doi:10.2533/chimia.2025.689 Chimia 79 (2025) 689–692

Educating Future Chemists in the Age of AI: A Digital Chemistry Course

Lauriane Jacot-Descombes[#a,b], Stefan P. Schmid[#a,b], and Kjell Jorner[*a,b]

Abstract: Artificial intelligence (AI) and machine learning (ML) are developing fast and are increasingly adopted in both chemical industry and academic research. With the projected role such tools will play in the future, for every chemist, these developments call for a fundamental and sound education for future generations of scientists in these areas. In this perspective, we describe the development of the course *Digital Chemistry* at ETH Zurich, which addresses these topics. In particular, we outline our approach to teaching ML and its applications in chemistry. We especially emphasize that the skills of understanding, applying and critically assessing ML models will be fundamental for future chemists. We hope that this article will serve as inspiration for educators in this field and help to enhance the teaching in this area of future significance.

Keywords: Critical thinking · Digital chemistry · Education · Machine learning

Lauriane Jacot-Descombes is a doctoral researcher at the Digital Chemistry Laboratory at ETH Zurich under the supervision of Prof. Dr. Kjell Jorner. After completing her MSc in chemistry, Lauriane joined the Digital Chemistry Laboratory, where her research focuses on the inclusion of chemical knowledge into data-driven methods for the prediction of chemical reactivity. In particular, she works on methods to generate transition state geometries at low computational costs and the inclusion of stereo-electronic descriptors calculated on transition states into machine learning models.

Stefan P. Schmid is a doctoral researcher at the Digital Chemistry Laboratory at ETH Zurich under the guidance of Prof. Dr. Kjell Jorner. After completing his MSc in chemistry, Stefan joined the Digital Chemistry Laboratory, where his research focuses on the optimisation of chemical reactions for the discovery and development of new catalytic methods. In particular, he works on Bayesian Optimisation algorithms that can be applied in chemical laboratories to efficiently optimise reactions and accelerate the development of novel catalysts.

Kjell Jorner is an Assistant Professor of Digital Chemistry at ETH Zurich. He completed his PhD in computational organic chemistry at Uppsala University, studying the effect of aromaticity on photochemical reactions. He then joined AstraZeneca UK as a Postdoctoral Fellow, building machine learning models for predicting the outcome of chemical reactions, followed by a postdoctoral fellowship at the University of Toronto, working on machine learning for molecular design with Alán Aspuru-Guzik. Since 2023, he has been leading a group at ETH, focusing on accelerating chemical discovery using digital tools.

1. Introduction

Artificial Intelligence (AI) and Machine Learning (ML) are arguably among the most headline-generating and fast-paced developments in this decade. With their advertised potential to revolutionise entire industries, these developments have also made their way into the chemical industry and academic chemical research. While the application of ML is being developed in the chemical industry, parallel efforts are ongoing in teaching future generations of chemists the knowledge and skills to apply ML to solve the challenges they will be facing in the future.[1–4] The skills that we regard as critical for a constructive application of ML in chemistry are both technical and cognitive.

As technical skills, we count, among others, the ability to know the fundamentals of coding to work with chemical data and being able to visualise it. Furthermore, the ability to understand, use, and potentially even train ML models is required, in such a way that the model output can be trusted, and its limitations understood. These technical skills should be complemented by the ability to use precise technical terminology, so that chemists can both correctly describe their work and have discussions with domain experts.

Acquiring the technical skills enables our students to write programs to work more efficiently with their data, a requirement in high demand both in industrial applications and academic research. By working with their data in a consistent and reproducible way, and understanding the associated challenges, they also become aware of important aspects of data handling, such as curating FAIR (findable, accessible, interoperable, and reusable) data in their own projects for better reproducibility and collaboration.

As for cognitive skills, we focus on skills that are shared across scientific disciplines and have been taught for a long time but are becoming even more important today. Since many ML models provide a prediction regardless of whether that prediction is ac-

*Correspondence: Prof. Dr. K. Jorner, E-mail: kjell.jorner@chem.ethz.ch
[#]Authors contributed equally. [a]Institute of Chemical and Bioengineering, Department of Chemistry and Applied Biosciences, ETH Zurich, CH-8093 Zurich, Switzerland; [b]NCCR Catalysis, Switzerland

tually reliable, critically assessing the output becomes paramount for responsible use. Such critical assessment requires knowledge about applicability domains (the chemical space where a model can be responsibly applied), model errors, and limitations, as well as explainability (why the model comes up with a certain prediction).

As the developments in AI are changing the landscape of chemistry, they are also transforming education.[1–4] Teaching students how to use AI can be supported by also employing AI in the course itself, using both predictive and generative models. Generative tools such as chatbots can be used to write code and explain scientific concepts, while predictive models can be used to obtain chemical properties. All these applications are encouraged throughout our course.

