Comparison of Bioacoustic Recognition Models in Finnish Conditions

Muuttolintujen kevät v4 | BirdNET 2.4 | BirdNET 3.0 | Perch 2.0

Foreword

Mobile applications that identify bird sounds have become widespread in recent years. They have brought bioacoustic recognition into everyday birdwatching and nature monitoring: a phone can suggest which species is calling while the user is still in the field. Many current applications rely on BirdNET 2.4 or on the local BSG Finland v4 model built on top of it and used in the Muuttolintujen kevät application.

This comparison places these established solutions alongside two newer models. BirdNET 3.0 represents the next generation of the BirdNET development line, while Google DeepMind's Perch 2.0 broadens the perspective beyond birds to other animals and environmental sounds.

The report describes how the models perform in a Finnish rural soundscape in midsummer, when bird vocal activity differs markedly from the intense song period of spring. At the same time, the comparison provides a practical view of how automated sound recognition and bioacoustic monitoring may develop over the next few years.

Established applications and emerging models in the same comparison

BirdNET 2.4 and BSG Finland v4 represent the practical foundation of many current recognition applications. BirdNET 3.0 and Perch 2.0 show how species coverage, recognition of other animal groups and analysis of the full soundscape are expanding.

BirdNET 2.4	BSG Finland v4	BirdNET 3.0	Perch 2.0

Executive summary

Bioacoustic recognition uses an AI model to analyse recorded sound. The model searches for features typical of species vocalisations, proposes a source species and assigns a confidence score. This allows animals to be monitored continuously without an observer being present.

Four bioacoustic recognition models of different generations were compared: the long-established BirdNET 2.4, the Finnish Muuttolintujen kevät v4 model built on top of it, and two newer models, the BirdNET 3.0 developer preview and Perch 2.0. The aim was not to determine which model produced the largest number of detections, but to assess how consistently the models identified species in a Finnish soundscape and what limitations appeared in practical use.

Across the two comparison days, Muuttolintujen kevät v4 appeared to be the most reliable option at the time of testing for automated baseline monitoring of Finnish birds. BirdNET 3.0 was still under development but looked highly promising as a global model. Perch 2.0 stood out for its broad recognition scope, while BirdNET 2.4 continued to identify much of the common species core.

Key interpretation

When interpreted correctly, bioacoustic recognition is a reliable way to monitor common species that vocalise repeatedly. Isolated detections should nevertheless be checked before being accepted into a species list.

Muuttolintujen kevät v4

Most reliable overall in Finnish conditions

BirdNET 3.0

Most promising global model; still in development

BirdNET 2.4

Oldest model in the test; common species still recognised

Perch 2.0

Broadest animal and environmental-sound coverage

Three identical devices - four models across two comparisons

26 June 2026		8 July 2026
Muuttolintujen kevät v4	Same location and soundscape	Muuttolintujen kevät v4
BirdNET 3.0	Each model on its own device	BirdNET 3.0
Perch 2.0	Only three devices available	BirdNET 2.4

BirdNET 3.0 and Muuttolintujen kevät v4 linked the two comparison days.

Aim of the comparison and test setup

The models use different sample lengths, confidence thresholds and methods for generating results. Detection counts produced from the same continuous vocalisation are therefore not directly comparable. The comparison emphasised species recurrence and whether the models detected the same species within the same hour.

Location and timing

The recording site was in a rural environment near the boundary of Lohja and Siuntio in southern Finland. The timing was challenging for recognition: easily identified territorial songs had already declined, while calls of fledglings and warning calls from adults were more prominent in the soundscape.

Processing settings

Model	Confidence threshold
BirdNET 3.0	68%
BirdNET 2.4	70%
Muuttolintujen kevät v4	80%
Perch 2.0	70%

50% overlap

All models processed samples with 50% overlap: the next sample began halfway through the previous one. Sample lengths varied by model, so a single duration cannot describe all models.

Thresholds were selected from practical experience so that the results remained broadly comparable without excluding too large a proportion of genuine detections.

Why were all four models not tested on the same day?

Each model ran on its own identical device, but only three devices were available. On 26 June the comparison included Muuttolintujen kevät v4, BirdNET 3.0 and Perch 2.0. On 8 July Perch 2.0 was replaced by BirdNET 2.4. BirdNET 3.0 and Muuttolintujen kevät v4 provided the link between the two comparisons. Because natural conditions differed between days, BirdNET 2.4 and Perch 2.0 cannot be compared directly.

Assessing reliability

Common and repeatedly vocalising species were already recognised quite reliably. A single isolated detection should usually be sent to a review queue, although during migration a fly-over bird or flock may produce only one detection sequence.

