Figure 7 Cont'd next page.

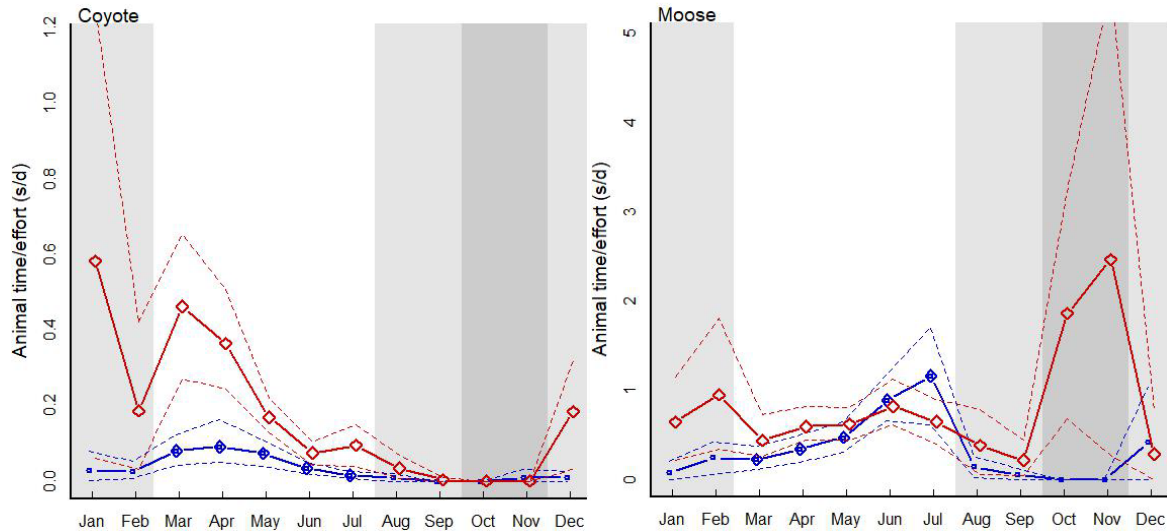


Figure 7. Catch per unit effort (seconds of species presence per day of camera time) of example species over time, for lured (red) and unlured (blue) deployments. Relatively few cameras were operating in the grey periods, and these were in non-representative geographic areas, especially during the darker grey period.

These seasonal changes in density estimates can confound results when our deployment and pick-up, or camera-failure, dates are not all the same. This is a concern for trend, as deployment dates change over time, and for habitat modeling, if more remote northern sites are more likely to have longer or over-winter deployments.

To correct for this problem, we could conceivably develop seasonal correction factors for each species. However, this would be complicated, and the current ABMI design does not support doing this because we don't have good samples from all seasons well-dispersed spatially and across habitats. In other words, this solution would suffer from the same confounding that it is trying to solve. Instead, we are currently using a simpler approach of truncating all cameras to a range of dates when most were operational. That common time period has to be the same for each year, to avoid confounding year and the habitat types sampled more in some years with different sampling windows. For the current camera density summaries, we use a window of March 22 to June 27, when > 90% of cameras were operating on any date (**Figure 8**). Animal records and camera operating time outside that window are omitted.

The main problem here is that we sacrifice a lot of data – 37.7% of cameras' operating times lie outside that window. The percentage of images of some example species that fell within this period is summarized in **Table 2**. Species detected mainly in summer, such as black bears and caribou, have a high percentage of observations in the common period, while species seen more commonly on snow in late winter have only a moderate percentage. Bison are represented mainly by one herd seen outside of the common period.

We also still have concerns about missed times within that common window for some cameras, such as off-grids that missed the early part of that window, and hence the snow pack where we see several species almost exclusively, or cameras that ended earlier and hence would have increased density estimates for species detected primarily in March and April (because those early-season records are not "diluted" by extra late-season days with few detections). Truncating to a common operating period is probably not a long-term solution, as the common period will get shorter, or have a lower

operating percentage, if deployment periods continue to vary over time. The best choice, of course, is to maintain consistent deployment times over the years and in different habitats.




