

Dissection of the molecular properties that distinguish primary and metastatic cancers may reveal new underlying biological drivers of aggressive disease and inform clinical stratification. The rapid increase in the size of molecularly profiled patient cohorts creates an opportunity to develop machine learning algorithms that interrogate these data for discovery and clinical application. Recently, deep learning models have achieved the state of the art performance in different fields including computer vision, speech processing, natural language processing, and bioinformatics. However, the superior performance of deep learning models typically comes with the downside of reduced interpretability. Here we introduce P-NET, an artificial neural network with biologically informed, parsimonious architecture that accurately predicts metastasis in PrCa patients based on their genomic profiles. In P-NET, each node encodes some biological entity and each edge represents a known relationship between the corresponding entities. P-NET can be used to simultaneously rank features, genes, and biological pathways based on their importance to the clinical classification. We applied P-NET to whole-exome sequencing data from 1012 primary and metastatic prostate cancers and validated our model on two independent validation sets of 130 primary samples and 95 metastatic samples. We compared P-NET performance to other models and visualized the ranked features. We believe that using such interpretable models can provide a better understanding of the molecular differences between cancer stages and help us generate novel biological hypotheses.