



# City Research Online

## City St George's, University of London

**Citation:** Quintanilla Portugal, R. L., Eder, M., Thurman, N. & Haim, M. (2025). Responsible AI in Journalism: A Design Thinking Approach to Audience-Centered Topic Selection for Migrant Communities. Paper presented at the The European Data & Computational Journalism Conference 2025, 8-10 Sep 2025, Athens, Greece.

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

**Permanent repository link:** <https://openaccess.city.ac.uk/id/eprint/35833/>

**Copyright and Reuse:** Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

# Responsible AI in Journalism: A Design Thinking Approach to Audience-Centered Topic Selection for Migrant Communities

Roxana L. Quintanilla Portugal<sup>\*1,2</sup>, Maximilian Eder<sup>1</sup>, Neil Thurman<sup>1</sup>, Mario Haim<sup>1</sup>

<sup>1</sup>Department of Media and Communication, LMU Munich, Germany

<sup>2</sup>Departamento Académico de Informática, UNSAAC, Perú

{roxana.portugal, maximilian.eder, neil.thurman, mario.haim}@ifkw.lmu.de

## ABSTRACT

Local newspapers have faced declining readership and advertising revenues, alongside growing skepticism from audiences with migration backgrounds. This undermines local journalism's democratic role. To address this, we conducted a participatory Design Thinking (DT) project with a German regional newsroom to develop a low-fidelity prototype of a responsible AI tool to improve engagement with migrant audiences. Two focus groups with individuals from immigrant backgrounds informed five DT workshops with newsroom staff and community partners, following the Stanford d.school framework. The resulting prototype includes three dashboards for topic selection, editorial support, and post publication feedback. Guided by a value-sensitive design approach grounded in transparency, engagement with audience perspectives, and responsive feedback mechanisms, this paper presents a practical solution for embedding responsible AI in local journalism workflows to support democratic and organizational sustainability.

## CCS CONCEPTS

• General and references → Cross Computing Tools and Techniques → Design

## KEYWORDS

Responsible Artificial Intelligence, Local Journalism, Design Thinking, Value-Sensitive Design, Migrant Audiences, Participatory Design, Topic Selection, Audience Engagement.

## 1 Introduction

The internet disrupted local newspapers' legacy business models, reducing advertising and subscription revenues. Editorial decisions are increasingly shaped by commercial pressures often leaving smaller groups—such as migrants—underrepresented [1]. Strengthening connections with these audiences serves both a democratic and economic purposes.

In collaboration with a regional newsroom in Germany, we carried out a participatory Design Thinking (DT) process to co-develop a prototype of a responsible AI tool aimed at improving engagement with migrant communities. Before the DT workshops, we conducted two focus groups with people from migrant backgrounds (e.g., Turkey, Syria) to explore their information needs and perceptions of local journalism. Participants reported consuming news mainly via social media and news aggregators (e.g., Google News), with little engagement with local journalism. While interested in municipal politics, they felt marginalized, noting negative biases and lack of content diversity. These concerns echoed findings by Ross Arguedas et al. [2] across disadvantaged communities in Brazil, India, the UK, and the US.

More inclusive editorial practices and audience-aware topic selection are essential to address these gaps. Yet newsrooms often lack the time and resource to engage deeply with underrepresented groups. AI tools offer promise: as Lin & Lewis [3] note, they can amplify diversity by improving multilingual access and translation. When designed to support accuracy, accessibility, diversity, relevance, and timeliness, AI can align editorial practices with democratic values. Poorly implemented, however, AI risks reinforcing biases [4], and further alienating these audiences.

To translate focus group insights into concrete software requirements, we conducted five DT workshops (October 2024 - January 2025) following the Stanford d.school framework. Journalists, editors, and product staff used participatory methods (e.g., LEGO Serious Play, Crazy 8s, paper prototyping) to examine newsroom practices and envision improvements. The newsroom staff interpreted the focus group results as indicating that shortcomings largely stemmed from topic selection and thus chose this stage as the focus for change. We also applied early goal modelling with the i\* framework [5][6] to capture newsroom motivations, dependencies, and constraints. This revealed a key challenge: limited structured knowledge about audiences and weak feedback mechanisms. Addressing this gap requires both technological support and balance between two priorities—on one hand, editors need actionable audience insights, repeatedly described in the workshops as “audience transparency”; on the other, audiences expect visible impact from their feedback, which they interpret as editorial accountability.