Interpreting detections in practical monitoring

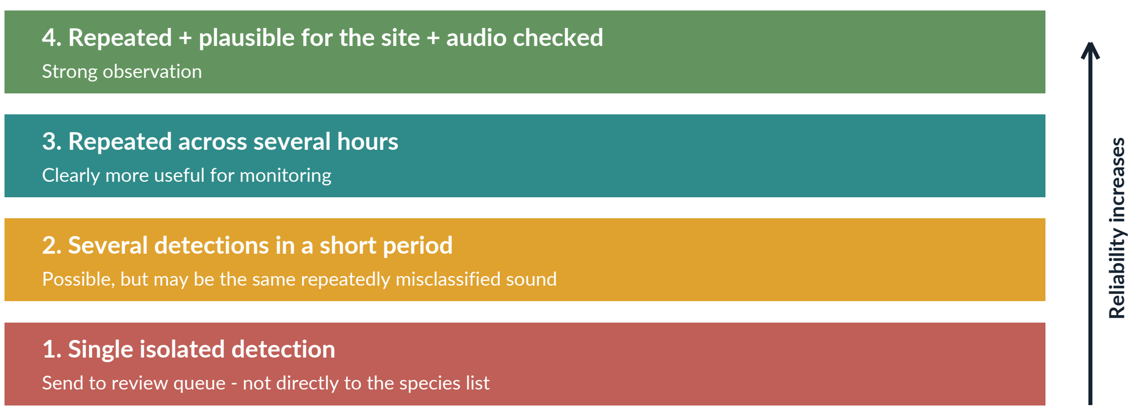

Detailed day-specific results are provided in a separate PDF appendix.

Model background, age and coverage

Model age refers here to the time since the tested model version was publicly released, as of July 2026. It does not refer to the age of the broader research programme, training data or developer organisation.

Model	Developer and age in the comparison	Coverage
Muuttolintujen kevät v4 / BSG Finland v4	A locally fine-tuned model created through the Bird Sounds Global collaboration. The Muuttolintujen kevät app was developed at the University of Jyväskylä. V4 was released in late 2025; about 7 months old.	263 Finnish bird species plus non-bird and human classes. Uses the BirdNET 2.4 backbone. [1-3]
BirdNET 2.4	Stefan Kahl, Connor Wood and Holger Klinck; Cornell Lab of Ornithology and Chemnitz University of Technology. The oldest model generation in the comparison; in use by 2023 at the latest.	6,522 bird species worldwide plus a place-and-time model. [4-6]
BirdNET 3.0 Preview 3.1	Mario Lasseck, Maximilian Eibl, Holger Klinck and Stefan Kahl. Released 15 June 2026; only a few weeks old during testing.	About 11,000 species, including more non-avian animal groups than before. [7-8]
Perch 2.0	Bart van Merriënboer, Vincent Dumoulin, Jenny Hamer, Lauren Harrell, Andrea Burns and Tom Denton; Google DeepMind and Google Research. Released 6 August 2025; about 11 months old.	14,597 animal species and 198 general sound-event classes. [10-11]

Breadth of the recognition space

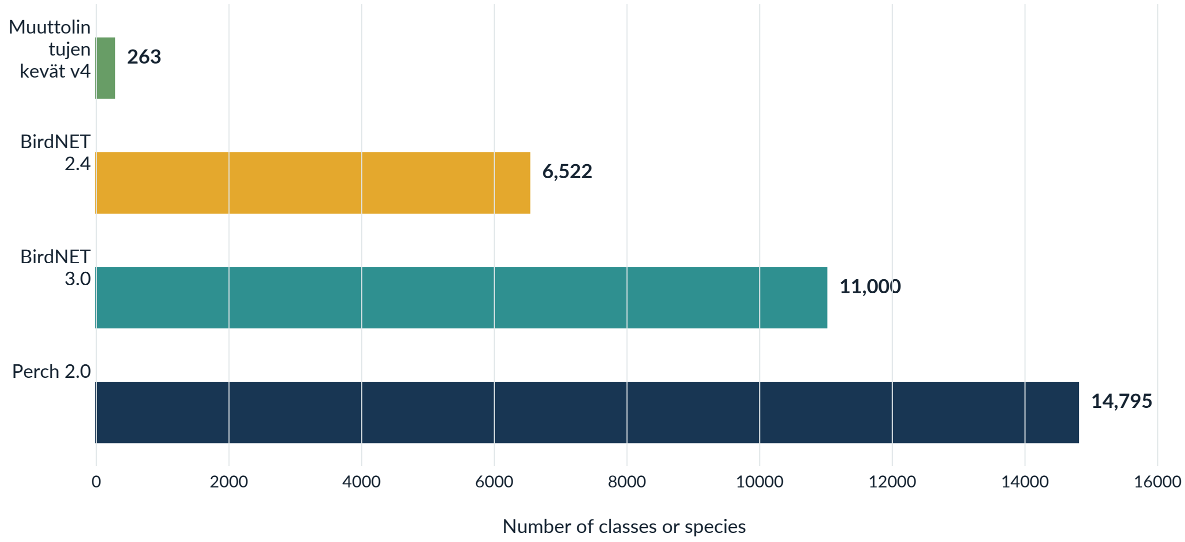

The class sets differ in size and content; the numbers should not be interpreted directly as an accuracy comparison.

