SIGIR 2025 Low Resource Environments Track Report

Catherine Chavula
University of Strathclyde
United Kingdom
catherine.chavula@strath.ac.uk

Charles Clarke
University of Waterloo
Canada
claclark@gmail.com

Douglas W. Oard
University of Maryland
USA
oard@umd.edu

Abstract

The 48th International ACM SIGIR Conference on Research and Development in Information Retrieval introduced a new low resource environments (LREs) track, dedicated to build an information retrieval (IR) research community among scholars from low- and middle-income countries and those addressing IR challenges associated with limited resources. Rather than traditional research paper submissions, the track invited presentation proposals focused on information access in the context of LREs. The track had a full day schedule and was held on 16 July 2025 in Padua, Italy. This report provides the track's objectives, activities, lessons learned, and recommendations. Primary outcomes include: (i) participants from low-resource environments valued the opportunity to attend and present at SIGIR, with many expressing intent to submit a full paper at the conference; (ii) a need was identified for alternative, community-building activities to share knowledge, pool resources, and learn collectively; and (iii) early action is required to address visa and logistical challenges that may hinder participants from low income countries to participate in the conference.

Date: 16 July 2025.

Website: https://sigir2025.dei.unipd.it/low-resource-environments-track.html.

1 Introduction

Low-resource environments are characterized by constraints that limit research, development, and deployment of information retrieval (IR) systems. These constraints may, for example, include lack of publicly available test collections, limited research funding, or inadequate computing infrastructure. While commonly associated with low-income countries, such conditions can also exist within underserved or marginalized communities in middle and high income countries. Low resource environments therefore encompass the places, people, platforms involved in information access. These settings present unique IR challenges that require innovative approaches to IR system design and evaluation.

The IR community and SIGIR has produced limited research addressing these challenges, with low representation from researchers working in such contexts. To begin addressing this gap,

the SIGIR 2025 conference introduced the low resource environments track. The track received financial support from SIGIR executive and University of Padua, to allow participants from low income countries to attend and present their work at the track. This report summarizes the track's purpose, activities, lessons learned, recommendations and its long-term impact.

2 Purpose and Objectives

The purpose of the track is to broaden the participation of IR researchers who are working **in** low resource environments, and build a community of researchers working **on** low resource environments in the field of IR. Specifically, the objectives of the low-resource environments track are:

- Broaden participation: Expand engagement of researchers working in or on low-resource contexts within the IR community, as described above.
- Provide a collaborative platform: Offer a space for researchers and practitioners facing resource constraints to connect and collaborate.
- Facilitate knowledge exchange: Encourage participants to share experiences, identify common challenges, and explore strategies to improve information access in low-resource settings.
- Promote community building: Strengthen long-term ties to support future partnerships, mentoring opportunities, and cross-regional research initiatives to promote research addressing low resource settings.

3 Scope of Topics

The scope of the track broadly encompasses IR challenges associated with low resource settings affecting places, people and platforms. The track invited submissions for presentations on emerging ideas or positions, challenges, methods, tools, techniques, or experiments that address IR issues in low-resource environments. Topics covered included:

- IR systems and applications: Exploration or evaluation of new IR applications for low-resource environments. For example, low-cost tools for low-income, underserved communities or rural populations.
- IR for sustainable development: Systems or technologies aligned with UNESCO's Sustainable Development Goals.
- Societal impact of IR systems: Studies on the societal implications and ethical concerns involving IR platforms and algorithms in low resource environments including issues of equitable information access.
- Building IR research networks: Initiatives and frameworks for establishing collaborations and communities to address low-resource environments.
- IR Dataset creation: Development and evaluation of new datasets for low-resource languages and other low-resource environments.
- Search and ranking: Research on retrieval models and ranking techniques for low-resource languages and other low-resource environments.

- Efficiency and scalability: Studies of search systems that focus on efficiency and scalability in resource-constrained environments.
- Cross-cutting research involving IR: Research bridging gaps between IR and related fields such as natural language processing, speech processing, image processing or video processing for low-resource languages and other low-resource environments.
- Evaluation: Innovative methods for assessing IR systems in low-resource environments.

The submission deadline was March 3, 2025. This date was intended to allow the track chairs to promote the track. The track was advertised through mailing lists to African researchers working on AI related topics, as well as CHI and other computing related communities.