\*Corresponding Author

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the authors.

*DataConf 2025, September, 2025, Athens, Greece*

© 2025 Copyright held by the authors.

The resulting low-fidelity prototype comprises three interconnected dashboards: (1) topic selection, with AI-enhanced filters highlighting stories relevant to migrant audiences; (2) editorial support, offering contextual suggestions and value-aligned feedback; and (3) post-publication feedback, tracking content performance and generating recommendations. Together, these dashboards illustrate how value-driven AI can help local journalism balance economic goals with its democratic mission and rebuild trust with underserved audiences.

The article is structured as follows: Section 2 outlines the DT process and methods; Section 3 reviews related work; Section 4 concludes with avenues for future research.

## 2 DT process for Audience-Centered Topic Selection.

### 2.1 Empathizing: Exploring Editorial Decision-Making

The first phase of the DT process explored how journalists select topics [10][11], and how these choices align with the expectations of migrant audiences [2].

In the initial exercise, participants used a visual association task, assigning cards [12]—visual prompts with symbolic images and metaphors—to political headlines from local newspapers (Figure 1). This aimed to uncover the emotional responses that news imagery can trigger, fostering empathy with readers’ perspectives.

Next, participants carried out a card-sorting task (Figure 2). They classified two articles according to established newsworthiness factors [10][11] and journalistic principles (e.g. German Press Council [13]). In a second round, the same method was applied using audience expectations derived from previous studies [2][14].

After the workshop, the identified quality dimensions (newsworthiness factors, journalistic principles, and audience demands) were modeled with the NFR framework [9] to analyze their interrelations (Figure 4). This task is complex, as it usually requires explanatory descriptions from transcripts, or extensive modeling. To support early ideation, we used ChatGPT (output available upon request) to generate example statements, not as validated claims but as prompts. For instance:



Figure 1: **Headline prompting reflection on political imagery’s impact**

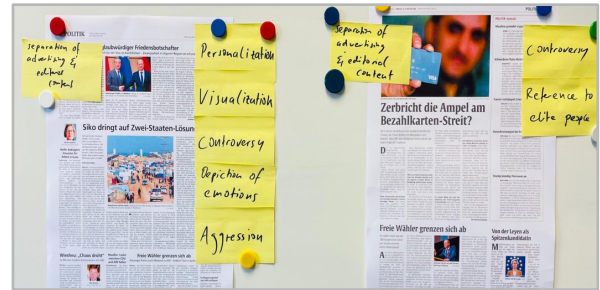


Figure 2: **Card-sorting of two political articles by principles and newsworthiness.**

*“Depiction of emotions has a high contribution to both impartiality and diversity, as emotional storytelling can enhance the understanding of different perspectives; it has a moderate contribution to accessibility, as emotional narratives can aid comprehension.”*

Such prompts have proven useful in early-stage design [15] and align with what Göpfert et al. [16] describe as argumentative discourse, making explicit the reasoning often left implicit in engineering problem-solving.

According to the NFR framework, the relationships in Figure 4 can be interpreted as follows: Depiction of emotions is satisfied only if satisfactorily supported by accessibility and authenticity. In turn, authenticity is achieved only when positively influenced by personalization and truthfulness. Under these conditions, authenticity contributes positively both to newsworthiness and to reach, showing how interconnected qualities shape topic selection.

This type of reasoning is necessary not only for machines to balance competing qualities, but also for tool designers to recognize the inherent complexity of developing value-driven systems for journalism.

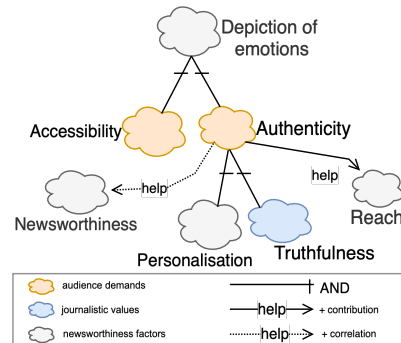


Figure 4: **Emotional Depiction and its Quality Dependencies in Topic Selection**

### 2.2 Defining: Mapping the Information Flow

The second phase of the DT process sought to define the core challenge of audience-centered topic selection. The workshop mapped the information flow—input, throughput, and output—and clarifying the responsibilities of newsroom stakeholders at each stage.