Timeline of the tested versions

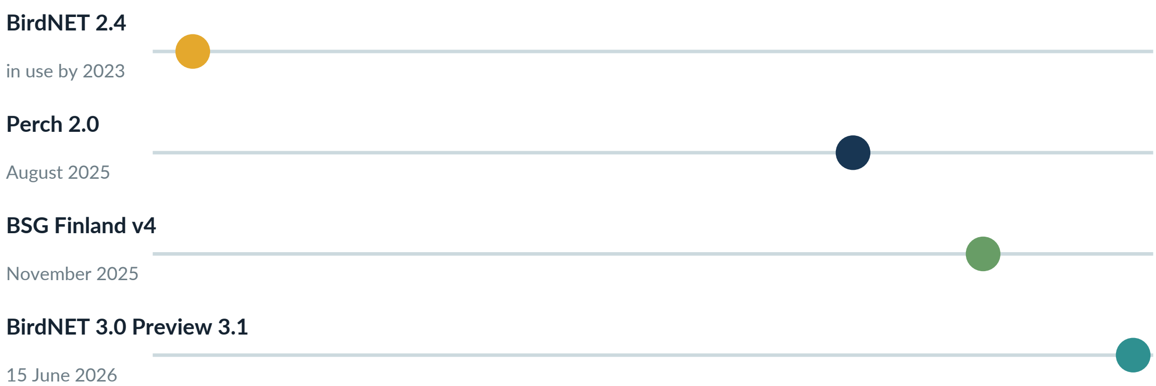

Muuttolintujen kevät v4

Muuttolintujen kevät v4, also known as BSG Finland v4, is a model specialised for Finnish conditions and created through the Bird Sounds Global collaboration. The Muuttolintujen kevät mobile application was developed at the University of Jyväskylä, and its current recognition set covers 263 Finnish bird species. [1]

Development of the Finnish model is part of the wider international Bird Sounds Global effort. The lead authors of the model-building publication include Patrik Lauha, Meeri Rannisto, Panu Somervuo and Mikko Heikkinen. [2-3]

Muuttolintujen kevät v4 is not a separate end-to-end network operating directly from raw sound

BirdNET 2.4 produces acoustic features, on top of which the Finnish classification head makes the final species decision.

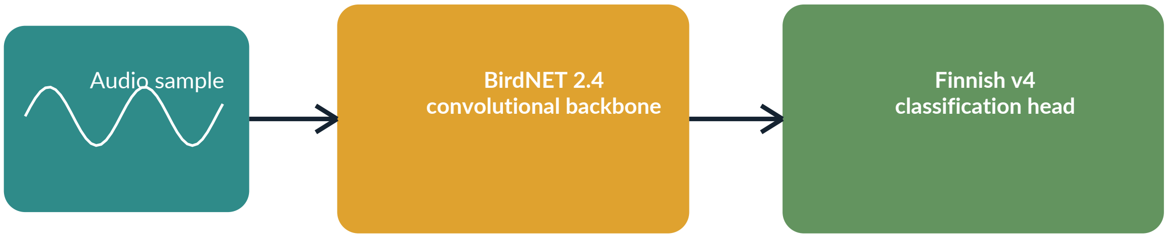

263 bird species + non-bird / human classes

The Finnish v4 model is a locally trained classification head on top of BirdNET 2.4.

Muuttolintujen kevät v4 is not a completely separate neural network that maps an audio waveform directly to a species result. The convolutional backbone of BirdNET 2.4 converts the audio sample into acoustic features useful for species recognition. A Finnish classification head trained on those features makes the final species decision. This is why using v4 also requires the BirdNET 2.4 base model. [3,5]

In this comparison, Muuttolintujen kevät v4 produced the most consistent overall results for local conditions. It identified Hooded Crow, Canada Goose and swallows or swifts better than the other models. Some Eurasian Jay calls were nevertheless included among the Hooded Crow results: one ordinary jay screech was classified as a crow with more than 90% confidence.

Strength and limitation

The local set of 263 bird species reduces implausible alternatives produced by global models. On the other hand, the Finnish model cannot recognise a species that is absent from its class list, and it does not provide equally broad coverage of other animal groups.

BSG Finland v4 is not the only local model built using this method. The Bird Sounds Global model family also contains versions targeted at other countries and regions. Only the Finnish model was used here, so this material cannot support conclusions about the performance of those other regional versions.

Possible future direction

If a comparable Finnish classification head can be trained on BirdNET 3.0 feature representations, the combination could benefit from both the newer global acoustic representation and local classification. This is a development hypothesis, not a result tested in this comparison.

BirdNET 2.4

The official release credits Stefan Kahl, Connor Wood and Holger Klinck as the authors of BirdNET 2.4. BirdNET is developed through collaboration between the K. Lisa Yang Center for Conservation Bioacoustics at the Cornell Lab of Ornithology and Chemnitz University of Technology. [4-6]

The model's class list contains 6,522 bird species. A separate occurrence model based on location and time of year can be used alongside it to reduce locally improbable alternatives. [6]

BirdNET 2.4 still identified a large share of the common-species core. Abundant and repeatedly vocalising species of forests, farmland and gardens often appeared similarly in its results and those of the other models.

