

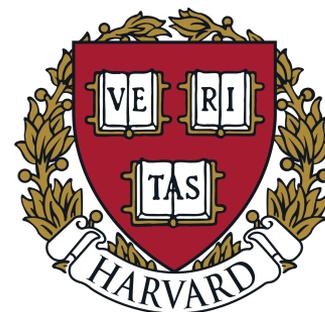
# Machine Learning for Drug Development

Marinka Zitnik

Department of Biomedical Informatics  
Broad Institute of Harvard and MIT  
Harvard Data Science Initiative

marinka@hms.harvard.edu

<https://zitniklab.hms.harvard.edu>



# Outline

- ✓ Overview and introduction
- ✓ Part 1: Virtual drug screening and drug repurposing
- ✓ Part 2: Adverse drug effects, drug-drug interactions
- ✓ Part 3: Clinical trial site identification, patient recruitment
- ✓ Part 4: Molecule optimization, molecular graph generation, multimodal graph-to-graph translation
- ✓ Part 5: Molecular property prediction and transformers

Demos, resources, wrap-up & future directions



# Datasets to facilitate algorithmic innovation

# Therapeutics are one of most exciting areas for computational scientists. However,

Retrieving, curating, and processing datasets is time-consuming and requires extensive domain expertise

Datasets are scattered around the bio repositories and there is no centralized data repository for a variety of therapeutics

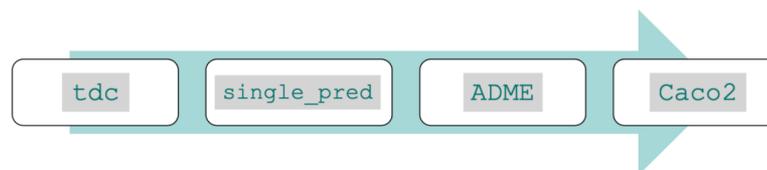
Many tasks are under-explored in AI/ML community because of the lack of data access



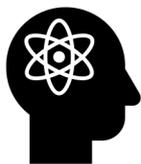
# THERAPEUTICS DATA COMMONS

- **Open-Source ML Datasets for Therapeutics:**
  - **Wide range of tasks:** target discovery, activity screening, efficacy, safety, manufacturing
  - **Wide range of products:** small molecules, antibodies, vaccine, miRNA
- **Numerous Data Functions:**
  - Extensive data functions
  - Model evaluation, data processing and splits, molecule generation oracles, and much more
- **3 Lines of Code:**
  - Minimum package dependency, lightweight loaders

```
from tdc.single_pred import ADME
data = ADME(name = 'Caco2_Wang')
splits = data.split()
```



# Our Vision for TDC



**Domain  
scientists**

Identify meaningful  
learning tasks



Design powerful  
ML models

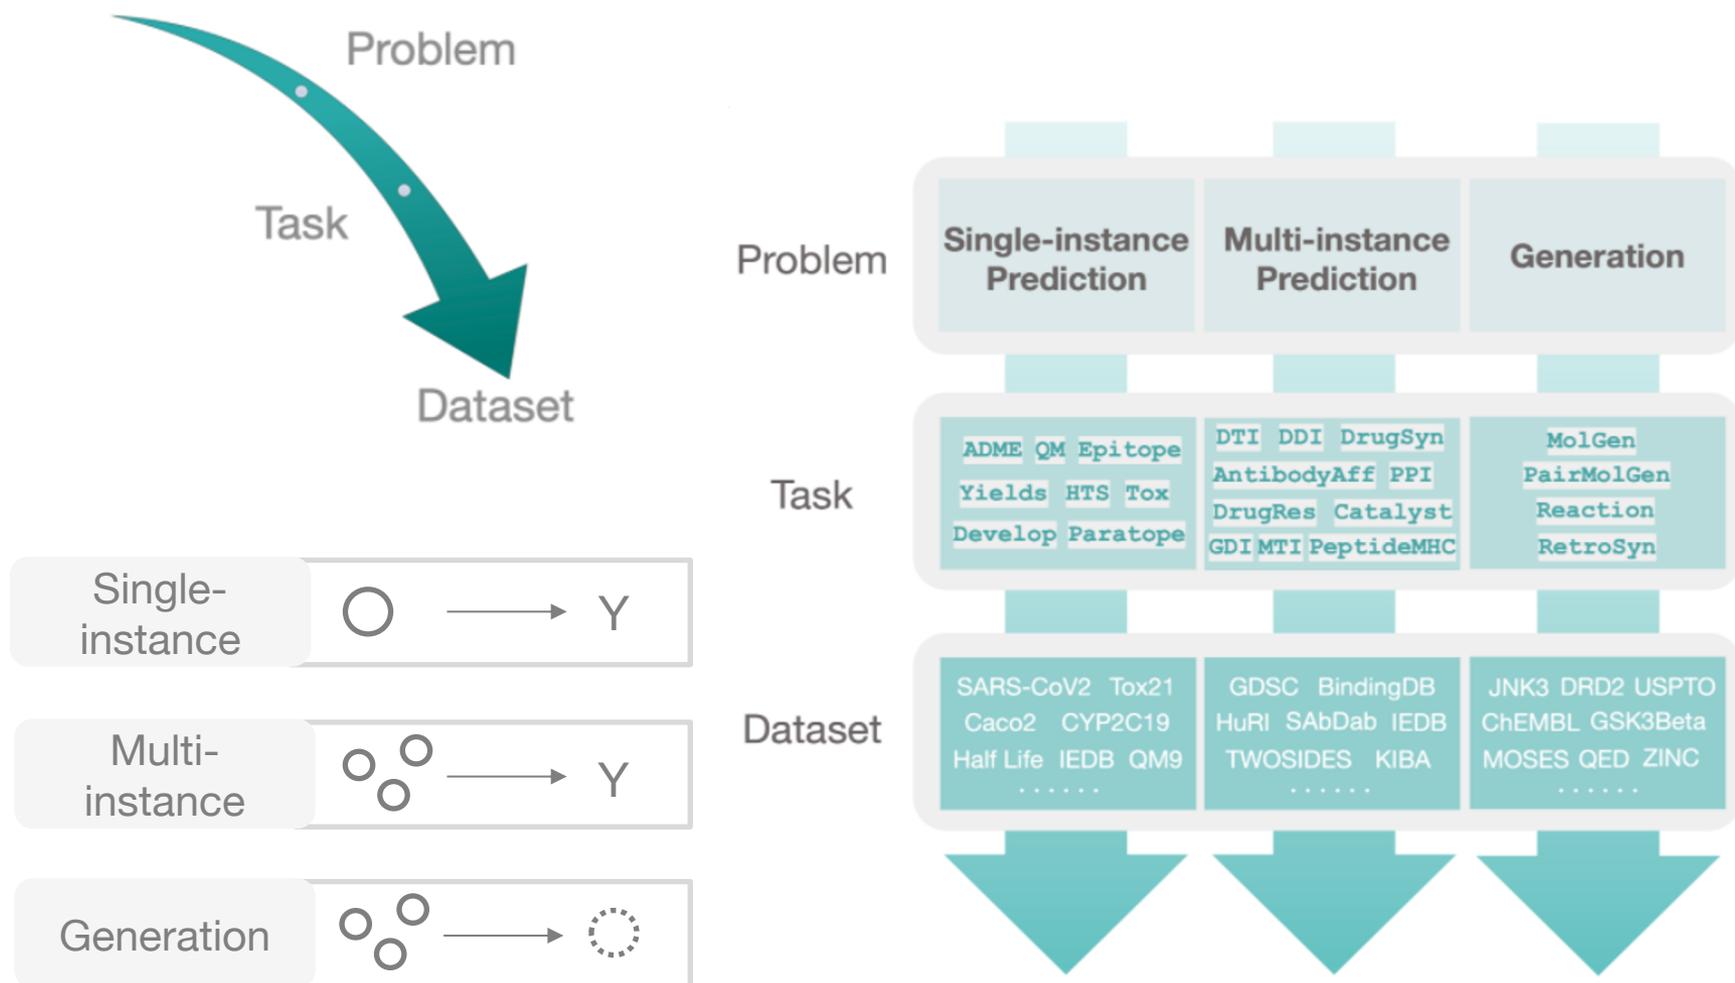


**ML  
scientists**



**Advancing algorithms for key therapeutics problems**

# Modular Structure of TDC



DATASET INDEX

Absorption

- Caco-2 (Cell Effective Permeability), Wang et al.
- HIA (Human Intestinal Absorption), Hou et al.
- Pgp (P-glycoprotein) Inhibition, Broccatelli et al.
- Bioavailability, Ma et al.
- Bioavailability F20/F30, eDrug3D
- Lipophilicity, AstraZeneca
- Solubility, AqSolDB
- Solubility, ESOL
- Hydration Free Energy, FreeSolv

ADME

Distribution

- BBB (Blood-Brain Barrier), Adenot et al.
- BBB (Blood-Brain Barrier), Martins et al.
- PPBR (Plasma Protein Binding Rate), Ma et al.
- PPBR (Plasma Protein Binding Rate), eDrug3D
- VD (Volumn of Distribution), eDrug3D

Metabolism

- CYP P450 2C19 Inhibition, Veith et al.
- CYP P450 2D6 Inhibition, Veith et al.
- CYP P450 3A4 Inhibition, Veith et al.
- CYP P450 1A2 Inhibition, Veith et al.
- CYP P450 2C9 Inhibition, Veith et al.

