Improving Small Molecule pK_a Prediction Using Transfer Learning with Graph Neural Networks

Fritz Mayr (0000-0002-6621-2108)¹⁺, Marcus Wieder* (0000-0003-2631-8415)¹⁺, Oliver Wieder (0000-0003-4967-7613)¹, Thierry Langer (0000-0002-5242-1240)¹

⁺contributed equally to this work; ¹Department of Pharmaceutical Sciences, Pharmaceutical Chemistry Division, University of Vienna, Althanstrasse 14, 1090 Vienna, Austria

***For correspondence:**
marcus.wieder@univie.ac.at (MW)

Abstract Enumerating protonation states and calculating micro-state pK_a values of small molecules is an important yet challenging task for lead optimization and molecular modeling. Commercial and non-commercial solutions have notable limitations such as restrictive and expensive licenses, high CPU/GPU hour requirements, or the need for expert knowledge to set up and use. We present a graph neural network model that is trained on 714,906 calculated micro-state pK_a predictions from molecules obtained from the ChEMBL database. The model is fine-tuned on a set of 5,994 experimental pK_a values significantly improving its performance on two challenging test sets. Combining the graph neural network model with Dimorphite-DL, an open-source program for enumerating ionization states, we have developed the open-source Python package `pkasolver`, which is able to generate and enumerate protonation states and calculate micro-state pK_a values with high accuracy.

Introduction

The acid dissociation constant (K_a), most often written as its negative logarithm (pK_a), plays a major role in the field of molecular modeling, as it influences the charge, tautomer configuration, and overall 3D structure of molecules with accessible protonation states in the physiological pH range. All these factors further shape the mobility, permeability, stability, and mode of action of substances in the body [1]. In case of insufficient or missing empirical data, the correct determination of pK_a values is thus essential to correctly predict the aforementioned molecular properties.

Authors and studies disagree on the exact percentage of drugs with ionizable groups, but a conservative estimate suggests that at least two-thirds of all drugs contain one or more ionization groups (in a pH range of 2 to 12) [2]. The importance of pK_a predictions for drug discovery has been widely recognized and has been the topic of multiple blind predictive challenges (most notable the Statistical Assessment of Modeling of Proteins and Ligands (SAMPL) series SAMPL6 [3], SAMPL7, [4] and ongoing SAMPL8 ¹ challenge.

Multiple methods have been developed to estimate pK_a values of small molecules, ranging from physical models based on quantum chemistry calculations [5, 6] and/or free energy calculations [7, 8] to empirical models based on linear free energy relationships using the Hammet-Taft equation or more data driven methods using quantitative structure-property relationship (QSPR) and machine learning (ML) approaches like deep neural network or random forest models [4, 9–11]. In general empirical methods require significantly less computational effort than their physics-based counterparts once they are parameterized but require a relatively large number of high-quality data points as training set [4].

In recent years, machine learning methods have been widely applied to predict different molecular properties including pK_a predictions. Many of these approaches learn pK_a values on fingerprint representations of molecules [10, 12]. The pK_a value of an acid and conjugate base pair is determined by the molecular structure and the molecular effects on the reaction center exerted by its neighborhood, including mesomeric, inductive, steric, and entropic effects [13]. Ideally, these effects should be included and encoded in a suitable fingerprint or set of descriptors. For many applications, extended-connectivity fingerprints (ECFPs) in combination with molecular features have proven to be a suitable and powerful tool to learn structure-property relationships [14, 15].

The emergence of graph neural networks (GNNs) has shifted some focus from descriptors and fingerprints designed by domain experts to these emerging deep learning methods. GNNs are a class of deep learning methods designed to perform inference on data described by graphs and provide straightforward ways to perform node-level, edge-level, and graph-level prediction tasks [16–18]. GNNs are capable of learning representations and features for a specific task in an automated way eliminating the need for excessive feature engineering ([19]). Another aspect of their attractiveness for molecular property prediction is the ease with which a molecule can be described as an undirected graph, transforming atoms to nodes and bonds to edges encoded both atom and bond properties. GNNs have proven to be useful and powerful tools in the machine learning molecular modeling toolbox [19, 20].

Pan et al. [21] have shown that GNNs can be successfully applied to pK_a predictions of chemical groups of a molecule, outperforming more traditional machine learning models relying on human-engineered descriptors and fingerprints, developing Mo1Gpka, a web server for predicting pK_a values. Mo1Gpka was trained on molecules extracted from the ChEMBL database [22] containing predicted pK_a values (predicted with ACD/Labs Physchem software²). Only the most acidic and most basic pK_a values were considered for the training of the GNN models.

The goal of this work was to extend the scope of predicting pK_a values for independently ionizable atoms (realized in Mo1Gpka) and develop a workflow that is able to enumerate protonation states and calculate the corresponding pK_a values connecting them (sometimes referred to as ‘sequential pK_a prediction’). To achieve this we implemented and trained a GNN model that is able to predict values for both acidic and basic groups by considering the protonated and deprotonated species involved in the corresponding acid-base reaction. We trained the model in two stages. First, we started by pre-training the model on calculated microstate pK_a values for a large set of molecules obtained from the ChEMBL database [22]. The pre-trained

¹<https://github.com/samplchallenges/SAMPL8>

²version 12.01, Advanced Chemistry Development Inc. 2010ACD/Labs

model already performs well on the two independent test sets used to measure the performance of the trained models. To improve its performance we fine-tuned the model on a small training set of molecules for which experimental pK_a values were available. The fine-tuned model shows excellent and improved performance on the two test sets.

