KI - Künstliche Intelligenz (2023) 37:99–105
https://doi.org/10.1007/s13218-022-00791-5

INTERVIEW

GeoAI as Collaborative Effort

Interview with Devis Tuia

Kai-Florian Richter[1] · Simon Scheider[2] · Devis Tuia[3]

Received: 9 December 2022 / Accepted: 13 December 2022 / Published online: 20 January 2023

Devis Tuia is an Associate Professor at Ecole Polytechnique Fédérale de Lausanne (EPFL), where he heads the Environmental Computational Science and Earth Observation (ECEO) laboratory. Devis is an expert in earth observation and remote sensing research, machine learning, and image processing. His current work includes interpretable deep learning in environmental modeling, human-machine interaction in remote sensing, and digital wildlife conservation, i.e., using imaging to automatize censuses and conservation efforts.

✉ Kai-Florian Richter
kai-florian.richter@umu.se

Simon Scheider
s.scheider@uu.nl

Devis Tuia
devis.tuia@epfl.ch

1 Department of Computing Science, Umeå University, 901 87 Umeå, Sweden

2 Department of Human Geography and Spatial Planning, University Utrecht, Princetonlaan 8a, Utrecht 3584 CB, Netherlands

3 Environmental Computational Science and Earth Observation laboratory, Ecole Polytechnique Fédérale de Lausanne, EPFL Valais Wallis, Rue de l'Industrie 17, Case postale 440, 1951 Sion, Switzerland

1 Personal Questions

KI: Devis, let's start with how you ended up where you are now. You have worked at various places around the world; mostly in Switzerland, but also in (or with) South America and groups in Spain, the US, and the Netherlands. You have held prestigious Swiss SNSF Ambizione and Professorships grants that funded your own position and research groups. Ever since your PhD studies (at least), you have worked on issues related to data extraction and processing, mostly in the context of remote sensing and imagery, employing and developing AI and machine learning techniques. How has your biography shaped your views on the field of AI?

In a way, I stumbled into it almost by mistake because, originally, I was studying geography and back then I was very unaware of machine learning and AI. And one day I ended up in an advanced statistics course and I realized that this is what I wanted to do in my life. And then little by little I started climbing that very big hill of catching up on the techniques side but, on the other hand, I had all the background in geospatial data and geospatial thinking that I was bringing along with me from my studies as a geographer. So, I boosted my profile with a lot of technical skills over the years, which came complementing this core that was more thematically oriented. And I think that is also what is shaping my view, especially now that I am more settled and I come back to my former loves of environmental science and environmental engineering. I am now bringing back all this background that I have been looking at for several years.

KI: Your work is located in an area between geographic information science (GIScience) and AI. In the last decades, there have been several important technological developments, which have also particularly influenced GIScience:

(1) an ever increasing availability of computing power; (2) the availability of large amounts of "geographic" data, both official (e.g., satellite image repositories) and from (geo-) social media and crowdsourcing, along with machine learning techniques for natural language processing (NLP), which has fuelled work in geo-parsing and geographic information retrieval (GIR); (3) deep learning techniques have been broadly adopted to extract knowledge from remote sensing images and other imagery sources. How do you think your career has echoed these developments?

I have always loved nature; studying nature, understanding nature in all its different forms, from the natural world to the built environment. Research without a nature component would not be for me. And yes, it is true that we now live in an an era of plenty. When I was doing my PhD, we had only few images we could work with and had free access to. Landsat[1] was just starting to become completely available. And then there was suddenly this kind of golden triangle that appeared between the data, the computing power, and the exciting problems. They all arrived together. And since Earth images lend themselves well to machine learning algorithms, the field became immediately very exciting. And as the cherry on top, there are all these other sources of data, such as social media, which made image analysis even more challenging. Now you can get first-hand opinions of people, their experiences and perception, data that we did not have before. In my opinion, this widened the palette of things that we can do with all this geospatial data. Because previously—and I don't mean this in a negative way—we were a bit forced to work on land cover because that is what we had. But now there is an almost infinite amount of problems that we can tackle with(in) remote sensing and the community is responding to this.

KI: Have the problems that the community addresses now always existed and you just could not address them previously? Or has the awareness of all this new data and possibilities shaped and opened new problems that the field had not even thought about before?

