

Projecting Climate Change Impacts on Wetland-Dependent Birds in the Prairie Pothole Region

John Ossanna: To start things off, please join me in welcoming Laura Thompson with the National Climate Change and Wildlife Science Center who will be introducing our speaker today. Laura?

Laura Thompson: Thank you, John. I'm very happy to introduce Helen Sofaer, a speaker at our center's webinar series. Helen is a quantitative ecologist who studies the impacts of global change, including climate change, land use change, and species invasions.

After receiving her PhD in Ecology from Colorado State University, she worked with Drs. Barry Noon and Susan Skagen on a North Central Climate Science Center funded project assessing climate change impacts in the Prairie Pothole Region.

She is currently a Mendenhall Postdoctoral Fellow at US Geological Survey's Fort Collins Science Center. Thank you, Helen, for being here.

Helen Sofaer: Thanks so much for having me today, and thanks to everyone for tuning in. As she said, I'm going to talk about work that my collaborators, Barry Noon and Susan Skagen, Valerie Steen, Curt Flather and others did, projecting climate change impacts on wetland-dependent birds in the Prairie Pothole Region.

Let me make sure I can advance this slide.

I want to talk about how we can use modeling to inform management in this era of global change. Many management decisions come down to prioritizing what to do with limited resources.

Where on the landscape to invest in conservation or management and which species to actively monitor or manage for. The reality of climate change is that these decisions have to be made despite uncertainty about the future.

One theme that cuts across all the work I'm going to present today is attempting to quantify and address the uncertainties that are inherent to any modeling effort.

I want to start with the spatial prioritization. In this context, I modeled the distribution of wetlands across the US part of the Prairie Pothole Region in the Dakotas and Montana. The goal was to project the impacts of climate change.

As many of you likely know, these wetlands are the most important breeding habitat for harvested waterfowl in the US. They also provide habitat for many other wetland-dependent species. This plot shows the average wetland densities along segments of the Fish and Wildlife Services' May pond count survey.

Densities are highest in a region called the Missouri Coteau, which is east of the Missouri River. The Prairie Pothole Region boundary is shown with that dotted line that's not the state boundaries.

I used wetland counts from this dataset to model the responses of the system to land use and land cover and also to weather. I did this with the random forest machine learning method and also with Bayesian spatial models.

In general, I'm not going to get into too much of the modeling details, but I'm happy to answer questions about that for anyone that's interested.

We know that the number of wetlands varies in space, but it also varies in time. This is a dynamic system. In wet years, there can be millions more ponds, which are shown in the dashed line, than in drought years.

Duck populations actually track that variation as well. That's the solid line on this graph. Because the system is so sensitive to weather, that's why there's so much concern that climate change could affect both the distribution and number of wetlands and in turn, the avian populations that rely on them.

The concern in this system is a general one. Basically, will the lands in the current network of conservation, the current reserve network, continue to provide suitable habitat in the future?

In particular, in the system, research by Carter Johnson and his colleagues suggested that the central portion of the Prairie Pothole Region, which is where the highest wetland densities are now, could become too dry to host productive wetlands in the future.

That work considered the potential effects of changes in temperature and precipitation on vegetative cycles in wetlands that basically underlie wetland productivity. In their models, the highly favorable conditions are shown in red.

With only changes in temperature and not compensatory increases in precipitation, they projected that the most productive habitats would move to the east. One of the assumptions in this work was that wetland basins are evenly distributed. Also, that climate change will occur in a uniform way across the landscape.

Those were things we wanted to consider in our analysis, both climate change and also realistic patterns of land use and land use change.

Oh, this one came out terribly. Because water balance is so key to the dynamics of the system, basically a major part of our work was understanding what the different climate models had to say about the projected future water balance.

We can't really see much here, but basically what this was trying to show was precipitation minus potential evapotranspiration. The red is drier conditions, and blue would be wetter conditions.