4 Organisation

Our call for participation invited researchers working on low-resource environments to submit presentation proposals in any convenient format. The call encouraged submissions from low income countries and was promoted in associated forums. However, for the camera ready submissions, the authors of accepted presentation proposals submitted a two-page abstract using ACM style format. Each proposal was evaluated based on its relevance to SIGIR, its focus on low-resource environments, potential to engage the SIGIR audience, and the identified strengths and weaknesses. Every submission received at least two reviews. The program committee was well balanced internationally, with strong representation from people who themselves were working in low-resource countries. The chairs reviewed these evaluations, and based on their discussion and their final consensus, a meta-review was prepared for each proposal.

5 Participation

In total, the track received 29 presentation proposals on topics related to low-resource languages, agriculture and food, healthcare, under-studied user groups, and efficiency in low-resource environments. Just over half of all submissions came from one of ten African countries: Cameroon, Ethiopia, Nigeria, Malawi, South Africa, Tanzania, Tunisia, Togo, Kenya, and Uganda. In total, the track received 29 presentation proposals on topics related to low-resource languages, agriculture and food, healthcare, under-studied user groups, and efficiency in low-resource environments. Just over half of all submissions came from one of ten African countries: Cameroon, Ethiopia, Nigeria, Malawi, South Africa, Tanzania, Tunisia, Togo, Kenya, and Uganda. There were also submissions from Asian countries, such as India, Nepal, and the Philippines. No submissions were received from South America. Submissions were also received from researchers working in higher-resource environments, including Italy, France, Germany, Ireland, The Netherlands, the United Kingdom, and the United States, some of which involved international collaborations between low and high income countries and others of which involved only authors from high-income countries.

Ten submissions were accepted, five of which were from African countries (Ethiopia, Malawi, Cameroon, Nigeria, and Tanzania), one from India, one from the Philippines, and three others (two from Italy and one from The Netherlands). Three of the five presenters from Africa were

able to travel to the conference; the other two were denied visas. The track chairs coordinated with the SIGIR chairs to allow them to participate online.

6 Schedule

The track featured a full-day schedule with ten presentations, five in the morning and five in the afternoon. Each presentation was 18 minutes, including questions. The morning session consisted of five papers from Africa, followed by a panel discussion involving those five speakers that focused on "The Future of Information Retrieval in Africa." The panel focused on community building, key challenges, and potential solutions. Presenters from outside Africa presented in the afternoon session. In the evening, seven participants continued these discussions over dinner, and focused on the track's future and building an African IR research community.

7 Lessons Learned: Positives

There are several positive experiences that came from running the track for the first time. A few have been outlined below:

1. Inclusive process for presentation proposals: Our decision to solicit presentation proposals without specific format requirements served us well by generating a substantial number of submissions from participants who had not previously contributed to SIGIR from a broad range of countries.
2. Number of submissions: The number of submissions was adequate to identify 10 presentations that the program committee felt would contribute well to the program.
3. Topics: Presenters from Africa focused on the use of IR in food and agriculture, low-resource languages, or in both. Presenters from elsewhere addressed efficient implementation (in two cases) or low-data problems involving specific populations (children, autism, or Augmented and Assistive Communication (AAC)).
4. Mentoring process: The track chairs each worked closely with the authors of 3 or 4 papers to help them revise their submissions, prepare their presentations, and navigate travel and visa applications. This proved to be very helpful for participants who were working in low-resource environments.
5. Support for registration and accommodation: SIGIR organizers assisted track presenters with conference registration and accommodation, and this eliminated the need for upfront payments and reimbursement requests for these expenditures.

8 Lesson Learned: Challenges

Some challenges were met while running the track stemming from visa issues and reimbursement policies.

1. Timing of visa applications: It took some time after our acceptances were announced to finalize the type of support that would be provided to track participants who were traveling

from low-income countries. Participants from low to middle income countries (Africa, India and the Philippines) required extensive visa support and had a considerably more extended timeline than did the more typical visa support that the conference provided to all participants. Participants seeking visas were required to have accommodations arranged before applying for their visas. Several participants requested intervention from the conference general chairs as the time of the conference approached, which was provided, and which was usually successful. Because visa processing is a national activity, the responsiveness of consulates and embassies for future SIGIR venues may differ, so it is difficult to extrapolate from this one experience. Overall, we can say that much more time is needed, and that the demands on the conference organizers may differ from those for more traditional attendees.

2. Participation in conference sessions: The track session in the afternoon was well attended (30 or so non-presenters and non-organizers), but the morning session with presentations and the panel from Africa was not (less than 10 non-presenters and non-organizers). The explanation for this seems to be that some of the afternoon presenters were well known to the SIGIR community, and (at least) the two efficiency-focused presentations in the afternoon were more typical of high-interest topics for a SIGIR audience.
3. Reimbursement: ACM reimbursement policies require recipients to pay expenses such as airline tickets, ground transportation, and meals (other than those provided by the conference) upfront and request reimbursement afterward. This is a substantial burden for people working in low-resource environments. The other key problem was that we sought to reimburse all out of pocket costs, but the ACM system did not have provisions for reimbursing some costs (e.g., meals and visa fees).

