Chapter 7

Supervision and Learning Attitudes

In philosophy, logic and psychology the word 'learning' is used to denote a variety of phenomena. In this chapter, as in the previous ones, we take the phrase 'to learn' to mean 'to acquire information in order to arrive at a certain (correct) conclusion'. In this we follow the lines of learning theory (see, e.g., Jain et al., 1999) where learning consists of a sequence of mind changes that should lead to a correct conclusion. All belief states visited on the way there, together with the correct one, are drawn from some given set of possibilities that constitutes the initial uncertainty range. To be more specific, let us try to describe in those terms the process of language learning. Agent L's (Learner's) aim is to 'arrive' at a grammar that correctly describes his native language. This scenario can be represented by a graph in which the vertices represent possible grammars and an edge from vertex s_1 to vertex s_2 labeled with α stands for a possibility of a mind change from grammar s_1 to s_2 that is triggered by the incoming information α. The properties of such a graph are determined by features of the initial range of possibilities, the language being learned, and the nature of agent's learning strategy. Then, the process of learning can be represented simply as a 'journey' of L through the graph until he finally reaches the correct grammar. Obviously, other inductive inference processes can also be described in this way.

The setting can be enriched by the presence of another agent, let us call her Teacher, T, who decides which data are presented to L, in which order, etc. In other words, we can introduce another player who supervises the process of learning and manipulates the data in order to influence the speed and accuracy of convergence. The analysis of the role of the teacher has been increasingly present in formal learning theory (e.g., see Angluin, 1987 for the *minimally adequate teacher* in learning from queries and counterexamples, and Balbach & Zeugmann, 2009 for recent developments in teachability theory).

In this chapter, we investigate the interaction between Learner and Teacher in a particular kind of *supervision* learning game that is played on a graph. Learner's information state changes while he moves around the graph, from one conjecture to another. Teacher, having a global perspective, knows the structure of the graph, and by providing certain information eliminates some initially possible mind changes of L. We are interested in the complexity of teaching, which we interpret in a similar way as in Chapter 6. Assuming the global perspective of Teacher, we identify the teachability problem with deciding whether the success of the learning process is possible. We interpret learning as a game and hence we identify learnability and teachability with the existence of winning strategies in a certain type of game. In this context, we analyze different Learner and Teacher attitudes, varying the level of Teacher's helpfulness and Learner's willingness to learn.

We interpret our learning game within the existing framework of *sabotage games* (Van Benthem, 2005). We start by recalling sabotage games and *sabotage modal logic*. Then, we explore *variations of the winning condition* of the game, providing sabotage modal logic formulae that characterize the existence of a winning strategy and their model-checking complexity (in other words: the complexity of deciding whether a player has a winning strategy in the game) in each case. Then we observe the asymmetry of the players' roles, and we allow Teacher to skip moves and we analyze how such removal of strict alternation of the moves affects our previous results. Finally, we recapitulate the work and discuss possible extensions.

With respect to the previous chapter, which also deals with computational complexity of teaching, we will now, in a way, take a step back. We will lose the detailed view on the content of epistemic states. Instead, we will gain the global picture of the process on conjecture change, and will be able to see the influence of intentional attitudes on the complexity of teaching.

7.1 Sabotage Games

As we already mentioned, our perspective on learning leads naturally to the framework of sabotage games. Sabotage games are useful for reasoning about various interactive processes involving random breakdowns or intentional obstruction in a system, from the failures of server networks to the logistics of traveling by public transport. We argue that it can also be interpreted positively, as some form of learning. But before we get to that, we first introduce the general, basic framework of sabotage games.

A sabotage game is played in a directed multi-graph, with two players, *Runner* and *Blocker*, moving in alternation, with Runner moving first. Runner moves by making a single transition from the *current* vertex. Blocker moves by deleting *any* edge from the graph. Runner wins the game if he is able to reach a designated goal vertex; otherwise Blocker wins.

To define the game formally let us first introduce the structure in which sabotage games take place, a directed multi-graph (see, e.g., Balakrishnan, 1997).

Definition 7.1.1. *A* directed multi-graph *is a pair $G = (V, E)$ where V is a set of vertices and $E : V \times V \to \mathbb{N}$ is a function indicating the number of edges between any two vertices.*

The sabotage game is defined in the following way.

