



Workshop summary

Machine learning for remote sensing (ML4RS) has emerged as an exciting area of research with the potential to address some of the most pressing global challenges, including conservation, food security, and disaster response. Remote Sensing offers the Machine Learning community a uniquely rich and structured view of the Earth - vast, multimodal, and temporally continuous data that challenge and advance ML methods for representation learning, generalization, and reasoning about complex real-world systems. Remote sensing data, collected from diverse instruments capturing Earth observations across various spatial, temporal, and spectral dimensions, offer unique research opportunities and challenges for the ML community. Unlike conventional data modalities, these datasets are high-dimensional, extremely multi-modal, and contain patterns at a multitude of spatial and temporal scales that pose inherent scientific challenges for learning representations. These characteristics often require specialized approaches in cross-cutting ML topics like self-supervised/semi-supervised learning, domain adaptation/generalization, and multi-modal learning/data fusion.

A dedicated research focus in ML4RS – fostered in large part by our past ICLR workshops and complementary workshops offered at other AI/computer vision conferences – has spurred a wealth of new research in topics such as foundation models, fine-tuning techniques, datasets, and implementations, all tailored to the unique characteristics and challenges that remote sensing data pose to machine learning. This explosion of work has come with growing pains. Proper benchmarking of new models within the scientific community and translating research insights into practical use in real-world applications are particularly difficult in ML4RS. Our workshop this year will focus on bridging these gaps while continuing to catalyze state-of-the-art research through discussion and feedback on early-stage research.

The special **theme of our workshop this year is “ML4RS: publication to practice.”** Reflecting this theme, this year’s workshop features a new “tutorials” track, with dedicated time during the in-person workshop for contributed talks from submitted tutorials, followed by hands-on “breakout intensives” led by the tutorial authors. We also integrate a new remote sensing track with a direct path to proceedings in conjunction with the IEEE Geoscience and Remote Sensing Letters (GRSL) journal. Finally, our invited speakers are asked to begin their research talks with “provocations” on the theme of “foundation models in ML4RS: are we there

yet?,” setting the stage for a productive (possibly spicy) panel discussion between our speakers, who span users and developers of foundation models.

By fostering collaborations between ML researchers and remote sensing domain experts, our workshop promises to break new ground in advancing both the methodologies and applications of ML for remote sensing, setting the stage for future advances in this field. Our workshop solicit contributions tackling problems including (but not limited to):

-  **Foundation models:** How can large models pretrained on unlabeled satellite data capture spectral, spatial, and temporal nuances to accelerate downstream tasks?
-  **Active learning & annotation efficiency:** How can we target limited labeling effort to maximize model performance?
-  **Imperfect evaluation data:** How can we assess model quality and detect artefacts with sparse or uncertain labels?
-  **Interpretability & generalization:** How can geospatial priors and physical models improve transparency and robustness?
-  **Benchmarking & impact:** How can evaluation metrics reflect real-world impact and enable fair model comparison?
-  **Accessibility & efficiency:** How can we democratize ML4RS through distributed, low-cost training and experimentation?
-  **Local vs. global models:** Should we aim for one global model or many local specialists, and can we identify scopes for transfer automatically?
-  **Precomputed embeddings:** How can we best use precomputed embeddings from foundation models, instead of requiring users to compute embeddings themselves?

Tackling these and related problems requires a confluence of perspectives from developers and users of ML4RS research across different sectors and sub-focuses of machine learning, which our workshop will enable.

Location-specific focus

This ML4RS edition will continue the tradition established by our 2023-2025 workshops of integrating local researchers and stakeholders. This year will emphasize and celebrate Brazil's vibrant community of researchers, practitioners, and students in machine learning and remote sensing. Past iterations integrated speakers and participants from the Rwanda Space Agency and CMU Africa (2023 in Kigali), the European Space Agency (2024 in Vienna), and NASA's international SERVIR Southeast Asia Hub and the Earth Observatory of Singapore (2025 in Singapore) to tackle regionally important challenges. This year's keynote speakers (one from Brazil's National Institute for Space Research and one from Google, US) and the confirmed invited speakers (from Brazil, Canada, and the US) span applied and methodological aspects of machine learning for remote sensing. This will spur rich discussion and facilitate connections between communities that might otherwise remain disjointed without our workshop. We will continue our tradition of advertising our workshop to local communities (Brazilian universities,

companies, and NGOs), with an emphasis this year on our tiny papers, remote sensing, and tutorials submissions track.