Figure 8. Percentage of cameras operating through the year.

Table 2. Percentages of images of selected species that were in the March 22 to June 27 “common period”.

Species	% images in common period
Caribou	83.9
Black Bear	79.2
WT Deer	64.2
Coyote	63.5
Lynx	55.5
Fisher	36.8
Bison	5.4

7. Lure effect

The assumption that animals are not attracted to or repelled by the camera implies that no lure or other attractant is used. However, results from lured deployments can be used if a factor is included to account for how much the lure increases the effective field-of-view for attracted species. ABMI chose to use lure at half their deployments, and no lure at the others. Lure is intended to attract more carnivores, which are generally rare and therefore subject to high uncertainty. However, long-term trend is a major objective of ABMI monitoring, and there is concern that lure formulations change over time, potentially ruining confidence in long-term trend estimates if only lured sites were used. The ABMI therefore uses four deployments at each site (separated by 600 m), with two lured and two unlured. In addition to hedging bets about whether data from lured or unlured sites will be more valuable in the

future, this paired design provides a strong way to estimate the effect of lure. Direct comparisons of lured and unlured deployments are not confounded by geographic location, year, or, with enough samples, microhabitats.

Using the paired design, I summarized the ratio of occurrences (presence/absence) at lured:unlured deployments, and the ratio of “abundance given presence” (the density when the species is present) (**Table 3**). Total density is the product of those two components. Large carnivores occur moderately more often at lured sites, and stay considerably longer when they do occur. Smaller carnivores occur more often at lure sites, but only stay somewhat longer, if at all. Overall, density estimates are inflated 2–10 times with lure than without for the carnivores. Ungulates, rodents and hares did not show significant effects of lures, although there was a tendency for the ungulates to have slightly lower densities at lured sites. When we do analyses that include estimated densities from lured deployments, we divide those densities by the appropriate factor for the species. Again, the ABMI direct paired comparison across a large number of sites is necessary to be able to do this correction effectively.

Table 3. Lure effect on presence/absence at a deployment and on abundance (duration in field-of-view) when present at a deployment, for the carnivores.

Species	Lured:Unlured	
	Presence/Absence	Abundance Presence
Black Bear	1.33	2.18
Grey Wolf	1.76	3.21
Coyote	1.27	3.06
Red Fox	2.12	3.42
Cougar	1.26	1.26
Lynx	1.70	2.28
Wolverine	1.29	3.50
Fisher	5.52	2.04
Marten	2.70	0.89
Small Weasels	2.67	0.84
Striped Skunk	1.57	1.51

Confirming our concerns about using uncorrected-for lure in trend estimates, several species have shown apparent changes in the lure effect over the three years of ABMI sampling (**Figure 9**; the three points left to right for each species are 2015, 2016 and 2017. Error bars are 90% CI’s, from bootstrapping with the site as the resampling unit). For example, the lure effect is stronger for bears, canids and wolverines in 2017 compared to 2016 and 2015. CI’s mostly overlap the other means, but results are consistent across the 7 species. Marten and fisher might shower a lesser lure effect in 2017, but sample sizes are small for those species. The geometric mean lure effect for the nine main carnivores was 3.3 in 2015, 3.0 in 2016, and 4.6 in 2017 (CI’s not calculated). The non-carnivores, which generally don’t have strong responses to lure, do not show any consistent trends over time, although there are anomalous individual years for a few species.




Figure 9. Lure effect for each main species. The three points per species are, left to right, 2015, 2016 and 2017. Error bars are 90% CI's. (Note: a value of 1 would mean no lure effect. We know that lures attract carnivores, so the significant differences from 1 are not surprising. The important point here is the comparison of years within a species.)