Excretion

- Half Life, eDrug3D
- Clearance, eDrug3D

DATASET INDEX

- BindingDB
- DAVIS
- KIBA

DTI

DATASET INDEX

- SARS-CoV-2 In Vitro, Touret et al.
- SARS-CoV-2 3CL Protease, Diamond.
- HIV

HTS

DATASET INDEX

- DisGeNET

GDA

DATASET INDEX

- GDSC1
- GDSC2

DrugRes

DATASET INDEX

- OncoPolyPharmacology

DrugSyn

DATASET INDEX

- Tox21
- ToxCast
- ClinTox

Tox

DATASET INDEX

- USPTO

Reaction

DATASET INDEX

- MOSES
- ZINC
- CHEMBL

MolGen

DATASET INDEX

- DRD2
- QED
- LogP

PairMolGen

DATASET INDEX

- USPTO-50K
- USPTO

RetroSyn

DATASET INDEX

- IEDB, Jespersen et al.
- PDB, Jespersen et al.

Epitope

DATASET INDEX

- miRTarBase

MTI

DATASET INDEX

- USPTO

Catalyst

DATASET INDEX

- HuRI

PPI

DATASET INDEX

- TAP
- SAbDab, Chen et al.

Develop

DATASET INDEX

- DrugBank Multi-Typed DDI
- TWOSIDES Polypharmacy Side Effects

DDI

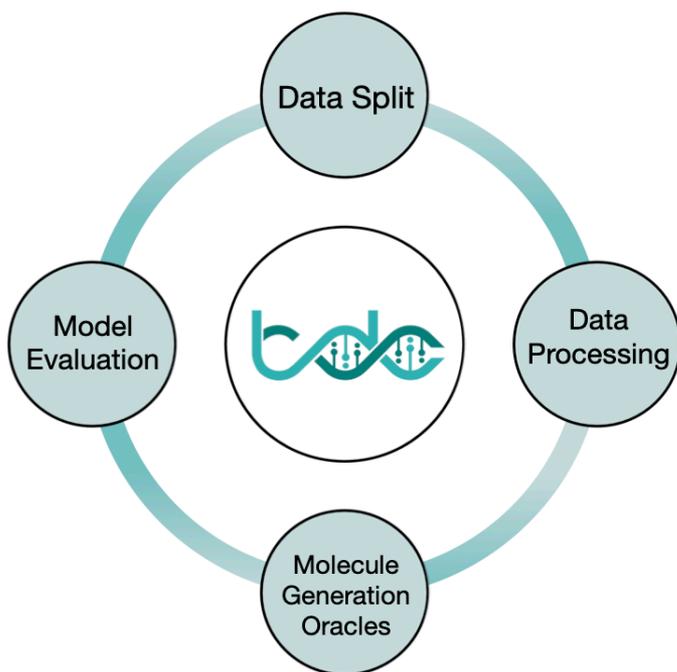
DATASET INDEX

- Buchwald-Hartwig
- USPTO

Yields

67 datasets spread over 22 learning tasks

# Data Functions to Support Your Research



## Model performance evaluators

### FUNCTION INDEX

#### Regression Metric

- Mean Squared Error (MSE)
- Mean Absolute Error (MAE)
- Coefficient of Determination ( $R^2$ )

#### Binary Classification Metric

- Area Under the Receiver Operating Characteristic Curve (ROC-AUC)
- Area Under the Precision-Recall Curve (PR-AUC)
- Accuracy Metric
- Precision
- Recall
- F1 Score

#### Multi-class Classification Metric

- Micro-F1, Micro-Precision, Micro-Recall, Accuracy
- Macro-F1
- Cohen's Kappa ( $\kappa$ )

#### Token-level Classification Metric

- Average ROC-AUC

## A variety of data splits

### FUNCTION INDEX

#### Data Split Overview

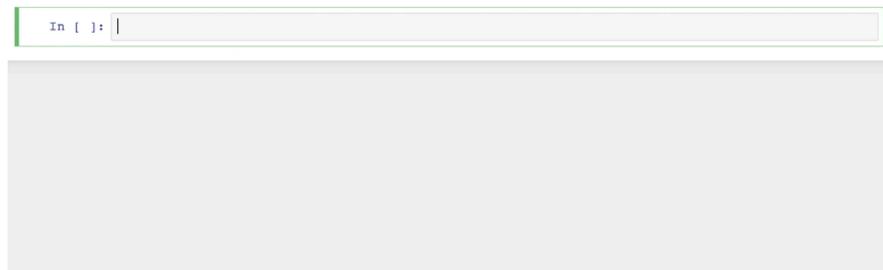
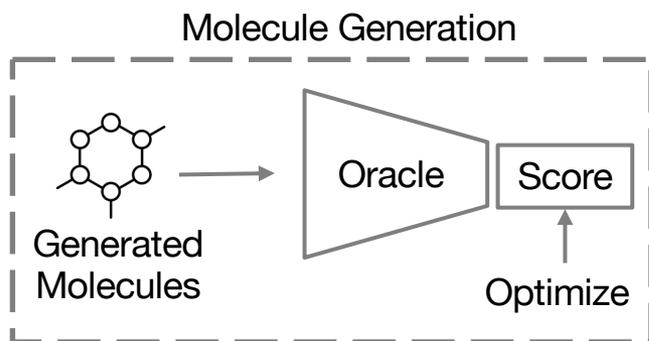
- Random Split
- Scaffold Split
- Cold-Start Split

## Data processing helpers

### FUNCTION INDEX

- Label Distribution Visualization
- Label Binarization
- Label Units Conversion
- Label Meaning
- Basic Statistics
- Data Balancing
- Graph Transformation for Pair Data
- Negative Samples for Pair Data
- From PubChem CID to SMILES
- From Uniprot ID to Amino Acid Sequence

