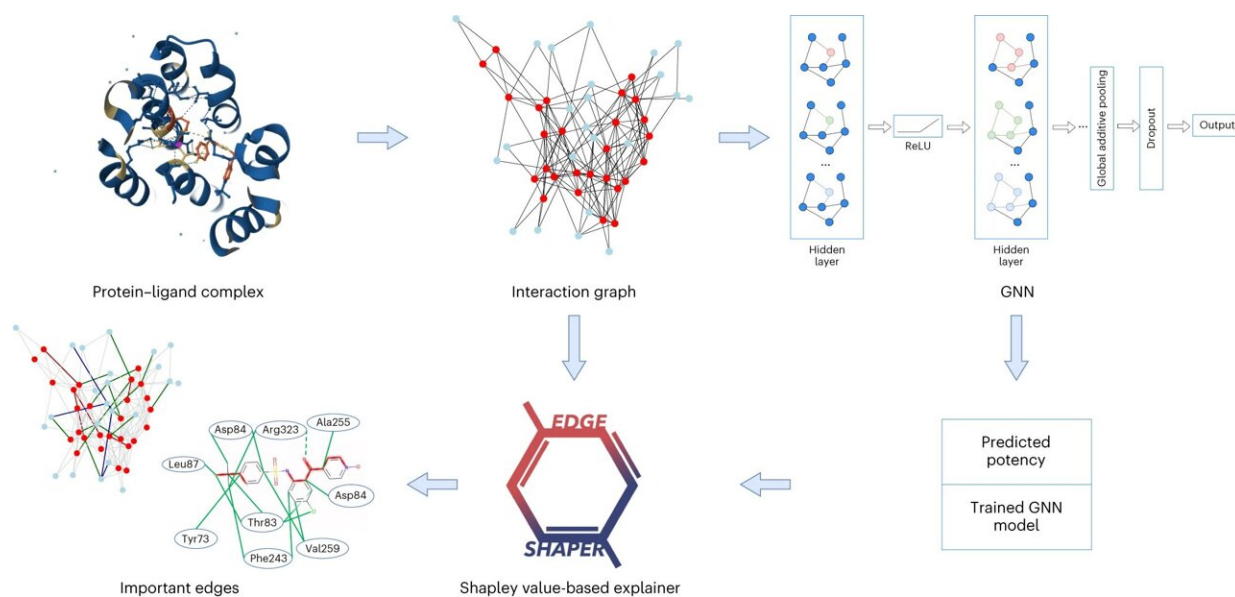


Artificial intelligence for drug discovery offers up unexpected results

November 13 2023, by Johannes Seiler



Rationalizing affinity predictions based on protein–ligand interaction graphs. The schematic representation summarizes the different stages of the analysis including the generation of interaction graphs from X-ray structures for training and testing a GNN to predict numerical affinity values, followed by the determination of edge importance for predictions and delineation of subgraphs determining the predictions. Credit: *Nature Machine Intelligence* (2023). DOI:10.1038/s42256-023-00756-9

Which drug molecule is most effective? Researchers are feverishly searching for efficient active substances to combat diseases. These

compounds often dock onto proteins, which usually are enzymes or receptors that trigger a specific chain of physiological actions.

In some cases, certain molecules are also intended to block undesirable reactions in the body—such as an excessive inflammatory response. Given the abundance of available [chemical compounds](#), this research is like searching for a needle in a haystack at first glance. Drug discovery, therefore, attempts to use scientific models to predict which molecules will best dock to the respective [target protein](#) and bind strongly. These potential drug candidates are then investigated in more detail in experimental studies.

Since the advance of AI, [drug discovery](#) research has also been increasingly using machine learning applications. As one "Graph neural networks" (GNNs) provide one of several opportunities for such applications. They are adapted to predict, for example, how strongly a certain molecule binds to a target [protein](#).

To this end, GNN models are trained with graphs that represent complexes formed between proteins and chemical compounds (ligands). Graphs generally consist of nodes representing objects and edges representing relationships between nodes. In graph representations of protein-ligand complexes, edges connect only protein or ligand nodes, representing their structures, respectively, or protein and ligand nodes, representing specific protein-ligand interactions.

