

How to aggregate deep ensembles?

Supervisors Pierre-Alexandre Mattei, Damien Garreau, Inria, Université Côte d’Azur.

Objective Ensemble methods have been increasingly popular in recent years, notably through deep ensembles [Lakshminarayanan et al., 2017, Fort et al., 2019]. Deep ensembles aggregate the outputs of several classifiers trained on the same dataset but with different initialisations. Empirically, these ensembles perform much better than their individual members. There are several ways of aggregating the individual predictions. The goal of this internship is to study the role of the aggregation method in the success of deep ensembles, and to see if more complex aggregation functions than simple average may be desirable. Other natural notions of centrality could be used, for instance the trimmed mean, the median, or the generalised mean (for a discussion, see *e.g.* Kuncheva, 2014, Section 5.3). An ambitious goal would be to prove theoretically in which situations which aggregation method is best, and whether or not it is worth learning this aggregation. While deep ensembles are typically used in a supervised framework, it would also be interesting to study ensembles of generative models (for instance VAEs, or diffusion models).

Skills required The internship will involve both theory (the candidate should have a strong background in statistics/machine learning) and implementation (prior knowledge of deep learning frameworks like Pytorch, JAX, or Tensorflow, is a plus). The balance between theory and practice will depend on the candidate’s skills and will.

Context of the internship This internship will be supervised by Pierre-Alexandre Mattei (Research scientist at Inria, Maasai team), and will involve collaborations with Damien Garreau (Associate prof. at Université Côte d’Azur, also a member of Maasai). We have already started working on this project, so the intern will benefit from our first advances.

The intern will join the Maasai team of Inria Sophia-Antipolis and Université Côte d’Azur, which is composed of 25 researchers in statistical and machine learning (web: <https://team.inria.fr/maasai/>). The team is part of the Institut 3IA Côte d’Azur <https://3ia.univ-cotedazur.eu/>, which offers a lot of opportunities (thesis offers, seminars & meetings with PhD students/postdoc in machine learning).

Duration: 6 months

Salary: approx. 550€ / month

PhD opportunities within the Maasai team may be pursued after the internship, to continue this work.

Contact To apply, please contact Pierre-Alexandre Mattei (pierre-alexandre.mattei@inria.fr).

References

Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep ensembles: A loss landscape perspective. *arXiv preprint arXiv:1912.02757*, 2019.

Ludmila I Kuncheva. *Combining pattern classifiers: methods and algorithms*. Wiley, 2014.

Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 2017.