I think it is both. The problems existed before and we did not have the means to address them. And at the same time, having now the means just unlocked new ideas and new possibilities that before were not present by lack of images, by lack of computational resources, or by lack of people from AI interested in our problems. This cross-skills type of communities is emerging now. I am lucky enough to work in at least two of these; digital ecology and geoAI. And I see similar patterns emerging in them: instead of blocking new ideas and ways of doing things, one tends to synergise with people from the domain, who may be rather non-technical in an AI sense. That is where the real magic happens, at least the way I see it.

KI: You have applied your work to 'classic' remote sensing questions, such as land cover classification (e.g., detecting blue ice in the Antarctic [1]), *but also and in particular in recent times, to other disciplines and domains, such as quality of life, sustainable development* [2], *and wildlife ecology* [3]. *What challenges do you see, or have you encountered in such interdisciplinary work? What are the benefits?*

The benefits are infinite in the sense that you can unlock something that is really useful for society. In our work in animal ecology, for example, I have the feeling I am doing my little part for at least monitoring and providing information to decision makers for the biodiversity crisis. And this is something I always wanted to do. But somehow I think I had to go through this learning phase of really mastering the tools first. And also at some point you need to get over your preconceptions. For some time, especially during my PhD, I was living in this world of "remote sensing because of remote sensing."

And then with time you start to understand that you can do so much more. But to get to that understanding you need to go through this first phase. And I think that moving around geographically helps a lot because at every new appointment I got exposed to communities that were completely different. For example, in Wageningen I was working a lot with animal ecologists, which I had not done in Zürich. And not because there are no animal ecologists in Zürich, it just did not happen. There, I was more close to people working on phenology. And also with the group of Ross Purves[2]; we were working on gazetteers and then ontologies. These experiences keep building up, and so you end up understanding the problems of a field and bringing in your own background. I think that without such discussions it just cannot work. And now at EPFL, it is even more so because I am part of the ALPOLE center[3], working on alpine and polar environments. I also work on coral reefs because we have a number of people who come with this eagerness of working together. And it is not like they say "just do your magic with machine learning." It is all about "can we do this together and learn from each other?"

KI: As a researcher you have gathered a lot of practical experience in doing geoAI and remote sensing research. What have been important practical insights for you? Have you come across some lessons learned? Are there "things you believe but cannot prove?"

1 https://landsat.gsfc.nasa.gov.

2 https://www.geo.uzh.ch/en/units/gco.html.

3 https://www.epfl.ch/schools/enac/research/environmental-engineering-institute-iie/alpole-en/.

Most importantly, you cannot make it alone. If you want to have real impact and really want to do something for society, you cannot do it only while sitting in front of a computer, you cannot do it only in the field. You need this synergy, getting the data, scaling it up, understanding the real implications and bringing it to practice. I realize how difficult this is. It is not done in a few months, you don't do it just in one project. It is something you create by working on several fronts, one being the scientific work, one being maybe a bit more lobbying. You need to get the right persons together.

For example, we have this project in the Red Sea for coral monitoring[4], and we are working with a team, where we have hydrologists, geneticists, ecologists, and us in AI. And we have a good part of the team dealing with diplomacy. Because when you go to the Red Sea, the political situation you work in is not easy. It is not simply doing some machine learning to see which corals are doing well. It is like a big pile of threads you are trying to disentangle. Of course all this work results in nice technical papers about computer vision, and it increases our knowledge about aquatic ecology, but it is also very much about practical questions, such as how to find a generator for the boat. All this interconnects and is complicated. And I believe that geoAI has a strong role to play here as a connector. I think the geographical information aspect almost serves as a kind of spider web where all these disciplines can finally connect with each other. Which is fantastic!

KI: So, what does that now mean for geoAI, which is about intelligence, right? What does this mean if we really want to have intelligence in a somehow automated fashion in a computer? Because it seems you need to bring together a lot of people with a lot of different competences to solve these problems, and researchers in AI may say "well, that is not what AI is about."

It is not AI, but I think there are different types of intelligence, as there are different types of people. And I think that human versus machine intelligence is also very different. But, actually, I am not even sure whether intelligence is the right word to use here. It is this capacity of cross-checking incredibly complex and high dimensional spaces that is difficult for us. And it is a kind of intelligence even though maybe you see it as just a computer program searching the data. But then combining all these different disciplines and perspectives gives you the boundary conditions of the model. It leads to the priors that you want to use in the model, so that it does not predict anything that would be absurd.