Definition 7.1.2 (Löding & Rohde 2003a). *A* sabotage game

$$\mathrm{SG} = \langle V, E, v, v_g \rangle$$

is given by a directed multi-graph (V, E) and two vertices $v, v_g \in V$. Vertex v represents the initial position of Runner and v_g represents the goal state (the aim of Runner).

Each match consists of a sequence of positions and is played as follows:

1. *the initial position $\langle E_0, v_0 \rangle$ is given by $\langle E, v \rangle$;*
2. *round $k+1$ from position $\langle E_k, v_k \rangle$ consists of:*
   (a) Runner moving to some v_{k+1} such that $E(v_k, v_{k+1}) > 0$, and then
   (b) Blocker removing an edge (v, v') such that $E_k(v, v') > 0$.
   The new position is $\langle E_{k+1}, v_{k+1} \rangle$, where $E_{k+1}(v, v') := E_k(v, v') - 1$ and, for every $(u, u') \neq (v, v')$, $E_{k+1}(u, u') := E_k(u, u')$;
3. *the match ends if a player cannot make a move or if Learner reaches the goal state, which is the only case in which he wins.*

In other words, Blocker removes an edge between two states v, v' by decreasing the value of $E(v, v')$ by 1. As we will see later, this description of the game based on the above definition of multi-graphs can lead to some technical problems when we want to interpret modal logic over these structures. Therefore, we will now present an alternative definition, which we later show to be equivalent with respect to the existence of a winning strategy[1].

Definition 7.1.3. *Let $\Sigma = \{a_1, \ldots a_n\}$ be a finite set of labels. A* directed labeled multi-graph *is a tuple $G^\Sigma = (V, \mathcal{E})$ where V is a set of vertices and $\mathcal{E} = (\mathcal{E}_{a_1}, \ldots, \mathcal{E}_{a_n})$ is a collection of binary relations $\mathcal{E}_{a_i} \subseteq V \times V$ for each $a_i \in \Sigma$.*

In the above definition, the labels from Σ are used to represent multiple edges between two vertices; $\mathcal{E}$ is simply an ordered collection of binary relations on V with labels drawn from Σ. Accordingly, the modified definition of sabotage game is as follows.

[1] In what follows we take the size of any multi-graph $G = (V, E)$ to be bounded by: $|V| + \max\{E(v, w) \mid v, w \in V\} \cdot |V^2|$.

Definition 7.1.4. *A labeled sabotage game*

$$\mathtt{SG}^{\Sigma} = \langle V, \mathcal{E}, v, v_g \rangle$$

is given by a directed labeled multi-graph $(V, \mathcal{E})$ and two vertices $v, v_g \in V$. Vertex v represents the initial position of Runner and v_g represents the goal state.[2]

Each match is played as follows:

1. *the initial position $\langle \mathcal{E}^0, v_0 \rangle$ is given by $\langle \mathcal{E}, v \rangle$;*
2. *round $k+1$ from position $\langle \mathcal{E}^k, v_k \rangle$ with $\mathcal{E}^k = (\mathcal{E}^k_{a_1}, \ldots, \mathcal{E}^k_{a_n})$, consists of:*
   (a) *Runner moving to some v_{k+1} such that $(v_k, v_{k+1}) \in \mathcal{E}^k_{a_i}$ for some $a_i \in \Sigma$, and then*
   (b) *Blocker removing an edge (v, v') with label a_j ($(v, v') \in \mathcal{E}^k_{a_j}$) for some $a_j \in \Sigma$.*

   The new position is $\langle \mathcal{E}^{k+1}, v_{k+1} \rangle$, where $\mathcal{E}^{k+1}_{a_j} := \mathcal{E}^k_{a_j} - \{(v, v')\}$ and $\mathcal{E}^{k+1}_{a_i} := \mathcal{E}^k_{a_i}$ for all $i \neq j$;
3. *The match ends if a player cannot make a move or if Runner reaches the goal state, which is the only case in which he wins.*

It is easy to see that both versions of sabotage games have the *history-free determinacy property*: if one of the players has a winning strategy then (s)he has a winning strategy that depends only on the current position. Then, each round can be viewed as a transition from a sabotage game $\mathtt{SG}^{\Sigma} = \langle V, \mathcal{E}^k, v_k, v_g \rangle$ to another sabotage game $\mathtt{SG}'^{\Sigma} = \langle V, \mathcal{E}^{k+1}, v_{k+1}, v_g \rangle$, since all previous moves become irrelevant. We will use this fact through the whole paper.

