

If a high quality subgroup is present, it is very likely that several variants of this subgroup exist, which are also evaluated as good subgroups (e.g., by adding a constraint  $A = v$  for some attribute  $A$  and a rare feature  $v$ ). Furthermore, deviation analysis typically focusses on a target variable rather than associations between any attributes, which offers some optimization potential [6]. If it is possible to transform the database such that association mining (e.g., via the Apriori algorithm) can be applied, the validation and ranking of the patterns found are merely a post-processing step [15]. Thus, it may happen that 198 7 Finding Patterns many of the subgroups in the beam are variations of a single scheme, which prevents the beam search from focussing on other parts of the dataspace—a small number of diverse subgroups would be preferable. A subset of diverse subsets can be extracted from the beam by selecting successively those subgroups that cover most of the data. Exhaustive searching is prohibitive unless intelligent pruning techniques are applied that prevent us from losing too much time with redundant, uninteresting subgroups. For nominal target variables, efficient algorithms from association rule mining can be utilized. If most of the subgroups in the beam are variations of one core subgroup, only a few diverse subgroups will be selected. Thereby subgroups cannot be rediscovered, but the method has to .(focus on different parts of the dataspace. 8