In the first session, participants built LEGO models to depict how topic selection operates in their daily work. This revealed new insights, for example “only the relevant parts of reality are shown to each target group, since not everyone wants all the available information” (Participant 1). The modeling was then refined iteratively: news factors were added in a second round, and later a more structured model emerged (Figure 5) addressing value tensions in light of the quality map (Figure 4).




Figure 5: LEGO model illustrating audience, transparency, and geographical focus

Afterwards, participants explained their models by answering guiding questions on challenges and solutions. Two key questions emerged: “What do our target groups want?” and “How do we reach different target groups?” To explore this further, they reflected on their department’s topic selection using established news values. The first question was linked to keywords such as data use, personalization, reach, and influence; the second to values like unexpectedness, continuity, unambiguity, reference to persons, and frequency. These challenges were then mapped across the information flow.

The workshops’ results were synthesized into a consolidated *i\** model [21], mapping key stakeholders (editorial staff, audiences, political actors, and paid content teams) and their goals. This revealed the absence of audience feedback mechanisms as a central barrier. To address it, we proposed a feedback loop—topics selected → articles produced → articles distributed → feedback collected → topic selection refined—which mirrors the news cycle but explicitly incorporates input from migrant audiences.

### 2.3 Ideation

The third phase of the DT process aimed to identify mechanisms for gathering audience feedback, focusing on two key questions: how to achieve “transparency about the audience” and how to gain “more perspectives from readers”. To address these, participants engaged in creative exercises combining individual reflection and collaborative exploration.

Using the Image Elicitation method, they assigned visual trigger cards to express perceptions of these questions. In discussing “transparency about the audience,” they highlighted *background knowledge, different challenges, diverse perspectives, unknowns, and emptiness*, underscoring both the

multiplicity of audience standpoints and the newsroom’s gaps in understanding them. Metaphors such as *different colors or tones* symbolized diversity, while the aspiration “*my view on life is understood and represented*” reflected the importance of inclusion and recognition (see Figure 6).

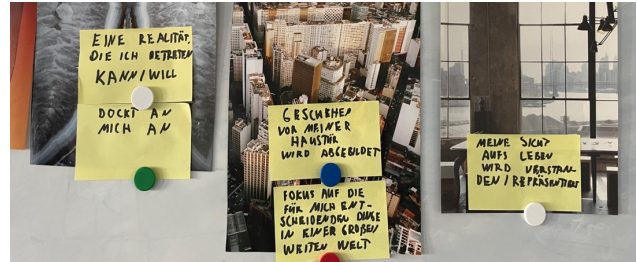


Figure 6: Visual trigger cards and keywords reflecting participants’ perceptions of audience

This was followed using the Crazy 8’s method, which encouraged participants to generate multiple ideas rapidly, and the LEGO Serious Play® method, which allowed them to model envisioned solutions. These exercises helped turn abstract concerns into concrete mechanisms.

Proposed mechanisms included reader surveys, direct contact with influential readers, analysis of diverse data sources, and observing the practices of other media. Participants also stressed overcoming language barriers and strengthening the newsroom’s digital presence. Emerging guiding principles were closeness, credibility, authority, and familiarity.

They also formulated new questions to better understand audience needs, such as:

*Do readers like our newspaper? What other media do they trust? What do they do after work? How do they spend their money? For what kinds of organizations do they work? With whom do they interact?*

Addressing these questions in the prototype would provide a stronger basis for profiling target audiences.

### 2.4 Prototype: Developing a low-Fidelity AI Tool

In the fourth phase of the DT process, participants created low-fidelity prototypes of an AI tool to support topic selection across the information flow. Building on audience-centered questions from the ideation phase, we prepared pre-designed metrics with the support of ChatGPT, which generated example prompts based on publicly available knowledge [15] (Figure 7). This use of the LLM served to stimulate participants’ critical evaluation rather than provide authoritative answers by illustrating how abstract audience questions could be translated into concrete, measurable elements. Data sources such as Google Analytics and social media were mentioned by participants themselves.