In the July comparison, BirdNET 2.4 also produced many species assignments that appeared only once or for a very short period. This is not unique to BirdNET 2.4; the same phenomenon occurred with every tested model. The results are useful when analysis is restricted to repeated detections and place-and-time filtering is used.

Relationship to the Finnish v4 model

Muuttolintujen kevät v4 uses the BirdNET 2.4 backbone that produces acoustic features. Agreement between the two models is therefore not as independent a confirmation as agreement between two completely separate model families.

Hooded Crow: BSG v4 performed best

The presence of Hooded Crow was verified from the original audio. Muuttolintujen kevät v4 identified the species correctly on both comparison days, but BirdNET 2.4 scores on 8 July remained below the threshold used in the analysis.

BSG v4 performed best for Hooded Crow, although Eurasian Jay calls also entered its crow class: one ordinary jay screech was labelled as a crow with more than 90% confidence. For the global models, the main problem appeared to be confidence scores that were too low.

BirdNET 3.0

The tested Developer Preview 3.1 was released on 15 June 2026. It is the next model generation in the same BirdNET project as BirdNET 2.4. The author teams overlap: Holger Klinck and Stefan Kahl are credited on both releases; the other Preview 3.1 authors are Mario Lasseck and Maximilian Eibl. The development model covers about 11,000 species and includes more non-avian animal groups than earlier BirdNET versions. [7-8]

The developers state that the model, species list and software will continue to change before the full release. At the time of testing, the model version was only a few weeks old. [7-8]

BirdNET 3.0 was the most promising global model in this comparison. Despite being a development version and covering far more species than the Finnish model, it recognised a large share of the same common Finnish species as Muuttolintujen kevät v4.

Hooded Crow (Corvus cornix), a common Finnish corvid. Photo: Petri Suutarinen-Maaheimo.

Most obvious observed weakness: Hooded Crow

Hooded Crow was verified from the original audio, but BirdNET 3.0 scores were too low for the species to pass a practical general threshold.

BSG v4 detected it best. BirdNET 2.4 and Perch 2.0 also missed the verified event at their applied thresholds, while BSG crow results contained some Eurasian Jay calls.

The material is insufficient to determine whether the issue resulted from training-data coverage, partial similarity between crow and other corvid calls, confidence calibration or a combination of these factors. Because BirdNET 3.0 was still in development, species-specific performance may yet be adjusted before release.

BirdNET GeoModel

GeoModel filters the results of the global acoustic model

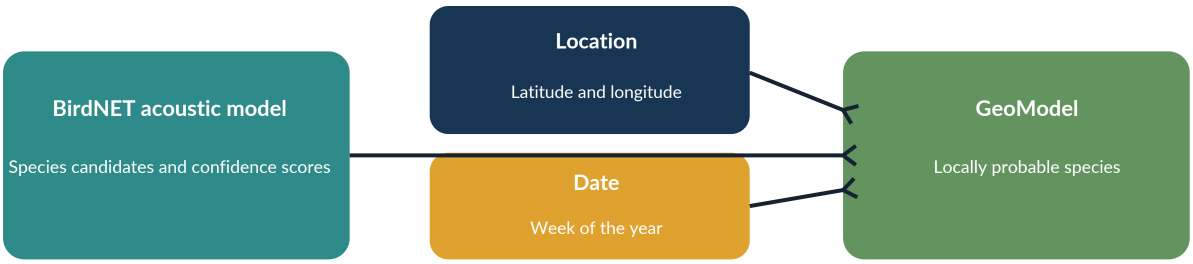

Note: a genuine rarity or unusual seasonal occurrence can be excluded by overly strict filtering.

GeoModel estimates which species are locally probable from the location and week of the year. [9]

GeoModel does not listen to the audio or perform the primary species recognition; it acts as a post-filter for BirdNET acoustic results. Overly strict filtering can nevertheless exclude a genuine rarity or a species occurring outside its usual season.

Perch 2.0

The authors of Perch 2.0 are Bart van Merriënboer, Vincent Dumoulin, Jenny Hamer, Lauren Harrell, Andrea Burns and Tom Denton. The model was developed in bioacoustics research at Google DeepMind and Google Research and released in August 2025. [10-11]

Perch 2.0 is a broad bioacoustic foundation model that provides ready-made recognition outputs and acoustic features for building new classification tasks. Its training data cover 14,597 animal species and 198 general sound-event classes. In addition to birds, the species coverage includes mammals, amphibians and insects. [10]

In the June comparison, Perch 2.0 identified much of the same common bird core as Muuttolintujen kevät v4 and BirdNET 3.0. The breadth of the model became especially visible when thresholds were lowered for selected classes.

At lower class-specific or sound-class-specific thresholds, the results could be searched for engine sounds and other interesting soundscape events. Perch 2.0 correctly identified a lawnmower through its Engine class. These results did not pass the normal general threshold because their confidence scores were lower than those of typical bird detections.

