Exercise 1. ***Moralisation: Converting DAGs to undirected minimal I-maps***

In the lecture, we had the following recipe to construct undirected minimal I-maps for $\mathcal{I}(p)$:

- Determine the Markov blanket for each variable x_i
- Construct a graph where the neighbours of x_i are given by its Markov blanket.

We can adapt the recipe to construct an undirected minimal I-map for the independencies $\mathcal{I}(G)$ encoded by a DAG G. What we need to do is to use G to read out the Markov blankets for the variables x_i rather than determining the Markov blankets from the distribution p.

Show that this procedure leads to the following recipe to convert DAGs to undirected minimal I-maps:

1. For all immoralities in the graph: add edges between *all* parents of the collider node.
2. Make all edges in the graph undirected.

The first step is sometimes called "moralisation" because we "marry" all the parents in the graph that are not already directly connected by an edge. The resulting undirected graph is called the moral graph of G, sometimes denoted by $\mathcal{M}(G)$.

Exercise 2. ***Moralisation exercise***

Consider the DAG G:

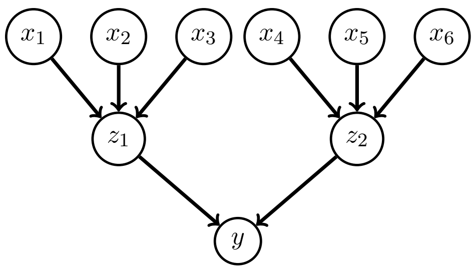

A friend claims that the undirected graph below is the moral graph $\mathcal{M}(G)$ of G. Is your friend correct? If not, state which edges needed to be removed or added, and explain, in terms of represented independencies, why the changes are necessary for the graph to become the moral graph of G.

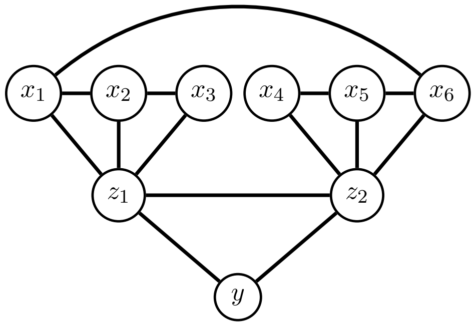

Exercise 3. ***Choice of elimination order in factor graphs***

We would like to compute the marginal $p(x_1)$ by variable elimination for a joint pmf represented by the following factor graph. All variables x_i can take K different values.

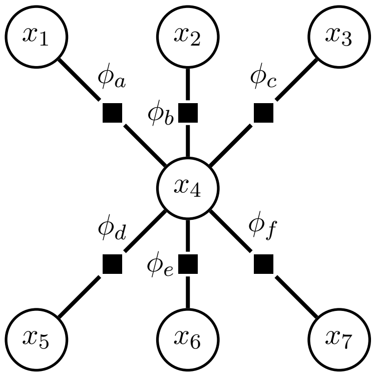

(a) A friend proposes the elimination order $x_4, x_5, x_6, x_7, x_3, x_2$, i.e. to do x_4 first and x_2 last. Explain why this is computationally inefficient.

(b) Propose an elimination ordering that achieves $O(K^2)$ computational cost per variable elimination and explain why it does so.