In this perspective, we outline how the course we created, *Digital Chemistry*, enables our students to obtain these technical and cognitive skills, and how they can do this with the help of AI tools. We focus on: (i) understanding how the ML methods work; (ii) using these algorithms in practice; and (iii) sensibly using and critically evaluating the output of the models. We describe how we teach these skills by modern pedagogical methods, *e.g.* interactive sessions. We hope that this article serves as inspiration for educators in this field.

2. Structure of the Course

The course that we developed at ETH Zurich is called *Digital Chemistry* and is taught by Prof. Dr. Kjell Jorner, together with his PhD researchers Lauriane Jacot-Descombes and Stefan Schmid. The course is listed as a master elective in four study programs: Biochemistry – Chemical Biology, Chemistry, Chemical and Bioengineering, and Computational Science and Engineering. It involves 6 ECTS credits and includes one lecture per week (2 × 45 min), one weekly exercise session (1 × 45 min) and a course project throughout the semester in groups of four students. The role of the lectures is to introduce the methods used in the field of digital chemistry and teach the theory behind them, while the exercises focus on how to write the code corresponding to the concepts taught in the lectures. The goal of the project is for the students to apply what they learn to a concrete research question.

3. Understanding the Methods

Given the diverse backgrounds of our students (the course is listed in four study programs, with students from bachelor to doctorate level), we decided that it was necessary to start with the fundamentals to get everyone on the same page. After laying the groundwork by giving an introduction to Python (the programming language of our choice), RDKit, and other important cheminformatics packages, we started putting this knowledge to work in applying ML to chemical problems. To give the students an introduction to predictive models, we start with linear regression. This method is already familiar to the students and can be used to teach them about important concepts such as hyperparameter optimisation, and how a model is trained in practice (*e.g.* gradient descent). Moreover, even such simple models allow us to introduce key concepts, such as applicability domain, extrapolation *vs.* interpolation, bias and variance, regularization, as well as the best practices of model training, selection, and evaluation. In the development of the course, we consciously decided to emphasise these concepts from the very beginning, as they build the framework for understanding and interacting responsibly with ML methods. Since linear models are still widely used in academic chemistry research, we include them both in the lectures and exercises so that the applications of the learned concepts are from real-world cases.

After introducing these important concepts to students, the complexity of the considered models is increased, swiftly introducing kernel methods and ultimately neural networks (NNs) as the foundation of modern ML developments. Throughout this stage of the course, we refer continuously to state-of-the-art research in the field that employs these methods, again underlining the utility of the learned concepts.[5] At the beginning of every topic, we pick recent publications, often from chemical companies, and when choosing these illustrative applications, we ensure that the different backgrounds of the students is accounted for by alternating the scientific fields (chemistry, biochemistry, material sciences, *etc*.) from which these publications originate.

The last parts of the course deal with teaching the students the necessary knowledge to understand the most recent developments in AI for chemistry, focusing on language and graph models, both predictive and generative. We continuously emphasise throughout these sections the importance of the concepts taught in the beginning by analysing the applicability domain, avoiding overfitting by model regularisation (*e.g.* dropout or early stopping) and statistical methods (*e.g.* cross-validation). In line with critically questioning the models, we also teach students interpretability and explainable AI (xAI) techniques (more detailed in Section 5) that can reveal which factors influence the model predictions. We also exploit the domain knowledge of our students, by inviting them to scrutinize the explanations given by models and argue why, potentially, the model came up with chemically counterintuitive outputs.

Since modern ML architectures are by no means easy to understand for someone with little experience, particular attention is paid to keeping students engaged throughout an entire lecture (2 × 45 min) and ensuring the students do not get lost along the way. To this end, we created interactive in-class exercises (*via* the Mentimeter platform[6]), this allows students to have discussions with each other and to challenge themselves to make sure that they understood the concepts. These interactive elements of the lecture, explicitly lauded multiple times in the student evaluations, are strategically placed around the middle point of every lecture hour to allow the students to catch a breath and recapitulate the most important concepts thus far. These in-class exercises include open-ended, multiple choice, and true/false questions. To gain an intuition on more complex models, the in-class exercises in the latter stages of the course focus on letting the students use visualisation tools that explain the discussed models and then discuss a question about the insights gained in the visualisation. Examples of the questions are shown in Fig. 1. Next to the in-class exercises, we also decided to convey messages *via* the use of deliberately humorous depictions of learnt concepts in the form of memes. Facilitated by the domain expertise of the teaching assistants, memes proved an effective and memorable technique for students.

