Guerrilla Badges for Reproducible Geospatial Data Science

Daniel Nüst, Lukas Lohoff, Lasse Einfeldt, Nimrod Gavish,
Marlena Götza, Shahzeib Tariq Jaswal, Salman Khalid, Laura Meierkort,
Matthias Mohr, Clara Rendel and Antonia van Eek

Institute for Geoinformatics (ifgi)
University of Münster
Münster, Germany
daniel.nuest@uni-muenster.de

Abstract

The building blocks of research are developing at an unprecedented pace. Data collection, analysis, interpretation, presentation, review, and publication take place completely on computers. The final product often is still a static document with only limited links to the underlying digital material, making transparency and reproducibility a challenge. In this work we apply the mechanism of badges to provide prominent connections to underlying analyses environments and important (meta-)data to readers of scholarly publications in geospatial data science. An API specification and implementation for a badge server provide extended and regular badges. The badges leverage recognition value for executability, licensing, spatial extent, and peer-review metadata – base information which either is or should be made available. We show a client-side integration method across many third-party platforms that allows evaluation of badges in realistic scenarios. The server and client enable an independent spreading of badges. The open source publication of all used software enables reproducibility and extensibility. The badges show potential to enhance interaction with scholarly works and should be evaluated with academic users in the future.

Keywords: badges, open science, data science, publication infrastructures, scholarly publication, research compendium

1   Introduction

The building blocks of research are constantly developing, though arguably at an unprecedented pace in the current age of digitisation. Data collection, analysis, interpretation, presentation, review, and publication take place completely on computers. However, the main outcome still is often a static document (e.g. an HTML or PDF file) resembling the traditional form of dissemination – the research paper. Thus Buckheit & Donoho (1995) postulated: *"An article about computational science in a scientific publication is not the scholarship itself, it is merely advertising of the scholarship."* The typical research paper provides only limited links to the underlying building blocks of the actual scholarship, such as input datasets, software/hardware environment, workflow code, or output data. Therefore reproducibility and reusability, both cornerstones of science, have been identified as important challenges in geospatial data science (Nüst et al., 2018; Konkol et al., 2018). Efforts to improve the publication of and access to data and software, e.g. establishing citation principles (Wilkinson et al., 2016; Katz & Chue Hong, 2018), exist and they work (Piwowar & Vision, 2013). Practicing Open Science (Ibanez, 2014) and the advantages of openness, transparency and reproducibility, e.g. efficiency and collaboration effects, are clear (cf. Markowetz, 2015). Research compendia, a term first used by Gentleman & Temple Lang (2007) and since then taken up and extended[1], are but one concept to package all buildings blocks of a piece of research. Nevertheless these practices are not common yet.

In this work we use the concept of badges to expose, not only advertise, the building blocks of scholarship. Badges are an established artefact in the software development community to visually highlight important pieces of information, exploiting a high recognition value. A user can quickly grasp the current version of a piece of software of interest, whether its test suite completes successfully or fails, or whether a tool is available in an established public repository for easy installation. Gaining these badges, and keeping them "green" in the case of tests, works as a motivator for developers. In science, the core medium to disseminate work between users, i.e. scientists, is the research paper. Badges for relevant building blocks behind research papers could benefit both users of this medium. Readers could quickly assess multiple or single publication items. Authors are encouraged to share more complete information (cf. Grahe, 2013) at the prospect of a larger readership and reuse. Relevance is specific to each reader, e.g. a paper may use data from the same area of interest or may contain transferable methods. For reproducible geospatial data science, we see the following questions as crucial for readers to decide if a work is interesting for them, e.g. in a literature study, and badges could help to answer them: *Is all code, data, and documentation openly available? Is a software environment documented so the results can be reproduced? What is the area of interest or data location?*

In the remainder of this work we first give a detailed background on badges in science. Then we present and discuss the first prototype of an API, server, and client implementation for creating and spreading badges on scholarly communication platforms.