Obviously, this is very fuzzy, but the takeaway point is each one of these plots refers to a different model. The names are jumbled, at least, in the version I can see, but some of the models projected overall drying and some projected overall wetting.

The models actually agree that it's going to get warmer, but some project enough increased rainfall to compensate for that. This climate variable was the most important predictor in our models. You can imagine that what we found was that the drier models predicted fewer wetlands and the wetter models predicted more.

Let's see. Oh, man. This one's blurry, too. The citation at the bottom is where these figures are from. What we found was that wetland densities were projected to decline overall, but not to shift in space.

I'll [laughs] walk you through what you would be seeing here. What I want to emphasize is that climate mattered and land use mattered, too. The top row here shows projections to the historical simulation of the climate models. It's the average across 10 climate models.

In the middle row are projections to future simulations. In the bottom row is a comparison between those so that that bottom one gets at the impact of climate change. You can see the red diagonal line cutting across the bottom outermost two figures.

The columns here are different land use scenarios. The left one is the current pattern, basically assuming the future mirrors present.

The one that's pretty similar to that on the far right stems from a collaboration with Ben Rashford, who's an agricultural economist and who developed projections of land use change under these same climate models.

Basically, the same patterns of those two outer scenarios are the most realistic. They both show that red line, basically, right where wetland densities are the highest now. That red corresponds to lower wetland densities.

That's where we get the understanding that they're projected to decline, whereas the upper two rows...Again, I apologize for this.

I don't know if you can make out that the most blue is in that same place, so the most wetlands are still in the same place. The middle column corresponds to a scenario with uniform land use and a uniform distribution of wetland basins.

That's actually most similar to the assumptions underlying Carter Johnson's work. I'll say also that I view this work as complementary to that, that managers will want to know about both wetland numbers, which is what our work informed and also wetland condition, which is what that work primarily informed.

Let's see if we can move on, [laughs] and, hopefully, get to some figures that are useful. Woohoo! All right. One of the approaches that I just want to mention briefly is that what we did is we captured the whole distribution of projected changes within climate models.

That's why we had separate predictions to each model's historical simulation and each model's future simulation. We call that the model space method.

I had the great fortune in doing this work through the Climate Science Center to link up with the climate scientists themselves, and in particular with Joe Barsugli, who taught me a ton about assessing climate change impacts.

The model space method, we call it that because it asks what happens to the ecological impact in the world of the climate model. In contrast, most ecological climate change studies are based on the Delta Method. They assume that extremes are going to change at the same rate as the mean.

This is baked into datasets like WorldClim and ClimateNA. The model space method allowed us to take more information from the climate model about how different parts of the distribution were projected to change.

This is important because extremes can often be projected to change faster than climate means. This figure shows the projected change in the hottest July day that would be expected in a 10-year period, just for one climate model. Don't read into the precise spatial patterns here.

For this model, the Delta Method effectively assumes that this measure of a climate extreme is going to change at the same rate as the mean of July maximum temperature.

The climate model actually projects a much faster rate of change in the extreme in a lot of the Midwest. As you can see, it's a difference of five degrees C or something like that. It's a pretty big deal. I learned so much from Joe that I wrote a paper about it.

It's about how can we do these climate change impact assessments in a way that matches the drivers of our study systems. And when we know that extremes often are important to ecological systems, how can we design climate change impact assessments to take more from the climate models in terms of what they project about changes in extreme events?

Now, I want to shift gears to talk about prioritization among species. I'm going to break this into three different questions and spend just a brief overview on each one.

The first is, are waterfowl good surrogates for other wetland-dependent birds? Next, I'll talk about which wetland-dependent birds are most vulnerable to climate change. Lastly, I'll talk about more recent work that asks how reliable are our methods for projecting range change.

Again, a theme that cuts across all three of these questions is how do we deal with uncertainty and especially the uncertainty that arises from the modeling methods themselves.