# Molecule Generation Oracles



GuacaMol



MOSES



Literature

## FUNCTION INDEX

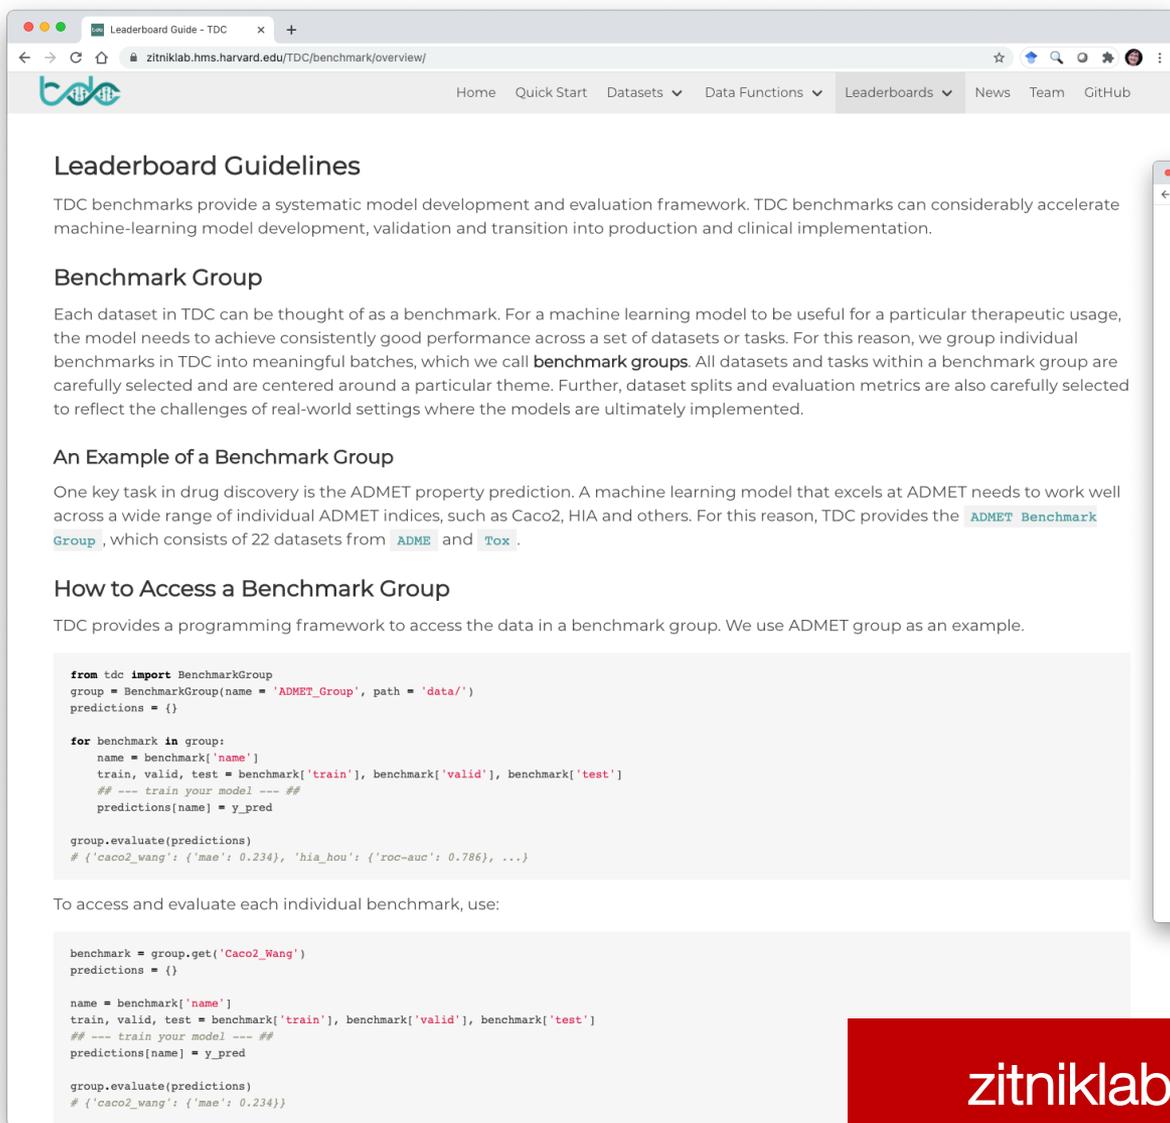
### Goal-oriented Oracles

- Glycogen Synthase Kinase 3 Beta (GSK3 $\beta$ )
- c-Jun N-terminal Kinases-3 (JNK3)
- Dopamine Receptor D2 (DRD2)
- Synthetic Accessibility (SA)
- IBM RXN Synthetic Accessibility (IBM\_RXN)
- Quantitative Estimate of Drug-likeness (QED)
- Octanol-water Partition Coefficient (LogP)
- Rediscovery
- Similarity/Dissimilarity
- Median Molecules
- Isomers
- Multi-Property Objective (MPO)
- Valsartan SMARTS
- Hop

### Distribution Learning Oracles

- Diversity
- KL divergence
- Frechet ChemNet Distance (FCD)
- Novelty
- Validity
- Uniqueness

# Leaderboards: Submit your Models



Leaderboard Guide - TDC

zitniklab.hms.harvard.edu/TDC/benchmark/overview/

## Leaderboard Guidelines

TDC benchmarks provide a systematic model development and evaluation framework. TDC benchmarks can considerably accelerate machine-learning model development, validation and transition into production and clinical implementation.

## Benchmark Group

Each dataset in TDC can be thought of as a benchmark. For a machine learning model to be useful for a particular therapeutic usage, the model needs to achieve consistently good performance across a set of datasets or tasks. For this reason, we group individual benchmarks in TDC into meaningful batches, which we call **benchmark groups**. All datasets and tasks within a benchmark group are carefully selected and are centered around a particular theme. Further, dataset splits and evaluation metrics are also carefully selected to reflect the challenges of real-world settings where the models are ultimately implemented.

## An Example of a Benchmark Group

One key task in drug discovery is the ADMET property prediction. A machine learning model that excels at ADMET needs to work well across a wide range of individual ADMET indices, such as Caco2, HIA and others. For this reason, TDC provides the [ADMET Benchmark Group](#), which consists of 22 datasets from [ADME](#) and [Tox](#).

## How to Access a Benchmark Group

TDC provides a programming framework to access the data in a benchmark group. We use ADMET group as an example.

```
from tdc import BenchmarkGroup
group = BenchmarkGroup(name = 'ADMET_Group', path = 'data/')
predictions = {}

for benchmark in group:
    name = benchmark['name']
    train, valid, test = benchmark['train'], benchmark['valid'], benchmark['test']
    ## --- train your model --- ##
    predictions[name] = y_pred

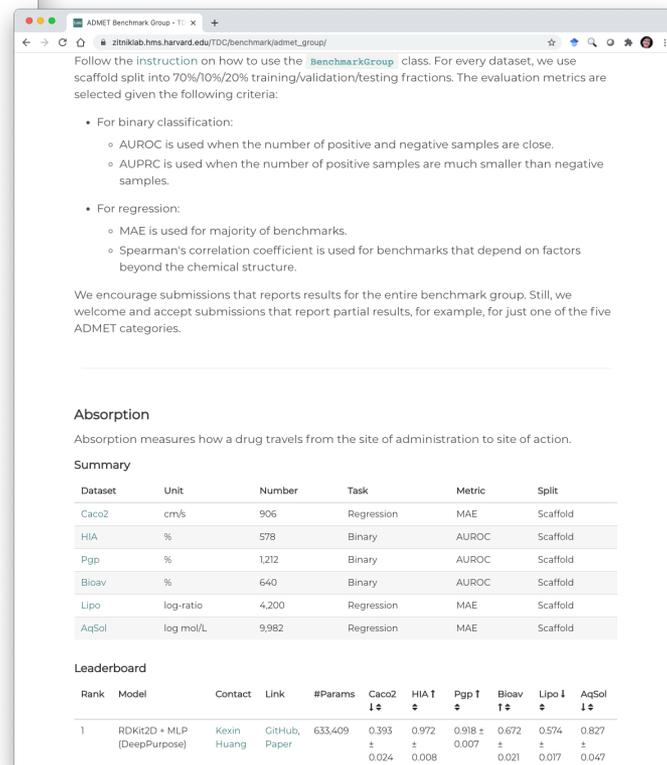
group.evaluate(predictions)
# {'caco2_wang': {'mae': 0.234}, 'hia_hou': {'roc_auc': 0.786}, ...}
```

To access and evaluate each individual benchmark, use:

```
benchmark = group.get('Caco2_Wang')
predictions = {}

name = benchmark['name']
train, valid, test = benchmark['train'], benchmark['valid'], benchmark['test']
## --- train your model --- ##
predictions[name] = y_pred

group.evaluate(predictions)
# {'caco2_wang': {'mae': 0.234}}
```



ADMET Benchmark Group - TDC

zitniklab.hms.harvard.edu/TDC/benchmark/admet\_group/

Follow the instruction on how to use the `BenchmarkGroup` class. For every dataset, we use scaffold split into 70%/10%/20% training/validation/testing fractions. The evaluation metrics are selected given the following criteria:

- For binary classification:
  - AUROC is used when the number of positive and negative samples are close.
  - AUPRC is used when the number of positive samples are much smaller than negative samples.
- For regression:
  - MAE is used for majority of benchmarks.
  - Spearman's correlation coefficient is used for benchmarks that depend on factors beyond the chemical structure.

We encourage submissions that reports results for the entire benchmark group. Still, we welcome and accept submissions that report partial results, for example, for just one of the five ADMET categories.

## Absorption

Absorption measures how a drug travels from the site of administration to site of action.