---

[1]   See https://research-compendium.science for a full list of recommendations and guidelines.

2 Related Work

Digital badges to show a specific accomplishment are popular accessories to "earn" in Free and Open Source Software (FOSS) development. They are awarded by platforms providing a service, or by third-party websites based on metadata available via APIs of said platforms. Developers must only follow common practices to provide the required information, e.g. structured metadata in a project description. They then include the badges in their documentation to show relevant bits of information to their users. Most badges are generated with current data each time they are loaded. They show a tuple of property name and value, and may use colour to visually distinguish property values. Badges can include icons, e.g. a logo, and are provided in different formats, e.g. vector (as SVG – Scalable Vector Graphics) or raster (PNG) graphics. *Shields.io* (https://shields.io/) is a popular badge service. It provides badges for example for software repositories (e.g. software version, number of downloads), license, popularity (download count, ratings), or build systems (e.g. status of automated tests). Shields.io also renders own information by providing text and styling information within a URL. Other badge services cover specific use cases, e.g. *MicroBadger* (http://microbadger.com/) provides images for container images published on Docker Hub.

Due to the high recognition value, badges have been picked up by platforms and groups in a scientific context, including several journals. These mostly show static content. The data repository *Zenodo* (http://zenodo.org/) and the journal *JOSS* (http://joss.theoj.org/) provide badges with the Digital Object Identifiers (DOIs) of records. The *ROpenSci* initiative uses them for different stages of its software review process (https://badges.ropensci.org/). The *Binder* project uses badges to advertise the availability of an interactive notebook for a project repository. Examples of these badges are shown in Figure 1.

Figure 1: Badges from (clockwise beginning at top left) JOSS, MicroBadger, Binder, Zenodo, and ROpenSci Review.

The Center for Open Science (COS, https://cos.io/) designed badges for acknowledging open practices in scientific articles (see https://osf.io/tvyxz). COS offers guidelines for incorporation into peer reviews and adding badges to documents. The badges are *Open Data*, *Open Materials*, and *Preregistration* of studies (see Figure 2) and are adopted by over a dozen of journals to date[2]. A study by Kidwell et al. (2016) reports a positive effect from the introduction of open data badges in the journal *Psychological Science*: After the journal started awarding badges for open data, more articles stating open data availability actually published data (cf. Baker, 2016). The COS badges are effective in promoting data publishing and show availability and transparency, but not geospatial aspects or reproducibility.

---

2   https://osf.io/tvyxz/wiki/5.%20Adoptions%20and%20Endorsements/

Figure 2: Full COS badges (from left to right: open data – blue, open materials – yellow, preregistered – red) in colour, reduced black-and-white versions for print output also exist.

Source: https://osf.io/tvyxz/wiki/home/.

Peng (2009, 2011) and Rowhani-Farid & Barnett (2018) report on the usage of badges in the journal *Biostatistics*: a set of "kite marks" led to a moderate increase in data sharing. Marks D and C are awarded if data respectively code is provided, and R if results were successfully reproduced during the review process (implying D and C). Figure 3 shows the usage of R on an article's title page.

Figure 3: *Biostatistics* kite mark R rendering (top right part of the page) in the PDF version of the paper Lee et al. (2009).

Source: https://academic.oup.com/biostatistics/article-pdf/10/3/409/17736633/kxp010.pdf.

The Association for Computing Machinery (ACM, https://www.acm.org/) provides a common terminology and standards for artefact review processes for its conferences and journals[3]. Three badges with several levels (see Figure 4) are awarded using specific criteria, e.g. the *Evaluated* badge means artefacts were made available to reviewers and it has levels *Functional* or *Reusable*. The ACM badges provide excellent information on reproducibility for human readers, but not on geospatial information and not across platforms. The Graphics Replicability Stamp Initiative[4] (GRSI) organises a community-driven additional evaluation process for computer graphics research. Its results are the basis for different badges for a number of journals and conferences, e.g. ACM's badges for *ACM Transaction on Graphics*.