We have implemented the training routine and prediction pipeline in an open-source Python package named `pkasolver`, which is freely available and can be obtained as described in the **Code and data availability** section. We also provide a ready-to-use Google Colab Jupyter notebook which includes the trained models and can be used to predict pK_a values for molecules without locally installing the package (for further information see the **Code and data availability** section) [23].

Results & Discussion

We will start by discussing the performance of the model on the validation set of the ChEMBL data set (which contains micro-state pK_a values calculated with `Epik` on a subset of the ChEMBL database) and the two independent test sets: the Novartis test set (280 molecules) and the Literature test set (123 molecules). This will be followed by a discussion of the fine-tuned model on its validation set, on both test sets, and on the ChEMBL data set. Finally, we will discuss the developed `pkasolver` package and its use cases.

Performance of the different predictive models is subsequently reported using the mean absolute error (MAE) and root mean squared error (RMSE). For each metric (MAE and RMSE) the median value from 50 repetitions with different training/validation set splits is reported and the 90% confidence interval is shown. To visualize training results a single training run was randomly selected and the results on the validation set plotted.

Pre-training model performance

The initial training of the GNN model was performed using the ChEMBL data set (microstate pK_a values calculated with `Epik`). Figure S.I.3 A shows the results of the best performing model on the hold-out validation set. The MAE and RMSE are 0.29 [90% CI: 0.28; 0.31] and 0.45 [90% CI: 0.44; 0.49] pK_a units showing an excellent accuracy across the reference pK_a values. The kernel density estimates (KDE) of the distribution of the reference and predicted pK_a values shown in Figure S.I.3 A highlights the ability of the GNN to correctly predict pK_a values throughout the investigated pH range.

The performance of the trained GNN model was assessed on two independent test sets: the Novartis and the Literature test set (both test sets are described in detail in the **Methods** section) [10]. The trained model performs well on both test sets with a MAE of 0.62 [90% CI: 0.57; 0.67] and a RMSE of 0.97 [90% CI: 0.89; 1.10] pK_a units on the Literature test set and a MAE of 0.82 [90% CI: 0.77; 0.85] and a RMSE of 1.13 [90% CI: 1.05; 1.21] pK_a units on the Novartis test set (shown in Figure S.I.2). The performance is comparable to the performance of `Epik` and `Marvin` on both test sets (shown in Table 1).

Fine-tuned model performance

While the performance on the test sets of the pre-trained model was already acceptable we were able to further increase model accuracy by fine-tuning the pre-trained model using a data set of experimentally measured pK_a values. The performance of the fine-tuned model on the validation set of the experimental data set is shown in S.I.3 B. The median performance of the fine-tuned model was improved from a RMSE of 0.97 [90% CI: 0.89; 1.10] to 0.82 [90% CI: 0.76; 0.88] pK_a units on the Literature test set and from a RMSE of 1.13 [90% CI: 1.05; 1.21] to 0.93 [90% CI: 0.85; 0.97] pK_a units on the Novartis test set (shown in Figure 1).

In order to avoid model performance degradation on the ChEMBL data set we randomly added molecules from the ChEMBL data set during the fine-tuning workflow. Adding molecules from the ChEMBL data set to restrict model parameters and avoid overfitting decreased the performance of the fine-tuned model on the ChEMBL data set only slightly (shown in Figure S.I.4). This was necessary since previous attempts without regularization showed decreased accuracy of the fine-tuned model in regions outside the limited pH range of the experimental data set while improving the performance on the test sets (details to the pH range of both the ChEMBL and experimental data set are shown in Figure S.I.5). An example of the performance of the fine-tuned model on the ChEMBL data set without regularization is shown in Figure S.I.6.

Model	Novartis data set		Literature data set	
	MAE	RMSE	MAE	RMSE
Random Forest ^{1,3}	1.15	1.51	0.53	0.76
ChemAxon Marvin (V20.1.0) ³	0.86	1.17	0.57	0.87
MolGpKa [21]	0.87 [0.77;0.97]	1.27 [1.08;1.45]	0.49 [0.40;0.65]	1.00 [0.56;1.53] ⁴
Epik ² [21]	0.83 [0.75;0.91]	1.16 [1.06;1.26]	0.58 [0.48;0.67]	0.92 [0.74;1.12]
pkasolver	0.71 [0.64;0.74]	0.93 [0.85;0.97]	0.52 [0.49;0.56]	0.82 [0.76;0.86]

Table 1. Performance of state-of-the-art knowledge-based approaches and commercial software solutions to predict pK_a values on the Novartis and Literature test sets are shown. For each data set, the mean absolute error (MAE) and root mean squared error (RMSE) is calculated. For MolGpKa, Epik, and pkasolver the median value and the 90% confidence interval are reported.

¹ used a random forest implementation with 1,000 estimators and the FCFP6 fingerprint. Values for the best performing random forest implementation are shown.

² Epik identified different protonation centers than were reported in the data sets for the Novartis data set for 26 out of 280 molecules. These molecules were excluded from the MAE and RMSE calculation for Epik.