8. Changes over time due to technology

A concern with long-term trend monitoring using any technology-mediated measure is that changes in the technology over time will generate spurious trends. More sensitive motion-detectors, for example, would lead to more time with animals in the field-of-view (although effective detection distances should also increase in compensation). Because of concern about those changes and the risk to credible trend estimates, ABMI has implemented paired comparisons of different camera types, using the two unlured deployments at each site. To date, all ABMI sampling has used the same make and model of camera. However, different batches of cameras were bought in 2013, 2014 or 2016. These were deployed in the pairwise design in pairs of 2013 versus 2014, or 2014 versus 2016 (there were more cameras bought in 2014 than other years). The comparison has only been done in 2017, so the cameras were 4, 3 or 1 years old at that point. (Note that this means that we are looking at differences in the

performance of cameras of different ages, not anything to do with yearly changes in the species themselves).

Although no year was radically different from its comparison year for any species, there were 5 of 14 species where the 2014 cameras had significantly lower total durations than the paired 2016 cameras (at $p < 0.1$) (**Figure 10** – the righthand values for these species are significantly < 1). Four of these 5 species were prairie or parkland species (coyote, red fox, elk, pronghorn), while the other was marten. No species showed a strong difference the other way (except porcupine with extreme CT's). The comparison of 2013 and 2014 was based on fewer paired sites, and showed no consistent differences. There were no paired comparisons of 2013 and 2016 cameras.

The 2014 versus 2016 results suggest some possible deterioration of cameras over time, or, alternatively, that the 2016 version of that camera model was just better, but the results are not overwhelming. It is worth continuing these comparisons, because substantial uncorrected changes with camera age and/or manufacture year would bias trends. Worse, the mere possibility of that bias would undercut confidence in our future trend results. Such designed comparisons will be far more important when different makes or models of cameras start being used.




Figure 10. Comparison of cameras bought in 2013 to cameras bought in 2014 (left point for each species), or cameras bought in 2014 to cameras bought in 2016. The value is the ratio of the species' total duration in cameras of the earlier year divided by its total duration in cameras of the later year, using only the cameras at the paired-comparison (unlured) sites. A significant difference from 1 would be a relevant result here, showing a difference in the cameras from the two years being compared.

9. Assumptions of the camera density index

There are a number of strong assumptions involved in using this measure to give absolute density of a species:

Assumption 1 The cameras are a random or otherwise representative sample of the area. As mentioned in the first section, density is calculated for the area in the fields-of-view of the cameras. To make inferences about density over a larger region, those fields-of-view need to represent the larger region (as is the case for any sampling). This means that cameras are not intentionally put along trails or other features that attract animals. Unfortunately, it also means that cameras should be placed in micro-habitats randomly. This is not possible logistically, because the cameras need some open space in front of them.

This practical limitation on camera micro-locations creates problems for interpreting habitat models when habitat types differ in amount of suitably open micro-habitats, if animals either favour or avoid those more open habitats. We are really surveying the use of suitably open micro-habitats within larger habitat types. Micro-habitat selectivity may also create uncertainty in trend, if there are substantial changes in habitat types and the amount of suitable open micro-habitats. Violations of this assumption do not intuitively seem like major potential sources of bias, but it isn't clear how we would evaluate their effect rigorously.

A more pervasive problem in the south of the province is that we cannot deploy cameras in active cultivated areas, so we use adjacent fence lines, leave strips or woodlot edges. Several species clearly favour those areas. This complicates interpretation of habitat use and trend in extensively cultivated regions, particularly as amount of those edge features changes over time. We would ideally be treating our design as a stratified sample, but we do not have the resources to estimate the areas of the strata we are sampling non-proportionally (or even to define them well) – i.e., we do not know the areas represented by the types of fence lines, woodlots, etc. that can accommodate cameras, and we do not track how those change over time. Interpreting trends from cameras in the south will require more substantive assumptions.

Cameras could potentially be deployed along game trails, if the design also included some cameras off of these trails, and the results were analysed as a stratified sample of the region. The difficulty again is in determining the area of the “game trail” and “not game trail” strata for whatever units the study is comparing (habitat types, years if the goal is to measure population change). Extremely detailed ecological information would be needed to know what constitutes a game trail – e.g., the first seismic line put into a forested area versus the hundredth, linear features in open areas, what happens if a natural game trail is clearcut or burned, etc.?

