Learning Materials at eXascale

AUTHORS:

Edoardo di Napoli, Jülich Forschungszentrum, Germany
Luigi Genovesi, The French Alternative Energies and Atomic Energy Commission (CEA), France
Davor Davidovic, Ruđer Bošković Institute (RBI), Croatia

CONTACT: e.di.napoli@fz-juelich.de, luigi.genovese@cea.fr, davor.davidovic@irb.hr
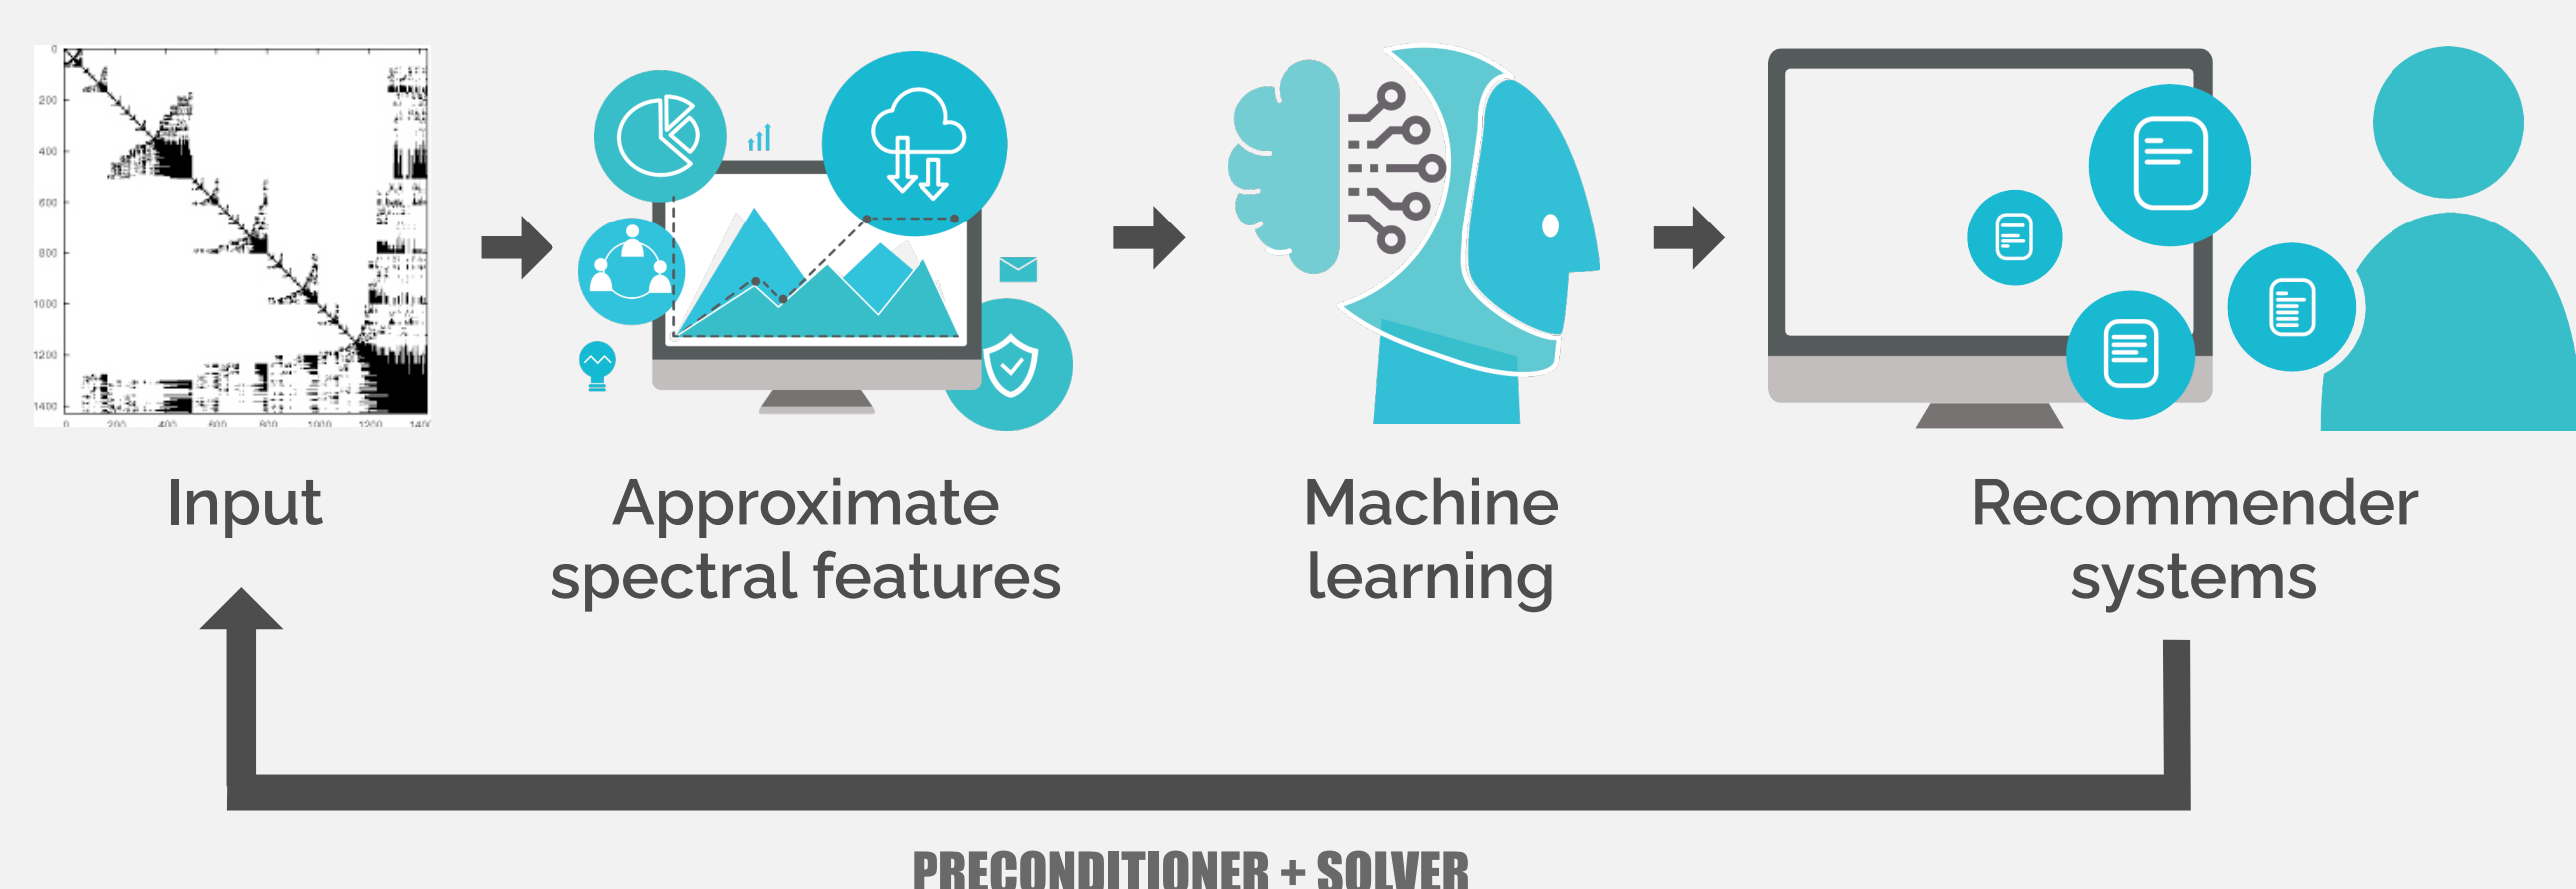
INTRODUCTION

CONTEXT: Simulating large molecular systems comprised of thousands of atoms is crucial across disciplines like solid-state physics, material science, chemistry and biology.

