The main task in this assignment is implementing a simple (toy) Datalog evaluation engine using the naive and semi-naive strategies. For the main part of the assignment, you may assume that the input data will be of modest size, permitting a main-memory implementation. Portions marked with one ⋆ are optional (extra credit) for COS 480 but required for COS 580; those marked with two ⋆s are extra credit for both. Please refer to the earlier assignments for general instructions, including those for packaging and submission, and use the class newsgroup for clarifications and discussions.

Input and Output The program reads from standard input and writes to standard output. The input consists of a sequence of interleaved period-terminated *statements* and `?`-terminated *queries*. Statements are interpreted as Datalog rules, including rules with empty bodies (i.e., facts) and produce no output. Queries are evaluated on the current database instance (IDB and EDB) and produce as output a listing of the queried relation with one tuple per line, the attributes on each line separated by a comma-space sequence. The syntax is illustrated below. Identifiers and literals (used for IDB and EDB relations, variables, and domain elements) follow the common lexical conventions of the C language, with the added convention that identifiers beginning with an uppercase letter denote constants while those beginning with a lowercase letter denote variables. Whitespace is optional.

```
Edge(A, B).
path(x, y) :- Edge(x, y).
Edge(B,C).Edge(A,C). Edge(B, A).
path(x,y):-path(x,y),path(z,y).
path(x,y):-Edge(x,z),Edge(z,y).
path(A,_)?
Edge(A,D). path(A,_)?
path(_,_)?
```

Special statements: evaluation mode and tracing The evaluation mode is naive evaluation by default and is changed by special input statements `evalmode naive.` and `evalmode seminaive.`. The program's tracing mode is turned on and off by the special input statements `trace on.` and `trace off.`. These special statements for evaluation mode-selection and tracing may occur anywhere regular statements and queries are permitted. In tracing mode, the program should write to standard output the values of all non-EDB relations after each iteration of the outer loop in the usual definition of naive and seminaive evaluation. The listing of the tuples of each such relation should be prefixed by the relation's name on a single line, with implicitly defined relations (such as the delta relations in seminaive evaluation) suitably named.

⋆ **External memory** Provide an external memory implementation. Include suitable documentation and tests in your submission to validate your implementation.

⋆⋆ **Magic sets** Implement the magic sets evaluation strategy and expose its functionality through a `evalmode magic.` statement and corresponding tracing mode output.