Even though the lecture makes explicit efforts to engage the students, the biggest benefit in learning about ML comes from the students putting it into practice themselves, as described in the following section.

4. Using Machine Learning

Complementing the lectures, which teach the theory of ML in chemistry, the students are provided with weekly exercises. These exercises were created in Jupyter Notebooks on Google Colab. Jupyter Notebook is a format combining explanatory text, images, equations, interactive code execution, and data visualisation, allowing the students to proceed through the exercise step by step, following the instructions and being challenged to fill in missing code. Google Colab provides computing resources in a controlled cloud environment, so that the students can easily execute code in their browser, overcoming the usual challenges of installing a local coding environment in a computer room or on each student's own computer.

The exercise notebooks follow the topics taught in the lectures and together provide examples of complete workflows of typical tasks performed in the field of ML for chemistry (Fig. 2). Stu-

Fig. 1. Two examples of the interactive in-class exercises. In the left one, the students comment on the code generated by a large language model-based chatbot (CoPilot). In the right one, they optimise the shot length of a catapult by tuning different parameters. The latter exercise takes place just before teaching the systematic methods such as Design of Experiment and Bayesian Optimisation and gives the students a tacit understanding of the need for efficient optimisation methods for complex problems.

dents thus learn the code corresponding to the methods introduced to them in the lectures. Moreover, as the exercises are based on recent papers in the field, the students see concrete examples of usage of these methods in real research applications.

To further deepen the students' ability to use digital chemistry techniques themselves, our course includes a mandatory graded project which the students work on in groups. Each group investigates a different research question, and we encourage the students to come up with their own project idea. The project consists in writing a research proposal, defining the milestones of the work to be done, writing reproducible code aligned with recognised best practices, handing in a report where the results are analysed, and creating a poster which they present to the whole class in a poster session. In addition to providing them with an opportunity to apply the concepts taught in the course and to confront themselves with common challenges of projects in digital chemistry, these activities develop project management and teamwork skills useful to any scientist. As the students come from different study programs, the project is also an opportunity for them to experience how their different backgrounds can be leveraged in an interdisciplinary team.

The use of ML extends beyond the predictive algorithms for chemistry; it is also a tool used to help in the process of writing code and text, through large language model-based chatbots or integrated code-completion assistants. Far from prohibiting their use, we encourage the students to integrate these tools into their coding and research workflow and teach them their limitations so that they can evaluate the output of the models critically.

5. Thinking Critically

The rapid and impactful advances in the field of AI have been accompanied by a wave of enthusiasm which can obscure a clear understanding of the actual scope and limitations of the underlying models. It is therefore important to teach the students how to critically assess the results and application of ML, and we included several pedagogical components in the course to this end. First of all, the students get an opportunity to participate actively during the lecture and reflect on the concepts with the interactive in-class exercises (as described in Section 3), which they are encouraged to discuss together. Some of these in-class exercises, as well as questions during the oral examination, specifically address the critical assessment of the output of large language model-based chatbots (Fig. 1). The students are presented with the text given by one of these AI chatbots and are asked to evaluate this output. They can then see that the chatbots can sometime make subtle mistakes that for example break the best practices in model validation, leading to unreliable metrics of how well the trained model will perform when confronted with new data.

The reasoning skills of the students are further developed with questions at the end of each exercise series, where they need to discuss the results obtained in the coding part of the task.

Dedicated chapters during the lectures also focus specifically on how to evaluate and interpret the ML results. We teach the students to evaluate the predictions through statistical methods such as cross-validation or bootstrapping a hold-out test set, taking into account the uncertainty of the model's predictions. Explainability methods such as SHAPley values,[7] counterfactuals,[8] or attention[9] scores are introduced, allowing the students to interpret which features influence the model predictions the most and to gain insight into the model's decision-making process. The concept of applicability domain is also introduced to assess whether new samples are within the scope where the model can make reliable predictions or whether the predictions should be flagged as potentially unreliable.