9 Recommendations

Based on both the successes and challenges of running the track, several recommendations have been proposed to improve future iterations

1. Continue the low-resource environments track, but focus on tasks that are of specific interest in low-resource environments. This focus would exclude low-data problems that are not motivated specifically by considerations in low-resource environments, and it would exclude work conducted on standard information retrieval tasks such as ranking or evaluation that are not applied to tasks specific to low-resource environments, regardless of where that work is conducted.
2. Consider including at least one track organizer who is themselves working in a low-resource environment at the earliest practical opportunity, perhaps considering past track participants for this role.
3. Encourage the creation of a SIGIR workshop on applications to low-data challenges. There was clear interest in our three afternoon sessions this year that addressed children, autism, and AAC, but that part of our program was something a SIGIR workshop could have addressed, with typical SIGIR attendees presenting to other typical SIGIR attendees. The extensive support structure for the low-resource track simply is not needed to bring together such a group.

4. Because the cost of running the track is high, and depends on the willingness of national authorities to issue the necessary visas, consider being selective about how often and where to run the track. Running the track every other year, for example, would be much better than never running it again.
5. Move things earlier that can be done earlier. In particular, ACM Exec commitments for travel support should be made before acceptance decisions are made, even though that would require estimating the need rather than working off actual data from acceptances, and if housing is provided and direct billed (as was done this year) that should be reserved as a block in advance based on the same estimate so that visa applications can start immediately upon acceptance. It would also be good to move the submission deadline earlier, both because it would allow more time for reviews and because it would allow more time for arranging visas. The usual concern about separating track submissions from the paper submission deadline is not currently a concern for people who are working in low-resource environments.
6. Some presenters asked if the track would be accepting full research papers in the coming iterations. We do not think that this is a good idea in the near term if the track, as we recommend, narrows to focus in the ways we have suggested. But a new workshop on low-data applications in which experienced SIGIR authors would be among the participants could certainly include such adoption.
7. Without clear ownership, there is a risk that focus on the track will diminish over time. Establishing an LRE steering committee, or assigning responsibility to an existing committee (such as the diversity committee or the executive), would provide oversight and advocacy to ensure continuity. A dedicated committee would also create a sustainable pipeline for future organizers.
8. To continue strengthening the LRE community, we propose the following approaches:
   (a) Encourage continued participation: Invite previous track participants to submit their work again to this track (and, of course, also to SIGIR more generally).
   (b) Pre-submission support: Organize one or more online sessions where people who might be interested can learn about what makes for a good submission.
   (c) Mentoring: Continuing the mentoring process that we used this year, which markedly improved both the abstracts and the presentation in several cases.
   (d) Encourage collaborative research: One of the interesting things we learned was that several of the accepted presentations involved collaborations between established IR researchers and researchers who are working in low-resource environments.
9. The section below on Community Awareness includes two additional recommendations.

10 Benefits of the Track

Through discussions with participants from Africa we learned that the track provided valuable skills for preparing submissions to ACM conferences, exposure to cutting-edge research topics, constructive feedback on their work, and opportunities to engage with the wider research community. Those participants reported that their research interests and activities are now more

closely aligned with the SIGIR community, and some plan to submit papers to future SIGIR conferences. Several of our presenters from Africa might now be invited as short paper reviewers, as a step towards expanding the geographic reach of the SIGIR program committee.

Although our audience was small in the morning, we did find that the presenters from Africa did change our thinking. In particular, we learned that there are at least three basic application areas that are of interest in these low-resource environments: (1) IR in low-resource languages, (2) food and agriculture, and (3) health care. We had no papers on the third, but we did have submissions on those topics. Essentially, all our presenters were aware of these three topics as central, yet our shared understanding emerged only through our track session. Of course, other topics (e.g., education, or cultural heritage) might also be of interest, but knowing these three can help us with shaping and advertising calls, and selecting the program committee members.

11 Impact: Building a community beyond the track

The vision of the track is to establish a community of researchers working on low-resource IR problems, with particular attention to cases in which that work is itself being done in low resource environments. The track has already nurtured an emerging community by providing participants with SIGIR experience and involving reviewers from low-income and lower-middle-income countries. This recognises their contributions and integrates them into the wider IR community. It is essential to continue developing an IR community focused on challenges in low-resource environments. Such a community would enable researchers to share expertise, pool resources, and create tools and datasets to support those starting work in these settings. This collaboration would also facilitate information sharing about grants, identifying collaborators, and collectively improving the quality of research outputs.