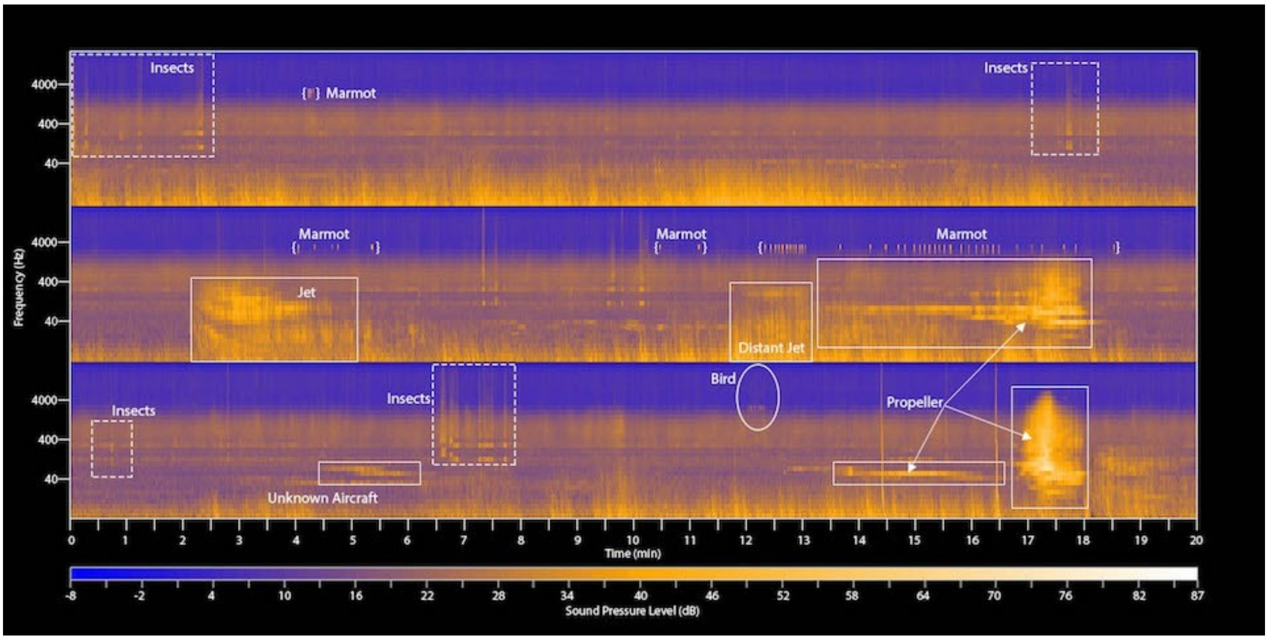

Example of a multi-source soundscape: insects, a bird, a mammal and aircraft are visible in the same spectrogram. Image: U.S. National Park Service, public domain.

A lower threshold is a search tool

Low-confidence outputs should not be interpreted directly as confirmed events. They can be used to locate interesting segments in a large audio archive and target review of the original recordings.

Perch 2.0 does not contain an internal place-and-time filter comparable to BirdNET GeoModel. In Finnish conditions, local plausibility therefore has to be assessed separately or through external filtering.

Hooded Crow was verified in the 26 June recordings, but Perch 2.0 did not produce an accepted crow detection at the applied threshold. This was a model miss in that test.

Beyond birds: other animal groups

BirdNET 3.0 and Perch 2.0 support recognition of animal groups beyond birds. Their class sets or training data include mammals, amphibians and insects, among others. Perch 2.0 also contains general environmental-sound classes. [7-8,10]

This makes it possible to use the same audio archive for nature monitoring beyond birds. Potential applications include activity periods of sound-producing insects, amphibian breeding activity, mammal presence and disturbance caused by vehicles or machinery.