Submission Tracks + Reviewing

Research papers (4 pages, non-archival). Our main call for papers will be research papers of up to 4 pages in length that describe new and ongoing/in-progress research, following our success with high-quality submissions in this track from past years. The paper PDFs will be made available on the website (See [Orals](#) & [Posters](#) from last year), but will be non-archival without proceedings or DOIs.

Short (tiny) papers (2 pages). We will welcome submissions that contribute findings of a smaller scope than the 4-page research papers, which we hope will help to extend engagement from local students and researchers from ICLR who may be new to the remote sensing domain. To support first-time authors, we will provide explicit guidance on what we expect from a tiny paper submission – for example, a modest in scope but complete experimental result, a fresh perspective with supporting evidence, a single theoretical result, or a proposal for work soliciting explicit feedback from the community. We will also provide explicit guidance to our reviewers on how to evaluate tiny papers in light of these expectations.

 **Tutorials track (Up to 4 pages).** Reflecting this year’s theme of bringing ML methods and practice closer together in remote sensing for machine learning, we will introduce a new “tutorials” track. The goal of this track is to expedite the use of new models, code libraries, datasets, and benchmark challenges – facilitating their use in both practical applications and comprehensive benchmarking in future research studies. Submitted tutorial proposals will detail a 15-minute video tutorial leading learners through a specific example (it is expected that this will correspond to one interactive Python notebook that can run locally on a laptop).

- All accepted tutorials will be posted publicly on our workshop website. We will give authors the option to post longer, more detailed tutorials as well.
- Authors of highly reviewed tutorial submissions will be invited to give (15-minute) spotlight presentations and lead breakout “intensives,” essentially holding office hours for extended tutorials during the workshop.

 *Reflecting our local focus, we will encourage tutorials with examples involving Brazil in their case studies, when feasible (for example, if the tutorial teaches users about a new, lightweight segmentation model for detecting tree crowns, the example in the tutorial notebook could focus on an area in Brazil). For some tutorials, such as existing benchmark datasets, this may not be feasible, and would not be counted against submissions during review.*

 **Remote Sensing Proceedings Track.** We are excited to introduce a new pathway for publishing remote sensing learning papers through a collaboration with a Special Stream of the IEEE Geoscience and Remote Sensing Letters (GRSL). Authors of ICLR ML4RS papers can opt to have their 4-page submissions evaluated as GRSL candidates, following the journal’s

review standards during the workshop. After the conference, accepted papers will enter a second, separate reviewing stage at GRSL, ideally with the same reviewers, ensuring a smooth and efficient transition toward publication. This dedicated remote sensing publication track provides an attractive incentive for researchers from the geoscience and remote sensing communities who may not typically attend ICLR. By offering a clear and recognized pathway toward GRSL publication, it allows these researchers to more easily justify their participation and travel to ICLR. The Editor in Chief of GRSL, R. Hänsch, is among the workshop organizers, ensuring a smooth and efficient publication process. While participation in this special stream is optional, it provides an excellent opportunity to combine open scientific exchange at the workshop with a valuable peer-reviewed publication (which may also help to warrant the costs for researchers for participating in an international event).

Highlighting contributed work: Accepted papers from all tracks are expected to present their work during our interactive poster sessions and will be encouraged to share materials online through our workshop website. A subset of the highest-reviewed submissions from each submission track will be invited for contributed talks to spotlight their work during the workshop.

AI-generated submissions: For all submission tracks, we will closely follow the posted ICLR requirements on AI authorship and publicize rules in our call for papers.

Avoiding conflicts of interest during reviewing: The review process will be double-blind, and we will ensure that no reviewers have conflicts of interest with the papers they are assigned to review. Possible conflicts of interest include submissions from past collaborators or submissions from the reviewer's organization.