The workshop used inquiry-driven data mapping, helping participants to familiarize themselves with the types of data needed for each metrics. For example, to answer “What do they do after work?”, they explored metrics such as those in Figure 7.

**Metrics:**

- Evening Traffic Share (ETS):**
  - Formula:**  $(\text{UsersBetween6:00PMand10:00PM} / \text{TotalUsers}) \times 100$
  - Description:** Percentage of website traffic occurring during post-work hours.
- Recreational Content Engagement (RCE):**
  - Formula:**  $(\text{InteractionsWithRecreationalContent} / \text{TotalInteractions}) \times 100$
  - Description:** Measures the level of engagement with topics like entertainment, hobbies, or lifestyle.
- Time Spent on Lifestyle Pages (TSLP):**
  - Formula:**  $\text{AverageTimeonLifestylePages} / \text{TotalTimeonSite}$
  - Description:** Tracks the time users spend on non-work-related content.

Figure 7: Metrics to answer the question “What do they do after work?”

These metrics were presented on pre-designed cards, one per question, enabling participants to evaluate relevance and note missing data. The cards illustrated how abstract questions could be operationalized into concrete data ready for integration into the prototypes.

In the low fidelity prototyping phase, participants sketched dashboard mockups that incorporated the selected metrics, visualizing how the prototype could support editorial decision-making. During this process, the prototype surfaced more specific functionalities—such as headline suggestions, archive integration, and feedback mechanisms—illustrating the importance of iterating on a prototype to refine and extend its design.

### 2.5 Testing

The final workshop tested and refined the low-fidelity prototype, which supported three dashboard stages: before

writing (topic selection), while writing (text editing), and after writing (feedback). Participants from the local newsroom, including editors and a product manager, engaged in two sessions.

In the first session, a cognitive walkthrough allowed participants to propose improvements. For topic selection, they stressed workflow transparency, topic prioritization, and audience data integration. For text-editing, they emphasized editorial criteria (e.g., relevance, virality), content validation, and adapting text to emotions, target groups, and political perspectives. For feedback, they highlighted the need to separate qualitative and quantitative input, linking suggestions to content management system (CMS) IDs, and integrating reader comments. Concerns included data reliability, language selection, and evaluating complex parameters.

The second session focused on usability. Participants suggested integrating the text-editing dashboard into the CMS, adding analytics and predictive features, and ensuring explainability of how AI suggests topics or target groups to build editorial trust. Figure 8 depicts the dashboard for topic selection.

While some features prototyped involve AI—such as an *AI Coach* suggesting alternative headlines—others focus on data mining. The prototype, available in [21] (German), has no confirmed implementation plans, though the partner newsroom expressed satisfaction in a LinkedIn post.

### 3 Embedding into related work

DT is increasingly applied in journalism and AI research to foster co-creation. Dimitrakopoulou and Lewis [17] highlight the value of *reflective listening*, and Miller and Rollnick [18] stress *structured summarization* to help participants clarify and deepen their input. In our workshops, having a moderator with expertise in both DT and value-sensitive design was crucial, as some activities alone did not achieve sufficient depth. Moderated *paraphrasing* helped participants articulate perspectives more fully and reflect on underlying values.

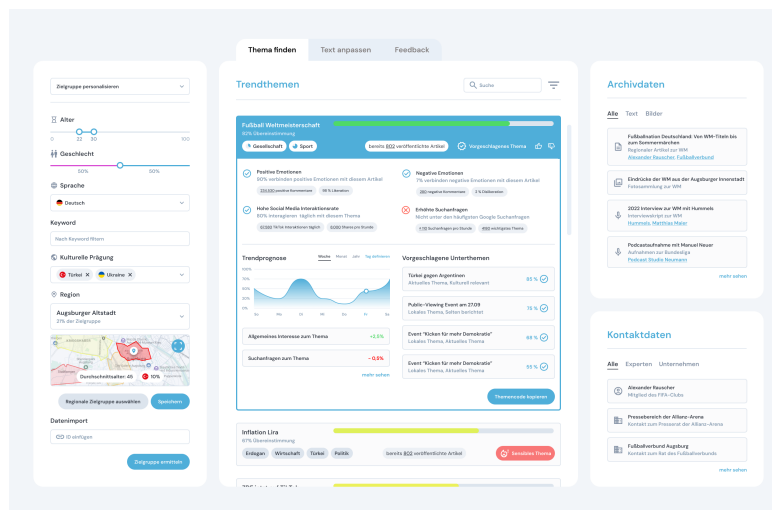


Figure 8: Final prototype of the Topic Selection Dashboard after two iteration rounds (see [21] for full details).

