	$sparse_1$	$moderate_1$	$dense_1$
job class	.08	.05	.03
department	.21	.15	.12
spatial	.366	.312	.272
`maxQ`	.50	.38	.33

Figure 2.13: Modularity values for partitions of $sparse_1$, $moderate_1$, and $dense_1$ graphs obtained via different methods. The modularity of the partition induced by "job class" falls well below the 0.3 threshold for a strong community structure, mentioned by Newman [80]. The last row shows modularity values for partitions obtained by using a greedy clustering algorithm, called `maxQ`, due to Clauset et al. [29]

graph's community structure) is the fraction of intra-community edges compared to the fraction of intra-community edges that the same node partition would have with a uniform random assignment of edges (see [81] for a precise definition). Modularity values upwards of 0.3 are said to indicate strong community structure [80]. Healthcare workers can be naturally partitioned by job class or by department. As shown in Figure 2.13, community structure induced in this way does not appear to be particularly strong. This is not surprising because healthcare workers in the same job class (e.g., nurses) are widely dispersed across multiple departments and departments are often composed of spatially dispersed units. One can do somewhat better by using spatial attributes. Specifically, for each healthcare worker u, define a *home location* $H(u)$ as the location of the computer in the hospital graph that u logs into most often. This maps each healthcare worker onto a vertex in the hospital graph and moreover establishes a metric space on the set of healthcare workers with the *distance* between healthcare workers u and v being the hop distance in the hospital graph between u