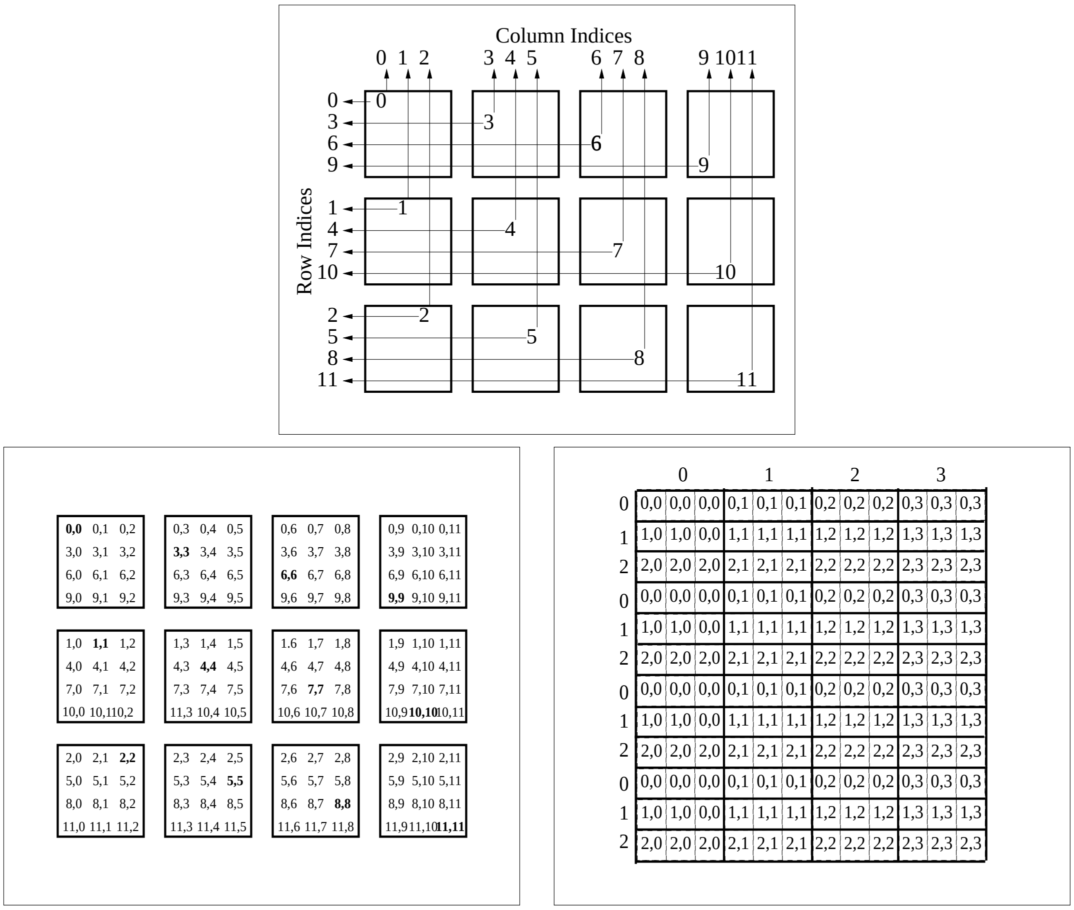

Figure 1.1
Inducing a matrix distribution from vector distributions. Top: Here each box represents a node of a 3×4 mesh. The sub-vectors of x and y are assigned to the mesh in column-major order. Thus the number in the box represents the index of the sub-vector assigned to that node. By projecting the indices of y to the left, we determine the distribution of the matrix row-blocks of A. By projecting the indices of x to the top, we determine the distribution of the matrix column-blocks of A. Bottom-left: The resulting distribution of the sub-blocks of A is given where the indices refer to the indices of the sub-blocks of A given in Equation (1.3.3). Bottom-right: The same information, except now from the matrix point of view. The figure shows the matrix partitioned into sub-blocks, with the indices in the sub-blocks indicating the node to which the sub-block is mapped.