It is easy to see that in *labeled* sabotage games, the label of the edge removed by Blocker is irrelevant with respect to the existence of a winning strategy. What matters is the number of edges that is left.

Observation 7.1.5. *Let $\mathtt{SG}^{\Sigma} = \langle V, \mathcal{E}, v_0, v_g \rangle$ and $\mathtt{SG}'^{\Sigma} = \langle V, \mathcal{E}', v_0, v_g \rangle$ be two labeled sabotage games that differ only in the labels of their edges, that is,*

$$\text{for all } (v, v') \in V \times V,\ |\{\mathcal{E}_{a_i} \mid (v, v') \in \mathcal{E}_{a_i}\}| = |\{\mathcal{E}'_{a_i} \mid (v, v') \in \mathcal{E}'_{a_i}\}|,$$

where $|\cdot|$ stands for cardinality. Then Runner has a winning strategy in $\mathtt{SG}^{\Sigma}$ iff he has a wining strategy in $\mathtt{SG}'^{\Sigma}$.

The existing results on sabotage have been given for the non-labeled version of the game. In what follows we show that the problems of deciding whether Runner has a winning strategy in sabotage games SG and $\mathtt{SG}^{\Sigma}$ are polynomially equivalent.

[2] We will sometimes talk about edges and vertices of $\mathtt{SG}^{\Sigma} = \langle V, \mathcal{E}, v, v_g \rangle$, meaning edges and vertices of its underlying directed (labeled) multi-graph $(V, \mathcal{E})$.

By doing this we establish that our modification of the definition makes only a slight difference and that the previous contribution is valid for our notion. We start by formalizing the two problems.

Definition 7.1.6 (SABOTAGE DECISION PROBLEM).

Instance *Sabotage game $\mathtt{SG} = \langle V, E, v_0, v_g \rangle$.*

Question *Does Runner have a winning strategy in SG?*

The Σ-SABOTAGE DECISION PROBLEM is very similar. The only difference is that it is concerned with slightly modified structures — labeled sabotage games.

Definition 7.1.7 (Σ-SABOTAGE DECISION PROBLEM).

Instance *Labeled sabotage game $\mathtt{SG}^{\Sigma} = \langle V, \mathcal{E}, v_0, v_g \rangle$.*

Question *Does Runner have a winning strategy in $\mathtt{SG}^{\Sigma}$?*

Theorem 7.1.8. SABOTAGE *and* Σ-SABOTAGE *are polynomially equivalent.*

Proof. The two problems can be polynomially reduced to each other.

(⇒) SABOTAGE can be reduced to Σ-SABOTAGE. Given a sabotage game $\mathtt{SG} = \langle V, E, v_0, v_g \rangle$, let m be the maximal number of edges between any two vertices in the graph, i.e.:

$$m := \max\{E(u, u') \mid (u, u') \in (V \times V)\}.$$

Then, we define the labeled sabotage game $f(\mathtt{SG}) := \langle V, \mathcal{E}, v_0, v_g \rangle$, where $\mathcal{E} := (\mathcal{E}_1, \ldots, \mathcal{E}_m)$ and each $\mathcal{E}_i$ is given by $\mathcal{E}_i := \{(u, u') \in V \times V \mid E(u, u') \geq i\}$.

We have to show that Runner has a winning strategy in SG iff he has one in $f(\mathtt{SG})$. The proof is by induction on n — the number of edges in SG, i.e., $n = \sum_{(v,v') \in V \times V} E(v, v')$. Note that by definition of f, $f(\mathtt{SG})$ has the same number of edges, i.e., $n = \sum_{i=1}^{i=m} |\mathcal{E}_i|$.

The base case

Straightforward. In both games Runner has a winning strategy iff $v_0 = v_g$.