Reviewing schedule: We will implement the following reviewing schedule. All tasks with quick turnaround will be covered by the workshop organizers.

- Submission Deadline: February 6, 2025 (Anywhere on Earth)
- Reviewer assignments to papers: February 7
- Reviewing period: February 7 - February 25 (2.5 weeks)
- Paper decisions: February 26 - February 28
- Acceptance Notification: March 1st, 2025 (in line with global workshop deadline)
- Camera Ready Submission Deadline: April 19, 2025 (Anywhere on Earth)
- Accepted papers shared on the ICLR website and the workshop website: April 23, 2025

Tentative schedule

Our proposed workshop agenda consists of three parts. In the first part of the day, we will focus on framing challenges around our theme of bridging the gap between publication and practice. We will hear from expert developers and end-users of remote sensing foundation models through invited talks and a debate/discussion about what is missing in the connection between publication and practice. In the second part of the day, we will highlight state-of-the-art advancements in ML research for remote sensing with a keynote, invited talks, and spotlight

paper presentations. The third part of the day will focus on putting research advancements into practice, featuring tutorials/demos from our new track.

★ Our proposed agenda features multiple sessions to foster **opportunities for discussion** among researchers, practitioners, and end-users, including a keynote-led debate, poster sessions, and interactive tutorials.

★ The poster sessions and tutorial “office hours” provide **inclusive, low-barrier opportunities for early-career and under-resourced researchers** to present their work and directly engage with senior experts.

start-end time	topic
9:00-9:15	Introduction and opening remarks
9:15-9:45	Keynote 1: Gilberto Camara, “Spatial Data Science for Sustainability: Exploring petabytes of satellite images to support greener and better societies”
9:45-10:30	Invited Talks from geospatial foundation model developers and end-users that open with “provocation” (3 x 15 minutes each) 1. Beth Tellman, “Are Geo-Foundational Models helping scientists or end-users make better flood maps? 2. Anthony Fuller, “We aren’t there but have travelled far via Self-Supervision, Architecture, and Data” 3. Daniela Szwarcman, “Developing Geospatial Foundation Models: the Prithvi-EO experience”
10:30-11:00	Coffee Break
11:00-11:30	Discussion/debate with morning keynote + invited speakers (topic: “foundation models - are we there yet?”)
11:30-12:00	Contributed Paper Spotlight Presentations: 3 x 10 minutes each
12:00-13:30	Lunch (opportunity to visit posters and have discussions)
13:30-14:00	Keynote 2: Christopher Brown, “The Ergonomics of Geospatial AI: Building Solutions for The People You’ll Never Meet”
14:00-14:30	Contributed Paper Spotlight Presentations: 3 x 10 minutes each
14:30-15:30	Coffee Break and Poster Session
15:30-16:20	Tutorials Spotlight Presentations (3 x 15 minutes each)
16:20-17:15	Tutorials Breakout: hands-on “intensives / office hours”
17:15-17:30	Best paper awards + closing remarks

Invited speakers

★ **Internationally-recognized speakers:** The proposed invited speakers are research leaders in ML for remote sensing and foundation models. All speakers have strong track records of impactful research and effective public presentation. Speakers span early-career and established researchers, international and local (Brazilian) researchers, and academic and industry affiliations.

★ We have **confirmed participation** from all five speakers listed below.

★ Our program features a **diverse array of speakers** spanning diverse career stages, geographies, ethnicities, genders, and institutions (industry/academia).

Gilberto Câmara (Keynote 1) is a distinguished Brazilian computer scientist and geoinformatician whose work bridges spatial data science, remote sensing, and artificial intelligence. He explores how machine learning can be applied to satellite image time-series analysis and geospatial big data—contributing to open-source tools (e.g., the sits R package) for Earth observation. His research agenda encompasses land-use change, environmental modeling, and the development of semantic models (geo-fields) to enhance the interpretation of AI outputs in spatial contexts. Through his advocacy of open data policies, his leadership in national and international programs (including as former director of INPE and former Secretariat Director of GEO), and his innovative integration of AI into geospatial research, he continues to shape how we monitor and understand environmental change.

