

The problem of determining the probability that a particular patient has a certain attribute (e.g., has a certain disease, or has a certain prospect of recovery) is of fundamental importance to medical diagnosis, prediction and treatment decisions. This is a single-case probability, rather than the frequency of disease or the frequency of recovery in the population at large. Single-case probabilities are often calibrated to estimates of frequencies in reference classes to which the individual of interest belongs. The major problem here is to determine an appropriate single-case probability when an individual belongs to several reference classes for which data is available, and where estimates of frequencies differ from reference class to reference class. This is the infamous reference class problem. Personalized medicine aims towards accurate predictions about the health of individuals, rather than towards predictions about average health-effects in large populations. Predictions are based on a wealth of molecular data called –omics. It is therefore one of the major promises of personalized medicine that it can solve the reference class problem. In this talk, I argue that this is not the case. Major research programs in personalized medicine are equally subject to the reference class problem.