The inductive case

(⇒) Suppose that Runner has a winning strategy in the game $\mathtt{SG} = \langle V, E, v_0, v_g \rangle$ with $n+1$ edges. Then, there is some $v_1 \in V$ such that $E(v_0, v_1) > 0$ and Runner has a winning strategy for all games $\mathtt{SG}' = \langle V, E', v_1, v_g \rangle$ that result from Blocker removing any edge (u, u') with $E(u, u') > 0$. Note that all such games SG′ have just n edges, so by the induction hypothesis Runner has a winning strategy in $f(\mathtt{SG}')$. But then, by Observation 7.1.5, Runner also has a winning strategy in all games $f(\mathtt{SG})'$ that result from removing an arbitrary edge from $f(\mathtt{SG})$. This is so because for any removed edge (u, u'), the only possible difference between $f(\mathtt{SG}')$ and $f(\mathtt{SG})'$ is in the labels of the edges between u and u' (in $f(\mathtt{SG}')$ the removed label was the largest, in $f(\mathtt{SG})'$ the removed label is arbitrary). Now, by definition of f, choosing v_1 is also a legal move for Runner in $f(\texttt{SG})$ and, since he can win every $f(\texttt{SG})'$, he has a winning strategy in $f(\texttt{SG})$.

($\Leftarrow$) Runner having a winning strategy in $f(\texttt{SG})$ means that he can choose some v_1 with $(v_0, v_1) \in \mathcal{E}_i$ for some $i \leq m$ such that he has a winning strategy in all games $f(\texttt{SG})'$ resulting from Blocker's move. Choosing v_1 is also a legal move of Runner in SG. Suppose that Blocker replies by choosing (v, v'). Let us call the resulting game SG$'$. By assumption and Observation 7.1.5, Runner also has a winning strategy in the game $f(\texttt{SG}')$ which is the result from Blocker choosing $((v, v'), E(v, v'))$. Since $f(\texttt{SG})' = f(\texttt{SG}')$, we can apply the inductive hypothesis.

($\Leftarrow$) Σ-SABOTAGE can be reduced to SABOTAGE. Given a labeled sabotage game $\texttt{SG}^\Sigma = \langle V, \mathcal{E}, v, v_g \rangle$ with $\Sigma = \{a_1, \dots a_m\}$, define the sabotage game $f'(\texttt{SG}^\Sigma) := \langle V, E, v, v_g \rangle$, where $E(v, v') := |\{\mathcal{E}_{a_i} \mid (v, v') \in \mathcal{E}_{a_i}\}|$.

Showing that Runner has a winning strategy in $\texttt{SG}^\Sigma$ iff he has one in $f(\texttt{SG}^\Sigma)$ is straightforward, and can be done by induction on the number of edges in $\texttt{SG}^\Sigma$, i.e., on $n := \sum_{a \in \Sigma} |\mathcal{E}_a|$.

Finally, let us observe that both f and f' that encode the procedures of transforming one type of graph to another, are polynomial, so the proof is complete. □

7.2 Sabotage Modal Logic

Sabotage modal logic (SML) has been introduced by Van Benthem (2005) to investigate the complexity of reachability-type problems in dynamic structures, such as the graph of our sabotage games. Besides the standard modalities, it also contains 'transition-deleting' modalities for reasoning about model change that occurs when a transition (an edge) is removed. To be more precise, we have formulae of the form $\blacklozenge\varphi$, expressing that it is possible to delete a pair from the accessibility relation such that φ holds in the resulting model at the current state.

Definition 7.2.1 (SML Language; Syntax). *Let* PROP *be a countable set of propositional letters and let Σ be a finite set of labels. Formulae of the language of sabotage modal logic are given by:*

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \Diamond_a\varphi \mid \blacklozenge_a\varphi$$

with $p \in$ PROP and $a \in \Sigma$. The formula $\boxminus_a\varphi$ is defined as $\neg\blacklozenge_a\neg\varphi$, and we will write $\Diamond\varphi$ for $\bigvee_{a\in\Sigma} \Diamond_a\varphi$ and $\blacklozenge\varphi$ for $\bigvee_{a\in\Sigma} \blacklozenge_a\varphi$.

The sabotage modal language is interpreted over Kripke models, that are here called *sabotage models*.