### Summary

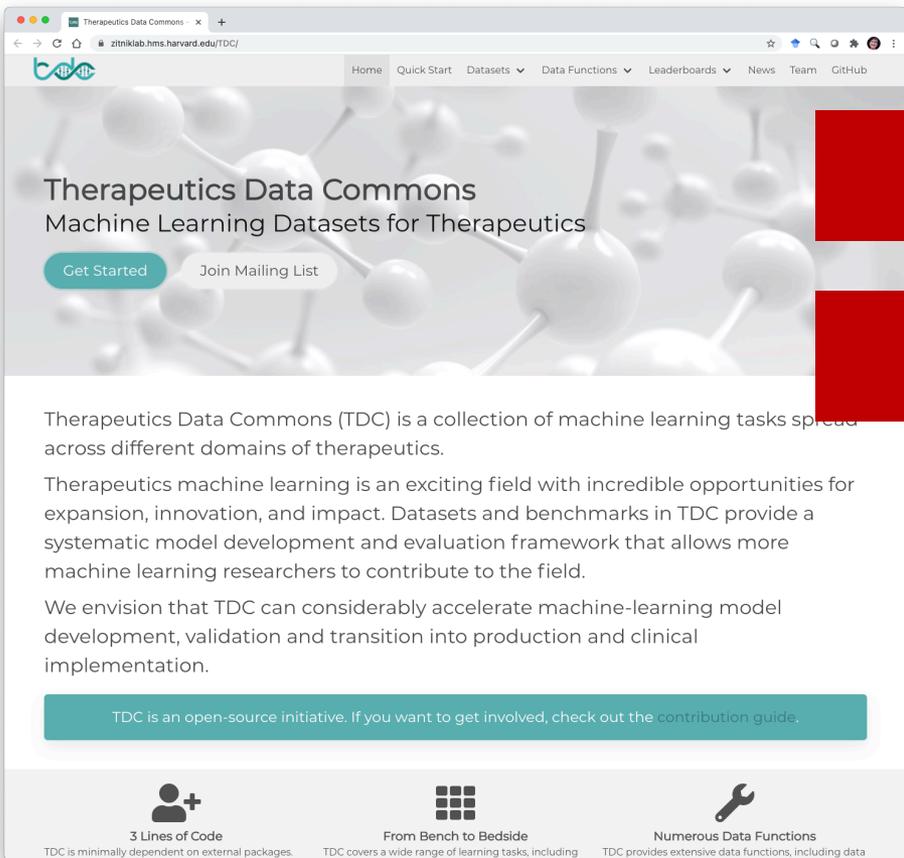
Dataset	Unit	Number	Task	Metric	Split
Caco2	cm/s	906	Regression	MAE	Scaffold
HIA	%	578	Binary	AUROC	Scaffold
Pgp	%	1,212	Binary	AUROC	Scaffold
Bioav	%	640	Binary	AUROC	Scaffold
Lipo	log-ratio	4,200	Regression	MAE	Scaffold
AqSol	log mol/L	9,982	Regression	MAE	Scaffold

### Leaderboard

Rank	Model	Contact	Link	#Params	Caco2 ↓	HIA ↑	Pgp ↑	Bioav ↑	Lipo ↓	AqSol ↓
1	RDKit2D + MLP (DeepPurpose)	Kevin Huang	GitHub, Paper	633,409	0.393 ± 0.024	0.972 ± 0.008	0.918 ± 0.007	0.672 ± 0.021	0.574 ± 0.017	0.827 ± 0.047

zitniklab.hms.harvard.edu/TDC

# You Are Invited to Join TDC! TDC is an Open-Source, Community Effort



The screenshot shows the homepage of the Therapeutics Data Commons (TDC) website. The header includes a navigation menu with links for Home, Quick Start, Datasets, Data Functions, Leaderboards, News, Team, and GitHub. The main content area features a large molecular structure graphic and the text "Therapeutics Data Commons Machine Learning Datasets for Therapeutics". Below this are two buttons: "Get Started" and "Join Mailing List". A paragraph of text describes TDC as a collection of machine learning tasks spread across different domains of therapeutics. Another paragraph explains that therapeutics machine learning is an exciting field with opportunities for expansion, innovation, and impact, and that TDC provides a systematic model development and evaluation framework. A third paragraph states the vision that TDC can accelerate machine-learning model development, validation, and transition into production and clinical implementation. At the bottom, there is a teal button that says "TDC is an open-source initiative. If you want to get involved, check out the contribution guide." The footer contains three icons with corresponding text: "3 Lines of Code" (TDC is minimally dependent on external packages), "From Bench to Bedside" (TDC covers a wide range of learning tasks, including), and "Numerous Data Functions" (TDC provides extensive data functions, including data).

[zitniklab.hms.harvard.edu/TDC](https://zitniklab.hms.harvard.edu/TDC)

[github.com/mims-harvard/TDC](https://github.com/mims-harvard/TDC)

## Tutorials

We provide a series of tutorials for you to get started using TDC:

Name	Description
101	Introduce TDC Data Loaders
102	Introduce TDC Data Functions
103.1	Walk through TDC Small Molecule Datasets
103.2	Walk through TDC Biologics Datasets
104	Generate 21 ADME ML Predictors with 15 Lines of Code
105	Molecule Generation Oracles

```
pip install PyTDC
```

# Demos, tools, and implementations

# DeepPurpose: Deep Learning Library for Compound and Protein Modeling

## DTI, Drug Property, PPI, DDI, Protein Function Prediction

<https://github.com/kexinhuang12345/DeepPurpose>

### Demos

Checkout 10+ demos & tutorials to start:

Name	Description
<a href="#">Dataset Tutorial</a>	Tutorial on how to use the dataset loader and read customized data
<a href="#">Drug Repurposing for 3CLPro</a>	Example of one-liner repurposing for 3CLPro
<a href="#">Drug Repurposing with Customized Data</a>	Example of one-liner repurposing with AID1706 Bioassay Data, training from scratch
<a href="#">Virtual Screening for BindingDB IC50</a>	Example of one-liner virtual screening
<a href="#">Reproduce DeepDTA</a>	Reproduce <a href="#">DeepDTA</a> with DAVIS dataset and show how to use the 10 lines framework
<a href="#">Virtual Screening for DAVIS and Correlation Plot</a>	Example of one-liner virtual screening and evaluate on unseen dataset by plotting correlation
<a href="#">Binary Classification for DAVIS using CNNs</a>	Binary Classification for DAVIS dataset using CNN encodings by using the 10 lines framework.
<a href="#">Pretraining Model Tutorial</a>	Tutorial on how to load pretraining models

and more in the [DEMO](#) folder!

DeepPurpose: a Deep Learning Library for Drug-Target Interaction Prediction, *Bioinformatics* 2020

# How can domain scientists interact with AI systems?

Gradio

57434.gradio.app

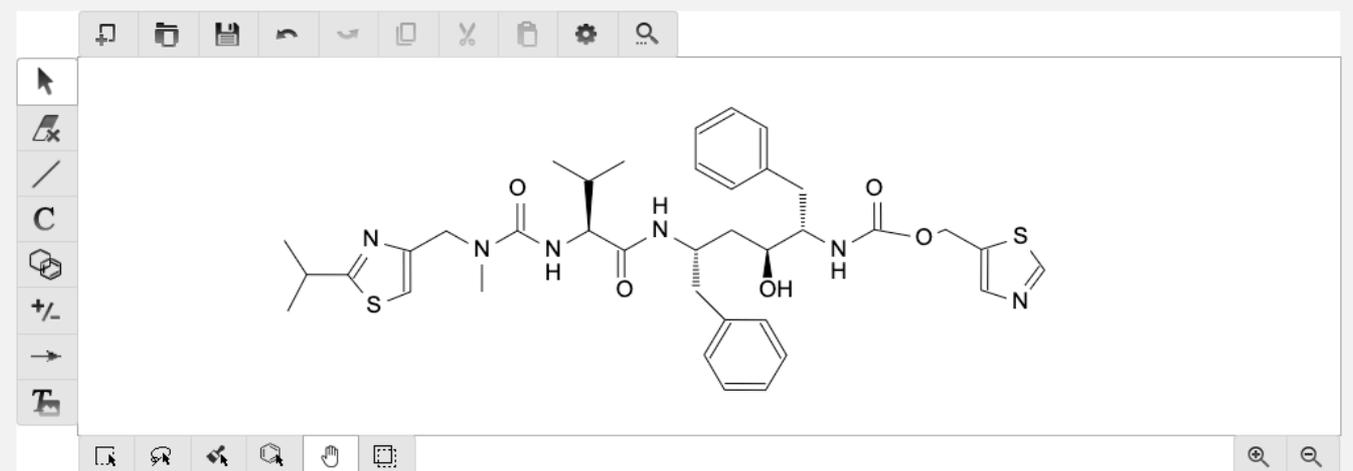
Live at <https://57434.gradio.app> Copy Link

## Interactive Molecular Design with Real-Time Binding Affinity and ADMET Prediction, powered by DeepPurpose

AMINO ACID SEQUENCE

```
SGFRKMAFPSPGKVEGCMVQVTCGTTTLNGLWLDVYVCPRHVICTSEDMLNPNYEDLLIRKSNHNFLVQAGNVQLRVIGHSMQNCVLKLVKVD  
TANPKTPKYKFVRIQPGQTFVSLACYNGSPSGVYQCAMPNFTIKGSFLNGSCGSVGFNIDYDCVSFCYMHMELPTGVHAGTDLEGNFYGP  
FVDRQTAQAAGTDTTITVNVLAWLYAAVINGDRWFLNRFTTTLNDFNLVAMKYNYEPLTQDHVDILGPLSAQTGIAVLDMCASLKELLQNGM
```

MOLECULE



The screenshot displays a web browser window with a Gradio interface. The browser address bar shows the URL '57434.gradio.app'. The main content area features a title 'Interactive Molecular Design with Real-Time Binding Affinity and ADMET Prediction, powered by DeepPurpose'. Below the title, there is a section for 'AMINO ACID SEQUENCE' containing a long string of amino acid codes. Underneath, a 'MOLECULE' section shows a 3D ball-and-stick model of a complex organic molecule. The molecule has a central backbone with various side chains, including a thiazole ring, a benzene ring, and a hydroxyl group. The interface includes a toolbar with icons for file operations and a search bar.