There are all these fashionable hybrid models coming now as this new wave, which is great because it gets back physics into machine learning. I have colleagues bringing back chemistry into machine learning, and not just develop machine learning models that solve all the problems in chemistry. Creating a loop between a brain powered science and a machine powered science, in my opinion, is what will unlock a lot of the areas that so far were very hard, as we just needed so much data to get at them. But since now we are able to introduce these fundamental laws of nature, we will hopefully need a lot less data and will be able to make more connections across disciplines.

2 AI

KI: Let's talk a little about AI more general. The connectionist and symbolic approaches so far have evolved largely as two parallel strands. Do you think any of these two approaches separately is able to achieve the goals of AI research, such as the idea of Strong AI, or even more modest ones, such as the substitution of human experts in the loop? Or will some form of hybrid approach be necessary?

I do not think it should be the objective of AI to replace humans to solve all the world's problems. The objective of AI—the way we define it—is to support decision making in a way that is as accurate as possible. People who take decisions need to have the facts. The more accurate the facts are and the more they reflect the real world, the better decisions can be taken.

And this opens the door to all the issues that have been shaping the questions we are doing here at my lab ECEO[5] at the moment, for example, explainable AI models and physics based models. So, models that mimic the way natural processes work or cognitive processes. It is difficult for a decision maker at the end of the chain to accept the results of a machine they do not understand or they are not able to query and question. That is, not only getting an answer from a query, but also understanding why the model answered the way it has. For me, this is the step that is needed to support the people I want to support and the way I want to support them.

KI: Following up, what do you think are some of the biggest and most relevant challenges in AI in general?

There are several. One would be to go back to the "small" in our models and the computational power they need instead of keeping on ever growing. We have reached a point where only a handful of people can follow. While those models will always be the best—the most accurate, if you want—this is simply not inclusive of all this brain power we have in the world. And I think this development is coming. There is this

[4] https://trsc.org/en/.

[5] https://eceo.epfl.ch.

cohort of people who try to make AI open, to go back to an AI that can be used by anyone.

A second big challenge is the fact that we are still way too reliant on the training data we use to train these machines. These models learn by looking at the same examples millions of times. In every training cycle they need to look at all the training data again and again, again and again. But this is not how we learn, right? We may see every example only once, or at least only a very limited amount of times. It just does not happen very often that as a human you see the exact same example twice. This seems like something worth exploring. How? I don't know. Our models learn the way they learn because right now it is the best way we have figured out. And if you were throwing different training examples at them in an online fashion, then different learning problems arise, especially since we want to be able to handle classes we have never seen before with one example or even with zero examples. As humans, this is something we are good at. Just now, we hide behind the fact that the recent language models are so big that they seem to be able to do it, but nobody really understands why. So, this seems to be something to dig in, to get to the next level of AI.

KI: Continuing on the data issues, the success of ML models, especially in deep learning, is based on preventing bias by increasing model variance. In consequence, such models become highly sensitive to the training data quality. This limits their robustness, as illustrated by various failures of deep learning models: for example, models, which recognize traffic signs for autonomous vehicles, can be easily hacked. And models for traffic object recognition may confuse pedestrians with road infrastructure, with possibly fatal consequences, as in the case of Uber in 2016. What do you think could be done to overcome this blind spot?

It is true that the fragility of these models has been shown several times by these adversary attacks. The Uber example is particularly tragic because someone died. But there are a lot of other examples as well, which raise plenty of ethical and discriminatory aspects. The biases of the models are, I believe, one of the big challenges when it comes to making these models robust and fair. I see a movement in AI here, but I am a bit tangential because I do not work on core AI topics. But I see that there are more and more voices asking for it. It may not be the majority at the moment, but I think people are getting more and more aware of this problem. Also it is right to claim that just because you are only an engineer programming your model you are still accountable for what this model does. There was this idea of the cheap way out by saying "it is the data, not me." But we cannot do that. As a society we need to be able to develop solutions that are right. And they are right because they have been thought through properly from the beginning till the end. Again, there are these voices, for example Timnit Gebru[6], fighting for this mindset. And they get a lot of heat for that.

KI: I was wondering about what you just said about the way how to make these models more robust, more accountable, or maybe to hold the people behind these models accountable. Would not one way of doing this relate to what you said earlier, namely to incorporate these people and their knowledge? For example, to integrate the a priori models, e.g., the physical models or the conceptual models, into the learning?