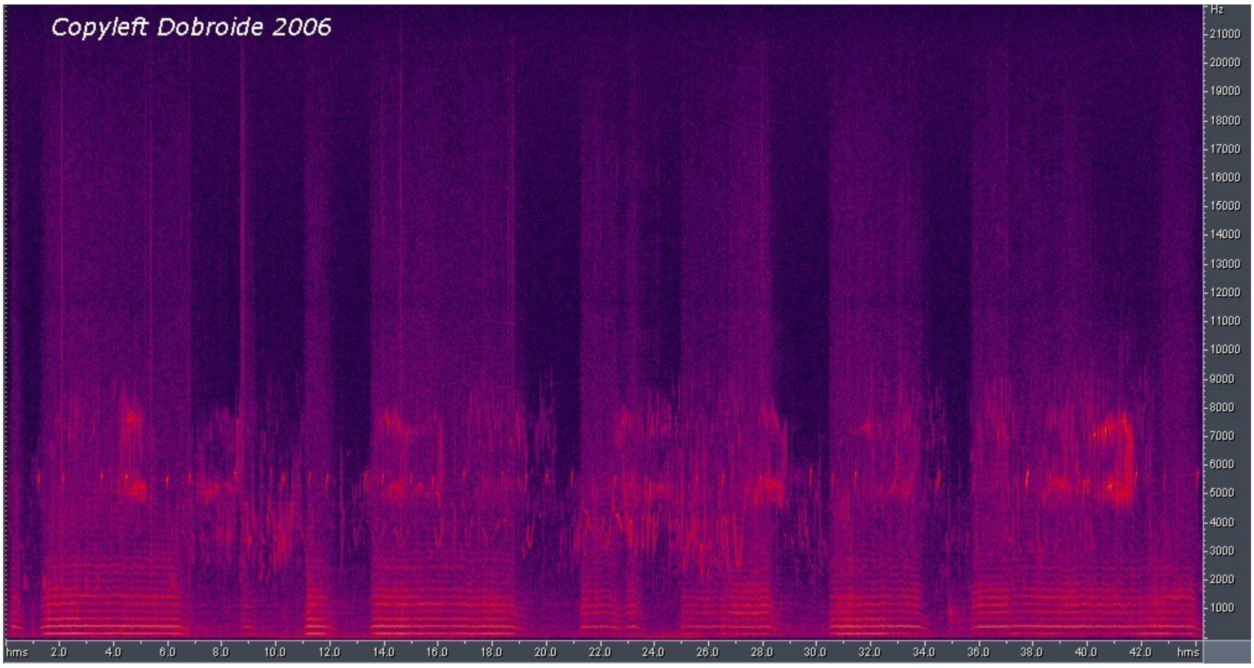

Bee flight sequences appear as acoustic traces in the spectrogram; bird song is also visible in the background. Image: dobroide / Wikimedia Commons, CC BY-SA 2.0.

Other possible uses

For sound-producing insects, acoustic monitoring can describe occurrence, daily and seasonal activity, and relative changes. Grasshoppers, bush-crickets and cicadas, for example, can be monitored automatically. However, the number of detections is not the same as the number of individuals: calling activity depends on temperature, weather, sex, behaviour and microphone detection range, among other factors. [12-13]

Vole population size cannot be estimated directly from an ordinary bird-monitoring recorder and model detection counts. Many small rodents produce high-frequency or ultrasonic sounds that require suitable microphones, and acoustic detections describe only animals that are vocalising. With specialised equipment, a species-specific detector and field calibration, acoustic activity could contribute to monitoring relative changes, but it would not by itself provide a reliable count of individuals.

General principle

An acoustic detection count becomes a population estimate only when species-specific call rate, detection probability, distance and the proportion of silent animals are known or estimated. [13]

Findings by species group

Results were also examined by species group. For owls, nocturnal singers, swallows and swifts, recurrence, time of day and checking the original audio were more informative than a single confidence percentage.

Owls: recognition varies with season

Season and call type affect owl recognition. In winter datasets, BirdNET and BSG v4 models have, in practical experience, often recorded owls accurately. A limitation is that the models do not distinguish male and female calls. In this summer dataset, owl candidates still had to be verified from the original recordings.

Species, 8 July	Models reporting it	Checked interpretation
Tawny Owl	BirdNET 2.4 + BirdNET 3.0	Included both genuine detections and misclassifications of human-made environmental sounds.
Long-eared Owl	BirdNET 2.4 + BirdNET 3.0	All checked detections were false.
Eurasian Eagle-Owl	BirdNET 2.4 + Muuttolintujen kevät v4	All detections in this dataset were false.

Conclusion: for owls, agreement between models, recurrence or a biologically plausible time of day did not by themselves provide sufficient confirmation. All checked Long-eared Owl and Eurasian Eagle-Owl detections were false in this material, while Tawny Owl included both true and false interpretations.

Nocturnal singers: European Nightjar stood out

European Nightjar appeared on 8 July in the results of both BirdNET 3.0 and Muuttolintujen kevät v4; BirdNET 2.4 did not report it. Muuttolintujen kevät v4 also reported Corn Crake on 26 June as a repeated nocturnal series. Both patterns were clearly more plausible than isolated owl hits.

Swallows and swifts: BSG v4 performed best

Muuttolintujen kevät v4 reported Common House Martin and Common Swift on 8 July, and Common Swift on 26 June. The detections were verified as correct. BirdNET 2.4 and BirdNET 3.0 did not produce accepted detections of these species, making BSG v4 clearly the most sensitive model for this group in the tested material.

Group-specific filtering

Owls: in this summer dataset, even repeated candidates were checked from the original audio. Nocturnal singers: support from more than one model or a coherent night-time sequence increased credibility. Swallows and swifts: the BSG v4 detections were correct and the model identified them best.

Data: hourly JSON results from 26 June and 8 July 2026, plus the day-specific results appendix.

Detailed data analysis

Hourly comparison did not use the number of detections as a measure of model quality because sample lengths and processing intervals differed. A table value indicates the number of hours in which a model produced at least one detection of the species. Because the recordings were made during the breeding season, calls, warning sounds and fledgling vocalisations were prominent alongside song.

The long tail of the data

Date and model	Species	Exactly one detection	Detected in at least 2 hours
26 June BirdNET 3.0	39	9	28
26 June BSG v4	44	11	27
26 June Perch 2.0	36	10	23
8 July BirdNET 3.0	52	17	32
8 July BSG v4	56	19	33
8 July BirdNET 2.4	35	10	22

Raw species count therefore does not directly describe accuracy. The larger result sets from BirdNET 3.0 and BSG v4 contained many species assignments that appeared only once.