³ values were obtained from Baltruschat and Czodrowski [10].

⁴ the reason for the large confidence interval is the incorrect prediction for a single molecule (Isomeric Smiles: CCNC) by MolGpKa with an error of 8.86 pK_a units.

To set the performance of the fine-tuned model in context we compare its performance with two recent publications investigating pK_a predictions using machine learning. In Table 1 the results are summarized for the methods presented in both Baltruschat and Czodrowski [10] and Pan et al. [21]. We extracted data from these publications where appropriate and recalculated values where needed. Pan et al. [21] split the reported results into basic and acidic groups making it necessary to recalculate the values reported there for MolGpKa and Epik, the values for Marvin were taken directly from reference [10] (reported values were calculated without confidence interval). The fine-tuned GNN model (called pkasolver in Table 1) performs on a par with the best performing methods reported there.

It is difficult to rationalize MAE/RMSE differences between different methods/models shown in Table 1) since training sets and methods are different. The small difference in performance between pkasolver and MolGpKa could be attributed to the transfer learning routine which added experimentally measured pK_a values. The random forest model was trained on significantly less data (only on the 5,994 pK_a values present in the experimental data set) than either pkasolver or MolGpKa yet performs comparably to both on the Literature data set while significantly worse on the Novartis data set. This might highlight the complexity of the Novartis data set, an observation previously made and investigated in Pan et al. [21].

Both Epik and Marvin perform well on both test data sets. It is surprising that pkasolver can slightly outperform Epik, even though its initial training was based on data calculated by Epik. We think this emphasizes the potential of transfer learning as used in this work and data-driven deep learning in general.

Sequential pK_a predictions with pkasolver

Combining the trained GNN model with Dimorphite-DL, a tool that identifies potential protonation sites, enabled us to perform sequential pK_a predictions. A detailed description of this approach is given in the Detailed methods section. We investigated multiple mono and polyprotic molecules for qualitative and quantitative (in reasonable bounds) agreement between prediction and experimental data. The results for the investigate systems were of excellent consistency. The list of molecules that we tested is included in the pkasolver repository.

We want to discuss the results generated by pkasolver on a challenging molecule: ethylenediaminetetraacetic acid, better known as EDTA. The protonation states and molecules are shown in Figure S.I.7. The correct protonation states are identified and the pK_a values for the 6 protonation states estimated are 1.49, 1.93, 2.28, 2.57, 6.05, and 9.44. Comparing this with the experimentally determined pK_a values of 0.0, 1.5, 2, 2.66, 6.16, and 10.24 the values are in good relative agreement [24]. Only the first and last pK_a estimates are noticeable too high and too low respectively. This highlights a limitation of pkasolver. The range of pK_a

Figure 1. The fine-tuned GNN model is able to predict the pK_a values of the Novartis and Literature test set with high accuracy. Panel **A** shows the performance of the fine-tuned model (initially trained with the ChEMBL data set and subsequently fine-tuned on the experimental data set) on the Literature test set. Panel **B** shows the performance of the same model on the Novartis test set. The solid red line in the scatter plot indicates the ideal behavior of the reference and predicted pK_a values, the dashed lines mark the ± 1 pK_a unit interval. Mean absolute error (MAE) and root mean squared error (RMSE) is shown, the values in bracket indicate the 90% confidence interval calculated from 50 repetitions with random training/validation splits. N indicates the number of investigated samples.

values present in the training data was between zero and 14, and predictions of pK_a values will stay in this range.

Conclusion

We have shown that GNNs can be used to predict micro-state pK_a values and achieve excellent performance on two external test sets. Training the GNN model in two stages with a pre-training phase using a large set of molecules with calculated pK_a values and a fine-tuning phase on a small set of molecules with experimentally measured pK_a values improves the performance of the GNN model significantly. This performance boost is especially noteworthy on the challenging Novartis test set (the RMSE was decreased from 1.18 [1.05;1.27] to 0.93 [0.85;0.97] pK_a units). A direct comparison with other software solutions and machine learning models on the two test sets shows that the fine-tuned GNN model performs consistently on a par with the best results of other commercial and non-commercial tools.

We have implemented `pkasolver` as an open-source and free-to-use Python package under a permissive licence (MIT licence). We provide an easy-to-use and accurate tool that enumerates protonation states and calculates micro-state pK_a values closing a gap in the computational chemistry toolbox. By providing the original training data, the fine-tuned model parameters, and training routines we encourage the community to build on this work.

Code and data availability

- Python package used in this work (release v0.2) and Colabs Jupyter notebook link: <https://github.com/mayrf/pkasolver>
- Data and notebooks to reproduce the plots/figures (release v0.1): <https://github.com/wiederm/pkasolver-data>

Author Contributions

Conceptualization: FM, OW, TL, and MW; Methodology: FM, OW, and MW; Software: FM and MW; Investigation: FM and MW; Writing-Original Draft: FM, OW, and MW; Writing-Review&Editing: FM, OW, TL, and MW; Funding Acquisition: MW, TL; Resources: TL; Supervision: MW, TL.

