Histopathology is critical for the diagnosis of many diseases, including cancer. These protocols typically require pathologists to manually evaluate slides under a microscope, which is time-consuming and subjective, leading to interest in machine learning to automate analysis. However, computational techniques are limited by batch effects, where technical factors like differences in preparation protocol or scanners can alter the appearance of slides, causing models trained on one institution or patient to fail when generalizing to others. Here, we propose a domain adaptation method that improves the generalization of histopathological models to data from unseen institutions, without the need for labels or retraining in these new settings. Our approach introduces an optimal transport (OT) loss, that extends adversarial methods that penalize models if images from different institutions can be distinguished in their representation space. Unlike previous methods, which operate on single samples, our loss accounts for distributional differences between *batches* of images. We show that on the Camelyon17 dataset, while both methods can adapt to global differences in color distribution, only our OT loss can reliably classify a cancer phenotype unseen during training. Together, our results suggest that OT improves generalization on rare but critical phenotypes that may only make up a small fraction of the total tiles and variation in a slide.