Figure 4: ACM badges, from left to right: *Artifacts Evaluated – Functional/Reusable* (pink/red), *Artifacts Available* (green, no evaluation), *Results Replicated* (light blue, artefacts provided by author), and *Results Reproduced* (dark blue, without author-supplied artefacts).

Source: https://www.acm.org/publications/policies/artifact-review-badging.

---

3   ACM policies: Artifact Review Badging, see https://www.acm.org/publications/policies/artifact-review-badging; Example article (badges rendered on landing page and PDF): https://doi.org/10.1145/3197517.3201397
4   http://www.replicabilitystamp.org/

These examples show the potential, diversity, and challenges in describing and awarding badges. This work explores novel badge types on an article's reproducibility and spatial area of interest, and an independent distribution mechanism, to contribute to the landscape of badges in scholarly publishing.

3 Geospatial Data Science Badges

3.1 Sources for Badge Information

The basis for useful badges are reliable data sources. For publications in geospatial data science, there are no established platforms or metadata protocols beyond regular article metadata, i.e. for code or data licenses or spatio-temporal extents of used datasets. Publication date, peer-review type (e.g. blind or double-blind), and license (i.e. if article is Open Access) are provided via online library reference APIs, namely Crossref (http://crossref.org/) or DOAJ (https://doaj.org/). Search terms are article title or DOI. The other properties are accessible via the prototypical o2r reproducibility service (Nüst, 2018). It provides access to the metadata of Executable Research Compendia (ERC) via the DOI of the related article. ERC contain all data and code used in a particular workflow and their creation process includes automatic extraction and user validation of metadata, including a spatio-temporal bounding box, for increased transparency and reproducibility. Both sources are used for geospatial research badges, but due to the prototypical state of the o2r API following examples rely on mock-up data.

Figure 5: Extended badges on an article page on DOAJ.org, integrated below the Abstract section.

### 3.2 Badge types, badge design, and an API

A RESTful[5] Application Programmer Interface (API) defines routes (i.e. URLs) to access badges of four types to answer the discovery questions in geospatial data science (see above):
- **executable**: code and runtime reproducibility
- **licence**: licensing information or all building blocks
- **spatial**: publication's geospatial area of interest
- **peerreview**: type of peer review

For these types we designed badges at two levels of detail, regular and extended. The extended badges contain a higher level of detail, while regular badges aggregate information to be suitable for search result listings, where they allow a visual comparison of hits, see Figure 5. Extended badges are more open in their design, while regular badges follow common badge styling.

The extended license badge has three categories (code, data and text), which are aggregated to simpler text ("open", "partially open", "mostly open", "closed") in the regular badge (see Figure 6).

Figure 6: Extended badges reporting different values for data, text and code licenses (left, middle) and regular badge (right, not to scale, based on Shields.io).
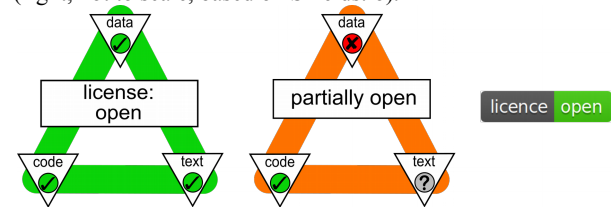
The extended spatial badge shows the bounding box as an interactive map, whereas the regular spatial badge only shows a suitable area name.

Badges for reproducibility, peer review status and license are colour coded to provide visual aids. They indicate for example (un)successful reproduction, a single-blind peer review process, or different levels of open licenses.

The API provides regular badges with HTTP GET requests, i.e. URLs following the pattern `/api/1.0/badge/:type/:doi`, where `1.0` is the API version, `:type` is the badge type, and `:doi` is the publication's DOI. Extended badges are returned when `/extended` is appended to the URL.

