Interpretable Tumor Growth Predictions

TRUST-AI in Healthcare

Johannes Koch, Tanja Alderliesten, and Peter A.N. Bosman

Introduction

Paraganglioma is a type of tumor in the head and neck region. These tumors are usually benign and are generally slow growing, making treatment potentially unnecessary. This is particularly relevant because treatment could have more risk than simply observing the tumor progression. However, paraganglioma can grow and cause severe (irreversible) complaints such as cranial nerve dysfunctions or hearing loss. In these cases, intervening timely is desirable to prevent possibly irreversible complaints. Moreover, the larger the tumor, typically, the larger the risk of the treatment. Hence, the current clinical approach is to "wait and scan", where the goal is to continuously observe and evaluate, together with the patient, how to best move forward and minimize the risk to the patient. In the initial follow-up years, the tumor development is closely monitored using medical imaging (e.g., MRI scans) and if a tumor is deemed low risk, the follow-up interval is increased. The future risk of growth is one of the deciding factors when it comes to determining if treatment needs to be considered or when to schedule the next follow-up moment.

Assessing the risk of tumor growth is however not trivial. Clinicians need to consider past tumor size measurements and other patient-specific factors together with their experience to make a statement about risk. To help clinicians, the TRUST-AI approach was used to develop interpretable AI models that predict tumor growth together with clinical experts. The potential benefits of accurate tumor growth predictions are two-fold, as on the one hand inevitable treatment can be done timely with decreased risk, and on the other hand potential complications and time and cost intensive follow-up can be avoided.

Having interpretable models is particularly relevant for applications in medicine, where the risk of wrong predictions can be high. Additionally, the EU AI act aims to ensure safe, transparent, traceable and non-discriminatory AI systems that are overseen by people, to prevent harmful outcomes.

Learning Interpretable Tumor Growth Predictions

To automatically learn interpretable growth prediction models, we have employed symbolic regression – a way to find small mathematical expressions that best describe the underlying data patterns, similar to physical laws but in this case for predicting the future tumor volume. Symbolic regression has two significant advantages compared to less transparent approaches such as deep learning: it does not demand massive datasets and, more importantly, yields small mathematical expressions that are easier for humans to read and reason about. In particular, to perform symbolic regression, we used modern model-based evolutionary computation techniques. Specifically, we employed GP-GOMEA (the Genetic Programming variant of the Gene-pool Optimal Mixing Evolutionary Algorithm), which is known to be capable of obtaining state-of-the-art results when used to perform symbolic regression [1-2].

To learn interpretable prediction models, we used retrospective patient data. First, medical images were segmented using modern image analysis techniques based on AI, so that tumor volume could be extracted from a large set of medical images automatically. This resulted in a dataset that, for each patient and each paraganglioma, contained a series of patient ages and associated tumor volumes [1].

To discover tumor-specific interpretable prediction models, a two-step symbolic regression approach was developed. The idea behind this was to find out whether there would be







generally speaking, different growth patterns in the data of all tumors of all patients. If so, it would be good to first extract these patterns, and then, for each tumor, use the most appropriate pattern to make tumor-specific predictions. To achieve this, multiple models containing free parameters were learned that describe a different pattern of growth, but by specifying values for the free parameters, can still be tuned to the data of a particular tumor of a particular patient. The first step requires a lot of computation, but this is common to all forms of learning in AI. Fitting growth patterns to individual tumor data, which constitutes the second step, is not computationally demanding. In this manner, it is possible to arrive at accurate, tumor-specific, prediction models that are fast to use and constitute small, interpretable, mathematical expressions [2].

A Use Case Specific User Interface

In practice, for a particular patient, the envisioned mode of operation includes re-fitting the different learned growth models to incorporate the latest acquired medical imaging data. The fitted model could then be used to make predictions about how the tumor will grow in the future. While this is key information, to really aid clinicians in making recommendations, these predictions need to be presented in proper context, which means having all relevant patient information available as well. This includes for instance the medical imaging scans that were used to derive tumor volume as well as patient characteristics. Importantly, from interactions with the clinical experts, it also became clear that beyond only having predictions of tumor growth as a singular value, it is important to have a notion of uncertainty associated with this prediction, so as to be able to gauge risk better. To this end, we developed an additional procedure around the fitting of the growth models to tumor measurements that incorporates uncertainties underlying the data and propagates these to the final predictions, enabling us to present predictions with associated uncertainty margins.



Figure 1: The predictions (solid lines) with associated uncertainty intervals (shaded areas) for two different tumors visualized in part of the developed user interface.

To support this, and to allow performing a first validation of the models in simulated decisionmaking scenarios, a custom user interface based on an open source medical image viewer has been developed. This user interface provides not only the tumor growth predictions as shown in Figure 1, but also relevant patient information. When it comes to making a decision, a clinical expert has to consider this additional information as well. For example, the tumor location and previous interventions may affect the available treatment options. Moreover, the patient's age, life expectancy, and their disposition towards treatment might already determine the decision without needing to consider future tumor growth. Compared to the current clinical approach, the developed tool for the first time provides all information needed to make treatment decisions in one interface.







In qualitative feedback sessions with clinical end-users, both the user interface and the prediction models showed promise for both supporting the making of decisions and patient-facing use, where the clinicians would use the tool to explain the tumor progression and discuss treatment options with patients.

Discussion and Conclusion

While we have already been successful in showcasing the potential of the overall approach, more patient data could be explicitly included in the development of the prediction models, such as biochemical screening and genetic data. While multiple models are currently already discovered using the new approaches to symbolic regression that we have developed in the TRUST-AI project, it is not yet studied which models are more representative of which groups of tumors and/or patients and which characteristics these tumors and patients have in common. Identifying this and subsequently incorporating more features into the prediction models directly, could make prediction models even more accurate for specific tumors and reduce the size of the uncertainty bands that are associated with tumor-growth predictions.

In this TRUST-AI use case, an auto-segmentation model, interpretable tumor growth predictions, and a customized user interface were developed and validated for paraganglioma, showing promising first results when evaluated together with clinical experts.

Acknowledgements

This research was funded by the European Commission within the HORIZON Programme (TRUST AI Project, Contract No.: 952060).

References

- [1] E. M. C. Sijben, J. C. Jansen, P. A. N. Bosman and T. Alderliesten, "Deep learningbased auto-segmentation of paraganglioma for growth monitoring," in Proc. SPIE 12929, Medical Imaging, 2024: Image Perception, Observer Performance, and Technology Assessment, 1292916; https://doi.org/10.1117/12.3006413.
- [2] E. M. C. Sijben, J. C. Jansen, P. A. N. Bosman and T. Alderliesten, "Function class learning with genetic programming: Towards explainable meta learning for tumor growth functionals," in Proceedings of the Genetic and Evolutionary Computation Conference – GECCO 2024, pages 1354-1362, 2024. https://doi.org/10.1145/3638529.3654145.
- [3] M. Virgolin, T. Alderliesten, C. Witteveen, P. A. N. Bosman; Improving Model-Based Genetic Programming for Symbolic Regression of Small Expressions. *Evol Comput* 2021; 29 (2): 211–237. doi: https://doi.org/10.1162/evco_a_00278.
- [4] W. La Cava, P. Orzechowski, B. Burlacu, F. de Franca, M. Virgolin, Y. Jin, M. Kommenda and J. Moore, "Contemporary Symbolic Regression Methods and their Relative Performance," in Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks, 2021.