In this system, there are five species of upland-nesting ducks that are the focus of a lot of current management. We wanted to quantify how well these ducks serve as surrogates for other species. What we're seeing here are the five species. Then on the map is what's often referred to as the thunderstorm map.

Basically, it's a map of the density of waterfowl breeding pairs. Again, you can see that red in the Missouri Coteau, which is this region that's been the focus of a lot of investment. I've overlaid

breeding bird survey routes in black, which were used to get abundance data for other wetland-dependent birds.

It makes sense that waterfowl could act as surrogates for other species because just as I showed earlier, duck populations track pond counts. So do the populations of some other wetland-dependent birds of interest.

In addition to that, the waterfowl can act potentially as umbrella species to the type of surrogate species because they use a variety of wetland habitats.

For example, they generally breed at highest densities in upland surrounding smaller wetlands. Then, when they're raising their broods, they move to larger, deeper wetlands that still have water later in the breeding season.

Because these species share this critical resource and a lot of the investments have basically been aimed at essentially getting easements to prevent draining of wetlands, that's something that benefits these other species. There has been work testing that and showing that.

A common method for selecting surrogates is to first cluster. Take the species you're interested in, in this case, wetland-dependent birds, cluster them into groups, and then select the representative from each group as a surrogate.

The goal here is that that can make the surrogate set be comprehensive in that all the species in the species pool are covered by one of the surrogates, i.e. grouped with it. Also, complementary in that each surrogate species represents different species in the species pool.

In practice, clustering is sensitive to the methods that are used. As part of this work, we generated species clusters based on trait data and abundance data. We used different sets of traits, and we used abundance data at different scales. For example, each breeding bird survey stop were summarized at a coarser scale.

Those are essentially the input data that go into a metric of dissimilarity and in turn into a hierarchical clustering method. Depending on the decisions that are used in the clustering, you can get a different set of relationships among species.

Here, we get essentially a tree depicting these relationships. Each letter at the tips of the tree represents a hypothetical species. In addition, the red lines represent the decision of where do we cut the tree, which essentially corresponds to how many groups do we want to create.

Each of these choices leads to a binary in versus out view of which species are grouped together, which species will be a surrogate for another. These results in groups can qualitatively vary a lot.

I'm not going to get into all these different groups. What we did in this work is we developed an ensemble across all these different methodological choices. One of the advantages of that is that it gave us a continuous perspective on the strength of surrogate relationships. At the same time, it encompasses all that uncertainty that arises from the input data and the clustering method.

We are visualizing here this ensemble in a network, and each species is a circle. I apologize for using the bird codes which only the birders will understand, but the five focal waterfowl are shown in gray, so they cluster together.

That's how we can interpret these network plots is that the closer the circles are to each other, the more similar the species were. In this case, only abundance data went into this particular network. This shows similarity in patterns of relative abundance.

This pattern that the five focal ducks clustered together was consistent across all of our both input trees and also showed up prominently in the ensembles. We see then this gradient in the strengths of the surrogate relationships, and the ducks have the weakest relationships with some songbirds.

There's some examples, like the sedge wren and the willow flycatcher in the upper left, and the strongest relationship here with the sora and also with other waterfowl. One of the advantages of making this continuous is that instead of having this binary view, it actually lets a manager decide what level of surrogacy is good enough, given the context and the resources.

Next, I want to talk about a work that was led by Valerie Steen, who did her PhD with Barry Noon and Susan Skagen and is now a postdoc with Morgan Tingley.

She did a nice study where she fits species distribution models to project range change in wetland-dependent birds in this region. She rigorously varies different modeling inputs and modeling decisions and looked at the implications for a range change.

Again here, there's a ton of uncertainty that just comes from the modeling methods as well as from the underlying climate model.

This is an example for the mallard, where for one set of decisions shown in the top row, each decision set is shown in a row. Basically, for different sets of decisions, you get similar current mapped predictions about the current distribution. The left-hand on top and bottom look pretty similar.