DeepPurpose: a Deep Learning Library for Drug-Target Interaction Prediction, *Bioinformatics* 2020

MolDesigner: Interactive Design of Efficacious Drugs with Deep Learning, *NeurIPS* 2020

ML for Drug Development - <https://zitniklab.hms.harvard.edu/drugml/> - Tutorial at IJCAI, Jan 6, 2021

# MolDesigner: Interactive Design of Drugs with Deep Learning

The screenshot displays the MolDesigner web interface. At the top, the browser address bar shows "Not Secure | deeppurpose.sunlab.org". The main content area is divided into several sections:

- AMINO ACID SEQUENCE:** A text box containing the sequence: "LGGSSVAIKITEHSWNADLYKLMGHFAWWTAFVTNVNASSEAFILGNYLGGKPREQIDGY VMHANYIFWRNTNPQLSSYSLPDMSKFPFLKLRGTAVMSLKEGQINDMILSLLSKGRLLI RENNRRVVISSDVLVNN".
- MOLECULE:** A central area showing a 3D ball-and-stick model of a complex organic molecule. The molecule features a central phosphorus atom bonded to a benzene ring, a nitrogen atom, and two oxygen atoms. It is also connected to a chain of atoms including a carbonyl group, a nitrogen atom, and a complex heterocyclic ring system with an amino group and a nitrile group.
- AFFINITY PREDICTION MODEL TYPE:** A dropdown menu set to "Daylight-AAC".
- ADMET PREDICTION MODEL TYPE:** A dropdown menu set to "MPNN".
- Buttons:** "CLEAR" and "SUBMIT" buttons are located below the model type dropdowns.
- CANONICAL SMILES:** A text box containing the SMILES string: CCC(CC)COC(=O)[C@H](C)NP(=O)(Oc1ccccc1)OC[C@H]2O[C@](C#N)([C@H]([C@@H]2O)O)c3n4c(cc3)C(N)=NC=N4
- BINDING AFFINITY (IC50):** A text box displaying "3896.74 nM".
- BINDING AFFINITY (PIC50):** A text box displaying "5.41".
- PREDICTED ADMET PROPERTY:** A table with two columns: "Property" and "Value".

Property	Value
Solubility	-3.57 log mol/L
Lipophilicity	0.68 (log-ratio)
(Absorption) Caco-2	-5.29 cm/s
(Absorption) HIA	77.32 %
(Absorption) Pgp	6.44 %
(Absorption) Bioavailability F20	75.45 %

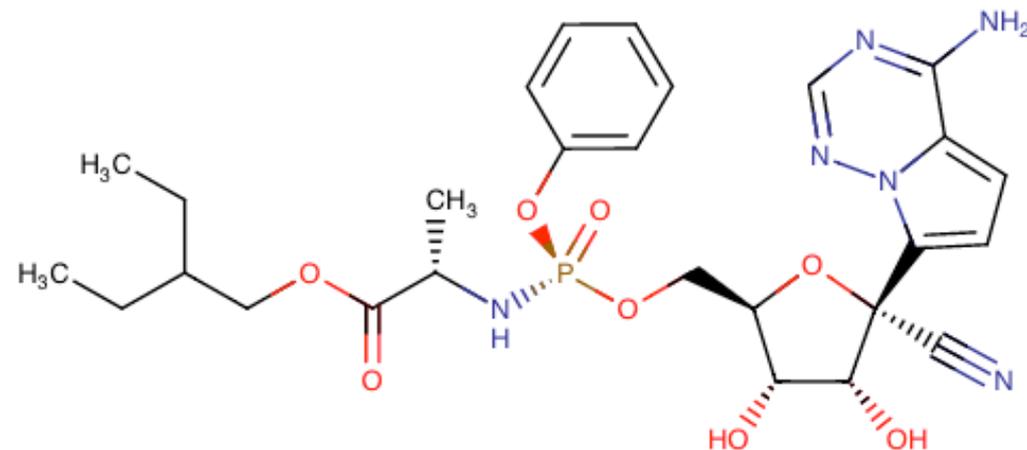
<http://deeppurpose.sunlab.org>

# DEMO: DRUG-TARGET INTERACTION PREDICTION

Drug: [Remdesivir](#) Remdesivir is indicated for the treatment of adult and pediatric patients aged 12 years and over weighing at least 40 kg for coronavirus disease 2019 (COVID-19) infection requiring hospitalization.

Target protein: [Replicase polyprotein 1ab](#). Multifunctional protein involved in the transcription and replication of viral RNAs

```
>lcl|BSEQ0052511|Replicase polyprotein 1ab
MESLVPGFNEKTHVQLSLPVLQVRDVLVRGFGDSVEEVLSEARQHLKDGTCGLVEVEKGV
LPQLEQPYVFIKRS DARTAPHGHVMVELVAELEGIQYGRSGETLGLVLPVHGEIPVAYRK
VLLRKNNGKAGGHSYGADLKSFDLGDDELGDPYEDFQENWNTKHS SGTRELMRELNGG
AYTRYVDNMF CGPDGYPLEC IKDLLARAGKASCTLSEQLDFIDTKRGVYCCREHEHEIAW
YTERSEKSYELQTPFEIKLAKKFDTFN GECPNFVPLNSIIKTIQPRVEKKKLDGFMGRI
RSVYVPVASPNECNQMCLSTLMKCDHCGETSWQTGDFVKATCEFCGTENLTKEGATTCGYL
PQNAVVKIYCPACHNSEVGP EHS LAEYHNESGLKTI LRKGGRTIAFGGCVF SYVVGCHNKC
AYWVPRASANIGCNHTGVV GEGSEGLNDNLEILQKEKVNINIVGDFKLN EEAIIILASF
SASTSAFVETVKGLDYKAFKQIVESC GNFKVTKGKAKKGAWNIGEQKSILSPLYAFASEA
ARVVRSIFSRTLETAQNSVRVLQKAAITILDGISQYSLRLIDAMMFTSDLATNNLVVMAY
ITGGVVQLTSQWLTNIFGTVYEKLPVLDWLEEFKEGVEFLRDGW EIVKFISTCACEIV
GGQIVTCAKEIKESVQTF FKL VNKFLALCADSIIIGGAKL KALNLGETFVTHSKGLYRKC
VKSREETGLMLPKAPKEIIFLEGETLPTEVLTEEVVLKTGDLQPLEQPTSEAVEAPLVG
TPVCINGLMLEIKDTEKYCALAPNMMVTNNTFTLKGGA PTKVTFGDDTVIEVQGYKSVN
ITFELDERIDKVLNEKCSAYTVELGTEVNEFACVVADAVIKTLQPVSELLTPLGIDLDEW
SMATYYLFDSEGEFKLASHMYCSFYPPDEDEEEGDCEEEEFEPSTQY EYGTEDDYQGKPL
EFGATSAALQPEEEQEEDWLDDDSQQT VGGQDGS EDNQT TTIQTIVEVQPQLEMELTPVV
QTIEVNSFSGYLKLTDNVYIKNADIVEEAKVKPTVVVNAANVYLKHGGVAGALNKATN
```



