Parallel uncertainty quantification with fused simulations

S. Wolf (Technical University of Munich), L. Seelinger (Heidelberg University), M. Bader (Technical University of Munich)

We develop a scalable workflow for parallel earthquake source inversion employing Bayesian methods.

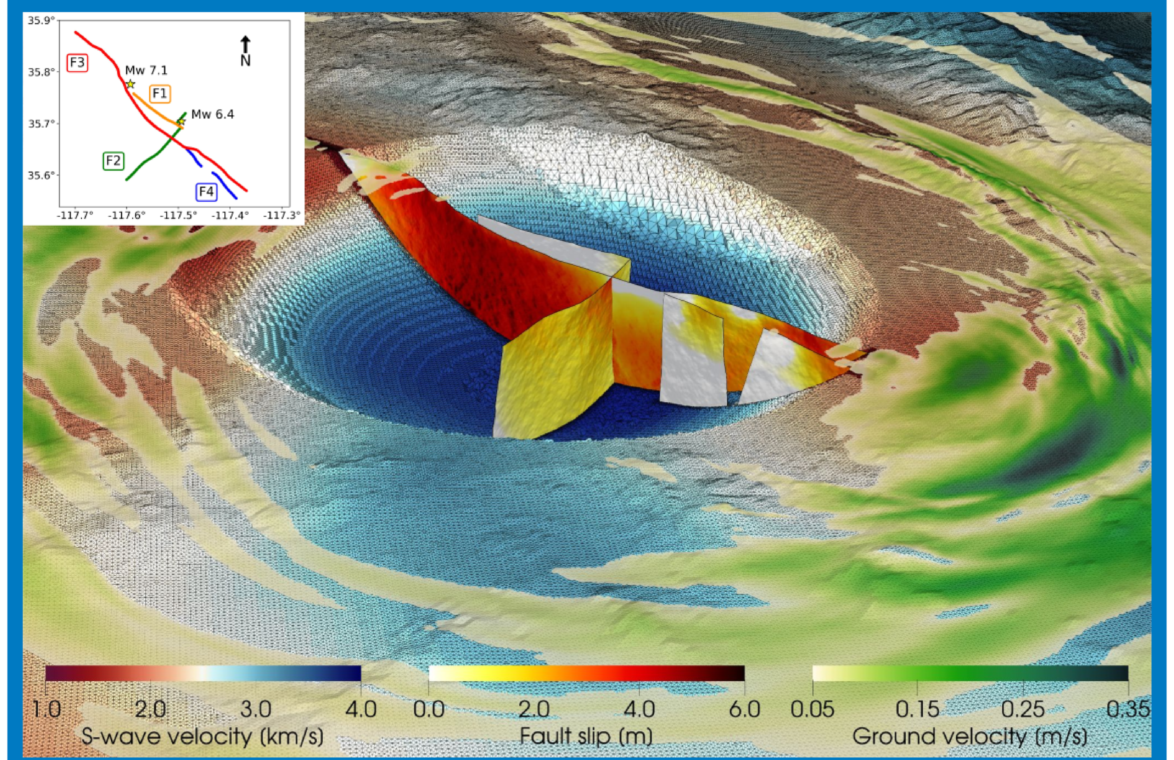

Figure 1: Ridgecrest scenario with complicated fault structure, topography and heterogeneous materials. Adapted from [3].

Earthquake simluations with SeisSol

SeisSol (`https://seissol.org`) is a well-established simulation software for earthquakes source dynamics and seismic wave propagation. SeisSol solves the hyperbolic PDE

$$\partial_t q + A\partial_x q + B\partial_y q + C\partial_z q = Eq.$$

Key features

Realistic Materials anisotropic elastic, isotropic elastic, poroelastic and viscoelastic materials with optional plastic deformation

Physics based sources dynamic rupture, rate-and-state friction

Geometric flexibility with tetrahedral meshes

Discontinuous Galerkin discretisation

High-order convergence polynomial basis functions + ADER time-stepping

Element local predictor corrector scheme facilitates parallelization.

Small matrix-matrix multiplications code-generator `YATeTo` with architecture specific backend for node-level performance

Parallelisation strategy

MPI + X based on mesh partitioning

OpenMP for CPU clusters

CUDA/SYLC for GPU offloading

Fused simulations

Compute N simulations simultaneously by adding another dimension to the solution *tensor*.