CHALLENGE: Density Functional Theory (DFT) calculations have emerged as a leading methodology due to their scalability and computational performance. Despite DFT's scalability, its computational demands limit its application to moderate system sizes.

MOTIVATION: Innovative algorithms are needed for computationally intensive tasks within DFT codes, such as linear system solvers and eigensolvers, demanding a delicate balance between precision and computational efficiency on massively parallel supercomputers. One promising path is the development of a recommender system optimized for linear system solutions, leveraging domain-specific information.
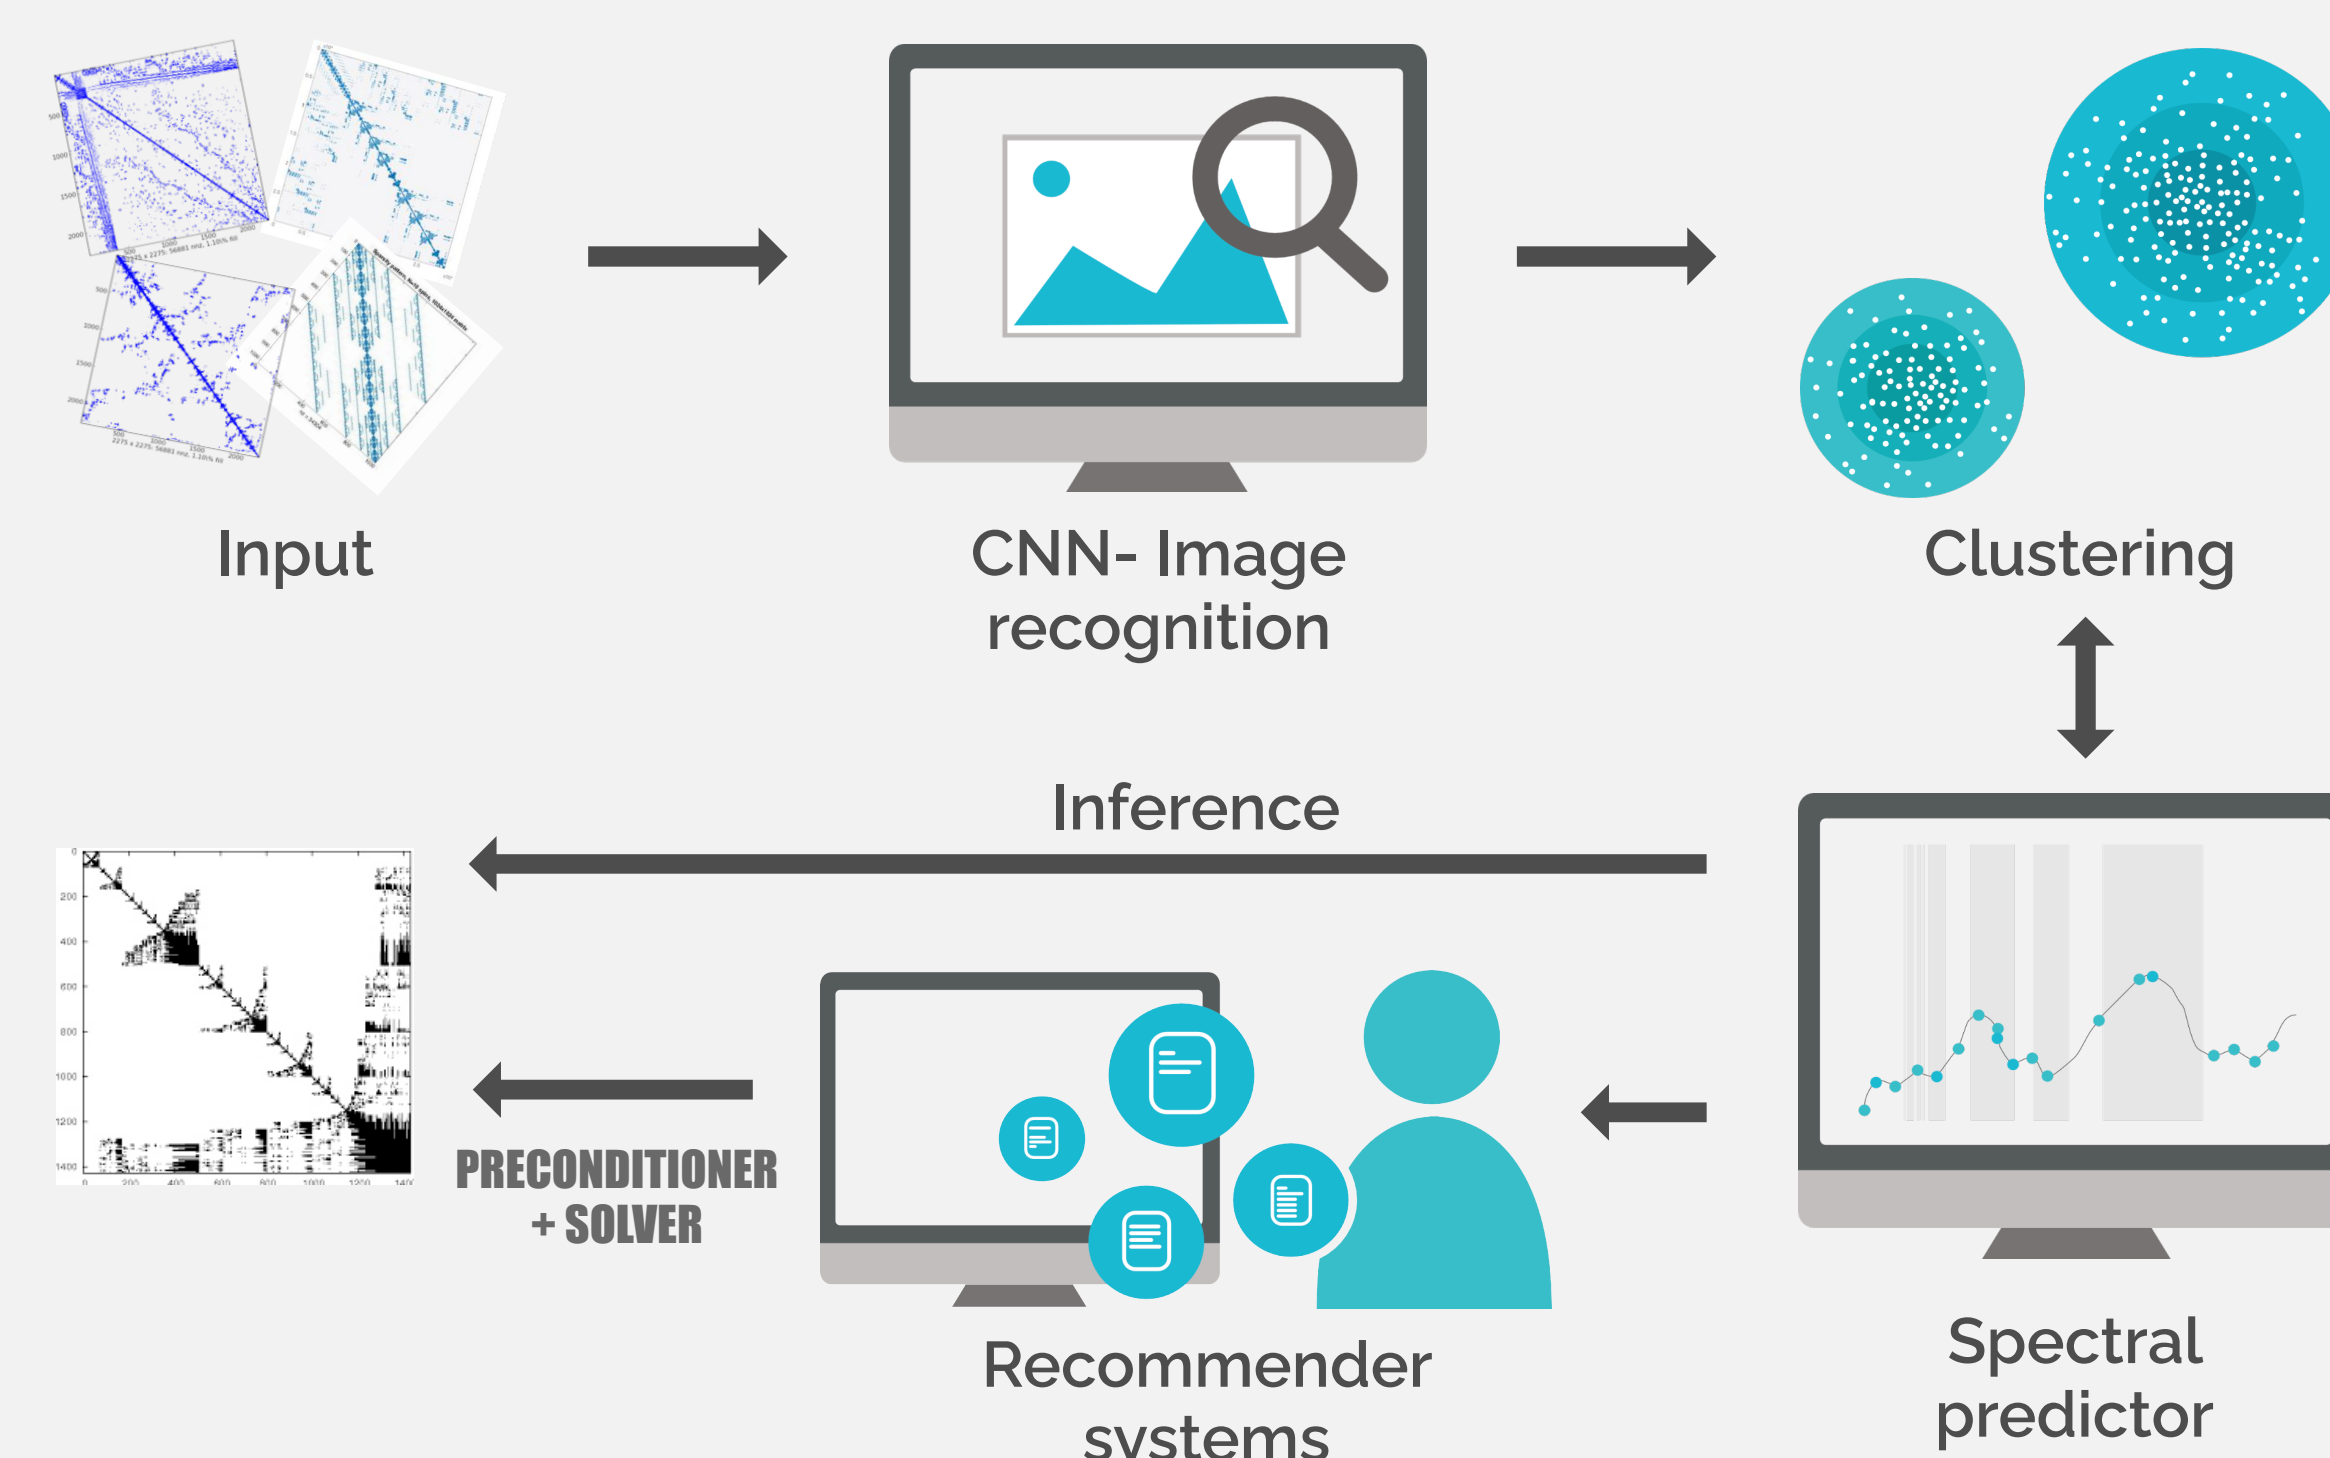
LIMITATIONS: Previous studies attempting to predict optimal sparse solvers have following limitations:
· Relying on spectral property computations (expensive), yielding improved predictions but at high computational cost, or
· Using "cheap" high-level features, sacrificing prediction accuracy.