Halskov and Lundqvist [19] emphasize defining a *design space* to organize workshop knowledge. We implemented such a space through the *i\** framework [5][6], modeling journalistic qualities, their operationalizations, and role-specific responsibilities and interdependencies. This enabled us to pinpoint challenges such as retrieving audience feedback transparently. Our approach aligns with Kolko [20], who notes that design often involves making decisions about things that do not yet exist, which requires *lateral thinking* and openness to emerging insights. Finally, a professional UX designer transformed participant-driven prototypes into a refined, implementation-ready version, increasing their practical applicability.

## 4 Conclusion

Designing responsible AI for journalism is challenging, as it requires balancing intertwined quality dimensions: traditional news factors, journalistic principles (e.g. those of the German Press Council), and audience-driven quality demands. These often conflict, making tensions hard to detect without structured methods. Support systems must combine multiple complementary approaches, as demonstrated in this study.

Our experience shows that DT can be applied flexibly, allowing participants to co-create both process and outcomes. Mapping insights into the newsroom's workflows proved crucial to identifying where information flows break down, while modeling social responsibilities and role dependencies with the *i\** framework [5][6] exposed gaps—particularly those linked to the reader—with no mechanisms to resolve inter-role challenges.

A key observation is that audience-centered systems—here based on newsroom staff's perceptions of audience needs—tend to be inherently data-driven. Maintaining traces of reader opinions is vital for explanation, accountability, and improvement, yet raises tensions with regulations such as GDPR. As a result, systems often rely on aggregated data, limiting the ability to detect, for instance, whether a reader belongs to local migrant communities. Achieving such identification would require multiple inference mechanisms, which not only reduce system efficiency but also risk violations of GDPR. Identifying actionable ways to enhance audience-centered topic selection while respecting privacy thus remains an open challenge and a priority for future work.

## ACKNOWLEDGMENTS

This Design Thinking project is part of the research project "Towards Responsible AI in Local Journalism", funded by the Volkswagen Foundation. We wish to thank Media Group pd for their cooperation in this Design Thinking project.

## REFERENCES

- [1] Lee Salter. 2016. Local journalism: The decline of newspapers and the rise of digital media. *European Journal of Communication* 31, 5 (October 2016), 606–609. <https://doi.org/10.1177/0267323116670811>

