

Realizing a Collaborative RDF Benchmark Suite in Practice

Piotr Sowiński, Maria Ganzha
Warsaw University of Technology, Poland

We wish there was an RDF benchmark that...

...anyone can edit

...allows for systematic reporting of benchmark results

...can be expanded with new datasets and tasks

...is fully machine-accessible with RDF metadata

...has clear and consistent documentation


...is super-fun and easy to use

RiverBench can do this with:


Open and public community processes



FAIR practices and Linked Open Data



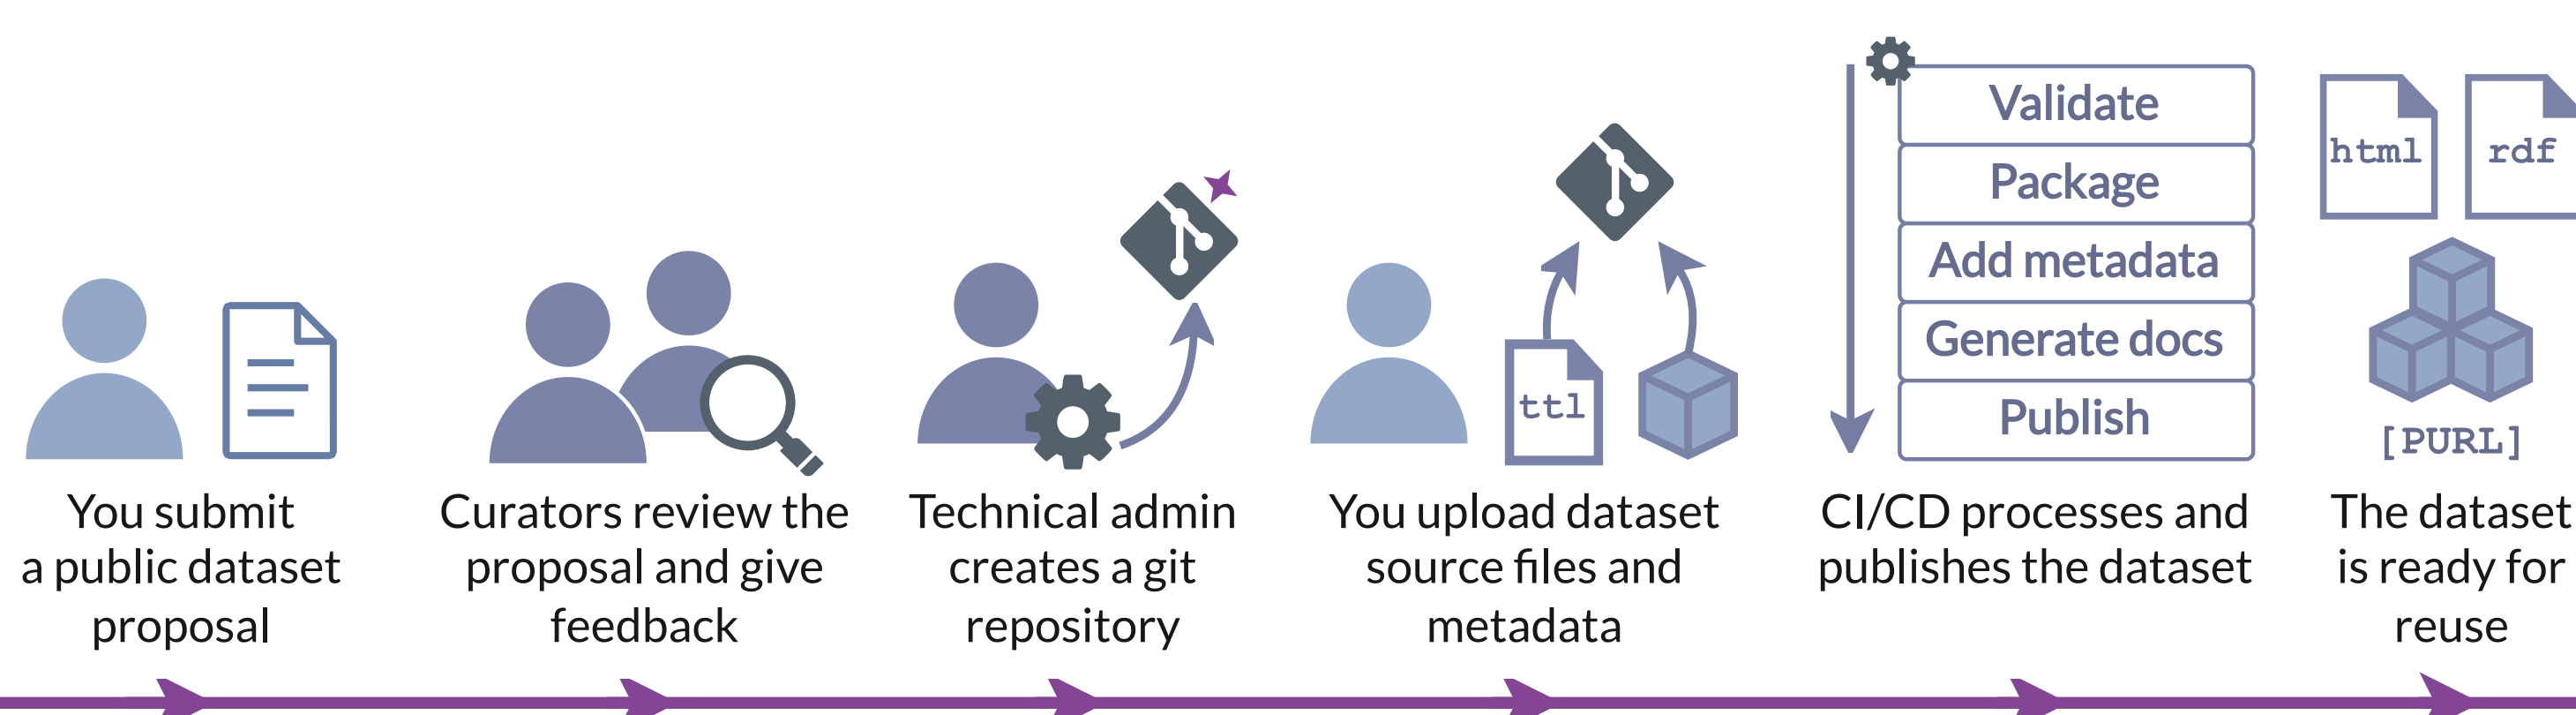
Seamless CI/CD automation



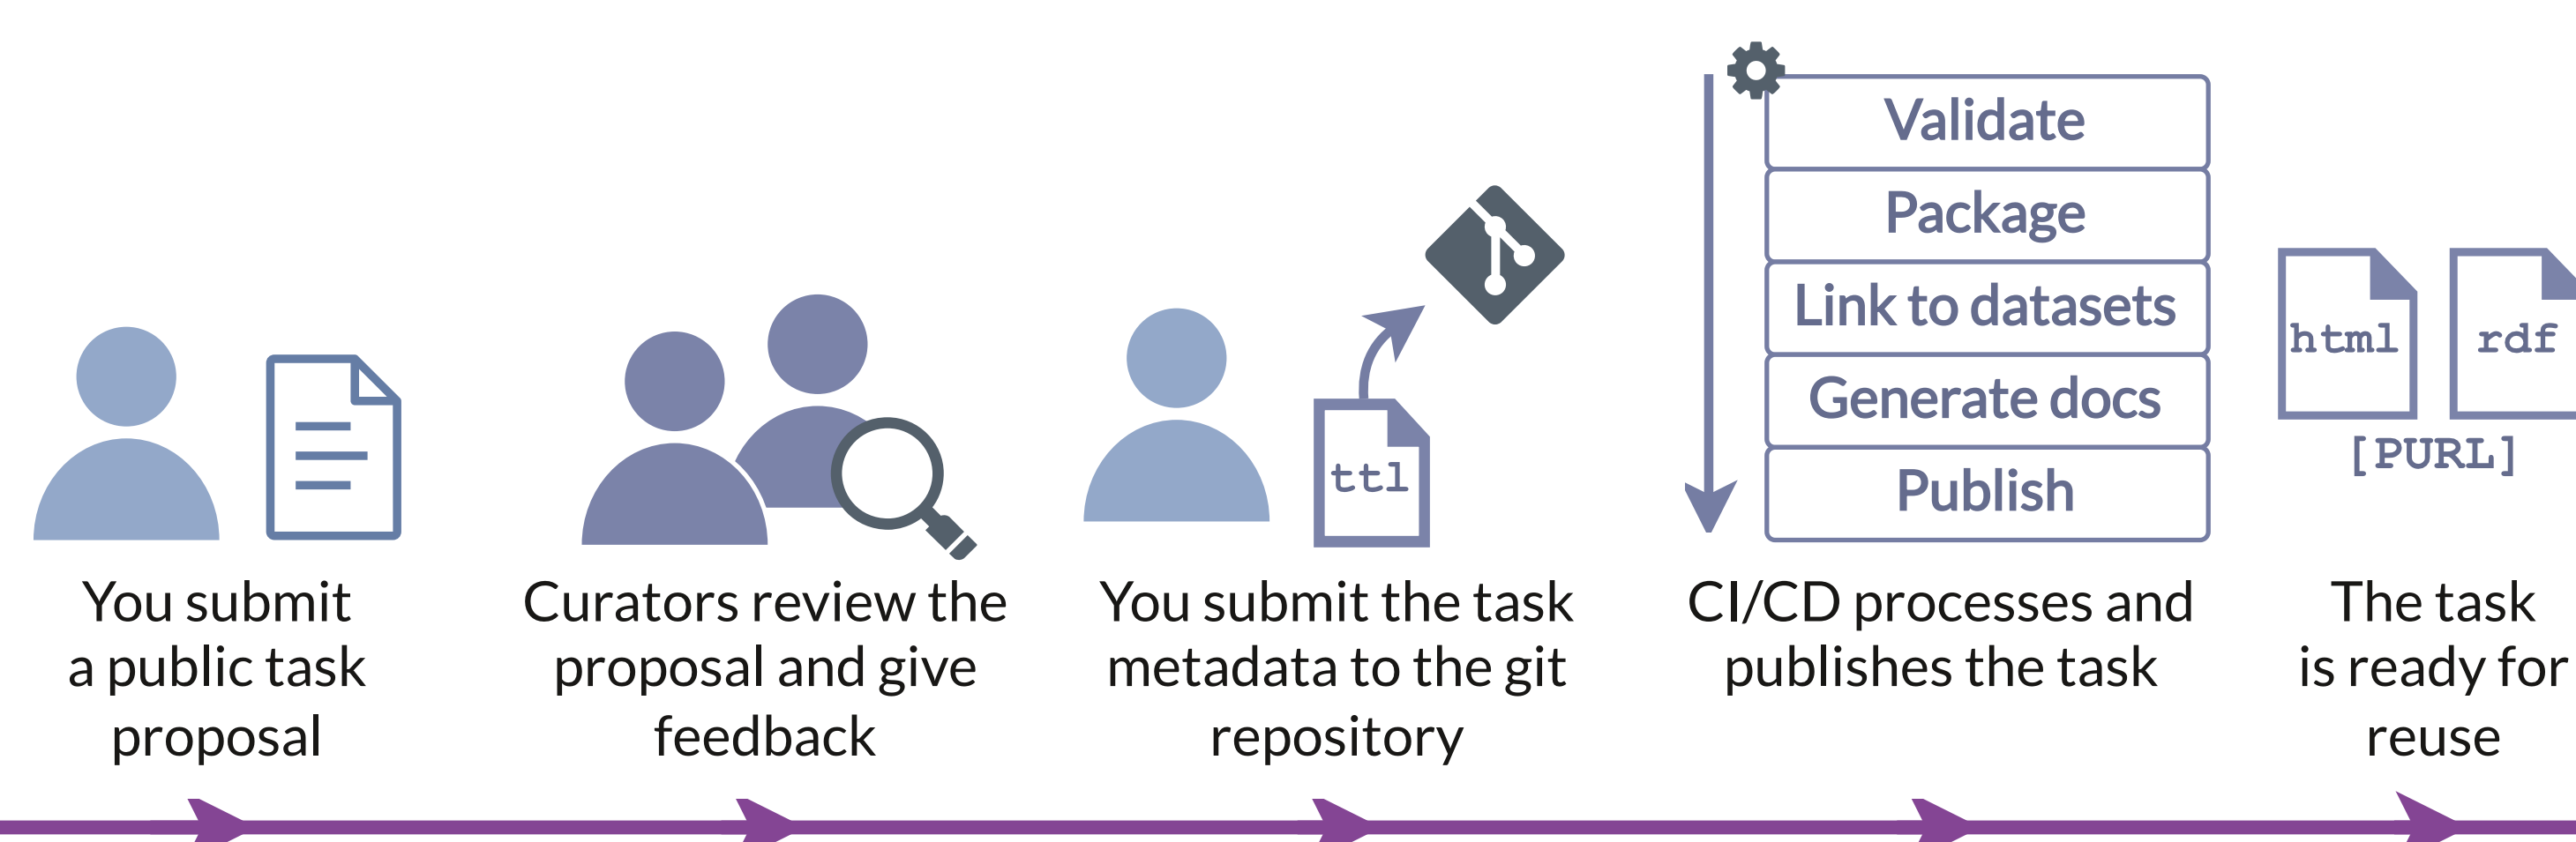
You! Yes, you! RiverBench is also yours



You can contribute new datasets:



And new benchmark tasks:

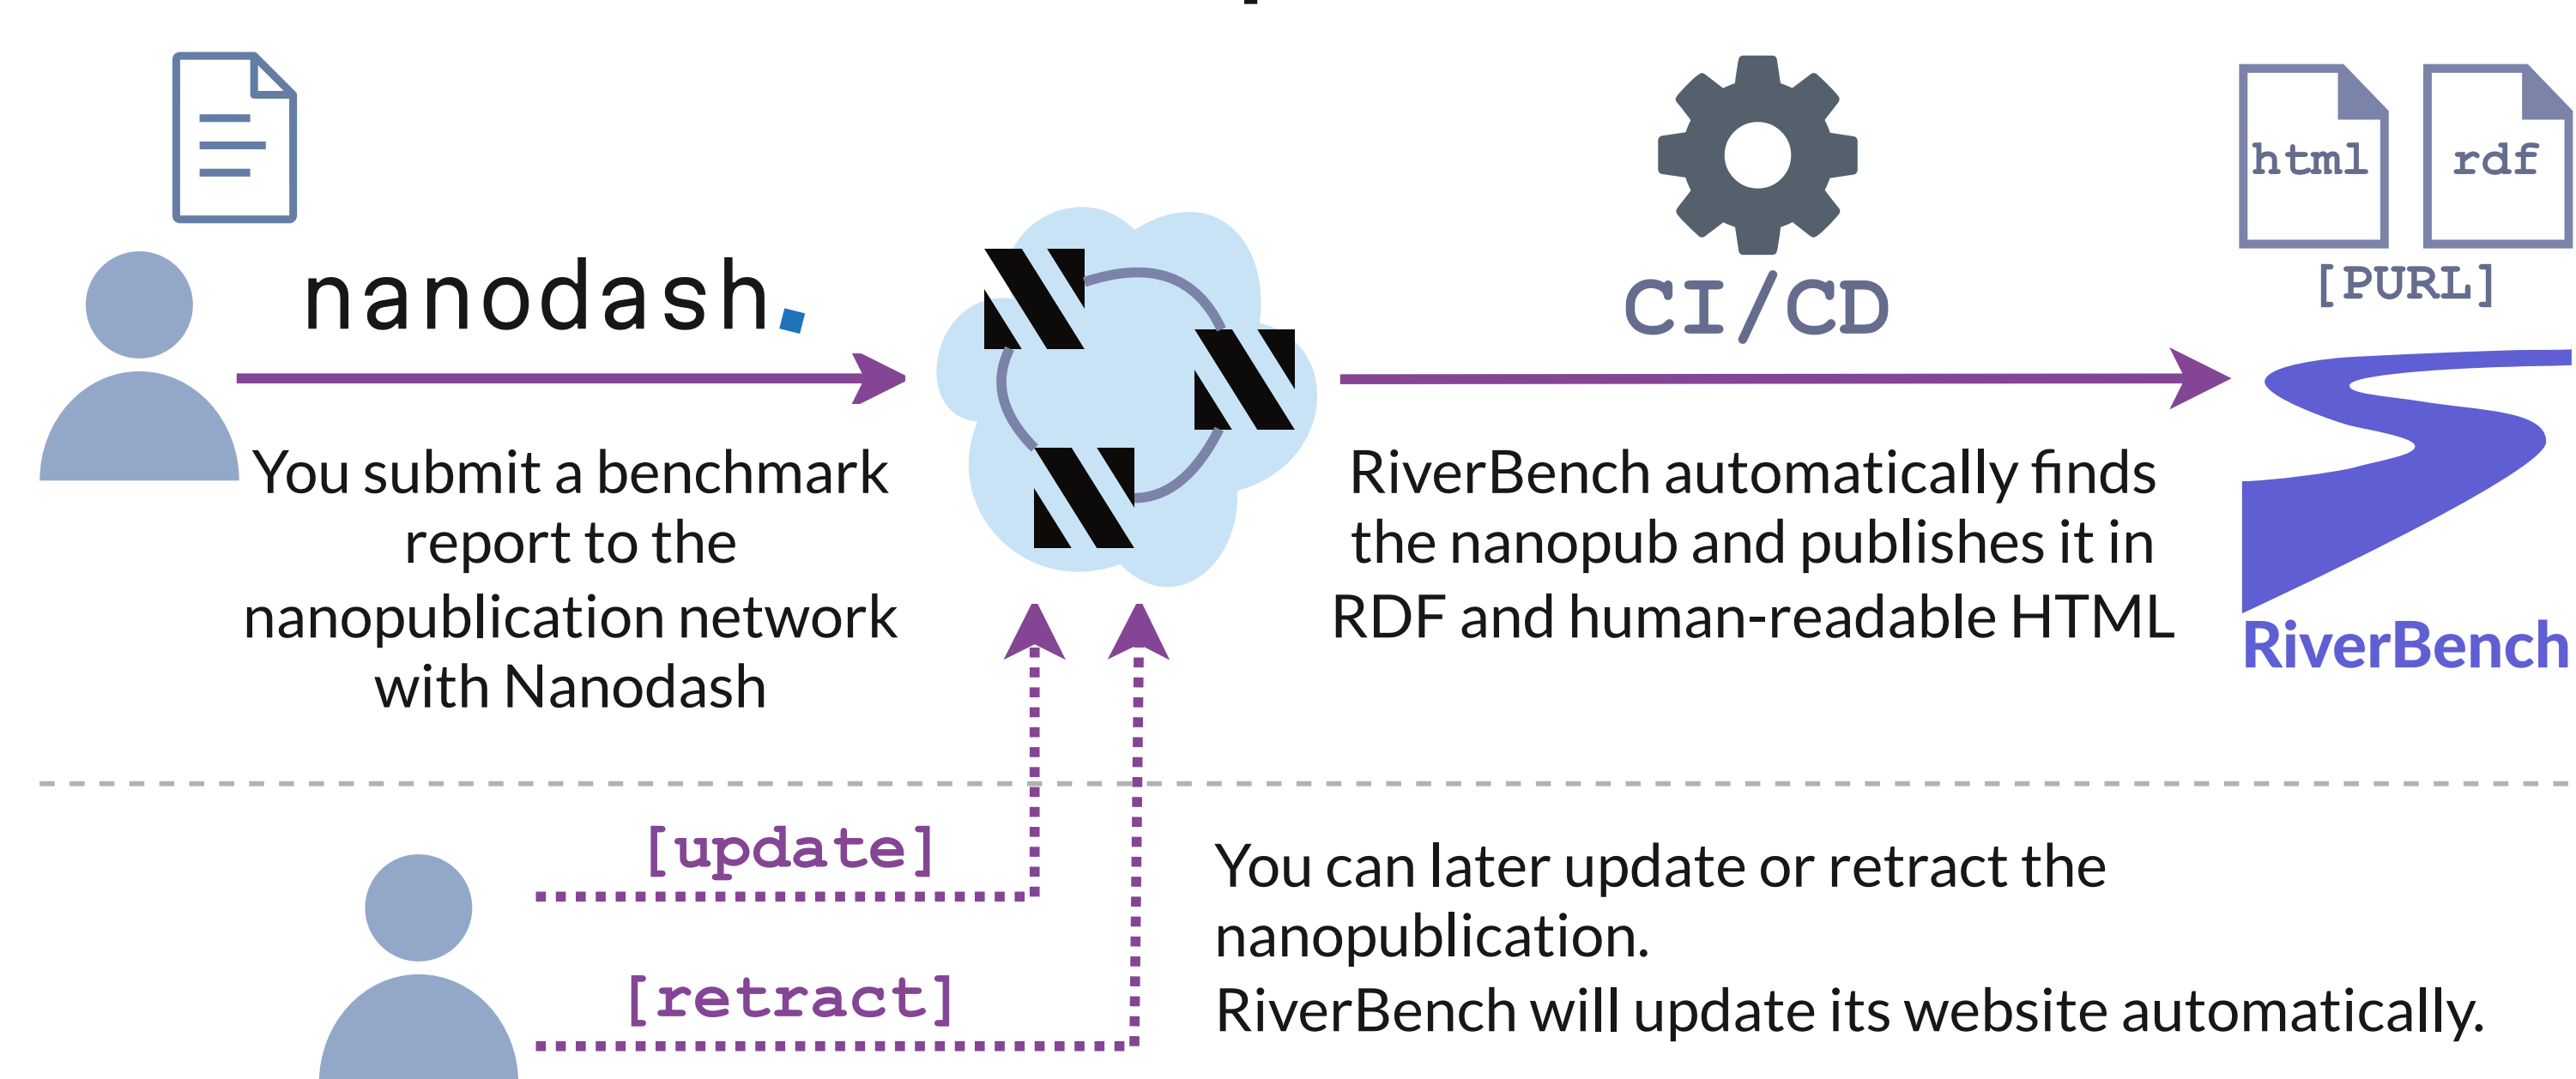


You can also edit anything easily:

[Edit this page](#) [Edit metadata](#)

Just press these buttons and be taken to a Turtle/Markdown editor

You can report your benchmark results with Nanopublications:



We want you to join RiverBench!

- Make your use case seen: add your dataset to the benchmark suite
- Report your benchmark results and gain visibility
- Propose an entirely new benchmark task
- Join the team of curators and decide the future of RiverBench



Today, RiverBench is:

- 13 diverse, real-life benchmark datasets, from IoT to popular music
- 9 benchmark tasks, from RDF compression to streaming RDF graphs over the network
- 11 benchmark run reports
- 4 community curators ensuring its high quality



Learn more:

<https://w3id.org/riverbench>



Hey, ask me for a demo! I can show it all to you live.
~ Piotr

Contact: piotr.sowinski.dokt@pw.edu.pl