Definition 7.2.2 (Löding & Rohde 2003b). *Given a countable set of propositional letters* PROP *and a finite set $\Sigma = \{a_1, \dots, a_n\}$, a* sabotage model *is a tuple $M = \langle W, (R_{a_i})_{a_i\in\Sigma}, Val\rangle$ where W is a non-empty set of worlds, each $R_{a_i} \subseteq W \times W$ is an accessibility relation and Val : PROP $\to \mathcal{P}(W)$ is a propositional valuation function. We will call a pair (M, w) with $w \in W$ a* pointed sabotage model.

To get to the semantics of sabotage modal language, we first have to define the model that results from removing an edge.

Definition 7.2.3. *Let $M = \langle W, R_{a_1}, \dots R_{a_n}, Val\rangle$ be a sabotage model. The model $M^{a_i}_{(v,v')}$ that results from removing the edge $(v, v') \in R_{a_i}$ is defined as*

$$M^{a_i}_{(v,v')} := \langle W, R_{a_1}, \dots R_{a_{i-1}}, R_{a_i} \setminus \{(v, v')\}, R_{a_{i+1}}, \dots R_{a_n}, Val\rangle.$$

Definition 7.2.4 (SML; Semantics). *Given a sabotage model*

$$M = \langle W, (R_a)_{a\in\Sigma}, Val\rangle$$

and a world $w \in W$, atomic propositions, negations, disjunctions and standard modal formulae are interpreted as usual. For the case of 'transition-deleting' formulae, we have

$$(M, w) \models \blacklozenge_a\varphi \text{ iff } \exists v, v' \in W \ ((v, v') \in R_a \ \& \ (M^a_{(v,v')}, w) \models \varphi).$$

One SML result that is of great importance to us is the SML model checking complexity (combined complexity model checking, see Vardi, 1982). We will use it to reason about the difficulty of our learning scenarios.

Theorem 7.2.5 (Löding & Rohde 2003b). *The computational complexity of model checking for SML is PSPACE-complete.*

7.3 Sabotage Learning Games

In this section, we reinterpret the sabotage game in the broader perspective of learning. We introduce variants of the winning condition of the game. For each variant, we will provide a sabotage modal logic formula characterizing the existence of a winning strategy. We also prove complexity results for model checking in each case. We will work with previously introduced labeled sabotage games, using the labeling of the edges to represent different kinds of information changes that take Learner from one state into another.

7.3.1 Three Variations

A sabotage learning game is defined as follows.

Definition 7.3.1. *A* sabotage learning game *(SLG) is a labeled sabotage game between* Learner *(L, taking the role of Runner) and* Teacher *(T, taking the role of Blocker). We distinguish three different versions,* SLGhe, SLGhu *and* SLGue. *Moves allowed for both players are those of the sabotage game. There is also no difference in the arena in which the game is played. However, the winning conditions vary from version to version (Table 7.1).*

Game	Winning Condition
SLGue	L wins iff L reaches the goal state, T wins otherwise.
SLGhu	T wins iff L reaches the goal state, L wins otherwise.
SLGhe	L and T win iff L reaches the goal state. Both lose otherwise.

Table 7.1: Sabotage Learning Games

The different winning conditions correspond to different levels of Teacher's helpfulness and Learner's willingness to learn. We can then have the cooperative case with H*elpful* Teacher and E*ager* Learner (SLGhe). But there are two other possibilities that we will be interested in: U*nhelpful* Teacher with E*ager* Learner (SLGue), and H*elpful* Teacher with U*nwilling* Learner (SLGhu).

Having defined the games representing various types of Teacher and Learner attitudes, we now show how sabotage modal logic can be used for reasoning about players' strategic powers in these games.

7.3.2 Sabotage Learning Games in Sabotage Modal Logic

Sabotage modal logic turns out to be useful for reasoning about graph-like structures where edges can disappear; in particular, it is useful for reasoning about sabotage learning games. In order to interpret the logic on our graphs we need to transform the arena of a labeled sabotage game into a sabotage model in which formulae of the logic can be interpreted. In fact, for each SLG we can construct a pointed sabotage model in the following straightforward way.

