
ROBERT MCCANN, University of Toronto
Optimal transport between unequal dimensions

In the last few decades, the theory of optimal transportation has blossomed into a powerful tool for exploring applications both within and outside mathematics. Its impact is felt in such far-flung areas as geometry, analysis, dynamics, partial differential equations, economics, machine learning, weather prediction, and computer vision.

The basic problem is to transport one probability density onto another, while minimizing a given cost $c(x,y)$ per unit transported. In the vast majority of applications, the probability densities live on spaces with the same (finite) dimension. After briefly surveying a few highlights from this theory, we focus our attention on what can be said when the densities instead live on spaces with two different (yet finite) dimensions. Although the answer can still be characterized as the solution to a fully nonlinear differential equation, it now becomes badly nonlocal in general. Remarkably however, one can identify conditions under which the equation becomes local, elliptic, and amenable to further analysis.