Minimally rigid graphs with many realizations via reinforcement learning

Jan Legersky

It is known that a graph is minimally rigid in the plane if and only if it is minimally rigid on the sphere, hence, one can speak about minimally rigid graphs without specifying in the plane or on the sphere. Algorithms that compute the number of complex planar or spherical realizations of a given minimally rigid graph are known. We focus on the task to find a minimally rigid graph with a given number of vertices that has as many realizations as possible. Previously, this has been done by exhaustive search or by extensive computations based on expert knowledge. We exploit the fact that a graph is minimally rigid if and only if it can be constructed using a sequence of 0- and 1-extensions. Such iterative constructions are suitable for reinforcement learning: in each round of the training process, a generation of graphs is constructed using a neural network that predicts extensions to be applied, and those leading to graphs with many realizations are used to train the network for the next round.
We design a model based on graph neural networks that optimizes a function from the set of all minimally rigid graphs with a given number of vertices to the real numbers. Our experiments yield the following results: while in the case of planar realizations we find faster the same graphs up to 18 vertices as in the previous works, we provide graphs with 15 to 17 vertices that have more spherical realizations.

This is joint work with Rodrigo Alves, Jan Rubeš and Oleksandr Slyvka.