Molecular structure of Remdesivir

Amino acid sequence of Replicase polyprotein 1ab

# How can domain scientists interact with AI systems?

Gradio

57434.gradio.app

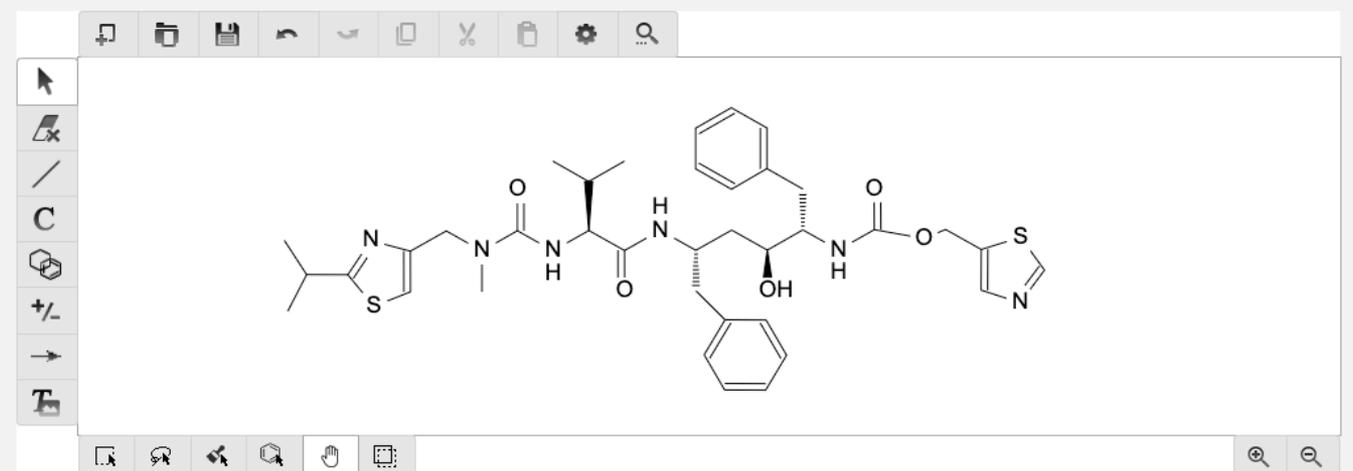
Live at <https://57434.gradio.app> Copy Link

## Interactive Molecular Design with Real-Time Binding Affinity and ADMET Prediction, powered by DeepPurpose

AMINO ACID SEQUENCE

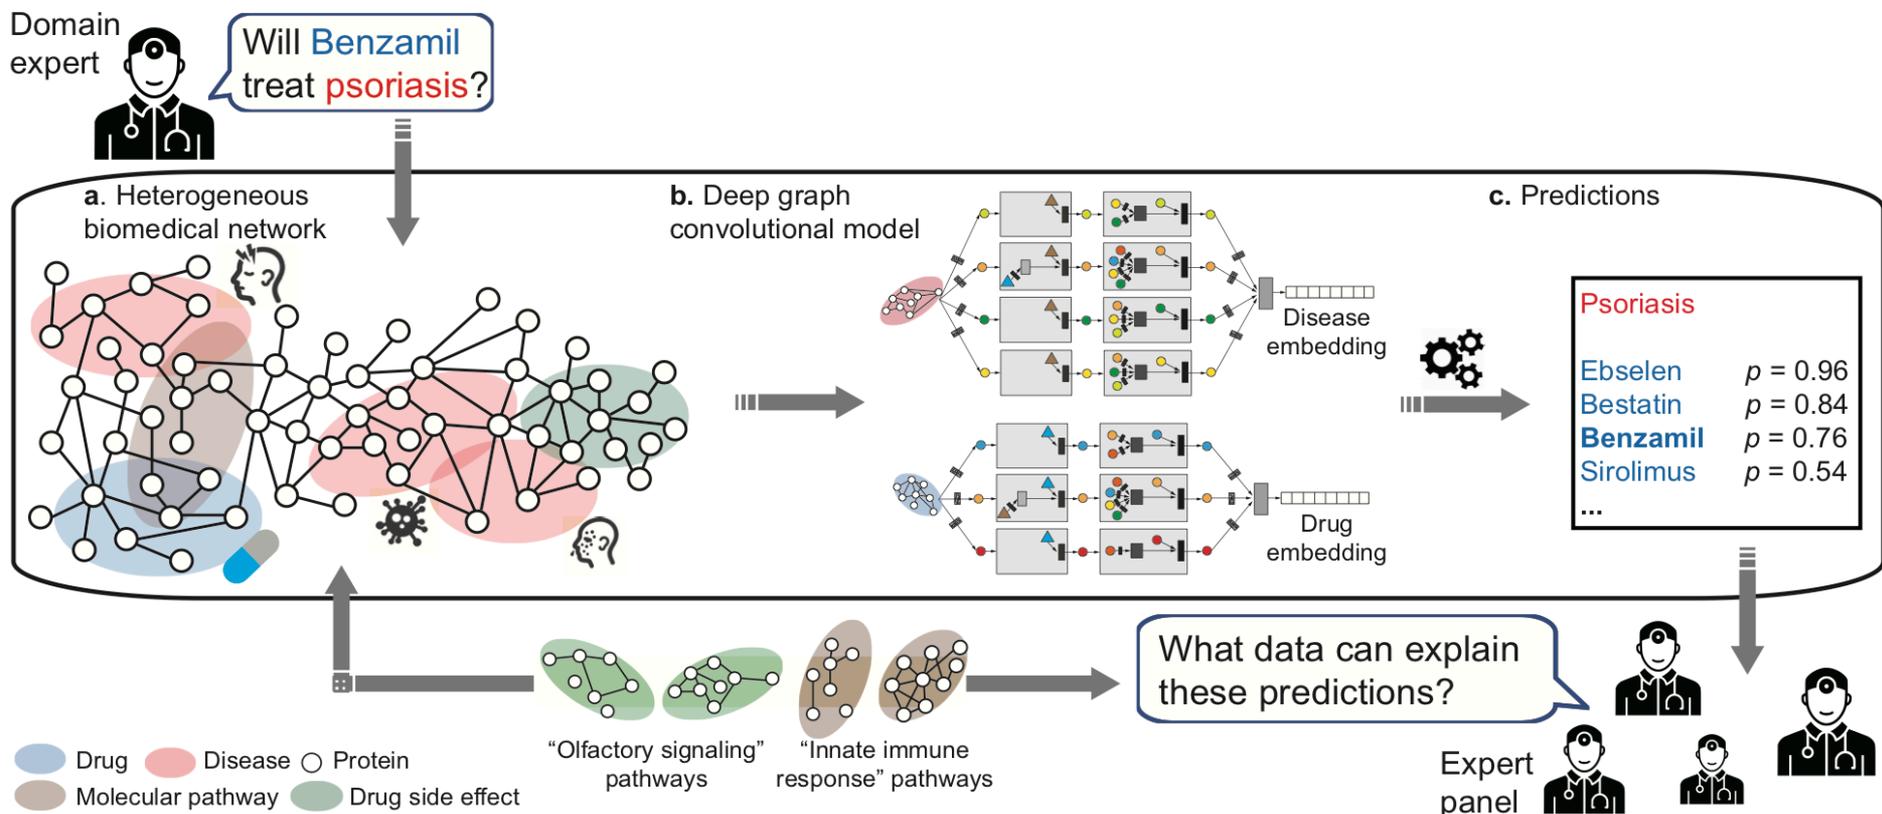
```
SGFRKMAFPSPGKVEGCMVQVTCGTTTLNGLWLDVYVCPRHVICTSEDMLNPNYEDLLIRKSNHNFLVQAGNVQLRVIGHSMQNCVLKLVKVD  
TANPKTPKYKFVRIQPGQTFVSLACYNGSPSGVYQCAMPNFTIKGSFLNGSCGSVGFNIDYDCVSFCYMHMELPTGVHAGTDLEGNFYGP  
FVDRQTAQAAGTDTTITVNVLAWLYAAVINGDRWFLNRFTTTLNDFNLVAMKYNYEPLTQDHVDILGPLSAQTGIAVLDMCASLKELLQNGM
```