Definition 7.3.2. *Let* $\mathtt{SG}^\Sigma = \langle V, \mathcal{E}, v_0, v_g \rangle$ *be a sabotage game and* $\mathcal{E} = (\mathcal{E}_a)_{a \in \Sigma}$. *The* pointed sabotage model $(M(\mathtt{SG}^\Sigma), v_0)$ *over the set of atomic propositions* $\textsc{Prop} := \{goal\}$ *is given by*

$$M(\mathtt{SG}^\Sigma) := \langle V, \mathcal{E}, \mathit{Val} \rangle,$$

where $\mathit{Val}(goal) := \{v_g\}$.

In the light of this construction, sabotage modal logic becomes useful for reasoning about players' strategic powers in sabotage learning games. Each winning condition in Table 7.1 can be expressed by a formula of SML that characterizes the existence of a winning strategy, that is, the formula is true in a given pointed sabotage model if and only if the corresponding player has a winning strategy in the game represented by the model.

Unhelpful Teacher and Eager Learner (SLGue) Let us first consider SLGue, the original sabotage game (Van Benthem, 2005). For any $n \in \mathbb{N}$, we define the formula $\gamma_n^{\mathtt{ue}}$ inductively as follows:

$$\gamma_0^{\mathtt{ue}} := goal, \qquad \gamma_{n+1}^{\mathtt{ue}} := goal \vee \Diamond \boxminus \gamma_n^{\mathtt{ue}}.$$

The following result is Theorem 7 of Löding & Rohde (2003b) rephrased for labeled sabotage games. We provide a detailed proof to show how our *labeled* definition avoids a technical issue present in the original proof.

Theorem 7.3.3. *Learner has a winning strategy in the* SLGue

$$\mathtt{SG}^\Sigma = \langle V, \mathcal{E}^0, v_0, v_g \rangle$$

if and only if $(M(\mathtt{SG}^\Sigma), v_0) \models \gamma_n^{\mathtt{ue}}$, *where* n *is the number of edges of* $\mathtt{SG}^\Sigma$.

Proof. The proof is by induction on n.

The base case

($\Rightarrow$) If L has a winning strategy in a game $\mathtt{SG}^\Sigma$ with no edges, then he should be already in the winning state, that is $v_0 = v_g$. Thus, $(M(\mathtt{SG}^\Sigma), v_0) \models goal$ and hence, $(M(\mathtt{SG}^\Sigma), v_0) \models \gamma_0^{\mathtt{ue}}$.

($\Leftarrow$) If $(M(\mathtt{SG}^\Sigma), v_0) \models \gamma_0^{\mathtt{ue}}$ then $(M(\mathtt{SG}^\Sigma), v_0) \models goal$. Since v_g is the only state where *goal* holds, then we have $v_0 = v_g$, and therefore L wins $\mathtt{SG}^\Sigma$ immediately.

The inductive case

($\Rightarrow$) Suppose that $\mathtt{SG}^\Sigma$ has $n+1$ edges, and assume L has a winning strategy. There are two possibilities: L's current state is the goal state (that is, $v_0 = v_g$), or it is not.

In the first case, we get $(M(\mathtt{SG}^\Sigma), v_0) \models goal$ and hence $(M(\mathtt{SG}^\Sigma), v_0) \models \gamma_{n+1}^{\mathtt{ue}}$. In the second case, since L has a winning strategy in $\mathtt{SG}^\Sigma$, there is some state $v_1 \in V$ reachable from v_0, i.e., for some $a_i \in \Sigma$ (that is, $(v_0, v_1) \in \mathcal{E}^0_{a_i}$) such that in all games $\mathtt{SG}^\Sigma_{(u,u'),a_j} = \langle V, \mathcal{E}^1_{(u,u'),a_j}, v_1, v_g \rangle$ that result from removing edge (u, u') from the relation labeled a_j, L as a winning strategy.[3]

All such games have n edges, so by inductive hypothesis we have

$$(M(\mathtt{SG}^\Sigma_{(u,u'),a_j}), v_1) \models \gamma_n^{\mathtt{ue}}.$$

for every edge (u, u') and label a_j. Now, the key observation is that each M-image of the game that results from L moving to v_1 and T removing edge (u, u') with label a_j, is exactly the model that results from removing edge (u, u') with a_j from the *model* $M(\mathtt{SG}^\Sigma)$.[4] Then, for all such (u, u') and a_j, we have

$$(M(\mathtt{SG}^\Sigma)^{a_j}_{(u,u')}, v_1) \models \gamma_n^{\mathtt{ue}}.$$

[3]The collection $\mathcal{E}^1_{(u,u'),a_j}$ is given by $(\mathcal{E}^0_{a_1}, \ldots, \mathcal{E}^0_{a_j} - \{u, u'\}, \ldots, \mathcal{E}^0_{a_{|\Sigma|}})$.