26 June 2026: clearest model differences

Species	BirdNET 3.0	BSG v4	Perch 2.0	Interpretation
Spotted Flycatcher	10	18	6	BSG v4 clearly the most sensitive.
Eurasian Siskin	8	11	3	Perch 2.0 detected the species less often.
Fieldfare	5	11	1	Perch 2.0 detected the species in only one hour.
Eurasian Blackcap	6	4	9	Perch 2.0 was strong for this species.
White Wagtail	3	4	0	Perch 2.0 missed a species supported by the other two models.
Common Crane	2	2	0	Likely miss by Perch 2.0.
Hooded Crow	0	13	0	BSG v4 detected it best; some Eurasian Jay calls were included.
Canada Goose	0	8	0	BSG v4's repeated detection sequence was verified as correct.
Eurasian Bittern	0	0	6	All detections were false.

Perch 2.0's Engine class correctly identified a lawnmower in the same recordings. Broad animal and environmental-sound classes, and Latin species names, are coverage features rather than errors in themselves.

Detailed data analysis: 8 July 2026

The July comparison showed clear species-specific differences between BirdNET 2.4 and the newer models. BirdNET 2.4 still recognised a strong core of common species, but its hourly coverage was much lower for several common midsummer vocalisers.

Species	BirdNET 3.0	BSG v4	BirdNET 2.4	Interpretation
Willow Warbler	6	5	0	BirdNET 2.4 missed it completely; both newer models detected it during the same five hours.
Eurasian Siskin	15	16	3	BirdNET 2.4 detected only a small part of the sequence.
European Robin	10	10	3	Hourly coverage of the older model was much lower.
Eurasian Blackcap	8	9	1	Clear under-detection by BirdNET 2.4.
Goldcrest	6	5	1	Clear under-detection by BirdNET 2.4.
Great Tit	9	10	2	The difference was also visible for a very common species.
Willow Tit	3	3	0	A likely miss, as both newer models detected it.
Spotted Flycatcher	5	14	4	BSG v4 was again the most sensitive.
Hooded Crow	0	18	0	BSG v4 detected it best; the others remained below threshold.

Verified as correct
BSG v4 Canada Goose and swallow/swift detections, plus Common Cuckoo, Common Gull, Great Black-backed Gull, Common Linnet and Vulpes vulpes (Red Fox).

Verified as false
All Eurasian Bittern, Eurasian Eagle-Owl and Long-eared Owl detections, plus Ortolan Bunting, Rook, Common Eider and Barnacle Goose.

Tawny Owl required checking each audio event. The material contained both genuine detections and misclassifications of human-made environmental sounds. In owls, agreement across models or recurrence alone did not provide sufficient confirmation.

Interpretation of model capabilities

BSG v4	Very sensitive for local species. Hooded Crow, Canada Goose and swallows/swifts were clear strengths, although some confusion between similar species occurred.
BirdNET 3.0	Balanced global model and more comprehensive than BirdNET 2.4 for the common midsummer species set. Hooded Crow scores were too low.
Perch 2.0	Recognised environmental sounds in addition to common birds. Finnish species results showed both misses and a need to review outputs from the broad class set.
BirdNET 2.4	Useful for common species, but hourly coverage for several species was lower than with the newer models in July.

Not every isolated detection was checked. In this material, a single detection should generally be treated as suspect, but during migration a fly-over bird or flock - especially a wader - may appear in only one detection sequence.

Model strengths and limitations

Model	Key strength	Clearest limitation or note
Muuttolintujen kevät v4	Most consistent overall in Finnish conditions. Detected Hooded Crow, Canada Goose and swallows/swifts better than the other models.	Cannot identify species outside its class list. Requires the BirdNET 2.4 backbone. Eurasian Jay calls were also included among the crow results.
BirdNET 2.4	Recognises a large share of common Finnish birds. Place-and-time filtering supports result screening.	Oldest model generation in the comparison. Isolated results require strong filtering. Verified Hooded Crow was absent from the 8 July results.
BirdNET 3.0	Highly promising global model given its development stage. Recognised much of the same common-species core as the Finnish model.	Hooded Crow scores were too low. The model and confidence calibration will still change before release.
Perch 2.0	Very broad coverage of animals and sound classes. The Engine class correctly identified a lawnmower.	No internal place-and-time filter. Unusual sounds often require class-specific thresholds. Verified Hooded Crow was missed on 26 June.

BSG v4 detected Hooded Crow best. BirdNET 2.4, BirdNET 3.0 and Perch 2.0 remained below their applied thresholds, while BSG crow results also contained some Eurasian Jay calls.

Limitations of the comparison

- • The comparison is based on two 24-hour periods at one site, so the results cannot be generalised to every season and habitat.
- • The four models were not all run on the same day. Differences between BirdNET 2.4 and Perch 2.0 cannot be fully separated from natural differences between the comparison days.
- • Not all original recordings or all isolated detections were checked manually.
- • A single detection usually needs verification, although a fly-over bird or flock during migration - particularly a wader - may produce only one detection sequence.
- • A formal accuracy test would require expert-annotated audio, the same files for all models and species-specific optimised thresholds.