Assumption 2 Animals are not attracted to or repelled by the cameras, and they do not change their behaviour (time spent in front of the camera) because of the cameras. This is clearly not true for some species, such as bears and moose that regularly investigate the camera and the pole we use to measure detection distance, even when the deployment does not have a lure. It is also violated when cameras are deployed in winter and animals follow the snowshoe tracks in front of the camera, which is frequently seen in the days after winter deployments. These effects increase density estimates. For habitat models or trend that only need relative density, we will have to assume that these effects do not vary substantially by habitat type or between years. Because we have images, we could measure whether species spend different amounts of time investigating the camera or unlured poles in different habitats and whether that changes over time. We could then potentially correct for the biasing effects in habitat models or trend estimates. We have not yet tried that to see how well it works.

Assumption 3 The method to estimate detection distances in different habitats assumes that the animal is certain to be detected within the 5m band. Some smaller animals can pass right in front of the camera undetected. This is a concern even using relative abundance for habitat models, if there is substantially higher non-detection in some habitat types, such as shrubby areas or habitats with abundant downed wood. It could also be a concern for trend, if there are large habitat changes that affect detectability near the camera. The assumption seems to be good for animals the size of coyotes or larger, but we only have a small trial that tested that.

Assumption 4 If camera deployment dates are not the same everywhere or standardized to a common period, the method also assumes that there is no major seasonal variation in detectability of a species. In other words, if the previous 3 assumptions are violated, those violations do not differ by time of year. The best approach is to ensure that cameras are deployed at the same time, or at least that they are all operating for a useably long common set of dates. Alternatively, we will need to find ways to adjust for some cameras (potentially in different habitats or years) sampling more at seasons when a species is more or less detectable.

Calibration of cameras, as technology changes or existing cameras age, is discussed separately (**Section 8**), but is anticipated to require serious effort. Lures, which violate assumption 2, can be accommodated if there is a rigorous design that allows their effects to be estimated (**Section 7**). For both of these corrections, it is important to remember that any uncertainty in the correction factor becomes uncertainty in the main scientific questions the cameras are meant to address. An imprecise correction for changing technology over time, for example, would add directly to the imprecision of any trend estimates. There is no point in putting in the sampling effort needed for a precise trend estimate if the effort is not put in to get a correspondingly precise calibration coefficient when technology changes.

10. Sampling distribution of species density

The distributions of density estimates of species at deployments are extremely skewed, with a large majority of deployments having none of a species, some deployments having low densities (one or a few individuals passing by) while a few have very high values (one or more individuals spending long periods in front of a particular camera) (**Figure 11**). Such distributions require large sample sizes to obtain precise estimates – for example, for yearly changes in density in a region or to use for estimating abundances in different habitat types. Ecologists' intuitive experience with adequate sample sizes may be very misleading with these kinds of distributions. Pilot studies and formal analysis of expected precision are recommended for large projects.

For habitat modeling, we find that the density estimates are best treated as a compound distribution of presence/absence – modeling how the 0 records differ from the non-0 records – and of abundance where the species is present – explaining variation in abundance where the species was recorded. We model presence/absence with the typical logit-linked binomial model, and the abundance-given-presence distribution with a log-normal distribution, which fits most of the species' distributions reasonably well (e.g., **Figure 10** bottom left). This compound distribution is the same as a zero-inflated log-normal distribution, but explicitly treating the two components separately allows more flexible modeling and critical examination of each component.