"How GNNs arrive at their predictions is like a black box we can't glimpse into," says Prof. Dr. Jürgen Bajorath. The chemoinformatics researcher from the LIMES Institute at the University of Bonn, the Bonn-Aachen International Center for Information Technology (B-IT), and the Lamarr Institute for Machine Learning and Artificial Intelligence in Bonn, together with colleagues from Sapienza University in Rome, has analyzed in detail whether graph [neural networks](#) actually learn protein-

ligand interactions to predict how strongly an active substance binds to a target protein.

The research is [published](#) in *Nature Machine Intelligence*.

How do the AI applications work?

The researchers analyzed a total of six different GNN architectures using their specially developed "EdgeSHAPer" method and a conceptually different methodology for comparison. These [computer programs](#) "screen" whether the GNNs learn the most important interactions between a compound and a protein and thereby predict the potency of the ligand, as intended and anticipated by researchers—or whether AI arrives at the predictions in other ways.

"The GNNs are very dependent on the data they are trained with," says the first author of the study, Ph.D. candidate Andrea Mastropietro from Sapienza University in Rome, who conducted a part of his doctoral research in Prof. Bajorath's group in Bonn.

The scientists trained the six GNNs with graphs extracted from structures of protein-ligand complexes, for which the mode of action and binding strength of the compounds to their target proteins was already known from experiments. The trained GNNs were then tested on other complexes. The subsequent EdgeSHAPer analysis then made it possible to understand how the GNNs generated apparently promising predictions.

"If the GNNs do what they are expected to, they need to learn the interactions between the compound and target protein and the predictions should be determined by prioritizing specific interactions," explains Prof. Bajorath. According to the research team's analyses, however, the six GNNs essentially failed to do so.

Most GNNs only learned a few protein-drug interactions and mainly focused on the ligands. Bajorath says, "To predict the binding strength of a molecule to a target protein, the models mainly 'remembered' chemically similar molecules that they encountered during training and their binding data, regardless of the target protein. These learned chemical similarities then essentially determined the predictions."

According to the scientists, this is largely reminiscent of the "Clever Hans effect." This effect refers to a horse that could apparently count. How often Hans tapped his hoof was supposed to indicate the result of a calculation. As it turned out later, however, the horse was not able to calculate at all, but deduced expected results from nuances in the facial expressions and gestures of his companion.

What do these findings mean for drug discovery research? "It is generally not tenable that GNNs learn chemical interactions between active substances and proteins," says the cheminformatics scientist.

Their predictions are largely overrated because forecasts of equivalent quality can be made using chemical knowledge and simpler methods. However, the research also offers opportunities of AI.

Two of the GNN-examined models displayed a clear tendency to learn more interactions when the potency of test compounds increased. "It's worth taking a closer look here," says Bajorath. Perhaps these GNNs could be further improved in the desired direction through modified representations and training techniques.

However, the assumption that physical quantities can be learned on the basis of molecular graphs should generally be treated with caution. "AI is not black magic," says Bajorath.

In fact, he sees the previous open-access publication of EdgeSHAPer

and other specially developed analysis tools as promising approaches to shed light on the black box of AI models. His team's approach currently focuses on GNNs and new "chemical language models."

"The development of methods for explaining predictions of complex models is an important area of AI research. There are also approaches for other network architectures such as language models that help to better understand how [machine learning](#) arrives at its results," says Bajorath.

He expects that exciting things will soon also happen in the field of "Explainable AI" at the Lamarr Institute, where he is a PI and Chair of AI in the Life Sciences.

More information: Mastropietro, A. et al, Learning characteristics of graph neural networks predicting protein–ligand affinities, *Nature Machine Intelligence* (2023). [DOI: 10.1038/s42256-023-00756-9](https://doi.org/10.1038/s42256-023-00756-9).
www.nature.com/articles/s42256-023-00756-9

Provided by Rheinische Friedrich-Wilhelms-Universität Bonn

Citation: Artificial intelligence for drug discovery offers up unexpected results (2023, November 13) retrieved 23 April 2024 from <https://techxplore.com/news/2023-11-artificial-intelligence-drug-discovery-unexpected.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.