Munich Health Foundation Model Symposium



Contribution ID: 48

Type: not specified

Unbalancedness in Neural Monge Maps Improves Unpaired Domain Translation

In optimal transport (OT), a Monge map is known as a mapping that transports a source distribution to a target distribution in the most cost-efficient way. Recently, multiple neural estimators for Monge maps have been developed and applied in diverse unpaired domain translation tasks, e.g. in single-cell biology and computer vision. However, the classic OT framework enforces mass conservation, which makes it prone to outliers and limits its applicability in real-world scenarios. The latter can be particularly harmful in OT domain translation tasks, where the relative position of a sample within a distribution is explicitly taken into account. While unbalanced OT tackles this challenge in the discrete setting, its integration into neural Monge map estimators has received limited attention. We propose a theoretically grounded method to incorporate unbalancedness into any Monge map estimator. We improve existing estimators to model cell trajectories over time and to predict cellular responses to perturbations. Moreover, our approach seamlessly integrates with the OT flow matching (OT-FM) framework. While we show that OT-FM performs competitively in image translation, we further improve performance by incorporating unbalancedness (UOT-FM), which better preserves relevant features. We hence establish UOT-FM as a principled method for unpaired image translation.

Primary authors: Mr KLEIN, Dominik; EYRING, Luca Vincent; USCIDDA, Théo
Co-authors: THEIS, Fabian; PALLA, Giovanni; KILBERTUS, Niki; AKATA, Zeynep
Session Classification: Break + Posters session

Track Classification: Poster