Reduce I/O overhead Read mesh only once.

No padding for SIMD Choose number of simulations as a multiple of the vector register width.

Reduce cache-transfers element-local mass/stiffness matrices only loaded once per N simulations

Scaling results

CoolMUC-2 812 compute nodes, each with two Intel Xeon E5-2697 v3 processors and 64 GB RAM

Improved strong scaling with fused simluations due to higher computational workload per element

Eight fused simulations optimal 33 % improvement over single simulations

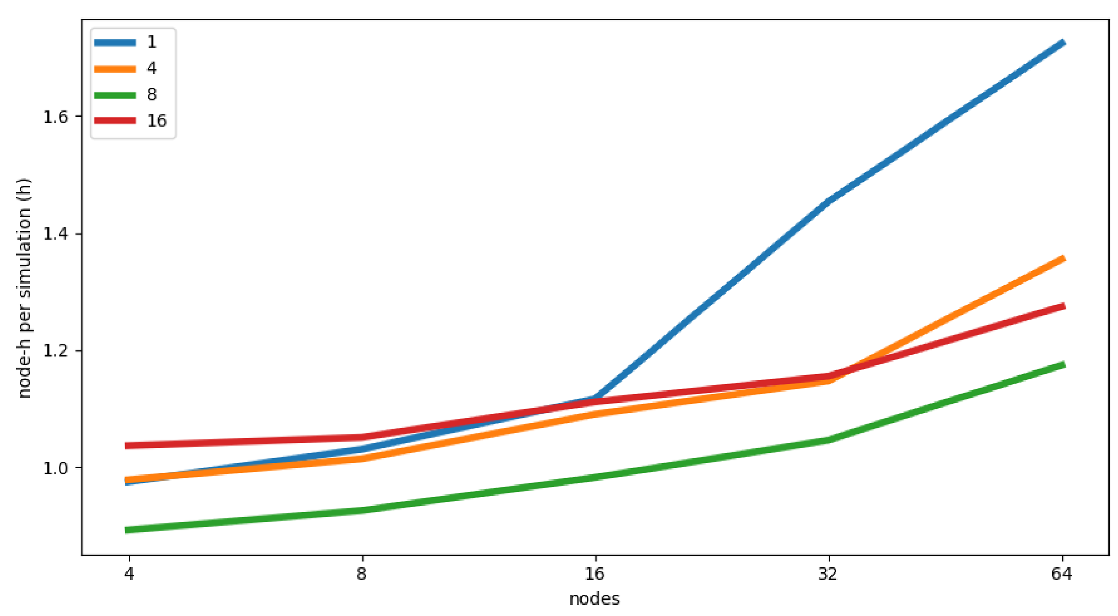

Figure 2: Computational effort per simulation for different number of fused simulations. The mesh contains 420000 elements. Using eight fused simulations, reduces the computational time by up to 33%.

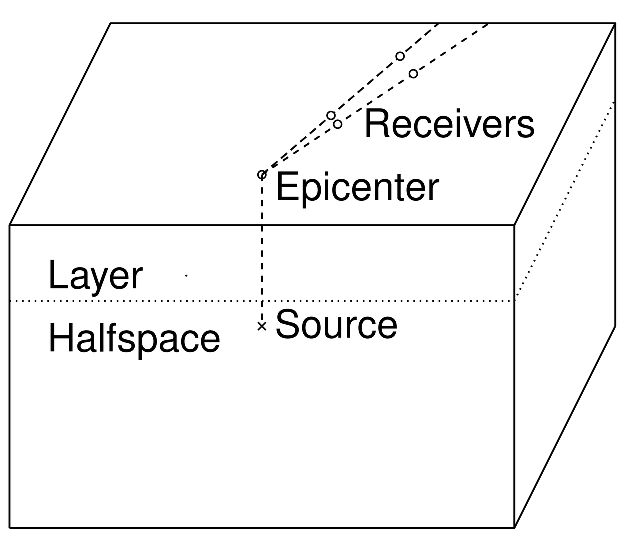

Figure 3: Geometry of the LOH1 scenario. The scneario features an elastic halfspace with a slow-velocity layer on top. A double couple source excites elastic waves, which are recorded at four surface receivers.