Christopher Brown (Keynote 2) is a Senior Research Engineer at DeepMind where he leads research at the intersection of AI and Earth observation. His work focuses on developing and applying large-scale AI models to address global environmental challenges. Christopher has led several planetary-scale and publicly-released projects, including AlphaEarth, Cloud Score+, and Dynamic World. His research aims to provide the scientific community with new tools to better monitor, model, and protect Earth's systems.

Beth Tellman is an Assistant Professor at the University of Wisconsin-Madison, in the Nelson School of Environmental Studies whose research addresses the causes and consequences of global environmental change on people, with a focus on flood risk and land use change. She engages in multiple disciplines and methods to “socialize the pixel” or understand the social processes behind environmental change captured in satellite image pixels and leverage satellite data to improve human well-being. She is a co-founder and Chief Science Officer of Floodbase, a public benefit corporation that leverages remote sensing to build flood parametric insurance products. She is a former Echoing Green Fellow, Fulbright Fellow, and won an NSF CAREER award (which was recently terminated!). Beth co-founded an NGO in Mexico, Umbela, to promote transformation from Global South Perspectives. Beth is passionate about co-producing knowledge with actors outside academia to achieve social impact. To learn more about her publications and projects see: <https://beth-tellman.github.io/>.

Anthony Fuller is a final-year PhD student at Carleton University and the Vector Institute, supervised by Jim Green and Evan Shelhamer. He has developed multiple multi-modal

pre-trained/foundation models for remote sensing and enjoys researching vision transformers and self-supervised learning.

Daniela Szwarcman is a Research Scientist in the Spatiotemporal Modeling group at IBM Research, where she began as a Ph.D. student intern in 2017. She holds a Ph.D. in Electronic Engineering with a specialization in Artificial Intelligence from PUC-Rio (2020). Her research focuses on developing foundation models for spatiotemporal data analysis, with applications in Earth observation, weather, and climate. She has contributed to award-winning projects, including NASA's Marshall Group Achievement Award and the American Geophysical Union Open Science Award.

Organizers and biographies

Our team includes experienced organizers of the Machine Learning for Remote Sensing workshop at ICLR (from its first 3 iterations in 2023-2025) as well as new team members (Bianca Zadrozny, local to Brazil) and leadership (Esther Rolf will lead the workshop organization for the first time, after being a co-organizer in 2025). The organizers span diverse career stages (PhD to senior researcher), geographies (Europe, North America, South America), ethnicities, genders, and institutions (industry, academia, and government). To ensure fair workload distribution and coverage of organizational activities, each member has defined their organizational role in the proposal stage (below).

1. **[Esther Rolf](#)** (first-time workshop lead 2026, co-organizer 2025-present) is an Assistant Professor in the computer science department at the University of Colorado, Boulder. Esther will be responsible for coordinating overall organization efforts, serving as the primary contact for ICLR, inviting speakers, and facilitating the workshop on-site.
2. **[NEW Bianca Zadrozny](#)** is a senior research manager at IBM Research Brazil, leading the Spatiotemporal Modeling group. Bianca will be responsible for coordinating local speakers and participants, including synergizing with local events and organizations.
3. **[Hannah Kerner](#)** (workshop lead 2025/2023, co-organizer 2023-present) is an Assistant Professor in the School of Computing and Augmented Intelligence at Arizona State University. Hannah will facilitate the timeline of workshop planning, ensuring key activities are planned and completed according to the ICLR timeline. She will also help coordinate reviewing efforts for submitted papers and awards.
4. **[Marc Rußwurm](#)** (workshop lead 2024, co-organizer 2024-present) is an Assistant Professor for Machine Learning and Remote Sensing at Wageningen University, Netherlands. Marc will be responsible for coordinating reviewing efforts, including recruiting reviewers, facilitating meta-reviewing, and making award decisions.
5. **[Evan Shelhamer](#)** (co-organizer 2025-present) is an assistant professor at the University of British Columbia in Vancouver and the Vector Institute in Toronto. Evan will be responsible for OpenReview administration and following ICLR announcements to align with ICLR dates and expectations (e.g., syncing and posting papers to OpenReview).