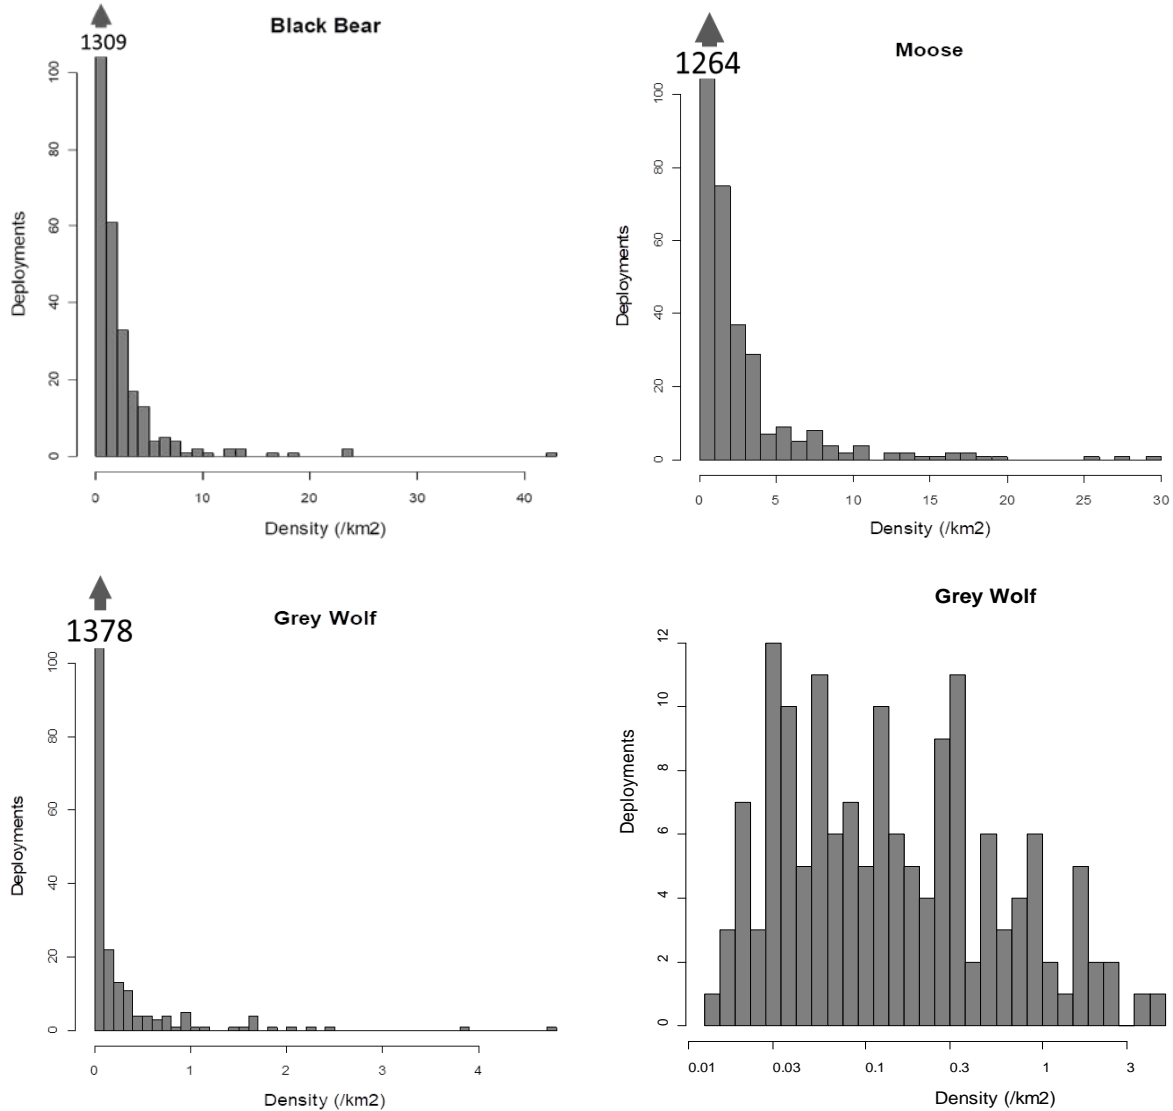


Figure 11. Distribution of density estimates across all deployments for three example species (top left and right, bottom left), and the distribution of log-transformed presence (non-0) densities for wolves (bottom right). Note that despite the extremely high densities at a few individual cameras, regional density estimates are reasonable, because they average in so many zeroes.