[4]In the original definition of a sabotage game this is not the case. In the game, the edges are implicitly ordered by numbers (the existence of an edge labeled with k implies the existence of edges labeled with $1, \ldots, k-1$); in the model, this is not the case. When we remove an edge from a *game* we always remove the one with the highest label, but when we remove an edge from a model we remove an arbitrary one: the operations of removing an edge and turning a game into a model do not commute.

It follows that $(M(\mathrm{SG}^\Sigma), v_1) \models \boxminus\gamma_n^{\mathrm{ue}}$ and therefore $(M(\mathrm{SG}^\Sigma), v_0) \models \Diamond \boxminus \gamma_n^{\mathrm{ue}}$, that is, $(M(\mathrm{SG}^\Sigma), v_0) \models \gamma_{n+1}^{\mathrm{ue}}$.

($\Leftarrow$) Suppose that $(M(\mathrm{SG}^\Sigma), v_0) \models goal \vee \Diamond \boxminus \gamma_n^{\mathrm{ue}}$. Then, v_0 is the goal state or else there is a state v_1 accessible from v_0 such that $(M(\mathrm{SG}^\Sigma), v_1) \models \boxminus\gamma_n^{\mathrm{ue}}$, that is, $(M(\mathrm{SG}^\Sigma)^{a_i}_{(u,u')}, v_1) \models \gamma_n^{\mathrm{ue}}$, for all edges (u, u') and labels a_j. By inductive hypothesis, L has a winning strategy in each game that correspond to each pointed model $(M(\mathrm{SG}^\Sigma)^{a_i}_{(u,u')}, v_1)$. But these games are exactly those that result from removing any edge from the game $\langle V, \mathcal{E}^0, v_0, v_g\rangle$ after L moves from v_0 to v_1. Hence, L has a winning strategy in $\langle V, \mathcal{E}^0, v_0, v_g\rangle$, the game that corresponds to the pointed model $(M(\mathrm{SG}^\Sigma), v_0)$, as required. □

Helpful Teacher and unwilling Learner (SLGhu) Now consider SLGhu, the game with helpful Teacher and unwilling Learner. We define γ_n^{hu} inductively, as

$$\gamma_0^{\mathrm{hu}} := goal, \qquad \gamma_{n+1}^{\mathrm{hu}} := goal \vee (\Diamond\top \wedge \Box\Diamond\!\!\!\!-\,\gamma_n^{\mathrm{hu}}).$$

In this case, Teacher has to be sure that Learner does not get stuck before he has reached the goal state — this is why the conjunct $\Diamond\top$ is needed in the definition of $\gamma_{n+1}^{\mathrm{hu}}$. We show that this formula corresponds to the existence of a winning strategy for Teacher.

Theorem 7.3.4. *Teacher has a winning strategy in the* SLGhu

$$\mathrm{SG}^\Sigma = \langle V, \mathcal{E}^0, v_0, v_g\rangle$$

if and only if $(M(\mathrm{SG}^\Sigma), v_0) \models \gamma_n^{\mathrm{hu}}$, *where n is the number of edges of* SG^Σ.

Proof. Similar to the proof of Theorem 7.3.3. □

Helpful Teacher and eager Learner (SLGhe) Finally, for SLGhe, the corresponding formula is defined as

$$\gamma_0^{\mathrm{he}} := goal, \qquad \gamma_{n+1}^{\mathrm{he}} := goal \vee \Diamond\Diamond\!\!\!\!-\,\gamma_n^{\mathrm{he}}.$$

Theorem 7.3.5. *Teacher and Learner have a joint winning strategy in* SLGhe

$$\mathrm{SG}^\Sigma = \langle V, \mathcal{E}^0, v_0, v_g\rangle$$

if and only if $(M(\mathrm{SG}^\Sigma), v_0) \models \gamma_n^{\mathrm{he}}$, *where n is the number of edges of* SG^Σ.