Then, these different decision sets, which include basically the climate model, the set of covariates, how they were thresholded from other modeling decisions, give different projections about how much range loss is expected.

This is a common feature of using species distribution models to project the impacts of climate change that often models that give pretty similar predictions in the current time period can be pretty divergent in terms of what they project for the future.

Valerie did this across a set of wetland-dependent birds, for each species, which, again, are these bird codes, there's this box plot that shows the variability in these different both input datasets, climate projections, and modeling decisions.

What they're showing variability in is the range change index, which is basically an index of how range size is going to change in the future where zero is no change and negative values are loss and positive values are gain.

We see there's a lot of variation within a species. The species here are ordered from the sora, which was projected to be the most vulnerable, i.e. to have the greatest loss on average, down to the great blue heron, which most of the decisions that's led actually to projections of gain.

One of the things that we wondered about is how stable is this ranking of vulnerability, how much can we rely on that ranking. For example, I showed you that some of the climate models project a wetter future and some of the climate models project a drier future.

It could be that the dry models are essentially contributing to these decision sets in which more loss is projected. Then the wetter models are on the right-hand side where they're associated with less loss or with gain.

If essentially varying these decisions gives a similar order of vulnerability, then we can be a lot more confident in that order. What we looked at is the vulnerability ranks across these decision sets. What Val found was that the sora was consistently projected to be the most vulnerable, and the great blue heron was consistently projected to be the least vulnerable.

The relative position of the other species depended on the method. The reason this is so important is it tells us at what level can we feel more confident in the information that we're providing to decision makers.

This highlights the coarseness of that level. Here, we can say that the sora rose to the top of this vulnerability assessment, but for most species there's so much modeling uncertainty that it's hard to rank them according to their vulnerability.

You might recall that the sora was one of the species that was expected to benefit most from management of waterfowl because it clustered pretty closely to those five focal ducks. That's both good news and also something to validate in future work.

The last piece of research I want to touch on is a more recent validation study asking how reliable range change projections are. Ideally, we'd want to see a nice positive relationship between projected and observed change in range size.

In a way, for these kinds of things we care probably more about these ranks. Can we be focusing on the right species and maybe less about the absolute numbers? These models are useful if they give us the right ranks.

This is work I did in collaboration with Curt Flather and Catherine Jarnevich. What we did is we estimated species distribution models based on Breeding Bird Survey data from the late '70s. Then we developed predictions to recent years.

What I'm showing is a map of the Breeding Bird Survey routes that went into this analysis. Actually, both what are labeled here as estimation and validation routes, both of those were used in estimation, and the validation is in time not in space.

These validation routes are routes that were surveyed in all six years that we used the data for. We knew how occupancy changed at those routes. We did this analysis basically in a way that if a species was present in one of the three years in each time period, then we counted that as present.

We only counted absences when the species was absent in all three years and when the route was surveyed in all three years. These validation routes were surveyed in all six years. There's 512 of those across the country.

What we did is for 190 songbird species, we estimated an ensemble of species distribution models based on biocovariables and pretty coarse metrics of land use. I want to highlight the contrast between this work and Val's work, where her study targeted the covariates towards things that were important for wetland-dependent birds, things -- for example, hydrological variables -- that we knew were important in the wetland model and things like that.

Whereas here, the covariates are totally generic across species, they're always the same. We had different sets, but they were always the same. They're not related to species ecology, this is a test of a more automated type of species distribution modeling.

Again, we estimated models using data from the late '70s, and then predicted to this recent time period.

What we found was that our models looked great in this recent time period. In terms of the model validation, it suggested that there's good transferability in time, which is a challenge for species distribution models. Actually, most validation studies often stop here.

What I'm showing is a histogram where the height of the bars represent the number of species, and AUC is the Area Under the Receiver-Operator Curve. It's one of the most common performance metrics in species distribution modeling. It's a metric of discrimination performance. How well does the model separate between the zeroes and ones?