Acknowledgments

MW is grateful for discussions with David Bushiri, John Chodera, Josh Fass, Nils Krieger, Magdalena Wiercioch, Steffen Hirte, Thomas Seidel, and the Tautomer Consortium, specifically Paul Czodrowski, Brian Radak, Woody Sherman, David Mobley, Christopher Bayly, and Stefan Kast. MW, OW, FM, and TL are grateful for the help of Gerhard F. Ecker and his group members who performed the reference pK_a calculations with `Epik` for the ChEMBL data set.

Funding

MW acknowledges support from an FWF Erwin Schrödinger Postdoctoral Fellowship J 4245-N28. FM and TL gratefully acknowledge funding by the NeuroDeRisk project (<https://www.neuroderisk.eu>), which has received funding from the Innovative Medicines Initiative 2 Joint Undertaking (IMI2 JU, https://european-union.europa.eu/institutions-law-budget/institutions-and-bodies/institutions-and-bodies-profiles/imi-2-ju_en) under Grant Agreement No. 821528. This Joint Undertaking receives support from the European Union's Horizon 2020 research and innovation program and the European Federation of Pharmaceutical Industries and Associations (EFPIA, <https://www.efpia.eu>).

References

- [1] Manallack DT, Prankerd RJ, Yuriev E, Oprea TI, Chalmers DK. The significance of acid/base properties in drug discovery. *Chem Soc Rev.* 2013; 42(2):485–496. <http://xlink.rsc.org/?DOI=C2CS35348B>, doi: 10.1039/C2CS35348B.
- [2] Manallack DT. The p K a Distribution of Drugs: Application to Drug Discovery. *Perspectives in Medicinal Chemistry.* 2007 jan; 1:1177391X0700100. <http://www.ncbi.nlm.nih.gov/pubmed/19812734><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC2754920><http://journals.sagepub.com/doi/10.1177/1177391X0700100003>, doi: 10.1177/1177391X0700100003.

- [3] **Işık M**, Rustenburg AS, Rizzi A, Gunner MR, Mobley DL, Chodera JD. Overview of the SAMPL6 pK_a challenge: evaluating small molecule microscopic and macroscopic pK_a predictions, vol. 35. Springer International Publishing; 2021. <https://doi.org/10.1007/s10822-020-00362-6>, doi: 10.1007/s10822-020-00362-6.
- [4] **Bergazin TD**, Tielker N, Zhang Y, Mao J, Gunner MR, Francisco K, Ballatore C, Kast SM, Mobley DL. Evaluation of log P, pK_a, and log D predictions from the SAMPL7 blind challenge, vol. 35; 2021. doi: 10.1007/s10822-021-00397-3.
- [5] **Selwa E**, Kenney IM, Beckstein O, Iorga BI. SAMPL6: calculation of macroscopic pK_a values from ab initio quantum mechanical free energies. *Journal of Computer-Aided Molecular Design*. 2018; 32(10):1203–1216. <http://dx.doi.org/10.1007/s10822-018-0138-6>, doi: 10.1007/s10822-018-0138-6.
- [6] **Tielker N**, Eberlein L, Güssregen S, Kast SM. The SAMPL6 challenge on predicting aqueous pK_a values from EC-RISM theory. *Journal of Computer-Aided Molecular Design*. 2018; 32(10):1151–1163. <http://dx.doi.org/10.1007/s10822-018-0140-z>, doi: 10.1007/s10822-018-0140-z.
- [7] **Zeng Q**, Jones MR, Brooks BR. Absolute and relative pK_a predictions via a DFT approach applied to the SAMPL6 blind challenge. *Journal of Computer-Aided Molecular Design*. 2018; 32(10):1179–1189. <http://dx.doi.org/10.1007/s10822-018-0150-x>, doi: 10.1007/s10822-018-0150-x.
- [8] **Prasad S**, Huang J, Zeng Q, Brooks BR. An explicit-solvent hybrid QM and MM approach for predicting pK_a of small molecules in SAMPL6 challenge. *Journal of Computer-Aided Molecular Design*. 2018; 32(10):1191–1201. <http://dx.doi.org/10.1007/s10822-018-0167-1>, doi: 10.1007/s10822-018-0167-1.
- [9] **Liao C**, Nicklaus MC. Comparison of nine programs predicting pK_a values of pharmaceutical substances. *Journal of Chemical Information and Modeling*. 2009; 49(12):2801–2812. doi: 10.1021/ci900289x.
- [10] **Baltruschat M**, Czodrowski P. Machine learning meets pK_a [version 2; peer review: 2 approved]. . 2020; p. 9.
- [11] **Rupp M**, Korner R, V Tetko I. Predicting the pK_a of Small Molecules. *Combinatorial Chemistry High Throughput Screening*. 2011; 14(5):307–327. doi: 10.2174/138620711795508403.
- [12] **Yang Q**, Li Y, Yang JD, Liu Y, Zhang L, Luo S, Cheng JP. Holistic Prediction of the pK_a in Diverse Solvents Based on a Machine-Learning Approach. *Angewandte Chemie - International Edition*. 2020; 59(43):19282–19291. doi: 10.1002/anie.202008528.
- [13] **Perrin DD**, Dempsey B, Serjeant EP. pK_a Prediction for Organic Acids and Bases. Dordrecht: Springer Netherlands; 1981. <http://link.springer.com/10.1007/978-94-009-5883-8>, doi: 10.1007/978-94-009-5883-8.
- [14] **Rogers D**, Hahn M. Extended-connectivity fingerprints. *Journal of Chemical Information and Modeling*. 2010 may; 50(5):742–754. <https://pubs.acs.org/doi/10.1021/ci100050t>, doi: 10.1021/ci100050t.
- [15] **Jiang D**, Wu Z, Hsieh CY, Chen G, Liao B, Wang Z, Shen C, Cao D, Wu J, Hou T. Could graph neural networks learn better molecular representation for drug discovery? A comparison study of descriptor-based and graph-based models. *Journal of Cheminformatics*. 2021; 13(1):1–23. <https://doi.org/10.1186/s13321-020-00479-8>, doi: 10.1186/s13321-020-00479-8.
- [16] **Wieder O**, Kohlbacher S, Kuenemann M, Garon A, Ducrot P, Seidel T, Langer T. A compact review of molecular property prediction with graph neural networks. *Drug Discovery Today: Technologies*. 2020; <https://www.sciencedirect.com/science/article/pii/S1740674920300305>, doi: <https://doi.org/10.1016/j.ddtec.2020.11.009>.
- [17] **Wu Z**, Pan S, Chen F, Long G, Zhang C, Yu PS. A Comprehensive Survey on Graph Neural Networks. . 2019 jan; .
- [18] **Zhou J**, Cui G, Hu S, Zhang Z, Yang C, Liu Z, Wang L, Li C, Sun M. Graph neural networks: A review of methods and applications. *AI Open*. 2020; 1(September 2020):57–81. doi: 10.1016/j.aiopen.2021.01.001.
- [19] **Gilmer J**, Schoenholz SS, Riley PF, Vinyals O, Dahl GE. Neural Message Passing for Quantum Chemistry. . 2017; doi: 10.1002/nme.2457.
- [20] **Deng D**, Chen X, Zhang R, Lei Z, Wang X, Zhou F. XGraphBoost: Extracting Graph Neural Network-Based Features for a Better Prediction of Molecular Properties. *Journal of Chemical Information and Modeling*. 2021; 61(6):2697–2705. doi: 10.1021/acs.jcim.0c01489.
- [21] **Pan X**, Wang H, Li C, Zhang JZH, Ji C. MolGpka: A Web Server for Small Molecule pK_a Prediction Using a Graph-Convolutional Neural Network. *Journal of Chemical Information and Modeling*. 2021; 61(7):3159–3165. doi: 10.1021/acs.jcim.1c00075.