This is one part of a longer-term strategy that started with the two AFIRM summer schools [Carterette et al., 2021], and indeed one of our five participants from Africa had participated in AFIRM so those two efforts have some degree of mutual reinforcement. Additional elements of the strategy should also be considered [Verberne et al., 2024]. For example, as networking improves there may be opportunities to continue what AFIRM started at lower cost (and with greater convenience) through online sessions, we might look for ways to build on the early work in building test collections for African languages (such as those created recently by Jimmy Lin or the earlier work in Mediaeval on speech retrieval) and SIGIR might also consider sponsoring an award for early career faculty who are working in low-resource environments. So the key question about impact is not in what the low-resource track achieves on its own, but rather where it fits in a broader engagement strategy.

It would also be worth considering whether greater attention to costs that are within the control of SIGIR might help to broaden participation. Registration fees are prohibitively high for researchers from low resource environments, and lodging expenses in many SIGIR venues are quite high as well. While we can not control expenses such as airfare, if special provisions for participants who are working in and traveling from low-resource environments could be made, it might be possible to cut total costs for such participants in half. That might be a useful step.

Of course, an even broader view would encompass what others have done, such as the growing set of faculty in Africa who have completed a Ph.D. in Europe, the role that institutions such as the University of Cape Town have played in similarly growing a network of IR faculty on that

continent, and how we can benefit from the activities of cognate research communities such as ACL and CHI in this space. In the longer term, we might also consider when it would be propitious to bring SIGIR itself to Africa.

Finally, we should recognize that this is a two-way relationship, as we as a community have much to learn from the work being done in low-resource environments as well. As we have noted, agriculture and health care are of particular interest in some low-resource environments, but of course those are also issues that the broader SIGIR community could benefit from engagement with. So creating workshops on those topics (or perhaps also on other topics such as support for governance and social justice) might result in richer interactions if organizers were to make special outreach efforts to people who are working in low-resource environments.

12 Community Awareness

Based on the poor attendance at the first track session, we believe that broader outreach and promotion within the SIGIR community is still needed. We therefore recommend including a plenary talk on research in low-resource environments on the second day of the main SIGIR program in order to increase the awareness of the types of work that will be presented in the low resource environments track, and then (as this year) schedule that track for the third day. This is a fairly expensive intervention that probably can be done only once, but if it is not done then it seems likely that attendance in the track session would again be low. If this is not feasible in 2026, then we would recommend that the low resource track be deferred to a future year in which such a plenary presentation would be possible.

We think that we might leverage the presence of researchers from low-resource environments at SIGIR in other ways as well. For example, this year we had a dinner for participants from low-resource environments, which was valuable, but if that dinner had been open to (and known about by) the full set of conference participants then there might have been even more productive interactions.

Financial Sustainability

Seven presenters from low-income and lower-middle-income countries were supported by grants to attend the conference in-person and present their work (although for visa reasons only five of those were able to come). As we noted above in the Recommendations section, this rather high cost might be a reason not to run a low-resource environments track every year. But if we do skip a year, we should do that intentionally, and not let the fact that we skipped one year become a reason to skip future years as well.

If there are years in which SIGIR chooses not to run a low-resource environments track, it would be worth considering whether there is some lower-cost engagement opportunity that could be crafted in that year. For example, perhaps those might be good years in which to encourage a workshop focused on limited data tasks, with special provisions for remote participation.

13 Summary

The inaugural low resource track was well received by researchers who were working in, and on, low-resource environments, and the track received a good number of submissions. Visa issues

were the main challenge. We have made a number of recommendations, but our most strident recommendation is that this be the first, but not the last, low-resource environment track. We believe that it is important to continue building and strengthening the community by creating opportunities for collaboration, sharing experiences, pooling resources, and providing collaborative learning opportunities to advance IR research in and on low-resource environments, and that this track can continue to contribute to that goal.

References

Ben Carterette, Hussein Suleman, and Douglas W. Oard. Report on the 1st acm sigir/sigkdd africa school on machine learning for data mining and search. *SIGIR Forum*, 53(1):3–13, March 2021. ISSN 0163-5840. doi: 10.1145/3458537.3458538. URL https://doi.org/10.1145/3458537.3458538.

Suzan Verberne, Hussein Suleman, Luca Soldaini, and Avijit Ghosh. Report on the sigir 2023 session on diversity, equity and inclusivity. *SIGIR Forum*, 57(2), January 2024. ISSN 0163-5840. doi: 10.1145/3642979.3642993. URL https://doi.org/10.1145/3642979.3642993.