As we see here, most species have an AUC above 0.9, which is considered generally excellent. This shows that, overall, the models are doing well on this validation data from the current years.

What we found was that the change in range size was not reliable at all. There is effectively no relationship between observed and predicted range change, and the vulnerability ranks. I'm better at darts than this right-hand [laughs] plot shows. It's pretty bad.

What we're looking at here is, on the left, we're looking at this plot of range change where -- just like in Val's work -- zero is no change in range size. Here, range size is represented by the number of occupied routes. It's not gridded cells across the landscape, it's just at the breeding bird survey routes.

Each point on this figure represents a species. The dots are bigger for the more common species, and they're smaller for the rarer species. The species that tripled its range size was actually the house finch -- this point on the far right -- and the models didn't predict that.

On this right-hand side plot, what we're seeing is a ranked version of this, where 1 is the most vulnerable and 190 is the least vulnerable. We're, again, looking at what was observed, in terms of loss of and gain of, in the number of routes that were occupied over this about 35-year period versus the prediction.

This was a head-scratcher for us, at first, because I just showed you in the last figure that we had this great overall performance. Yet, here I'm showing there's basically a total inability to predict which species will be most vulnerable.

I should mention, this includes land use change, at least, coarsely. We expect that might have been the greatest driver of distributional changes on the past few decades. Of course, we know more about how land use has changed in the past 30 years than we know about...When we try to start looking forward, there's a lot more uncertainty.

Basically, to explain this mismatch between good overall performance and yet terrible ranking of vulnerability, what we found was that there was actually poor prediction where change occurred. This collapses across all species, across all the validation routes.

What we see is that there's four potential categories we're referring to here as the "root status." A species could have gained -- could have not been there historically, and been there now -- could have lost occupancy on the route, could have always been there, or never been there.

You can see that most places are always or never occupied. There's a relatively small number of routes that change status. If you look within the always vertical bar here, it's mostly green, which corresponds to an always predicted status. Basically, the models are getting that right.

Similarly, if you look within this biggest bar, the never bar, it's mostly blue. That corresponds to this never predicted status, so the models are getting that right, too.

The models are doing well for places where the species is always or never there. When we look at the places where range expansion or contraction was observed, the models actually do badly at that. They usually, again, predict that the species was always or never there.

This is [laughs] a bummer, because the gain and loss is what we care about the most. We want to know where are the locations that are vulnerable to range loss. We want to know where are the locations that are becoming potentially suitable in the future. Precisely, what we want to know is what the models do badly.

It raises challenges for our ability to both understand relative vulnerability and also to link that vulnerability back to spatial planning, to say, "OK, where in the landscape are we going to change our management strategy to benefit a species that we expect to be vulnerable?"

Actually, I want to stop there to leave plenty of time for questions and challenge both the research and the management communities to think about, "How do we support these prioritization decisions under climate change, land use change? What kinds of information and ask, "What kind of resolution and level of confidence can we, as researchers, provide?"

Given our limitations, from a manager's perspective, what would be useful, despite the uncertainty? The fact remains that somebody still has to make hard decisions about where to invest resources and with what goal.

What my work tries to get at is to understand the uncertainty behind the science that can support those decisions. Then, hopefully, move towards supporting the decisions despite the uncertainty or acknowledging, in some cases, that we might not be able to say much with confidence.

With that, I'd like to acknowledge funding from the North Central Climate Science Center and also from the USGS Mendenhall Postdoc Fellowship and thank the many collaborators, without whom this work would not have been possible.

John: I would like to thank Helen and USGS, once again, for continuing this webinar series. Love having you guys as a partner. I'd like to thank everyone who was able to view the webinar.

[audio skips]

Helen: I showed a few different examples, thinking about, "What can we say with confidence?" It would be interesting to hear people chime in. I guess I'll answer...I saw one question pop up here. The question is, "Do you think it's possible...?"