- [22] **Gaulton A**, Bellis LJ, Bento AP, Chambers J, Davies M, Hersey A, Light Y, McGlinchey S, Michalovich D, Al-Lazikani B, Overington JP. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Research*. 2012 Jan; 40(D1):D1100–D1107. <https://academic.oup.com/nar/article-lookup/doi/10.1093/nar/gkr777>, doi: 10.1093/nar/gkr777.
- [23] **Bisong E**. Google Colaboratory. In: *Building Machine Learning and Deep Learning Models on Google Cloud Platform* Berkeley, CA: Apress; 2019.p. 59–64. http://link.springer.com/10.1007/978-1-4842-4470-8_7, doi: 10.1007/978-1-4842-4470-8_7.
- [24] **Latscha HP**, Klein HA, Linti GW. Analytische Chemie: Chemie-Basiswissen III. Chemie-Basiswissen, Springer; 2004. <https://books.google.pt/books?id=xVJ0WtmKMHQC>.
- [25] **Davies M**, Nowotka M, Papadatos G, Dedman N, Gaulton A, Atkinson F, Bellis L, Overington JP. ChEMBL web services: Streamlining access to drug discovery data and utilities. *Nucleic Acids Research*. 2015; 43(W1):W612–W620. doi: 10.1093/nar/gkv352.
- [26] **Shelley JC**, Cholleti A, Frye LL, Greenwood JR, Timlin MR, Uchimaya M. Epik: A software program for pKa prediction and protonation state generation for drug-like molecules. *Journal of Computer-Aided Molecular Design*. 2007; 21(12):681–691. doi: 10.1007/s10822-007-9133-z.
- [27] **Greenwood JR**, Calkins D, Sullivan AP, Shelley JC. Towards the comprehensive, rapid, and accurate prediction of the favorable tautomeric states of drug-like molecules in aqueous solution. *Journal of Computer-Aided Molecular Design*. 2010; 24(6-7):591–604. doi: 10.1007/s10822-010-9349-1.
- [28] **Baltruschat M**, Czodrowski P. Machine learning meets pKa [version 2; peer review: 2 approved]. . 2020; p. 9.
- [29] **Liao C**, Nicklaus MC. Comparison of Nine Programs Predicting p K_a Values of Pharmaceutical Substances. *Journal of Chemical Information and Modeling*. 2009 Dec; 49(12):2801–2812. <https://pubs.acs.org/doi/10.1021/ci900289x>, doi: 10.1021/ci900289x.
- [30] RDKit, Open-Source Cheminformatics;. <http://www.rdkit.org>.
- [31] **McNaught AD**, Wilkinson A, of Pure IU, Chemistry A, of Chemistry (Great Britain) RS. IUPAC Compendium of Chemical Terminology. International Union of Pure and Applied Chemistry;. <https://books.google.at/books?id=I2LojwEACAAJ>.
- [32] **Xu K**, Jegelka S, Hu W, Leskovec J. How powerful are graph neural networks? 7th International Conference on Learning Representations, ICLR 2019. 2019; p. 1–17.
- [33] **Fey M**, Lenssen JE, Fast Graph Representation Learning with PyTorch Geometric; 2019. <http://arxiv.org/abs/1903.02428>, cite arxiv:1903.02428.
- [34] **Paszke A**, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Kopf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In: Wallach H, Larochelle H, Beygelzimer A, d'Alché-Buc F, Fox E, Garnett R, editors. *Advances in Neural Information Processing Systems* 32 Curran Associates, Inc.; 2019.p. 8024–8035. <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [35] **Ioffe S**, Szegedy C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. CoRR. 2015; abs/1502.03167. <http://arxiv.org/abs/1502.03167>.
- [36] **Loshchilov I**, Hutter F. Decoupled weight decay regularization. 7th International Conference on Learning Representations, ICLR 2019. 2019; .
- [37] **Ropp PJ**, Kaminsky JC, Yablonski S, Durrant JD. Dimorphite-DL: An open-source program for enumerating the ionization states of drug-like small molecules. *Journal of Cheminformatics*. 2019; 11(1):1–8. <https://doi.org/10.1186/s13321-019-0336-9>, doi: 10.1186/s13321-019-0336-9.