3.3 Implementation

Badger implements the **server-side API**. It queries public APIs to elicit metadata and provide the aforementioned badges types. It uses two badge generation methods: (a) internally created SVG-based badges, and (b) redirects to Shields.io, where the information is encoded in the Shields.io-URL, which is generated on the fly. All badges can be requested at specific size and pre-rendered as a PNG image for compatibility. The process for generating the executability badge for a paper "Global Air Quality and Pollution" from *Science* identified by the DOI 10.1126/science.1092666 is as follows.

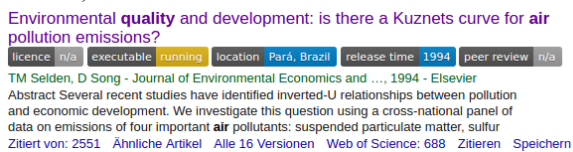5 https://en.wikipedia.org/wiki/Representational_state_transfer

1. A client calls the URL `https://badger/api/1.0/badge/executable/10.1126%2Fscience.1092666`.
2. *Badger* queries the o2r API for ERC connected to the given DOI. If it exists, it queries for the latest reproduction result status (a "job" in the API).
3. Depending on the result (success, running, or failure) specified in the job metadata, *Badger* generates a Shields.io-URL with green, yellow or red colour and matching property value.
4. The client displays a badge: `executable yes`.

If an extended badge is requested, *Badger* generates an SVG graphic or an embeddable HTML snippet. For the spatial badge it converts coordinates into textual information, i.e. country and if available district or place name, using the Geonames API[6], see Figure 7. When *Badger* does not find data for a certain DOI, it returns a grey "not available" - badge, see outermost badges "license" and "peer review" in Figure 7. Such a null result, e.g. "no spatial data included", can be equally helpful during discovery.

Badges are most successful when they are widely used and consequently quickly recognised by users. Though a desirable and more sustainable approach, it is unrealistic that (competing) publishers agree on a common badge system and design. Therefore we took an unusual approach to augment existing platforms for discovery of papers using a Chrome browser extension[7], similar to Unpaywall browser extension[8]. The *Extender* implements **client-side badge integration**. It inserts badges into search results or article pages using client-side browser scripting, also known as userscripts, on several websites including DOAJ (https://doaj.org/), Google Scholar (https://scholar.google.de/), PLOS (https://www.plos.org/), Microsoft Academic (https://academic.microsoft.com/), and ScienceDirect (http://www.sciencedirect.com/),

For each article displayed on these websites, either in a search result listing or dedicated article pages, *Extender* parses the DOI from the page's HTML code, requests badges from *Badger*, and inserts them into the page. The parsing and insertion is tailored to each supported website. Figure 7 shows an exemplary result. When the DOI is not directly provided, *Extender* queries the Crossref API with the paper title and uses the returned DOI if the result is unambiguous. *Extender* also inserts controls for filtering search results using badge values and for selecting displayed badges as shown in Figure 8.

Figure 7: Regular badges integrated into Google Scholar search result listing between title and authors (partial screenshot).

4   Discussion

A badge server for scholarly publications has the potential to improve discovery workflows for scientists by aggregating information, including underlying spatial data, with a high recognition value. It can enable identification of related work and reusability – an important aim of reproducibility. It also demonstrates badges as a mean to communicate more complex information compared to existing approaches.

An independently operated *Badger* and client-side integration with *Extender* may be favourable to a complex process of establishing a single set of badges across all involved stakeholders. This "guerrilla" approach allows to bring a new concept onto researcher's computers beyond a specific research project's own software or websites in a secure and reliable manner. It can also facilitate long-term studies, because users are exposed during their regular work and not only in a lab setting. However, the realised userscript integration into websites is less stable than an actual integration in platform APIs would be, because any UI change or code change may break the userscript. The Open Source nature of *Badger* and *Extender* allows research domains to adopt criteria to their needs. As a further effect, they may foster improved research practices regarding publication of data and code, and reproducibility.