What can still be concluded?

Two parallel test days revealed the benefit of local classification, the broader use cases of global models, species-specific differences and the continuing need to check original audio.

Conclusions

Across the two comparisons, Muuttolintujen kevät v4 appears to have been the most reliable model at the time of testing for automated baseline monitoring of Finnish birds. Recognition of common and repeatedly vocalising species is already quite reliable when justified confidence thresholds are used and results are interpreted species by species.

The strength of the Finnish v4 model is the combination of BirdNET 2.4's acoustic-feature backbone and a classification head trained on Finnish data. The local species set reduces the number of false alternatives available to a global model. At the same time, the model cannot recognise species outside its own class list or offer equally broad recognition of other animal groups.

BirdNET 3.0 was the most promising global model in the comparison and represents the next generation in the same development line as BirdNET 2.4. Despite its development status and much larger species set, it identified a large share of the same Finnish common-species core as Muuttolintujen kevät v4.

BSG v4 detected Hooded Crow clearly better than the other models. BirdNET 2.4, BirdNET 3.0 and Perch 2.0 remained below the applied thresholds. Some Eurasian Jay calls nevertheless entered the BSG crow results, so good sensitivity did not eliminate confusion between the species pair.

BirdNET 2.4 remains useful for common and repeatedly detected bird species, although its age is visible relative to newer model generations. Perch 2.0 offers the broadest range of applications in the comparison and recognises other animals and environmental sounds in addition to birds, including the correctly identified lawnmower in this material.

Overall result

When interpreted correctly, bioacoustic recognition is a reliable way to monitor common species that vocalise repeatedly. Isolated, short or locally unusual detections should be checked from the original recordings. During migration, however, a single detection may represent a genuine fly-over bird or flock.

Recommended practical workflow

- • Use a threshold justified for the species and model rather than one universal threshold whenever possible.
- • Send isolated detections to a review queue; during migration, remember the possibility of genuine fly-over birds.
- • Use recurrence, continuity across several hours, and information about location and season.
- • Retain unfiltered results for later searches for rare species and other animal groups.
- • Check important, rare and detections on which the models disagree from the original audio.

References

[1] University of Jyväskylä. Muuttolintujen kevät - Research for JYU Mobile. Accessed 11 July 2026. https://www.jyu.fi/fi/tutkimus/muuttolintujen-kevat
[2] Lauha, P., Rannisto, M., Somervuo, P., Heikkinen, M. et al. 2025. Bird Sounds Global - model builder: An end-to-end workflow for building locally fine-tuned bird classifiers. DOI: 10.22541/au.176599468.88578746/v1.
[3] Lauha, P. 2025. BSG_classifier_builder: Train locally fine-tuned bird sound recognition models. GitHub. Accessed 11 July 2026.
[4] Kahl, S., Wood, C. M., Eibl, M. & Klinck, H. 2021. BirdNET: A deep learning solution for avian diversity monitoring. Ecological Informatics 61, 101236. DOI: 10.1016/j.ecoinf.2021.101236.
[5] Kahl, S., Wood, C. & Klinck, H. 2025. BirdNET Model V2.4. Zenodo. DOI: 10.5281/zenodo.15050749.
[6] BirdNET Team. birdnet: A Python library for identifying bird species by their sounds. GitHub. Accessed 11 July 2026.
[7] Lasseck, M., Eibl, M., Klinck, H. & Kahl, S. 2026. BirdNET V3.0 model developer preview, Preview 3.1. Zenodo. DOI: 10.5281/zenodo.20703646.
[8] BirdNET Team. BirdNET V3.0 Developer Preview. GitHub. Accessed 11 July 2026.
[9] Kahl, S., Mauermann, M., Lasseck, M., Wood, C. & Klinck, H. 2026. Spatiotemporal species occurrence prediction for post-filtering BirdNET acoustic detections. GitHub. Accessed 11 July 2026.
[10] van Merriënboer, B., Dumoulin, V., Hamer, J., Harrell, L., Burns, A. & Denton, T. 2025. Perch 2.0: The Bittern Lesson for Bioacoustics. arXiv:2508.04665. DOI: 10.48550/arXiv.2508.04665.
[11] Google Research. Perch: A bioacoustics research project. GitHub. Accessed 11 July 2026.
[12] Do Nascimento, L. A. et al. 2024. Time and habitat structure shape insect acoustic activity in the Amazon. Scientific Reports / open-access publication.
[13] Marques, T. A. et al. 2013. Estimating animal population density using passive acoustics. Biological Reviews 88:287-309.

Image sources and licences

- Cover and Hooded Crow photograph: Petri Suutarinen-Maaheimo. Original photograph. The cover image was cropped and graphic elements were added.
- Soundscape spectrogram: U.S. National Park Service, Mount Rainier soundscape, public domain in the United States.
- Bee spectrogram: dobroide, Wikimedia Commons, CC BY-SA 2.0.
- All other diagrams and infographics were created for this report.