Input → Approximate spectral features → Machine learning → Recommender systems

PRECONDITIONER + SOLVER

RESEARCH OBJECTIVES

· Constructing a state-of-the-art recommendation system to accelerate and scale up the solution of large-scale sparse linear systems, enabling Materials Science research on exascale clusters.

· Applying novel recommendation systems to enhance the weak scaling of BigDFT simulations for large molecular systems with tens of thousands of atoms.

Current limitations of the BigDFT code:

· Number of iterations of the outer loop is 15 – 30
· The average number of iterations increases with the system size
· The larger the system is, the closer to an experimental set up and more representative of a real-world scenario it becomes.

NOVEL APPROACH

Combining AI with Linear Algebra & Material Simulations information

Input → CNN- Image recognition → Clustering → Spectral predictor → Recommender systems → PRECONDITIONER + SOLVER (Inference)

METHODOLOGY

CONSTRUCTING A COMPREHENSIVE DATABASE: "Developing an extensive database of matrices representing systems in Materials Science, specifically focusing on electronic structure calculations using Density Functional Theory (DFT)."

OPTIMIZING SPARSE MATRIX REPRESENTATION FOR ANALYSIS: "Identifying suitable sparsity patterns of sparse matrices for further analysis using machine learning, particularly Deep Learning methods."
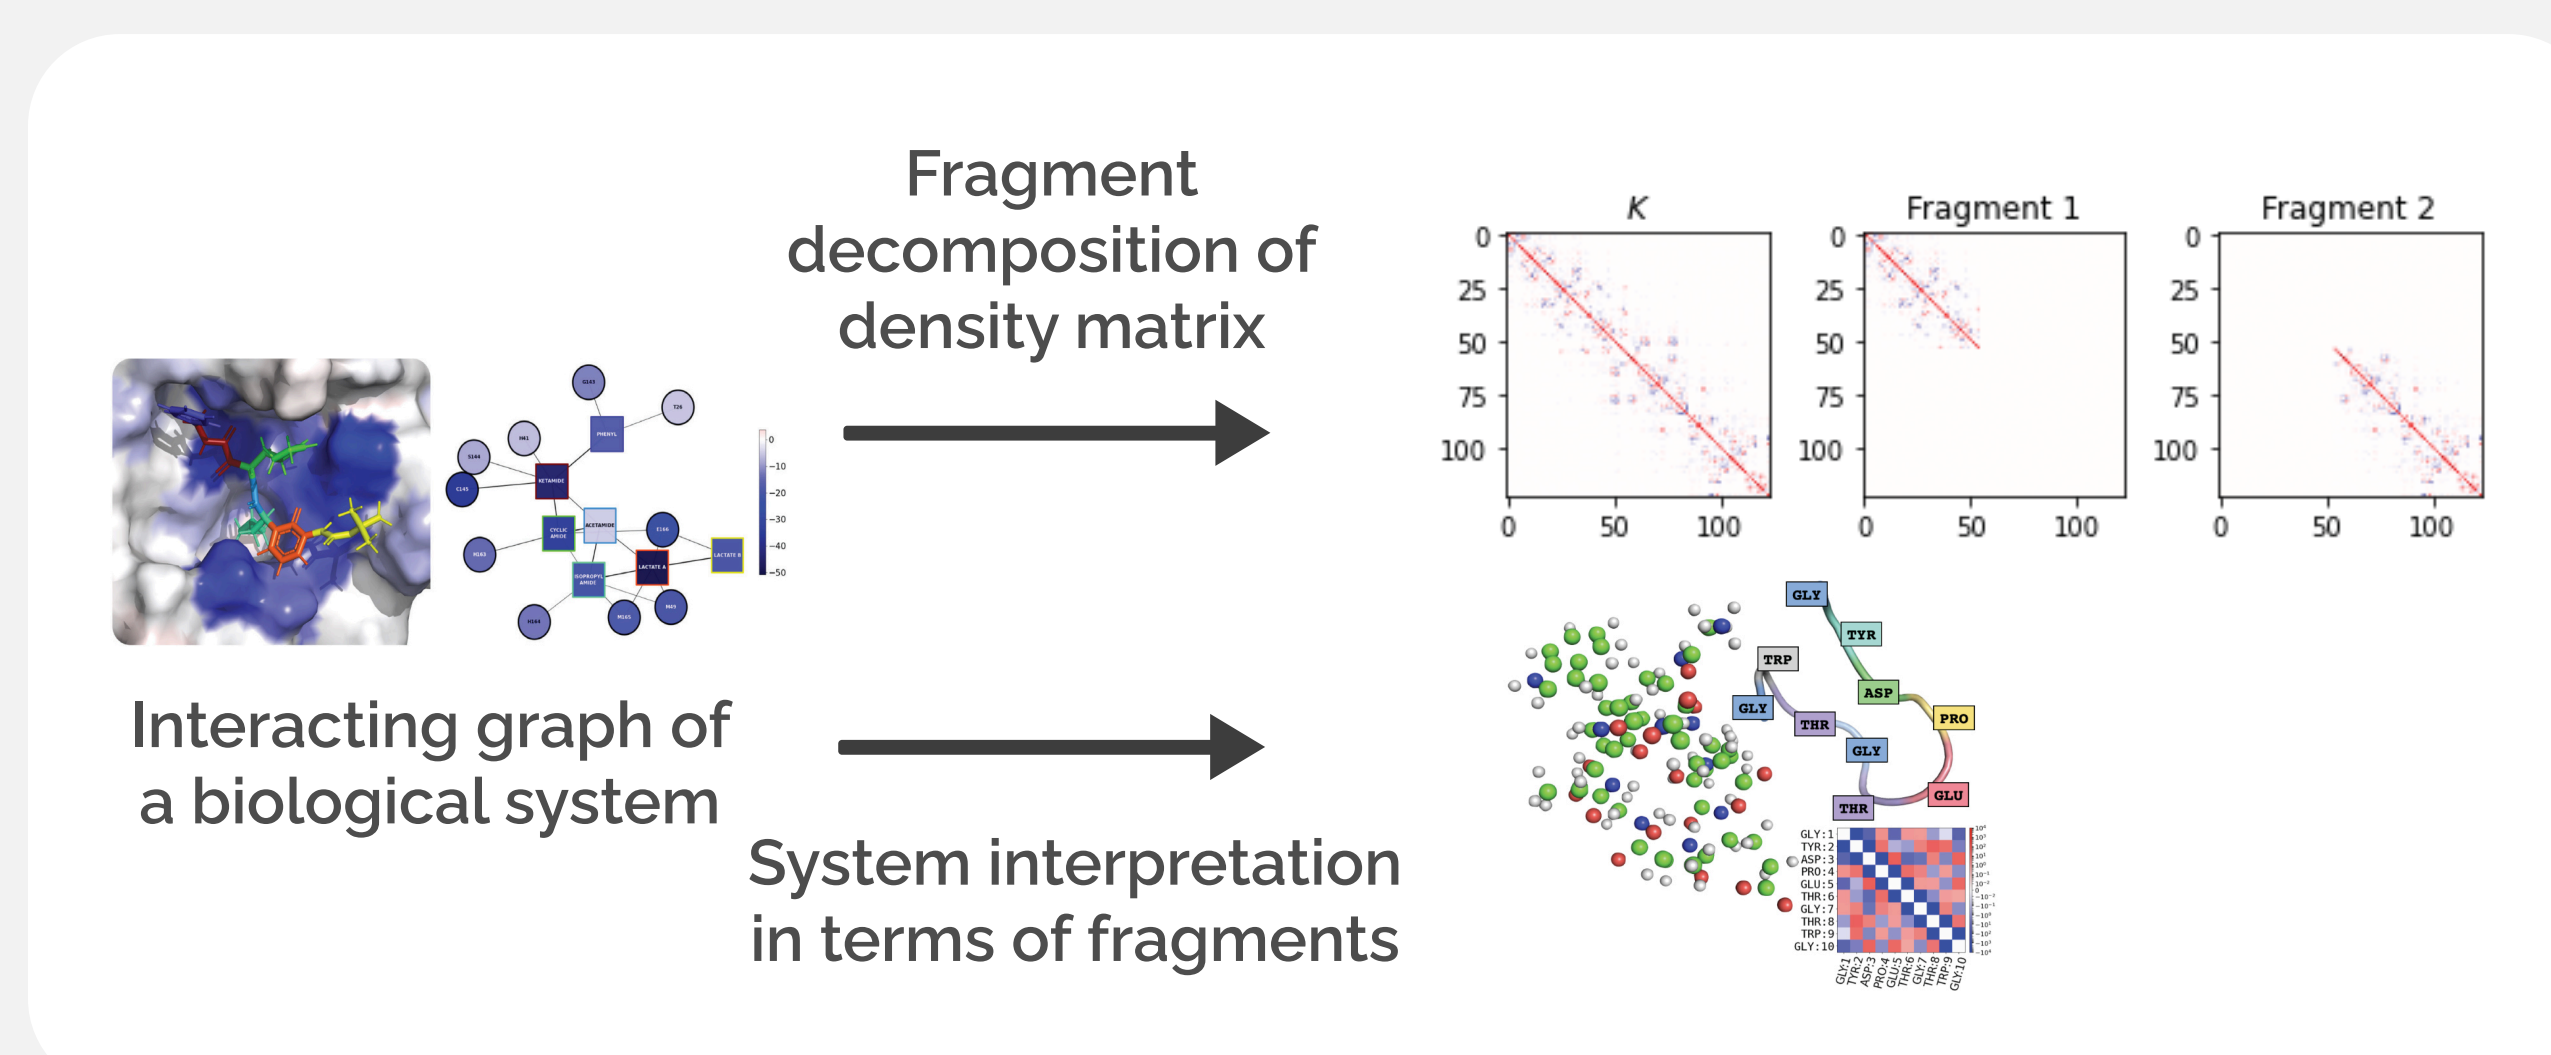
PREDICTING MATRIX SPECTRAL PROPERTIES: "Utilizing existing Deep Learning models, originally applied in image recognition, to predict spectral properties of matrices."

BUILDING A RECOMMENDER SYSTEM FOR PRECONDITIONER AND SOLVER: "Building a recommender system based on surrogate space and predictors to optimize the selection of preconditioner and solver combinations for solving linear systems derived from the DFT domain."

What we want to achieve

· Increase system size while maintaining limited number of outer iterations necessary to converge it.

Partitioning the systems into chemically meaningul fragments and describing their interactions - > resulting in a graph-like view of the systems

Interacting graph of a biological system → Fragment decomposition of density matrix

Interacting graph of a biological system → System interpretation in terms of fragments

Creation of curated database

Starting point: Existing database of density matrices of some 600 real-world scenarios.

Build a comprehensive database of density matrices by combining fragmentation approach and machine learning to characterize the nature of the interactions between the fragments. The sparse matrices should be provided either on the basis of the Support Functions, or in aggregated form, either on the atoms or on the fragments.

CURRENT STATUS OF THE PROJECT

The project commenced on February 1st, 2024. A comprehensive database comprising approximately 600 existing density matrices that describe real-world molecular systems has been compiled. The development environment for the predictor learning phase has been established utilizing the supercomputer facilities at the Juelich Supercomputing Centre in Germany.