VAEL: Bridging Variational Autoencoders and Probabilistic Logic Programming

Eleonora Misino[1,*], Giuseppe Marra[2] and Emanuele Sansone[2]

[1]University of Bologna, Bologna, Italy
[2]Katholike University Leuven, Leuven, Belgium

Keywords
Neuro-symbolic AI, Generative Model, Variational Autoencoder, Probabilistic Logic Programming

Neuro-symbolic learning has gained tremendous attention in the last few years [1, 2, 3, 4] as such integration has the potential of leading to a new era of intelligent solutions, enabling the integration of deep learning and reasoning strategies (e.g., logic-based or expert systems). While a lot of effort has been devoted to devising neuro-symbolic methods in the discriminative setting [5, 6, 7], less attention has been paid to the generative counterpart. An ideal generative neuro-symbolic framework should be able to encode the available small amount of training data into an expressive symbolic representation and to exploit complex forms of high level reasoning on such representation to generate new data samples. This is far from the actual state-of-the-art, where neuro-symbolic methods [8, 9, 10] have been mostly applied on generative tasks requiring only spatial-reasoning. As a motivation for this work, consider a task where a single image of multiple handwritten numbers is labeled with their sum. Suppose that we want to generate new images not only given their addition, but also given their multiplication, power, etc. Common generative approaches, like VAE-based models, have a strong connection between the latent representation and the label of the training task (i.e., the addition) [11, 12]. Consequently, when considering new generation tasks that go beyond the simple addition, they have to be retrained on new data.

In our work [13], we tackle the problem by providing a novel generative neuro-symbolic solution, named VAEL. Besides standard latent subsymbolic variables, our model exploits a probabilistic logic program to define a further structured representation, which is used for logical reasoning. The entire process is end-to-end differentiable. Once trained, VAEL can solve new unseen generation tasks by (i) leveraging the previously acquired knowledge encoded in the neural component and (ii) exploiting new logical programs on the structured latent space. Our experiments provide support on the benefits of this neuro-symbolic integration both in terms of task generalization and data efficiency. To the best of our knowledge, our work is the first to propose a general-purpose end-to-end framework integrating probabilistic logic programming into a deep generative model.

NeSy 2023, 17th International Workshop on Neural-Symbolic Learning and Reasoning, Certosa di Pontignano, Siena, Italy
*Corresponding author.
eleonora.misino2@unibo.it (E. Misino); giuseppe.marra@kuleuven.be (G. Marra); emanuele.sansone@kuleuven.be (E. Sansone)

Acknowledgements

GM is funded by the Research Foundation-Flanders (FWO-Vlaanderen, GA No 1239422N). This research is funded by TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215. ES is partially funded by the KU Leuven Research Fund (C14/18/062) and the Flemish Government (AI Research Program).

References

[1] T. R. Besold, A. d. Garcez, S. Bader, H. Bowman, P. Domingos, P. Hitzler, K.-U. Kühnberger, L. C. Lamb, D. Lowd, P. M. V. Lima, et al., Neural-symbolic learning and reasoning: A survey and interpretation, arXiv preprint arXiv:1711.03902 (2017).

[2] L. De Raedt, S. Dumančić, R. Manhaeve, G. Marra, From statistical relational to neuro-symbolic artificial intelligence, in: IJCAI, 2020.

[3] H. Kautz, The third ai summer, https://roc-hci.com/announcements/the-third-ai-summer/, 2020.

[4] Y. Bengio, G. Marcus, Ai debate, https://montrealartificialintelligence.com/aidebate/, 2020.

[5] R. Manhaeve, S. Dumancic, A. Kimmig, T. Demeester, L. D. Raedt, DeepProbLog: Neural Probabilistic Logic Programming, in: NeurIPS, 2018, pp. 3753–3763.

[6] K. Yi, J. Wu, C. Gan, A. Torralba, P. Kohli, J. B. Tenenbaum, Neural-symbolic vqa: Disentangling reasoning from vision and language understanding, in: NeurIPS, 2018.

[7] P. Minervini, S. Riedel, P. Stenetorp, E. Grefenstette, T. Rocktäschel, Learning reasoning strategies in end-to-end differentiable proving, in: ICML, 2020.

[8] J. Jiang, S. Ahn, Generative Neurosymbolic Machines, in: NeurIPS, 2020.

[9] R. Feinman, B. M. Lake, Generating New Concepts with Hybrid Neuro-Symbolic Models, in: CogSci, 2020.

[10] N. Gothoskar, M. Cusumano-Towner, B. Zinberg, M. Ghavamizadeh, F. Pollok, A. Garrett, J. B. Tenenbaum, D. Gutfreund, V. K. Mansinghka, 3DP3: 3D Scene Perception via Probabilistic Programming, in: NeurIPS, 2021.

[11] D. P. Kingma, S. Mohamed, D. J. Rezende, M. Welling, Semi-supervised Learning with Deep Generative Models, in: NeurIPS, 2014, pp. 3581–3589.

[12] T. Joy, S. M. Schmon, P. H. S. Torr, S. Narayanaswamy, T. Rainforth, Capturing Label Characteristics in VAEs, in: ICLR, 2021.

[13] E. Misino, G. Marra, E. Sansone, VAEL: Bridging variational autoencoders and probabilistic logic programming, in: A. H. Oh, A. Agarwal, D. Belgrave, K. Cho (Eds.), Advances in Neural Information Processing Systems, 2022. URL: https://openreview.net/forum?id=0xbP4W7rdJW.