Bayesian inverse problems with the Markov-Chain Monte-Carlo method

MCMC sampling

Bayesian inverse problems allow us to solve inverse problems in a probabilistic manner, i.e. find the probabilty distribution of an unknown parameter θ, which fits the data $\hat{y}$. Build a Markov chain of samples θ_i:

Draw new proposal θ' from proposal distribution $q(\theta_i, \cdot)$.

Evaluate forward model $y' = G(\theta')$.

Compute proposal likelihood by comparing y' with data $\hat{y}$.

Accept/Reject Accept $\theta_{i+1} = \theta'$ with certain probability otherwise keep previous sample $\theta_{i+1} = \theta_i$.

Stationary distribution $P(\theta|\hat{y}) \propto P(\hat{y}|\theta) \cdot \pi^0(\theta)$

Generalized Metropolis-Hastings algorithm

Metropolis-Hastings is sequenatial can only draw new sample after Θ' has been accepted or denied.

Generalized Metropolis-Hastings enables parallel model execution [1].

Draw N samples and evaluate forward model.

Compute stationary distribution.

Accept $K \leq N$ samples based on stationary distribution.

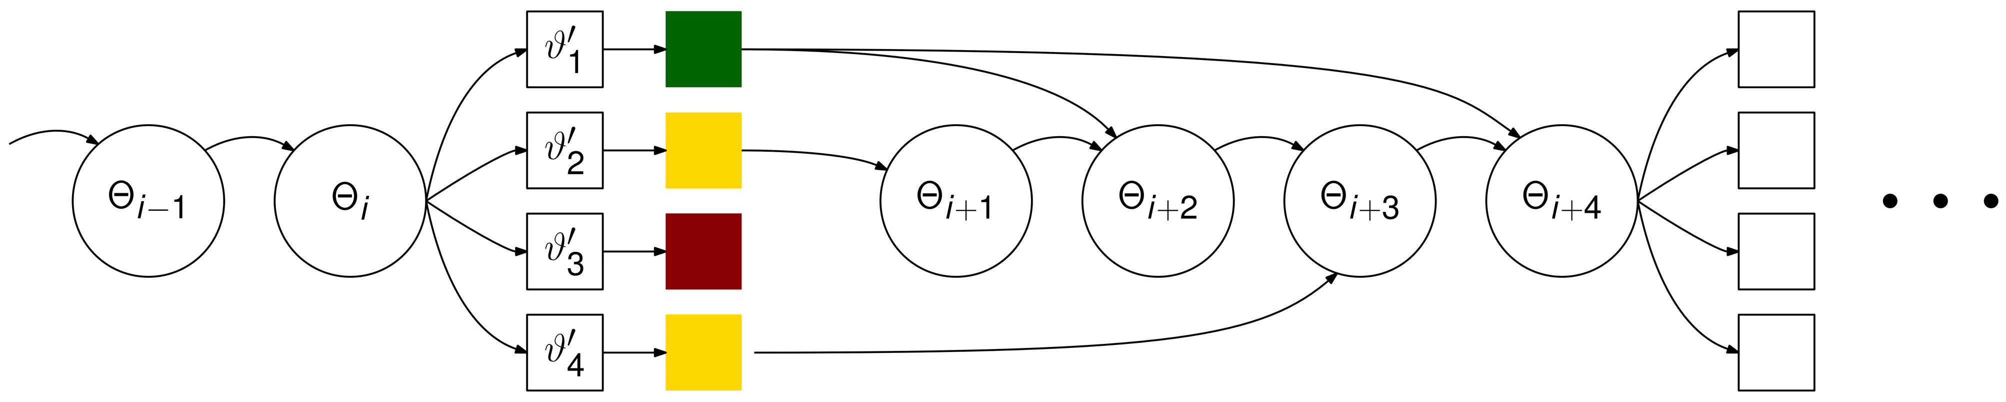

Figure 4: Sketch of the GMH algorithm for $N = K = 4$. Based on the sample Θ_i, four new proposals ϑ'_i are drawn. The forward model is executed and in comparison to the data, a likelihood is assigned to each propoasal (indicated by the colors). Based on this likelihood estimator, four new samples $\Theta_{i+1}, \ldots, \Theta_{i+3}$ are added to the Markov Chain.

Performance of the Generalized Metropolis-Hastings algorithm

Benchmarking GMH

We test the performance of the GMH algorithm on a simple ODE test case with two unknown model parameters.