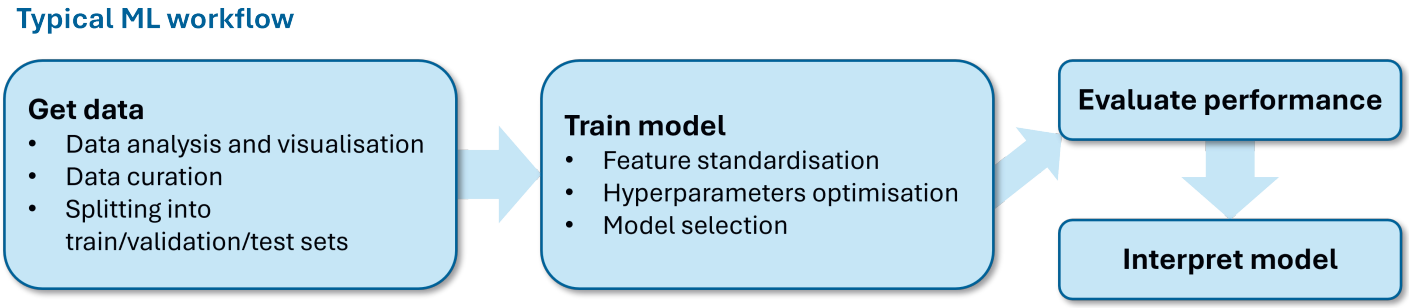

Fig. 2. Typical machine learning (ML) workflow in digital chemistry. The overall steps are to first curate the data, then to optimise and select the model, and lastly to evaluate its performance.

The students further sharpen their critical thinking skills in the poster presentation, where they explain and answer questions about their projects. Moreover, their ability to analyse their results reflectively carries a greater weight in the grading of the project than the model performance.

6. Conclusions

In this perspective, we have outlined the development of a Master level course in Digital Chemistry at ETH Zurich. We structured it around the technical and cognitive skills we are convinced will be required for future generations of chemists. We focus on understanding the model architectures and critically assessing the model outputs, the model limitations, and the applicability domain. Our lectures use interactive exercises to reemphasise important concepts and allow students to perform formative self-evaluation. Additionally, the practical implementation of the concepts is aided by pedagogical Jupyter notebooks that are discussed in guided exercise sessions.

These example cases also serve as a foundation for the students to apply their knowledge and skills in a practical project of their interest, experiencing how ML can be used to tackle their scientific problems, potentially on their own data. The projects are then presented in a poster session. Some of these projects have even led to peer-reviewed scientific publications.[10] By presenting their work, the students need to both critically assess their models and clearly communicate their workflow to their fellow students and teachers (domain experts), providing them experience for future collaborations in academia, industry, and beyond.

By formulating the necessary skills for future chemists and showcasing how the developed course addresses the teaching of these skills, we hope that this article serves as inspiration to other educators in this domain.

Acknowledgements

The authors thank all the students of the Digital Chemistry course at ETH Zurich in FS 2024 and FS 2025. K. J. thanks Jan Cvengros and Robert Grass for very helpful discussions throughout the first two years of the course. This publication was created as part of NCCR Catalysis (grant number 225147), a National Centre of Competence in Research funded by the Swiss National Science Foundation.

Author Contributions

L. J. D., S. P. S., and K. J. wrote and revised the article.

Received: September 2, 2025

[1] F. M. Fung, M. Lederbauer, Y. S. L. Choo, T. Gehring, K.M. Jablonka, K. Jorner, P. Schwaller, M. B. Sullivan, A. Volkamer, M.S. Sigman, L. Kuangbiao, C. Windle, *Chem* **2024**, *10*, 3519, https://doi.org/10.1016/j.chempr.2024.10.010.
[2] Y. Du, C. Duan, A. Bran, A. Sotnikova, Y. Qu, H. Kulik, A. Bosselut, J. Xu, P. Schwaller, *ChemRxiv* **2024**, https://doi.org/10.26434/chemrxiv-2024-h722v.
[3] M. Lederbauer *CHIMIA* **2025**, *79*, 174, https://doi.org/10.2533/chimia.2025.174.
[4] S. Berber, M. Brückner, N. Maurer, J. Huwer, *J. Chem. Educ.* **2025**, *102*, 1445, https://doi.org/10.1021/acs.jchemed.4c01033.
[5] P. Raghavan, A. J. Rago, P. Verma, M.M. Hassan, G.M. Goshu, A. W. Dombrowski, A. Pandey, C.W. Coley, *J. Am Chem. Soc.* **2024**, *146*, 15070, https://doi.org/10.1021/jacs.4c00098.
[6] Mentimeter, https://www.mentimeter.com/ (accessed 2025–08-27).
[7] L. S. Shapley, in 'Contributions to the Theory of Games II', H. Kuhn, A. W. Tucker, Eds., Princeton Univ. Press: Princeton, **1953**, 307, https://doi.org/10.1515/9781400881970–018.
[8] G. P. Wellawatte, A. Seshadri, A. D. White, *Chem. Sci*. **2022**, *13*, 3697, https://doi.org/10.1039/D1SC05259D.
[9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, in 'Advances in Neural Information Processing Systems', Eds. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett, Curran Associates, Inc., **2017**, *30*, 6000.
[10] Y. Zimmermann, L. Sieben, H. Seng, P. Pestlin, F. Görlich, *Npj Sci. Food* **2025**, *9*, 122, https://doi.org/F10.1038/s41538–025-00474-z.