Everything connects. Making the models interpretable, including physical priors, but also having a way of "stamping" our data saying that what the model is about to see is legit. All this needs to happen together to avoid some day having a catastrophe. I can only imagine some of the scenarios where such a catastrophe could happen, and I really hope it does not. I do not want to go into too much details here to keep others from misusing what I say. We need to make this technology right and make sure it cannot be misused. With today's AI systems, there are so many other biases that can enter. And these models discriminate people, negate things, you name it. And we really need to avoid this to happen in any possible way.

KI: Coming back to your work, would it be fair to say that your research in remote sensing and data processing is data-driven? On the other hand, the word 'semantic' appears in several of your paper titles, and your current work includes "making remote sensing accessible to everyone! Developing algorithms for human machine interaction" and "open the black box: interpretable deep learning in environmental modeling," which likely requires some form of explicit knowledge modeling. How do these two strands go together in your work?

Let's start with the word "semantic." There may be a mismatch of what we mean by "semantic" in different fields. Such mismatches are often the wall that prevents good collaborations; people use the same words with different meanings. In computer vision, "semantic segmentation" means to classify every pixel of an image to belong to a certain class. That is why the word "semantic" is all over the paper titles. We are not talking about ontologies; it is just saying "this pixel is part of a tree, and that is part of a house." It is probably somehow misusing the word "semantic" if I imagine how a philosopher would understand the term.

Now, in our work on interpretable AI and also visual question answering (VQA), we try to bring in a little bit more than just the name of a class. Because considering 'semantics' as used above just means that you know the class, you know what something looks like in the spectral domain (for instance the material it has been made of), and

[6] https://www.dair-institute.org/about.

you try to generalize from that. But we know that image understanding is much more complex than this. For example, we know that a car is usually on a street, and at the same time a car is conceptually closer to a boat than to a tree, even though it is located spatially more often closer to a tree than to a boat. But in terms of semantics, they are both vehicles and they would be closer in a knowledge graph. I think all these considerations start to enter the world of at least remote sensing. I cannot really speak for the geosciences in general.

We are also very active now in seeing how we can integrate these different types of semantic knowledge that appear on several levels; the material level where it really depends on how the light is refracted. And then this type of reasoning over objects being associated to another because they may be spatially close or they can use a certain level of space and not another. And the third level that is even more abstract where you state that a certain object can perform this action with these other objects, and it cannot be co-located with these other ones because they have a similar function. And all these types of different reasoning are required to work across different levels and different fields; in a way, they are related to different disciplines. And, thus, we come back to the initial point that we cannot do it alone. For example, these rules, these semantics if you want, for animal movement and how to discover them are with the ecologists. And before talking to them, for me it was all just "let's fly a drone and detect the animals by the color of their fur." But it is much more complicated than this, which I do understand a little bit better now. And I got there because I worked with ecologists.

KI: In this context, how do you define explainable AI? What are its pros and cons? To what extent do you think explainable AI is necessary for the further development of AI?

Maybe a simple way to put it is that understanding what a model is doing is always important? In machine learning, we somewhat lost track of this because models start to become more and more complex, but at the same time work very well. And people are excited because we manage to solve problems we could not solve before, but this comes with a price, which is losing an understanding of what is happening inside the model. It is not anymore transparent like in linear regression, where you know immediately which variables are important. With a neural network it is much more complicated.

Maybe we have to go back and dissect all these complex models to understand what they are doing. Our quest is importantly also a quest for knowledge. I want to understand processes because I am interested in how they work. For example, when we work on detecting blue ice, it is not only in order to detect blue ice and find meteorites, it is also about understanding issues related to climate change and the effects of global warming. And if you do not see how all the input parameters affect the final output, i.e., the prediction, you do not control the system. And if you do not control the system, you cannot simulate, you cannot foresee possible futures. And furthermore, with much more practical implications, when I give a model to an expert decision maker, and my model does not explain what the important variables are and how it reached a result, the expert won't be able to believe the results and stand behind them. I, as the developer, may trust this model, but it will never be used. It does not matter that predictions are 95% accurate. They will prefer a 70% model just because they understand it, and I perfectly understand this. So, we need to build this bridge between a model and explaining its results. This is very important.

KI: So, in some ways we may have partly traded science for engineering? Or we have traded performance for understanding? But it is important, for example for accountability, to being able to explain decisions, right?