6. [**Gabriel Tseng**](#) (co-organizer 2023-present) is a Ph.D. student at McGill / Mila - Quebec AI Institute and Research Scientist at Ai2. Gabriel will be responsible for maintaining the workshop website.
7. [**Ronny Hänsch**](#) (co-organizer 2024-present) leads the Machine Learning Team in the Signal Processing Group of the SAR Technology Department at the Microwave and Radar Institute of the German Aerospace Center (DLR). Ronny will be responsible for coordinating reviewing efforts, including recruiting reviewers, facilitating meta-reviewing, and making award decisions.
8. [**Hamed Alemdohmamad**](#) (co-organizer 2023-present) is an Associate Professor in the Graduate School of Geography at Clark University and also the Director of the Center for Geospatial Analytics. Hamed will be responsible for fundraising and coordinating travel/registration support applications.

Workshop Audience

Anticipated audience size

Based on the 2023-2025 editions of our workshop, we expect an attendance of approximately 50-100 participants. We expect a much larger audience online who will benefit from the materials that we host in perpetuity on our workshop website.

In past workshops, we observed that most of our participants stayed and remained engaged the entire day. This facilitates deep conversation and collaboration with lasting outcomes from our workshop. We also know that many participants alternate between workshops at ICLR on the same day, so we structured our workshop agenda to ensure participants will find meaningful collaboration and outcomes even when attending for part of the day.

Attracting an audience for the workshop

We will build on the promotion and outreach channels established in the last 3 years and will provide a 2026 website consistent with previous [ML4RS@ICLR websites](#). We have contacts with newsletters of Remote Sensing Societies, such as [Geoscience and Remote Sensing eNewsletter](#), [ISPRS e-Bulletin](#), and the [IAPR-TC7 newsletter](#). We will also promote the workshop in relevant Google groups, including affinity groups such as the Women in Machine Learning Group (e.g., as [in this post](#)). Most of the organizers have active Twitter/X profiles with a combined number of roughly 9.3k followers, with even greater reach from the organizers' institutions. We will also ask our diverse network of program committee members to promote the workshop in their networks. We will coordinate outreach to local institutions and organizations through our local organizing team member (Bianca Zadrozny) and speakers (e.g., Gilberto Camara) to promote local engagement in Rio and South America more broadly.

These efforts will ensure we draw a diverse audience including under-represented, under-resourced, and budding researchers to engage with the ML for remote sensing community at our workshop.

Virtual access to workshop materials and outcomes

Given that our workshop will be in-person, we will offer several ways to engage with the workshop for those unable to physically attend. During the in-person workshop, we will do the following to ensure our remote participants can engage:

- For the invited/contributed talks and panels, we will have the content synchronous on slideslive, as in past years. This includes the contributed tutorials talks.
- We will use online polls with QR codes to solicit audience feedback throughout the workshop, including: (i) compiling audience questions for talks/panels, (ii) voting for best poster awards.

Some aspects of our workshop are designed to facilitate conversation and co-learning are thus less conducive to a hybrid format. For these, we will offer:

- For the tutorials, we will require that accepted tutorials have a short (~15 minute) video and accompanying notebook which we will publicize on our conference website. In addition, for remote participants to be able to engage at a deeper level akin to our in-person “breakout intensives,” will encourage tutorial creators to publicly post any additional materials (e.g., deeper-dive notebooks, or a longer form video going through those) online, hosted or linked on our website. We will also encourage tutorial creators to host online “office hours” for their tutorials, either concurrent with the in-person workshop (if tutorials have a 2+ person author team) or at a different time.
- To enable rich remote engagement, we will post papers and posters for all accepted submissions on the official ICLR website (as we have in past years) so they are viewable by all in person and remote attendees, and we will encourage them to post them on our publicly available workshop website. Note that all submissions to this workshop are nonarchival.