MOLECULE

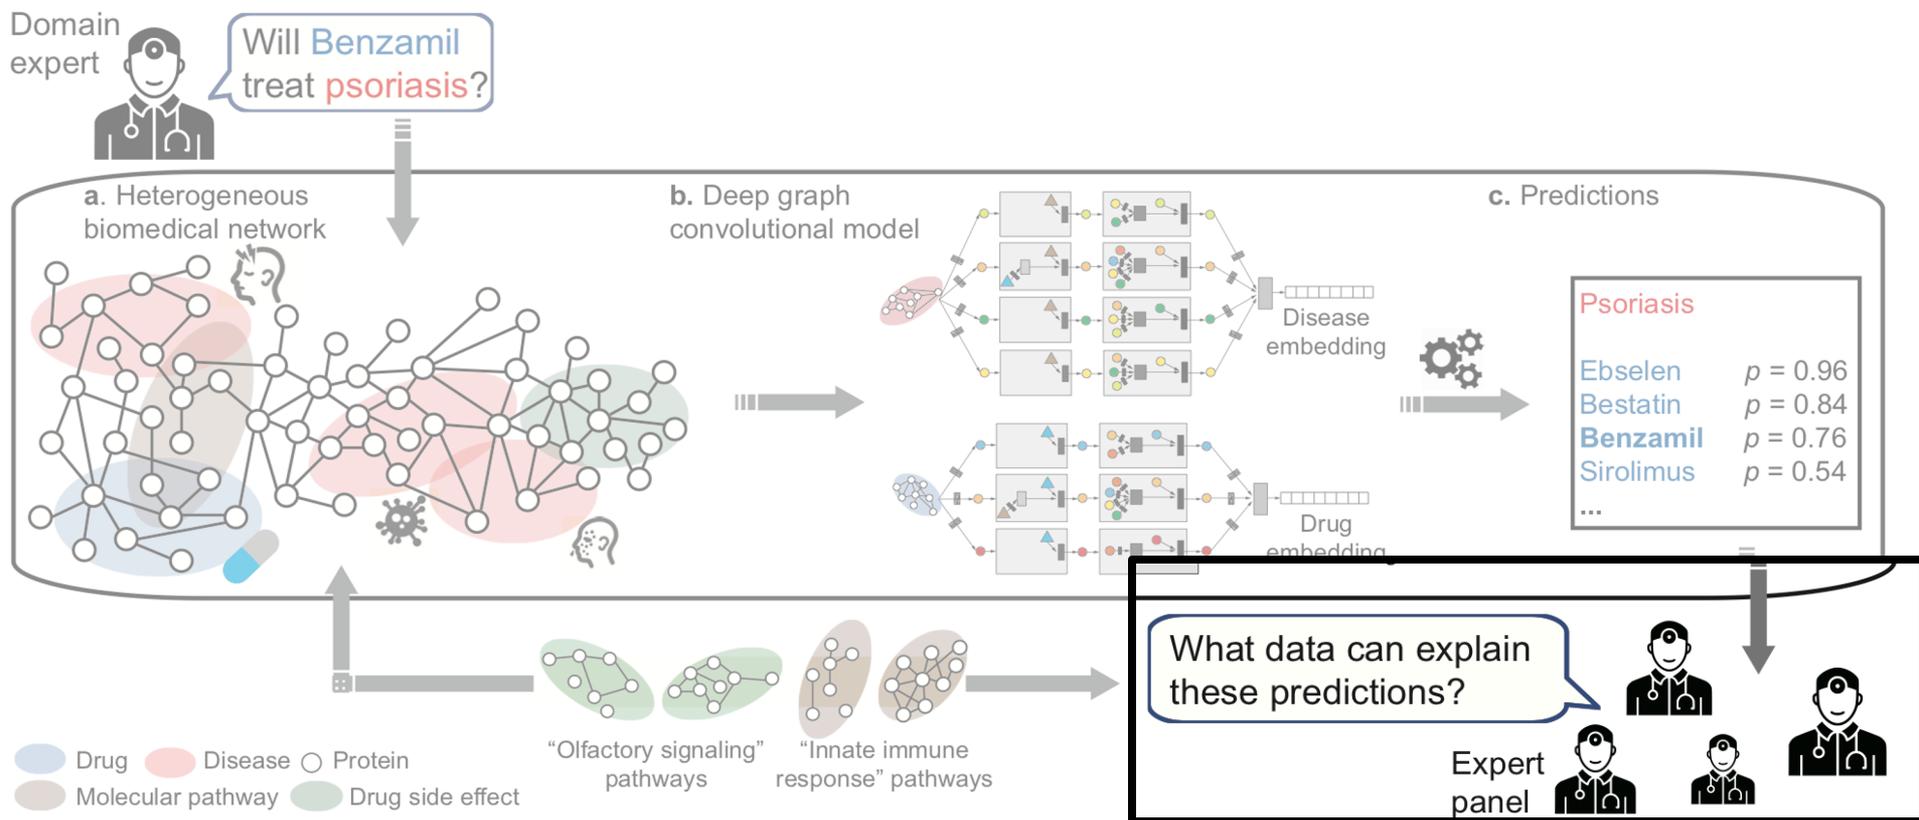


DeepPurpose: a Deep Learning Library for Drug-Target Interaction Prediction, *Bioinformatics* 2020  
MolDesigner: Interactive Design of Efficacious Drugs with Deep Learning, *NeurIPS* 2020

# Automating Science



# Automating Science



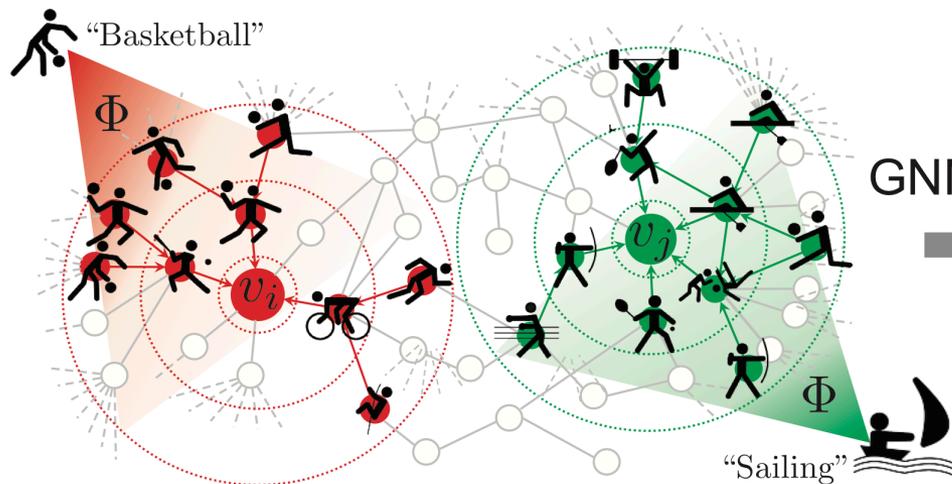
# How to explain predictions?

Key idea:

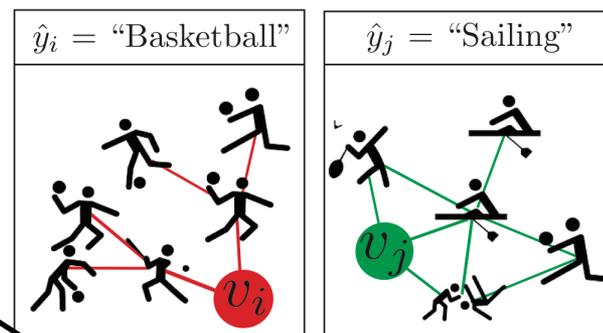
- Summarize where in the data the model “looks” for evidence for its prediction
- Find a small subgraph **most influential** for the prediction



GNN model training and predictions



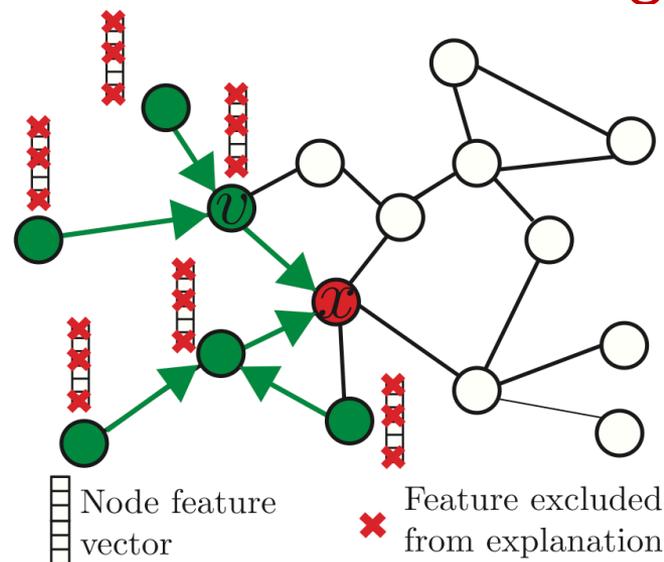
Explaining GNN's predictions



Approach to generate explanations for graph neural networks based on **counterfactual reasoning**

# GNNExplainer: Key Idea

- **Input:** Given prediction  $f(x)$  for node/link  $x$
- **Output:** Explanation, a small subgraph  $M_x$  together with a small subset of node features:
  - $M_x$  is most influential for prediction  $f(x)$
- **Approach:** Learn  $M_x$  via **counterfactual reasoning**
  - **Intuition:** If removing  $v$  from the graph strongly decreases the probability of prediction  $\Rightarrow v$  is a good counterfactual explanation for the prediction