Vary $N = 1, \ldots, 20$, $K = 1, \ldots, N$ to obtain 1000 samples.

$K << N$ inefficient high execution time, waste too many samples.

Acceptance ratio increases with N.

$K << N$ generates more independent samples.

Use $K = N$ with small N for best ESS per time ratio.

Note that in this benchmark test, the speed-up of using fused simluations was not as high as it was for SeisSol.

Performance metrics

Acceptance Ratio describes how many proposals are accepted as samples into the chain.
High acceptance ratio: exploration of the parameter space is insufficient. Low acceptance ratio: waste a lot of (expensive) samples.

Effective sample size subsequent samples of a Markov Chain are dependent, but we want to generate independent samples. The effective sample size (ESS) characterizes how many (approximately) independent samples have been drawn.

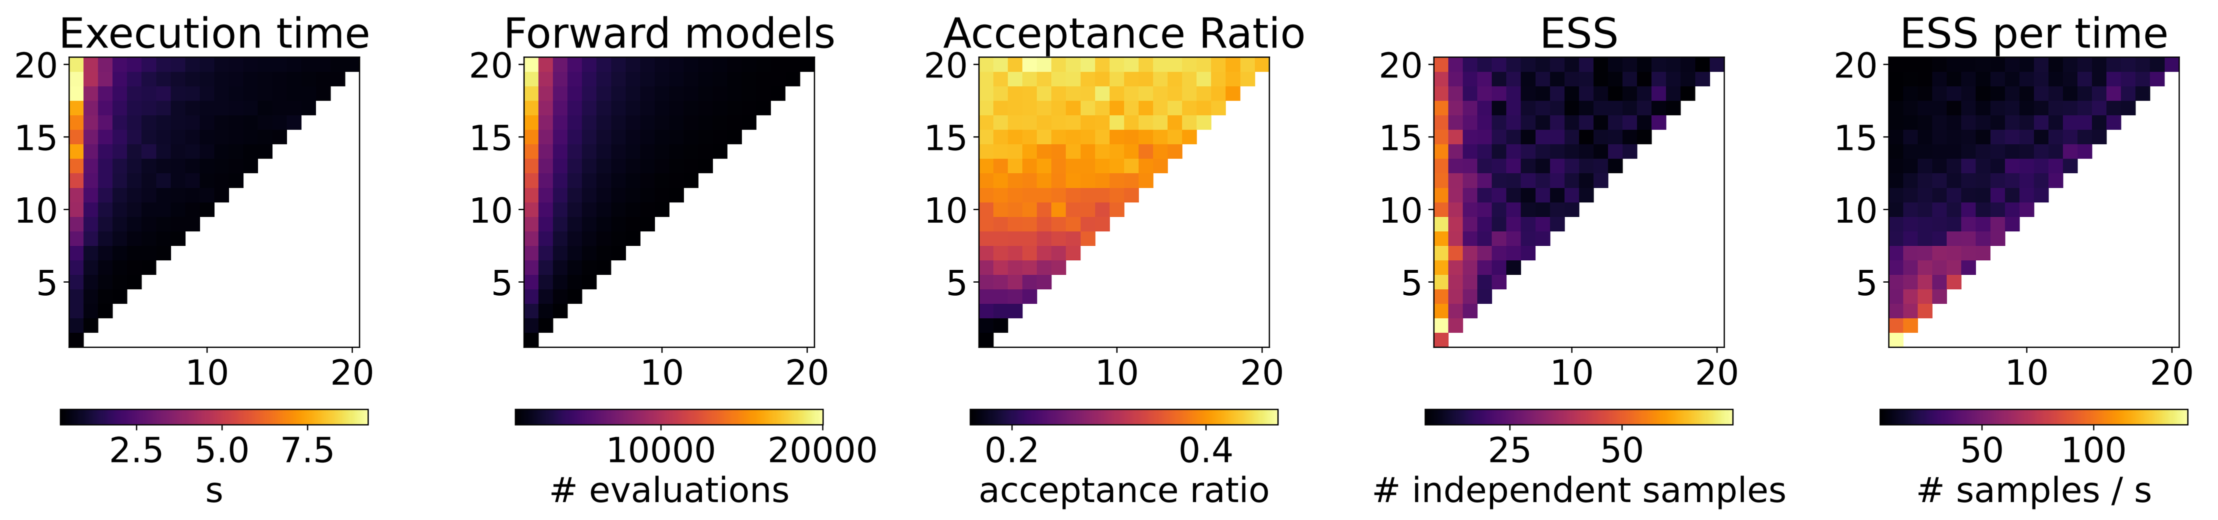

Figure 5: Performance metrics of the GMH algorithm for an ODE testcase. The vertical axis denotes the number of fused simluations (N) and the horizontal axis denotes the number of accepted samples (K).