I believe machine learning needed to gain respect of the disciplinary sciences somehow, so this was a necessary phase. 10 to 15 years ago people from the natural sciences did not trust machine learning results and were not searching for synergies with data science. And to gain this respect machine learning needed to make that choice of prioritizing performance and showing that they could crack important problems with learning algorithms. But I think people now are smart enough to understand that it is time to revert and to build models together with others who are experts in their domain. Now we are taken seriously, so let's do it together. The mistake would be to reject this because we think we are so much better than the others (which we are not).

3 GeoAI

KI: How would you define geoAI?

To which of the following two definitions would you agree more? "geoAI is…"

- "*…AI for geographic information*."
- "*…geographic information for AI*."

You know what I am going to say after all we have talked about before. You cannot have only one or the other, we need to have both. The title of my keynote talk on ecology is "machine learning supporting ecology supporting machine learning" because it is a cycle. It cannot be one-directional because we would miss all this knowledge that one domain can bring to the other. Of course AI will accelerate animal ecology. There is all this technology available now. We have been collecting camera traps data for years. We can finally process it. On the other hand, a machine learning model can benefit so much from rules about diffusion and movement that are there in ecology that the model knows nothing about. It works, but remains a waste of resources to ask the model to learn things that everybody knows anyway by looking at millions of examples. So, the only possible way forward here is symbiosis.

KI: And is there a particular approach in the field of geoAI which inspires your work?

I am always looking for new exciting things, I am a curious person. For example, I was very excited the first time I saw something about question answering. When I saw this, it really inspired me to show that there is more to it than just taking some data and trying to get an answer out of it in a pure forward, non-human like pipeline. This also tied in nicely with all this work we had done on humans in the loop. When I first encountered this QA work, it was about supporting visually impaired people to navigate in the real world. They would ask questions via their mobile phone together with a picture, and they would get the answer. You could help people cross the street this way, for example. This gave me an idea of how useful this can be, if transported to the Earth science domain.

The second inspiration is all this work on crowdsourcing that has become very popular. All these big portals, such as ebird[7], where there really is a link between people who enjoy nature, practitioner experts, and data scientists. Here, I saw that triangle I talked about before really happening. Especially the work of Serge Belongie[8] is very inspiring for me. It was a way of trying to integrate knowledge of people into models that would be top notch machine learning aimed at answering relevant conservation questions. For me this is the real inspiration. The work that we are doing in the lab now moves a little bit into this direction, always having multiple aspects of this triangle in there.

KI: Coming back to our previous discussion and related to aspects of crowdsourcing, one of the biggest challenges of current data-driven geoAI approaches seems to lie in their dependence on high-quality data collections, especially with human experimental subjects. For example, collecting large amounts of high quality answers to highly specific questions in geographic Question Answering (geoQA), getting enough experts to label maps or images, obtaining enough user interactions, and collecting specific text corpora is very costly. On the other hand, social media data and data from crowdsourcing projects often lack the required quality. How can we overcome this bottleneck for current data-driven geoAI?

I think there are different requirements on data. As long as you know what the quality of your data is, you can use it for what it is worth, but not more. For example, if you have corpora that you know are curated versus much larger corpora that are not, you know that in principle you can trust one more than another, but also that the other has the required size to be able to do more. To really make the most out of your data, you need to find a balance between the two. This has been used a lot over time, for example, in the ebird project I mentioned before. But there, the tasks that are crowdsourced are really simple, such as "is there a bird or not?" That is a task everybody can solve. And if you have more complex tasks, then you have to involve a network of more expert people. But there will likely never be perfect crowdsourced data, in fact, there will never be any perfect data at all.

KI: Moreover, the know-how and (cartographic) knowledge required for dealing with spatial information is non-trivial. In particular, it goes beyond knowing the geometry or just having the data. If this is true, isn't the reduction of geoAI to knowledge extraction from geodata generating a problem? Doesn't this run the risk of underestimating the difficulty of the task, because it underestimates what it means to interpret geodata?

I believe it is our job to make very clear for others where the difficulty of a task lies. Because for us it may be obvious, but many people do not think about it. For example, think about how concepts are non-stationary in space: if we call something a mountain in Switzerland it is not the same as if we call something a mountain in Ireland[9], or in the Netherlands. Where I used to live, there they call the "Wageningen Mountain" a mountain, but it is only 30 meters high!