Detailed methods

Data set generation and pre-processing

Four different data sets were used in this work: the ChEMBL data set, the experimental data set, the Novartis data set and the Literature data set.

The ChEMBL data set used for pre-training was obtained from the ChEMBL database using the number of Rule-of-Five violations (set to a maximum of one violation) as filter criteria [22, 25]. For each of the molecules, a pK_a scan for the pH range between zero and 14 was performed using the Schrodinger tool *Epik* [26, 27] (Version-2121-1). The sequential pK_a scan indicated for 320,800 molecules one or multiple protonation state/s, resulting in a total of 729,375 pK_a values. For each pK_a value, *Epik* further indicated the protonation center using the atom index of the heavy atom at which either a hydrogen is attached or removed.

To perform transfer learning we obtained a second data set with experimental pK_a values. This data set (subsequently called 'experimental data set') was developed by Baltruschat and Czodrowski [28] and can be acquired from their GitHub repository³. For a detailed description of the curating steps taken to generate this data set, we point the reader to the **Methods** section of [28]. The experimental data set consists of 5,994 unique molecules, each with a single pK_a value and an atom index indicating the reaction center. Some of the molecules had to be corrected to obtain their protonation state at pH 7.4 (examples shown in S.I.1).

To test the performance of the models, two independent data sets were used, which were provided and curated by Baltruschat and Czodrowski [28]. The Literature data set contains 123 compounds collected by manual curating the literature. The Novartis data set contains 280 molecules provided by Novartis [29]. For each molecule, a pK_a value and atom index indicating the reaction center was provided. To avoid training the model on molecules present in the Literature or Novartis data set we filtered the ChEMBL data set using the InChIKey and canonical SMILES strings of the neutralized molecules as matching criteria. 50 molecules were identified and removed from the ChEMBL data set. All checks were performed using RDKit [30].

Enumerate protonation states during training/testing

The goal of calculating microstate pK_a values is to find the pH value at which the concentration of two molecular species is equal. To do this efficiently, we provide as input the protonated and deprotonated molecular species of the acid-base pair for which we want to calculate the pK_a value (the Brønsted acid/base definitions are used here and subsequently [31]). This approach enables a consistent treatment of acids and bases with a single data structure (the acid-base pair).

This workflow made it necessary that we generate the molecular species at each protonation state starting from the molecule at pH 7.4 by removing or adding hydrogen to the reaction center (which was calculated by Marvin for the experimental, Novartis, and Literature data set and *Epik* for the ChEMBL data set). We do this by sequentially adding hydrogen atoms from highest to lowest pK_a for acids (i.e. at pH=0 all possible protonation sites are protonated) and removing hydrogen atoms from lowest to highest pK_a value for bases on the structure present at pH 7.4 (at pH=14 all possible protonation sites are deprotonated).

This approach presented challenges for the ChEMBL data set for which sequential pK_a values and reaction centers were calculated with *Epik*. *Epik* calculates the micro-state pK_a value on the most probable tautomeric/mesomeric structure. This leads to potential protonation states that require changes in the double bond pattern and redistribution of hydrogen. Since we do not consider tautomeric changes to the molecular structure in the present implementation such tautomeric changes can introduce invalid molecules in either the sequential removal or addition of hydrogen atoms. Whenever such molecular structures were encountered we removed these protonation states from further consideration. Additionally, we used RDKit's sanitize function to identify cases for which protonation state changes introduce invalid atom valences. In other cases in which the protonation state change on a mesomeric structure introduces valid yet improbable molecular structures (e.g. protonating the oxygen in an amide instead of the nitrogen) we keep these structures. This reduced the number of molecules and protonation states in the ChEMBL dataset to 286,816 molecules and 714,906 protonation states. The distribution of pK_a values for the ChEMBL and experimental data set is shown in Figure S.I.5.