The current API design lacks a transparent process (akin to ACM or COS badges though theirs are manual) to award the reproducibility and geospatial badges. The provenance of badges (i.e. who awarded it, to what, using which criteria) would be crucial in a scholarly setting to establish trust. It could be made accessible with interactive badges, e.g. clicking on a badge opens a pop-up with background information, but also for other services if the information behind the badges is exposed in a structured form via the API, supplementary to the mere images. The current approach could be extended in these directions leveraging SVG's features for interaction, and content-type negotiation for alternative representations. The novelty of ERC and the o2r reproducibility service is an issue, because three badge types rely solely on their existence. Only a wider uptake of ERCs or ERC-like metadata in other platforms, e.g. geospatial properties in publication metadata, can mitigate this.

We see automatically generated and independently spread badges as a promising supplement to the inspection-based badges by COS or ACM and as a way to expose still underused properties of publications' geospatiality and reproducibility. The biggest risks are fragmentation and establishing the trustworthiness of sources for badge information, both due to the distributed approach for defining, creating, and inserting badges.

Figure 8: Filtering search results using badge values.
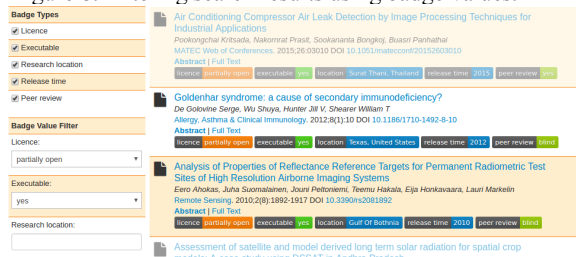
5   Future Work

While the prototypes show the technical feasibility, the most important next step is a user study to evaluate the design and content of the novel badges concerning the goal of improving user experience during discovery, and to learn more about the motivation, requirements, and preferences of involved user groups. The study should investigate potential effects on willingness to publish research compendia and elaborate on trust. It could potentially draw parallels to mechanisms behind other badges, e.g. organic food labels. Such a study can inform the further development of badges, e.g. interactive features, visual design, and regarding transparency (see Discussion). Technical measures can be taken to improve the experience, such as client- or server-side caching, and *Extender* can be reimplemented as a WebExtension[9] to make it available for other browsers such as Firefox. After solving technical and usability-related challenges, a real adoption by the scientific community requires an involvement of more stakeholders and individual early adopters, e.g. funding agencies and a leading journal or conference. Together these institutions can initiate a lighthouse project and a public discourse about the content and scope of badges, so that iterative improvements can make the badges more useful, even for larger user groups beyond the geospatial community.

Software and Data Availability

The implementations of *Badger* and *Extender* including Docker image tarballs, a docker-compose configuration, test data, and instructions for local evaluation (see file README.md) are published on Zenodo (Lohoff & Nüst, 2018). The source code projects are on GitHub at https://github.com/o2r-project/o2r-badger respectively https://github.com/o2r-project/o2r-extender.

Author Contributions

D. N. conceived the idea, supervised the development, and wrote the paper. L. L. developed software and drafted the paper. L. E., N. G., M. G., S. T. J. S. K., L. M., M. M., C. R. and A. v. E. designed badges and developed software. All authors approved the final version of the paper.

Acknowledgements

---

9   https://wiki.mozilla.org/Add-ons/Terminology
10   https://groups.google.com/forum/#!topic/reproducible-research/AP0k_xi69AA/discussion

References

Baker, M. (2016) Digital badges motivate scientists to share data. *Nature News*. doi:10.1038/nature.2016.19907.