Finding the source position of the LOH1 scenario with GMH sampling

Model setup

Unknown source position c.f. fig. 3

Ground truth $\hat{\theta} = (0, 0, 2000)^T$

Data $\hat{y}$ semianalytic solution of receiver seismograms

Implementation

Use SeisSol with fused simulation, $N = 8$.

Fused GMH kernel patch of MUQ library [2].

Compare result with data $\hat{y}$ using $\|\cdot\|_1$ norm.

Runtime 21 h on 32 nodes to collect 640 samples

Results

$\mathbb{E}(\theta)$ slightly off with significant offset in x/y direction: $(-471.4, 733.7, 2232)^T$

Receiver 2 and 4 match well.

Receiver 1 and 3 are troubled in particular v_3.

Accepance ratio = 11 %, ESS = 8

Upcoming work

Improve acceptance ratio and ESS.

More unknowns e.g. source orientation, frequency

Realistic scenario dynamic rupture, topography

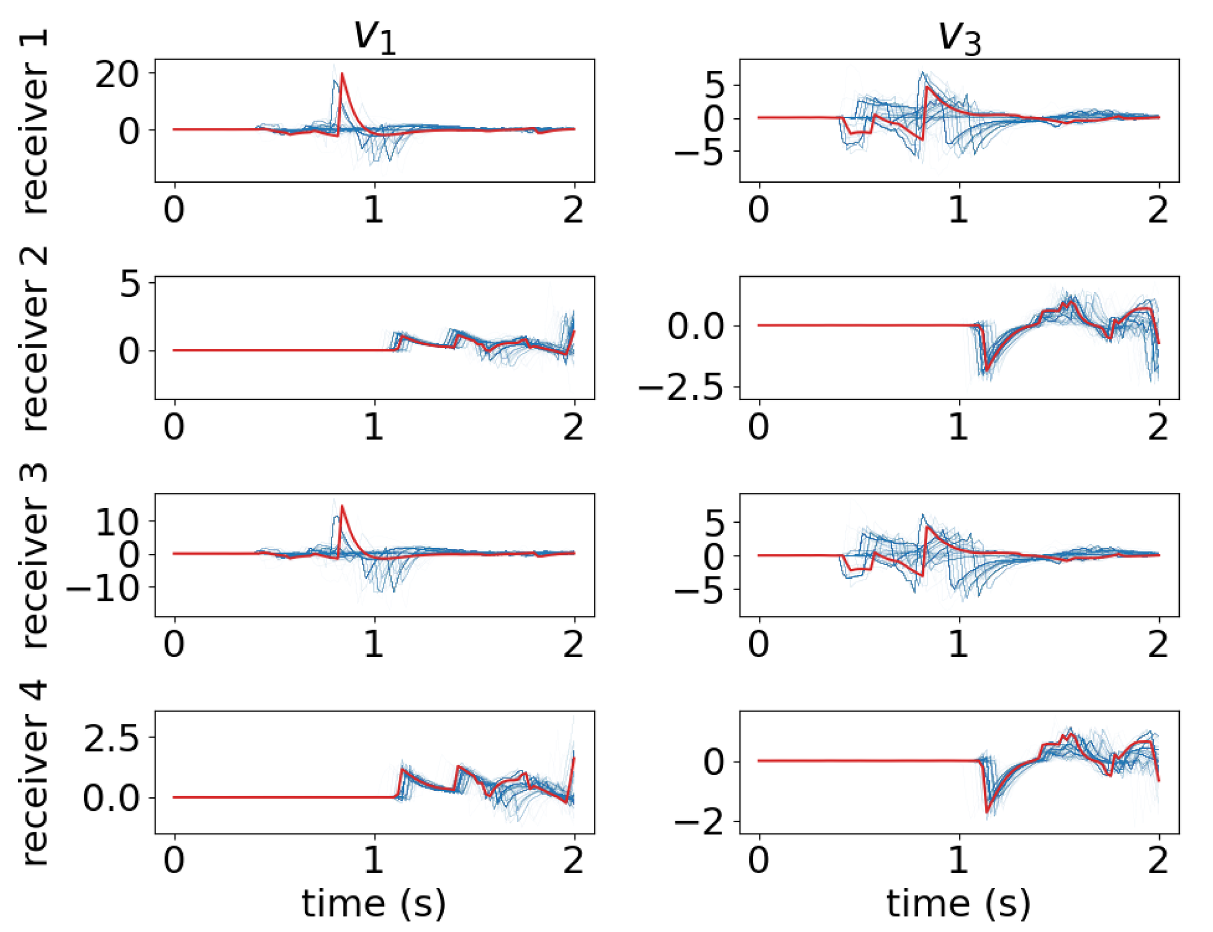

Figure 6: Comparison of the reference solution (red) and the MCMC samples (blue).