³<https://github.com/czodrowskilab/Machine-learning-meets-pKa>

Figure 2. Panel A shows the general workflow used to train the GNN on microstate pK_a values for a single molecule. During the training and testing phase, each molecule was provided in the structure dominant at pH 7.4 with atom indices indicating the protonation sites and corresponding pK_a values connecting them. In the **Enumeration of protonation states** phase we generate the protonation state for each pK_a value. The molecular species for each of the protonation states are then translated in their graph representation using nodes for atoms and edges for bonds, with node feature vectors encoding atom properties in the **Graph representation** phase. In the **pK_a prediction** phase graphs of two neighboring protonation states are combined and used as input for the GNN model to predict the pK_a value for the acid-base pair (using the Brønsted-Lowry acid/base definition [31]). The architecture of the GNN model is shown in detail in panel B. For a pair of neighboring protonation states two independent GIN (graph isomorphism network) convolution layers and ReLU activation functions are used for the protonated and the deprotonated molecular graph to pass information of neighboring atoms and achieve the embedding of the chemical environment of each atom [32]. The output of the convolutional layer is summarized using a global average pooling layer, generating the condensed input for the multilayer perceptron (MLP). To add regularization and to prevent co-adaptation of neurons a dropout layer was added.

Training and testing with PyTorch geometric

We use PyTorch and PyTorch geometric (subsequently abbreviated as PyG) for model training, testing, and prediction of pK_a values on the graph data structures [33, 34].

Graph data structure

A graph G is defined as a set of nodes V and edges E connecting the nodes. Each node $v \in V$ has a feature vector x_v , which encodes atom properties like element, charge, number of hydrogen, etc. as well as the presence of particular SMARTS patterns as a one-hot-encoding bit vector (all atom properties are shown in Table S.I.1). The adjacency matrix A defines the connectivity of the graph. A is defined as a quadratic matrix with $A_{uv} = 1$ if there is an edge between node u and v and $A_{uv} = 0$ if there is no edge between node u and v .

We used RDKit to generate a graph representation of the molecule with atoms represented as nodes and bonds as edges (in coordinate list format⁴ to efficiently represent the spares matrix).

Graph neural network architecture

To predict a single pK_a value the graph neural network (GNN) architecture takes as input two graphs representing the conjugated acid-base pair as shown in Figure 2. Figure 2 B shows the high-level architecture of the used GNN.

There are three phases to predict a pK_a value from a pair of molecular graphs. The first stage involves recurrently updating the node states using GIN (graph isomorphism network) convolution layers and ReLU activation functions [32]. We used 3 GIN layers with an embedding size of 64 bits each to propagate information throughout the graph and update each node with information about the extended environment. In the second stage, a global average pooling is performed to produce the embedding of the protonated and deprotonated graph, resulting in two 32 bit vectors. Concatenating the two 32 bit vectors produces the input for the third stage, the multilayer perceptron (MLP) with 3 fully connected layers (each with an embedding size of 64). To add regularization and to prevent co-adaptation of neurons a dropout layer randomly zeros out elements of the pooling output vector with $p = 0.5$ during training. Additionally, batch normalization is applied as described in [35].

GNN model training

Before each training run the ChEMBL and experimental data set were shuffled and randomly split in training (90% of the data) and validation set (10% of the data). To ensure that we can reproduce these splits the seed for each split was recorded.

The mean squared error (MSE) of predicted and reference pK_a values on the training data set was calculated and parameter optimization was performed using the Adam optimizer with decoupled weight decay regularization [36] as implemented in PyTorch.

Model performance was evaluated on the validation set and the model with the best performance was selected either for fine-tuning or further evaluation on the test data sets. The performance on the evaluation data set was calculated after every fifth epoch and the corresponding weights were saved. The learning rate for all training runs was dynamically reduced by a factor of 0.5 if the validation set performance did not change within 150 epochs (validation set performance threshold was set to 0.1).

Pre-training of the GNN was performed on the ChEMBL data set with a learning rate of 1×10^{-3} and a batch size of 512 molecules for 1,000 epochs. Fine-tuning was performed using the experimental data set with a learning rate of 1×10^{-3} and a batch size of 64 molecules for 1,000 epochs. All parameters of the GNN models were optimized during fine-tuning. To avoid overfitting to the experimental data set we added to each batch of the fine-tuning data set a randomly selected batch (1024 molecules) of the pre-training data set in the gradient calculation phase.

To calculate the confidence intervals of the model performance, pre-training and fine-tuning were repeated 50 times, each with a random training-validation set split resulting in 50 independently fine-tuned models.

⁴https://docs.scipy.org/doc/scipy/reference/generated/scipy.sparse.coo_matrix.html

282 Sequential pK_a value prediction with `pkasolver`

283 We use `Dimorphite-DL` to identify the proposed structure at pH 7.0 and all de-/protonation sites for a given
284 molecule [37].

285 Then we iteratively protonate each of the proposed de-/protonation sites generating a molecular pair
286 consisting of the protonated and deprotonated molecular species (in the first iteration the deprotonated
287 molecule is the molecule at pH 7.0). For each of the protonate/deprotonated pairs a pK_a value is calculated.
288 The protonated structure with the highest pK_a value (but below pH 7.0) is kept and the protonation site is
289 removed from the list of possible protonation sites. This is repeated until either (1) all protonation sites are
290 protonated, (2) no more valid molecules can be generated, or (3) the calculated pK_a values are outside the
291 allowed pK_a range.