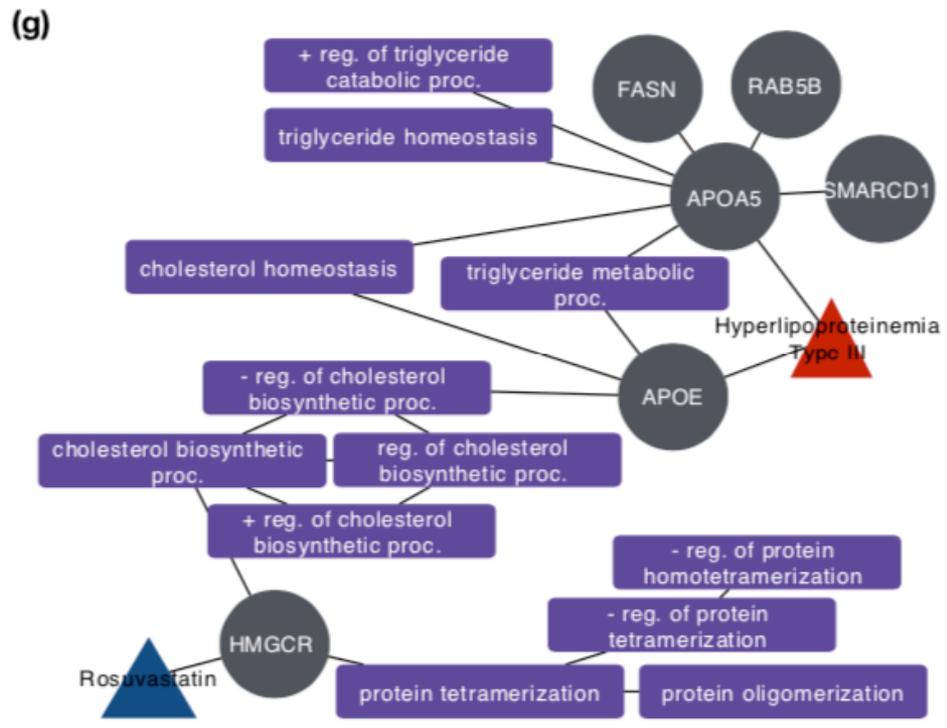
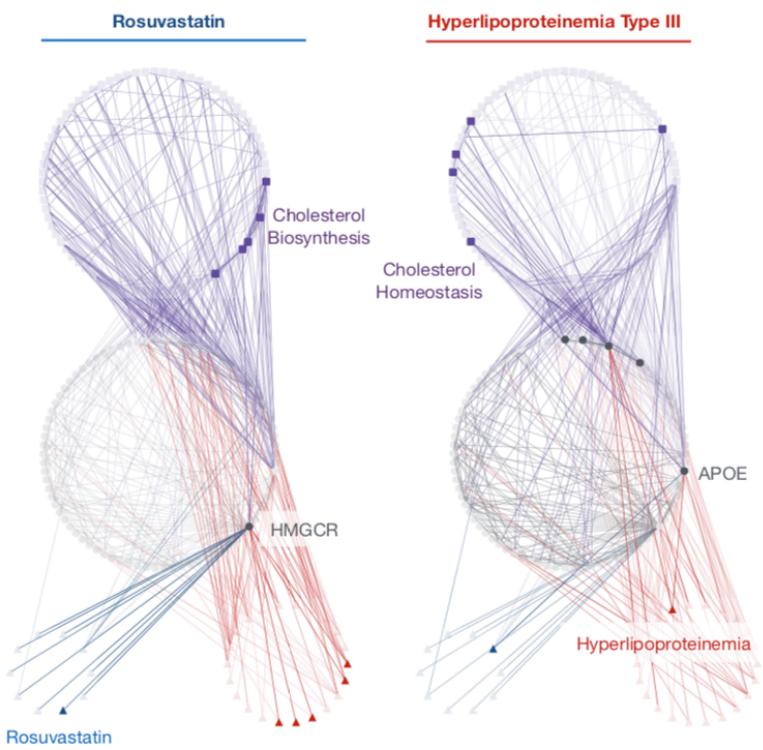


# Examples of Explanations

”Will **rosuvastatin** treat **hyperlipidemia**? What is the disease treatment mechanism?”

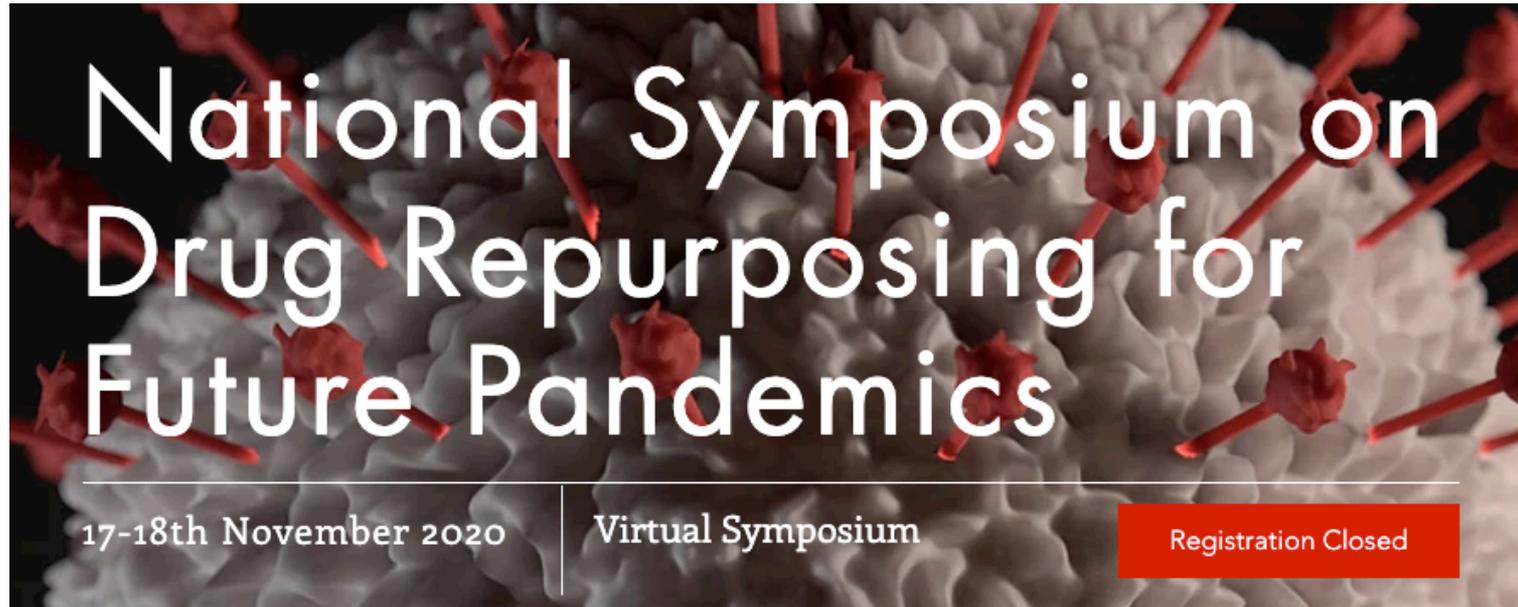


- Functional Pathways (n = 9,798)
- ▬ system
- ▬ organ
- ▬ tissue
- ▬ cell
- Proteins (n = 17,660)
- ▲ Drugs (n = 1,661)
- ▲ Diseases (n = 840)



# Open challenges and future directions

# Learn about Therapeutics ML!



<https://www.drugsymposium.org>

Videos from the presentations are now publicly available to everyone through the Symposium Video Channel

# Open Challenges

- **Disconnected, uncoupled biomedical knowledge:**
  - Challenge: Need to combine data in their broadest sense to close the gap between research and patient data
- **Diverse mechanisms of drug action:**
  - Challenge: Need to consider diverse mechanisms through which a drug can treat a disease
- **Novel drugs in development, emerging diseases:**
  - Challenge: Need to learn and reason about never-before-seen phenomena
- **Datasets for a variety of therapeutics tasks:**
  - Challenge: Need datasets and benchmarks to accelerate ML model development, validation and transition into production and clinical implementation

# Outline

- ✓ Overview and introduction
- ✓ Part 1: Virtual drug screening and drug repurposing
- ✓ Part 2: Adverse drug effects, drug-drug interactions
- ✓ Part 3: Clinical trial site identification, patient recruitment
- ✓ Part 4: Molecule optimization, molecular graph generation, multimodal graph-to-graph translation
- ✓ Part 5: Molecular property prediction and transformers
- ✓ Demos, resources, wrap-up & future directions



<https://zitniklab.hms.harvard.edu/drugml>