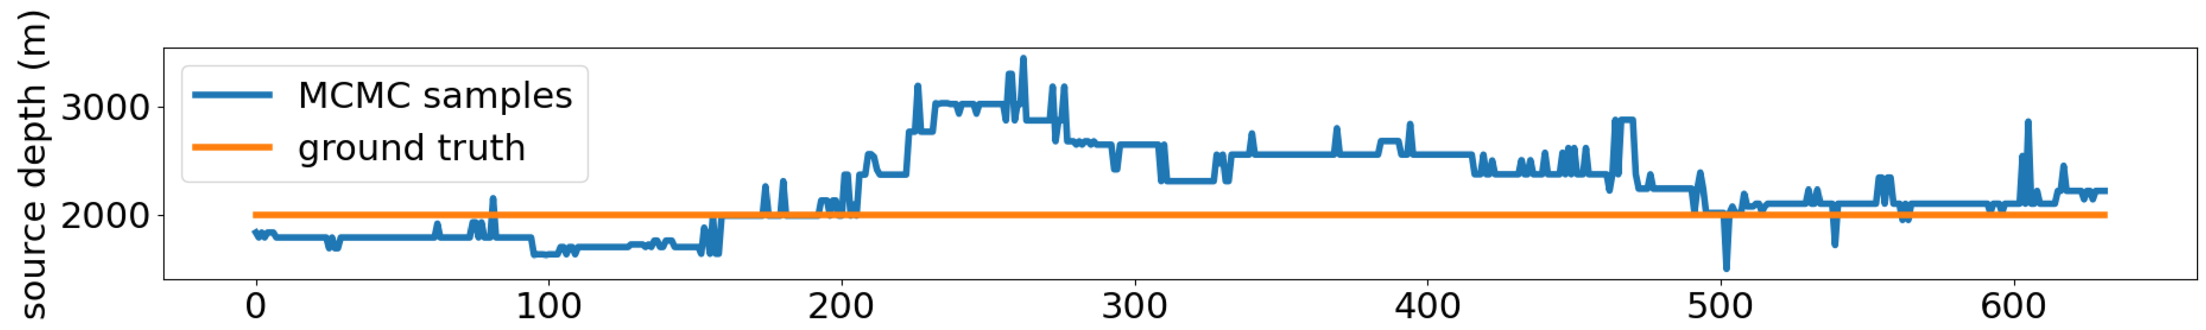

Figure 7: Source depth samples gathered during the MCMC inversion of the LOH1 example.

References

[1] B. Calderhead. "A general construction for parallelizing Metropolis-Hastings algorithms". In: *Proceedings of the National Academy of Sciences* 111.49 (Dec. 2014). Publisher: Proceedings of the National Academy of Sciences, pp. 17408–17413. DOI: 10.1073/pnas.1408184111.

[2] M. Parno et al. *MIT uncertainty quantification (MUQ) library.* 2014.

[3] N. Schliwa, A.-A. Gabriel, and Y. Ben-Zion. "3D Modeling of near-field ground motions and deformation in dynamic rupture simulations of the 2019 Ridgecrest earthquake including fault zone and fault roughness effects". Poster Presentation. 2023 SCEC Annual Meeting (Palm Springs, CA). Sept. 10, 2023.

Acknowledgement

The research leading to these results has received funding from Horizon Europe (ChEESE-2P grant No. 101093038).

Github