Sponsorship

In recent years, we have received sponsorships from Microsoft (2023, 2024, 2025), Astraea (2023), GRSS (2023, 2024), ESA (2024), Hyperplan (2024), ISPRS (2024), Google DeepMind (2025), the SIG Spatial Informatics Group (2025), and AI2 (2025). In 2026, we plan to build on this strong network to secure new sponsorships that will fund best paper and poster awards, offer financial support to students with limited resources, and enable the participation of invited speakers who might not typically attend ICLR. Specifically, we aim to cover conference registration costs for authors of accepted tiny papers from local universities, following the

successful model we implemented with Carnegie Mellon University in Kigali, Rwanda in 2023, which funded up to 100 students and had a lasting impact on the students involved.

Diversity commitment

We are committed to ensuring our workshop features diverse individuals, perspectives, and contributed research.

★ Our **organizing team and invited speakers** span diverse career stages (PhD students to senior researchers), geographies (Europe, North America, South America), ethnicities, genders, and institutions (industry, academia, and government).

★ We offer **multiple tracks** for individuals to present their work at our workshop with different levels of maturity required (research papers, TinyPapers, or tutorials). In our 2025 workshop, we offered the TinyPaper track for the first time as a pathway for researchers to submit work that may not be mature enough for acceptance on the research papers track. We found that authors and reviewers both struggled to differentiate between what qualified as an acceptable contribution for a TinyPaper vs. a research paper. To alleviate this roadblock, in our 2026 workshop we will provide specific guidance on what makes an effective TinyPaper to facilitate submission and ensure constructive reviewer feedback.

★ Our **audience recruitment strategy** includes outreach to affinity groups (such as Women in Machine Learning, LatinX in AI, Black in AI, etc). We will also advertise to local institutions and organizations through our local organizing team member (Bianca Zadrozny) and speakers (e.g., Gilberto Camara) to promote local engagement in Rio and South America more broadly.

Relation to Previous Workshops

Differentiation from previous ML4RS workshops at ICLR

As mentioned in detail in the introduction above, this year's workshop includes a new track for tutorial submissions (including invited tutorial talks and subsequent hands-on practice led by tutorial presenters) and a focus on debate/provocation during invited talks and panel discussion. Furthermore this ML4RS, like each edition, incorporates local expertise from the host country and city for ICLR, so this time we have invited Brazilian geoinformatician Gilberto Câmara.

Differentiation from complementary workshops at other conferences

We are well-connected to the CVPR EarthVision workshop series through common organizers H. Kerner and R. Hänsch and to the ICPR workshop Pattern Recognition in Remote Sensing (PRRS), co-organized by M. Rußwurm and R. Roscher. We believe that this ML4RS workshop placed at a machine learning and representation learning venue nicely augments the EarthVision and PRRS workshops for relevant environmental problems beyond pattern recognition, vision, and high-resolution imagery. For instance, we aim to include more topics such as time series to capture the dynamic nature of our planet and its shift with climate change, to fundamentally discuss transfer learning approaches, or how to mitigate domain shift on remote sensing data. The non-archival option of our ICLR workshop also provides a

complementary opportunity for researchers to present early-stage and in-progress innovative research with a higher probability for acceptance than EarthVision, which has a low acceptance rate due to its archival place in the CVPR proceedings. This opportunity is especially important for early-career researchers and those from under-resourced or less research-intensive institutions who have not yet “broken into” the mainstream ML conference space.

Our workshop audience and topics also have some intersection with the Climate Change AI (CCAI) workshop, frequently hosted at ICLR and other conferences. CCAI spans a broad set of topics related to climate change, within which remote sensing is just one subtopic. CCAI’s remote sensing contributions typically focus on application of remote sensing and AI to topics in climate change, whereas our workshop focuses on broad applications of remote sensing and ML and foundational methodologies for modeling remote sensing data with ML. We will work with CCAI to cross-promote intersecting workshop topics and ensure schedules do not create audience conflicts.

Other complementary workshops include Terrabytes (ICML 2025), Sustain EO (ICCV 2025), CV4EO (WACV 2025), and GeoCV (WACV 2025). These workshops are synergistic with ours, but ours differs significantly in our local geographic focus, international audience, tutorials track, and focus on lively discussion and debate, among other aspects.