292 To enumerate all deprotonated structures we start again with the structure at pH 7.0 and start to iter-
293 atively deprotonate each of the proposed de-/protonation sites. Here, we always keep the deprotonated
294 structure with the lowest pK_a value that is above 7.0.

295 The pK_a values are calculated using 25 of the 50 fine-tuned GNN models. For each protonation state, the
296 average pK_a value is calculated and the standard deviation is shown to enable the user to identify molecules
297 or protonation states for which the GNN model estimates are uncertain.

298 We provide a ready to use implementation of `pkasolver` to predict sequential pK_a values in our GitHub
299 repository (for further information see the **Code and data availability** section).

300 **Supplementary Information**

Atom features

atomic number:	1, 6, 7, 8, 9, 15, 16, 17, 33, 35, 53
formal charge	-1, 0, 1
hybridization:	1, 2, 3, 4
total number of Hs:	0, 1, 2, 3
explicit number of Hs:	0, 1, 2, 3
aromatic tag:	True, False
total valence:	1, 2, 3, 4, 5, 6
total degree:	1, 2, 3, 4
ring:	True, False
amide center atom:	True, False
SMARTS pattern:	one-hot-encoding of common acid-base pair substructures obtained from https://github.com/mcs07/MolVS

Table S.I.1. List of one-hot-encoding of atom features used for the node feature vector deposited in the node feature matrix X .

Figure S.I.1. Protonation state errors in the experimental data set. This exemplary selection shows molecules from the experimental data set provided by Baltruschat and Czodrowski [10] for which the protonation state provided does not correspond to the state at pH 7.4. For examples 1, 2, 4 and 5 with experimental pK_a values below 7.4 protonation at the reaction center would result in highly unlikely pentavalent nitrogen. For example 3 and 6 with pK_a values above 7.4 deprotonation at the reaction site can not be performed because of the lack of a suitable hydrogen. These error were corrected during our data preparation.

Figure S.I.2. Performance of the pre-trained GNN model on the Novartis and Literature test set is shown. 50 training runs with different training/validation splits were performed and for each training run the best model was selected based on its performance on the validation set (shown here is a single, randomly selected training run). Panel **A** shows the performance of the GNN model on the Literature data set. Panel **B** shows the performance of the GNN model on the Novartis data set. The solid red line in the scatter plot indicates the ideal behavior of the reference and predicted pK_a values, the dashed lines mark the ± 1 pK_a unit interval. Mean absolute error (MAE) and root mean squared error (RMSE) are shown, the values in bracket indicate the 90% confidence interval calculated from 50 repetitions with random training/validation splits. N indicates the number of investigated samples.

Figure S.1.3. Performance of the pre-trained and fine-tuned models are shown on the respective validation sets. 50 training runs with different training/validation splits were performed and for each training run the best model was selected based on its performance on the validation set (shown here is a single, randomly selected training run). Panel **A** shows the validation set performance of the best GNN model trained on the ChEMBL data set. Panel **B** shows the validation set performance starting from the same pre-trained model after fine-tuning on the experimental training set. The solid red line in the scatter plot indicates the ideal behavior of the reference and predicted pK_a values, the dashed lines mark the ± 1 pK_a unit interval. Mean absolute error (MAE) and root mean squared error (RMSE) are shown, the values in bracket indicate the 90% confidence interval calculated from 50 repetitions with random training/validation splits. N indicates the number of investigated samples.

Figure S.I.4. The accuracy of the fine-tuned GNN model only decreases slightly when molecules from the ChEMBL data set are used for regularization. 50 fine-tuning runs with different training/validation splits were performed, each initialized using the parameters of 50 pre-training runs, and for each training run the best model was selected based on its performance on the validation set. In order to generate a single plot we selected randomly a single fine-tuning run and generated the scatter plot with the best performing model on the validation set. The solid red line in the scatter plot indicates the ideal behavior of the reference and predicted pK_a values, the dashed lines mark the ± 1 pK_a unit interval. Mean absolute error (MAE) and root mean squared error (RMSE) are shown, the values in bracket indicate the 90% confidence interval calculated from 50 repetitions with random training/validation splits. N indicates the number of investigated samples.

Figure S.I.5. The pK_a distribution of ChEMBL and experimental data set are shown.

Figure S.I.6. The performance of the fine-tuned GNN model on the ChEMBL data set is shown. In contrast to the results obtained with the fine-tuned models shown in Figure S.I.4 the models shown here did **not** use regularization. The performance of the GNN model decreased significantly on the ChEMBL data, shifting pK_a values above 12 and below 2. The solid red line in the scatter plot indicates the ideal behavior of the reference and predicted pK_a values, the dashed lines mark the ± 1 pK_a unit interval. Mean absolute error (MAE) and root mean squared error (RMSE) are shown, the values in bracket indicate the 90% confidence interval calculated from 50 repetitions with random training/validation splits. N indicates the number of investigated samples.

Figure S.I.7. Results are shown for a sequential pK_a prediction using pkasolver for ethylenediaminetetraacetic acid (EDTA). For each protonation state the base-acid pair is shown and the consensus prediction for the pK_a value with the standard deviation is shown. The protonation site is highlighted for each protonation state.