John: Sorry. Do you think it's possible to use your earlier results to map climate change refugia for the prairie potholes? Sorry.

Helen: I'd say, "At what scale?" I'd say that the inputs to the pond models are already at a relatively coarse spatial scale, so they correspond to segments -- like the number of wetlands over a segment that was flown. At the same time, they're useful, but not at this kilometer here versus that kilometer there.

Probably, we tried to map refugia using GCM projections at a finer level than...There's real information within a GCM. In a way, it can be almost more useful to look at what's on the landscape now. Where is there an inversion now, and what can that tell us about where refugia might be? Things like that.

In the prairie potholes, what that corresponds to is, firstly, where are there areas where there are wetlands that, essentially, it's often referred to as wetland complexes that encompass a variety of depths?

The shallower wetlands can provide important habitats, especially for species like pintail, which are of concern because they've been declining. Of course, deeper wetlands are more likely to have water there and under drier conditions.

By looking at wetland complexes and bet hedging, in terms of water depth as much as in space, that's probably a useful method that actually aligns pretty closely with the current management strategy.

The other thing that popped out of this work was the prairie Coteau was a region, I'd say, that that was where there was both good alignment of current habitats with less negative projections for the future. That could be another area, generally.

Then within that general area, managers, of course, have to go and make decisions about what land and wetlands are in good condition.

Let's see. John Gross asks, "If some of the results aren't pessimistic -- since I assume all models are equal -- surely some are better than others?"

I want to believe these results are pessimistic. That paper actually just came out this past week online. If you look, there's this long supplement where we tried all these different things. Basically, we're trying to save it like, "I want to be able to do this."

We tried things like using the optimal calibration for each time period. Which, of course, you would never know how the optimal threshold for moving from continuous to binary predictions, would change in the future. You never know that. For this validation study, we did know that. We tried things like that, which helped a little bit, helped for common species.

In terms of this assumption that all the models are equal, the model for each species is actually an ensemble across several methods. Each method was weighted by...It's AUC. There's some attempt built into this ensemble to encompass variation model quality, like models that had an AUC lower than 0.7 didn't contribute to the ensemble.

Thinking about the work that Val showed, where some models can have similar predictions to the current time period, and yet divergent predictions for the future. It can be hard to tell which model's better, if they're all correlative models. It's ideal if we can have more of a mechanistic sense of the drivers and build careful models for a specific species of interest.

For this type of correlative modeling, sometimes it can be hard to use validation in withheld data to separate among these models that then lead to pretty different outcomes. Actually, that's true of the clustering, as well.

We validated our clusters, and we get a similar r-squared variant explained from different clusters. Yet, those clusters look different, in terms of which species are grouped together.

That's, again, one of the advantages of ensembling is that, when we're not sure which method is better, a priority. We can't make that decision up front, we can encompass these outputs from multiple models.

Let's see, "Could I see this methodology being applied to other regions of patchy, ephemeral wetlands, like the playa lakes?"

I'd say, absolutely. In the last few years, honestly, there are opportunities to do cool modeling. There's now remotely-sensed information. I think it's on a monthly level?

Google Earth Engine has surface water products, so there's the potential to use things like that to say not just, "Where are the wetlands in May?" which is all we did but, "How long are these wetlands holding water on the landscape?"

In the playa lakes area, they're often smaller wetlands. We know that there's some minimum time that these wetlands have to have water to line up with the phenology of breeding. There's this opportunity to scale this up pretty broadly and make some of those connections.

John: Thank you again, Helen. If we have anyone who types one while I'm saying my goodbyes. [laughs] I would like to, once again, thank everyone who participated. Thank you, everyone who threw out questions out there. Thank you, Helen and USGS, for continuing and moving this webinar series forward.

Helen: Thanks, everyone, for tuning in.

John: Have a good one.

Helen: Alright.

Transcription by CastingWords