Buckheit, J. B. and Donoho, D. L. (1995) WaveLab and Reproducible Research. In: Antoniadis, A., Oppenheim, G. (eds.) *Wavelets and Statistics*, 55–81. Lecture Notes in Statistics 103. Springer New York. doi:10.1007/978-1-4612-2544-7_5.

Gentleman, R. and Temple Lang, D. (2007) Statistical Analyses and Reproducible Research. *Journal of Computational and Graphical Statistics,* 16, no. 1, 1–23. doi:10.1198/106186007X178663.

Grahe, J. E. (2014) Announcing Open Science Badges and Reaching for the Sky. *The Journal of Social Psychology*, 154, no. 1, 1–3. doi:10.1080/00224545.2014.853582.

Ibanez, L., Schroeder, W. J., and Hanwell, M. D. (2014) Practicing Open Science. In: Stodden, V., Leisch, F., Peng, R. D. (eds.) *Implementing Reproducible Research*, 241–280. Chapman & Hall/CRC The R Series. CRC Press, 2014. https://osf.io/emvbz/.

Katz, D. S., and Chue Hong, N. P. (2018) Software Citation in Theory and Practice. In: Davenport, J. H., Kauers, M., Labahn, G., Urban, J. (eds.) *Mathematical Software – ICMS 2018*, 289–96. Lecture Notes in Computer Science. Springer International Publishing. doi:10.1007/978-3-319-96418-8_34.

Kidwell, M. C., et al. (2016) Badges to Acknowledge Open Practices: A Simple, Low-Cost, Effective Method for Increasing Transparency. *PLOS Biology* 14(5):e1002456. doi:10.1371/journal.pbio.1002456.

Lee, D., Ferguson, C and Mitchell, R. (2009) Air pollution and health in Scotland: a multicity study. *Biostatistics*, Volume 10, Issue 3, Pages 409–423, doi:10.1093/biostatistics/kxp010.

Markowetz, F. (2015) Five Selfish Reasons to Work Reproducibly. *Genome Biology* 16:274. doi:10.1186/s13059-015-0850-7.

Lohoff, L., and Nüst, D. (2018) Reproducibility package for: Badges for Geoscience Containers. *Zenodo*. doi:https://doi.org/10.5281/zenodo.1199271.

Nüst, D. (2018) Reproducibility Service for Executable Research Compendia: Technical Specifications and Reference Implementation. *Zenodo*. doi:10.5281/zenodo.2203843.

Konkol, M., Kray, C. and Pfeiffer, M (2018) Computational Reproducibility in Geoscientific Papers: Insights from a Series of Studies with Geoscientists and a Reproduction Study. *International Journal of Geographical Information Science*, 1–22. doi:10.1080/13658816.2018.1508687.

Nüst, D., Konkol, M., Pebesma, E., Kray, C., Schutzeichel, M., Przibytzin, H. and Lorenz, J. (2017) Opening the Publication Process with Executable Research Compendia. *D-Lib Magazine*. doi:10.1045/january2017-nuest.

Nüst, D., Granell, C., Hofer, B., Konkol, M., Ostermann, F. O., Sileryte, R. and Cerutti, V. (2018) Reproducible Research and GIScience: An Evaluation Using AGILE Conference Papers. *PeerJ,* 6, e5072. doi:10.7717/peerj.5072.

Peng, R. D., 2009. Reproducible research and Biostatistics. *Biostatistics*, Volume 10, Issue 3, Pages 405–408. doi:10.1093/biostatistics/kxp014.

Peng, R. D. (2011) Reproducible Research in Computational Science. *Science*. 334 (6060): 1226–27. doi:10.1126/science.1213847.

Piwowar, H. A., and Vision, T. J. (2013) Data Reuse and the Open Data Citation Advantage. *PeerJ*, 1:e175. doi:10.7717/peerj.175.

Wilkinson, M. D., et al. (2016) The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Scientific Data*, 3:160018. doi:10.1038/sdata.2016.18.