Another example are the language models that we use; in creating these data sets we have been taking some shortcuts because we had to, but we always try to talk openly about it in our papers. For example, if I start a question by "is it," the answer will be "yes" 99.8% of the time. These are biases that are in everybody's data sets, but not everybody talks about it. And when you talk about it, the reaction is often "oh, your model doesn't work." For example, we have done an experiment with the VQA system by randomizing the image part. Simply put, we gave on purpose a random image to the system, an image that had nothing to do with the question we were trying to answer, and still got 70% accuracy. The system gets just over 80% when seeing the right images. The message here is our model was not really using the visual information. Instead, it relied on language biases and we need to be very careful about such effects.

KI: How does this issue of non-trivial knowledge required for understanding and interpreting (results of) geographic information processing inform and challenge the aims of your current work, namely "remote sensing for everyone?"

7 https://ebird.org/home.

8 https://www.belongielab.org.

9 Example attributed to Ross Purves.

This big project really started with this issue in mind, namely: can we offer something that enables a journalist to get an answer about space? And without having ever used Python or not even knowing what a neural network is? I wanted to see if we could get to a point where we can use language as an interface, because it is an interface that everybody can use. Next came all the restrictions you may need to make. How much can the geography differ? Should this work all over the world or just in the country you trained the model in? Are the thematic areas constrained? Do you allow to ask questions only about forests or buildings? Since the more unconstrained the problem, the more challenging VQA becomes.

As you can see, there are many questions, but you need to get there one step at a time. We started with a simplified problem with a subset of themes, limited geographies and a language construction, which is really rather simple. It is only an expert system that creates sentences, always the same way. But these first steps then give you the knowledge that allows you to build the next system, which is a little better. And maybe one day we can also answer questions about themes that were not in the training set. Who knows?

KI: On the issue of data and data quality, you have released the data along with some code for your visual question answering for remote sensing (RSVQA) system. What was the motivation behind this?

It felt like the right thing to do, to promote a research problem that is interesting and give others the means to start developing. We were not required to do it, but if you want the community to be interested in your work and want them to challenge what you did and improve on what you did, we had to give out something. We decided to publish all the data we created, all the scripts to reproduce all the models we had in the paper, hoping that people will pick it up and improve—destroy—what we do in terms of performance.

I am always a bit curious about labs that do not do this. I can understand that if you are a company you want to have an advantage. But we are funded by public money. So I believe we need to give back something to society somehow. And for me, not sharing code and data is a mistake because it cripples the community, and I have the impression that the message is getting through. There is more and more out there, in fact maybe even a little bit too much. It would be good to have some kind of quality control as well. But I think every data set shared is good. And not just because if you do not share your data set, you are not allowed to publish your paper, because sometimes it seems like that is the reason people share. But really, sharing is the only way to move forward, so others can improve on what you do. And it is ok if people improve on what you do, and then your model is the one performing worst in the next paper. In fact, this is what should happen every time, right?

KI: Finally, what advice would you give a young researcher just entering the field of geoAI?

I would say be as curious as you can. Do not let people box you into a small field. And know about the problem that you are looking at, not just about the techniques or the GIS protocols to solve it. Spend the necessary time for the extra mile talking to experts, which means spend a lot of time understanding each other. Because if this is not done, the project will not work. You can work all the hours you want, train all the models you want. You will not get to where you want to get.

KI: Devis, thank you so much for this very inspiring interview!

References

1. Tollenaar V, Zekollari H, Tuia D, Kellenberger B, Rußwurm M, Lhermitte S, et al (2022) What determines the location of Antarctic blue ice areas? A deep learning approach. In: EGU General Assembly 2022. EGU22-1294. Vienna, Austria; Available from: https://doi.org/10.5194/egusphere-egu22-1294.
2. Persello C, Wegner JD, Hänsch R, Tuia D, Ghamisi P, Koeva M et al (2022) Deep learning and earth observation to support the sustainable development goals: Current approaches, open challenges, and future opportunities. IEEE Geosci Remote Sens Mag 10(2):172–200. https://doi.org/10.1109/MGRS.2021.3136100
3. Tuia D, Kellenberger B, Beery S, Costelloe BR, Zuffi S, Risse B et al (2022) Perspectives in machine learning for wildlife conservation. Nat Commun 13(1):792. https://doi.org/10.1038/s41467-022-27980-y