- [2] Amy Ross Arguedas, Sayan Banerjee, Camila Mont'Alverne, Benjamin Toff, Richard Fletcher, and Rasmus Kleis Nielsen. 2023. *News for the powerful and privileged: How misrepresentation and underrepresentation of disadvantaged communities undermines their trust in news*. Reuters Institute for the Study of Journalism. <https://doi.org/10.60625/RISJ-JQNY-T942>
- [3] Bibo Lin and Seth C. Lewis. 2022. The one thing journalistic AI just might do for democracy. *Digital Journalism* 10, 10 (November 2022), 1627–1649. <https://doi.org/10.1080/21670811.2022.2084131>
- [4] Nicholas Diakopoulos, Hannes Cools, Charlotte Li, Natali Helberger, Ernest Kung, Aimee Rinehart. 2024. *Generative AI in journalism: The evolution of newswork and ethics in a generative information ecosystem*. The Associated Press. <https://doi.org/10.13140/RG.2.2.31540.05765>
- [5] E.S.K. Yu. 1997. Towards modelling and reasoning support for early-phase requirements engineering. In *Proceedings of ISRE '97: 3rd IEEE International Symposium on Requirements Engineering*, 1997. IEEE Comput. Soc. Press, Annapolis, MD, USA, 226–235. <https://doi.org/10.1109/ISRE.1997.566873>
- [6] Xavier Franch, Julio Cesar Sampaio Do Prado Leite, Gunter Mussbacher, John Mylopoulos, and Anna Perini (Eds.). 2024. *Social modeling using the i\* framework: Essays in honour of Eric Yu*. Springer Nature Switzerland, Cham. <https://doi.org/10.1007/978-3-031-72107-6>
- [7] Batya Friedman, David G. Hendry, and Alan Borning. 2017. A survey of value sensitive design methods. *FNT in Human-Computer Interaction* 11, 2 (2017), 63–125. <https://doi.org/10.1561/11000000015>
- [8] Roxana L. Q. Portugal, Bartosz Wilczek, Maximilian Eder, Neil Thurman, & Mario Haim. 2023. "Design thinking for journalism in the age of AI: Towards an innovation process for responsible AI applications." *Proceedings of the Joint Computation + Journalism European Data & Computational Journalism Conference*, Switzerland. <https://openresearch.surrey.ac.uk/99947866602346>
- [9] Lawrence Chung, Brian A. Nixon, Eric Yu, and John Mylopoulos. 2000. *Non-functional requirements in software engineering*. Springer, Boston
- [10] Tony Harcup and Deirdre O'Neill. 2001. What is news? *Galtung and Ruge revisited*. *Journalism Studies* 2, 2 (May 2001), 261–280. <https://doi.org/10.1080/14616700118449>
- [11] Nora Denner, Thomas Koch, and Sascha Himmelreich. 2018. News selection within customer magazines: A quantitative survey among editors-in-chief in Germany. *Journalism Practice* 12, 7 (August 2018), 888–900. <https://doi.org/10.1080/17512786.2017.1343092>
- [12] Trigger Tool. 2023. Creative visual prompts for participatory workshops. Retrieved May 13, 2025, from <https://www.trigger-tool.de>
- [13] German Press Council. [Online]. Available: <https://www.presserat.de/en.html>. [Accessed: May 20, 2025].
- [14] Thomas R. Schmidt, Lisa Heyamoto, and Todd Milbourn. 2019. The social construction of media trust: An exploratory study in underserved communities. *Journal of Applied Journalism & Media* 8, 3 (November 2019), 257–271. [https://doi.org/10.1386/ajms\\_00001\\_1](https://doi.org/10.1386/ajms_00001_1)
- [15] Kevin Ma, Daniele Grandi, Christopher McComb, and Kosa Goucher-Lambert. 2024. Exploring the capabilities of large language models for generating diverse design solutions. *arXiv*. <https://doi.org/10.48550/ARXIV.2405.02345>
- [16] Jan Göpfert, Jann M. Weinand, Patrick Kuckertz, and Detlef Stolten. 2024. Opportunities for large language models and discourse in engineering design. *Energy and AI*, 17, 100383 (September 2024). <https://doi.org/10.1016/j.egyai.2024.100383>
- [17] Dimitra Dimitrakopoulou and Seth C. Lewis. 2023. The generative dialogue framework and the pursuit of better listening by journalists: A design-centered approach for more constructive conversations with audiences. *Digital Journalism* 11, 3 (March 2023), 547–568. <https://doi.org/10.1080/21670811.2022.2075415>
- [18] Miller, W.R. and Rollnick, S., 2012. *Motivational interviewing: Helping people change*. Guilford press.
- [19] Kim Halskov and Caroline Lundqvist. 2021. Filtering and informing the design space: Towards design-space thinking. *ACM Trans. Comput.-Hum. Interact.* 28, 1 (February 2021), 1–28. <https://doi.org/10.1145/3434462>
- [20] Jon Kolko. 2018. The divisiveness of design thinking. *interactions* 25, 3 (April 2018), 28–34. <https://doi.org/10.1145/3194313>
- [21] Roxana L. Q. Portugal, Kilian Frey, and Maximilian Eder. 2025. Prototyping Responsible AI for Local Journalism – Audience-Centered Topic Selection. In *The European Data & Computational Journalism Conference 2025 (DataJConf 25)*, Athens. Zenodo. <https://doi.org/10.5281/zenodo.17013207>