# Assignment3 - Agents and Search

### Problem 3.1

For each of the following *agent architectures*, develop a *PEAS* description of the task *environment*.

Additionally, characterize the *environments* of these *agents* according to the properties discussed in the *lecture*. Where not "obvious", justify your choice with a short sentence.

Finally, choose suitable designs for the agents.

- 1. Robot soccer player
- 2. Internet book-shop agent (that is: an agent for book shops that stocks up on books depending on demand)
- 3. Autonomous Mars rover
- 4. Mathematical theorem prover
- 5. First-person shooter (Counterstrike, Unreal Tournament etc.)

#### Problem 3.2

Explain the difference between agent function and agent program. How many agent programs can there be for a given agent function?

## Problem 3.3

Explain the commonalities of and the differences between the performance measure and the utility function.

## Problem 3.4 (Tree Search in ProLog)

Implement the following tree search algorithms in Prolog:

- $1. \ BFS$
- 2. DFS
- 3. *IDS*Iterative Deepening (with step size 1)

Remarks:

• In the *lectures*, we talked about *expanding* nodes. That is relevant in many AI applications where the tree is not built yet (and maybe even too big to hold in memory), such as game trees in move-based games or decision trees of agents interacting with an environment. In those cases, when visiting a node, we have to expand it, i.e., compute what its children are.

In this problem, we work with smaller trees where the search algorithm receives the fully expanded tree as input. The algorithm must still visit every node and perform some operation on it — the search algorithm determines in which order the nodes are visited.

In our case, the operation will be to write out the label of the node.

• In the *lectures*, we worked with goal nodes, where the search stops when a goal node is found. Here we do something simpler: we *visit all the nodes and operate on each one* without using a goal state. (Having a goal state is then just the special case where the operation is to test the node and possibly stop.)

Concretely, your submission **must** be a single *Prolog* file that extends the following *implementation*:

```
% tree(V,TS) represents a tree.
% V must be a string - the label/value/data V of the root node
\% TS must be a list of trees - the children/subtrees of the root node
% In particular, a leaf node is a tree with the empty list of children
istree(tree(V,TS)) :- string(V), istreelist(TS).
% istreelist(TS) holds if TS is a list of trees.
% Note that the children are a list not a set, i.e., they are ordered.
istreelist([]).
istreelist([T|TS]) :- istree(T), istreelist(TS).
% The following predicates define search algorithms that take a tree T
% and visit every node each time writing out the line D:L where
% * D is the depth (starting at 0 for the root)
% * L is the label
% dfs(T) visits every node in depth-first order
dfs(T) :- ???
% bfs(T) visits every node in beadth-first order
bfs(T) :- ???
%itd(T):- visits every node in iterative deepening order
itd(T) :- ???
```

Here "must" means you can define any number of additional predicates. But the predicates specified above must exist and must have that arity and must work correctly on any input T that satisfies istree(T). "working correctly" means the predicates must write out exactly what is specified, e.g.,

0:A 1:B

for the depth-first search of the tree tree ("A", [tree("B", [])]).