Proof. L and T have a joint winning strategy if and only if there is a path from v_0 to v_g. From left to right this is obvious. From right to left, if there is such a path, then there is also one without cycles[5], and a joint winning strategy is the one that follows the path and at each step removes the edge that has just been used (it is essential that L moves first). The theorem follows by observing that γ_n^{he} expresses the existence of such path. □

The above results for the three scenarios are summarized in Table 7.2.

[5]If this is not the case, i.e., if it is essential that L uses a path twice, removing used edges could cause L to be stuck somewhere away from the goal.

Game	Winning Condition in SML	Winner
SLGue	$\gamma_0^{\mathrm{ue}} := goal,\quad \gamma_{n+1}^{\mathrm{ue}} := goal \vee \Diamond \boxminus \gamma_n^{\mathrm{ue}}$	Learner
SLGhu	$\gamma_0^{\mathrm{hu}} := goal,\quad \gamma_{n+1}^{\mathrm{hu}} := goal \vee (\Diamond\top \wedge (\Box\Diamond\!\!\!\!-\,\gamma_n^{\mathrm{hu}}))$	Teacher
SLGhe	$\gamma_0^{\mathrm{he}} := goal,\quad \gamma_{n+1}^{\mathrm{he}} := goal \vee \Diamond\Diamond\!\!\!\!-\,\gamma_n^{\mathrm{he}}$	Both

Table 7.2: Winning Conditions for SLG in SML

7.3.3 Complexity of Sabotage Learning Games

We have characterized the existence of a winning strategy in our three versions of SLGs by means of sabotage modal logic formulae. In this section, we investigate the complexity of deciding whether such formulae are true in a given pointed model, i.e., the complexity of checking whether there is a winning strategy in the corresponding game.

By Theorem 7.2.5, the model checking problem of sabotage modal logic is PSPACE-complete. This gives us PSPACE upper bounds for the complexity of deciding whether a player can win a given game. In the three cases, we can also give tight lower bounds.

Unhelpful Teacher and eager Learner (SLGue) For SLGue, which can be identified with the standard sabotage game, PSPACE-hardness is shown by reduction from QUANTIFIED BOOLEAN FORMULA (Löding & Rohde, 2003b).

Theorem 7.3.6 (Löding & Rohde 2003b). SLGue *is PSPACE-complete.*

Helpful Teacher and unwilling Learner (SLGhu) Whereas at first sight, SLGhu and SLGue might seem to be duals of each other, the relationship between them is more complex due to the different nature of the players' moves: Learner moves locally by choosing an state accessible *from the current one*, while Teacher moves globally by removing *an arbitrary edge*. Nevertheless, we can show PSPACE-hardness for SLGhu. In the proof we will use the QUANTIFIED BOOLEAN FORMULA (QBF) problem, known to be PSPACE-complete.

Definition 7.3.7 (QUANTIFIED BOOLEAN FORMULA PROBLEM).

Instance *Let φ be an instance of* QBF, *i.e., a formula:*

$$\varphi := \exists x_1 \forall x_2 \exists x_3 \ldots Q x_n \psi$$

where Q is $\exists$ for n odd, and $\forall$ for n even, and ψ is a quantifier-free formula in conjunctive normal form.

Question *Is φ satisfiable?*

Theorem 7.3.8. SLGhu *is PSPACE-complete.*

Proof. From Theorem 7.2.5 and Theorem 7.3.4 it follows that SLGhu is in PSPACE. PSPACE-hardness of SLGhu is proved by showing that the QUANTIFIED BOOLEAN FORMULA (QBF) problem, can be polynomially reduced to SLGhu.[6]

We will construct a directed game arena for $\mathtt{SLGhu}_\varphi$ such that Learner has a winning strategy in the game iff the formula φ is satisfiable.

The ∃-gadget. Figure 7.1 represents the situation in which Learner chooses the assignment for x_i when i is odd. This part corresponds to assigning the value to an existentially quantified variable. Learner starts in A, and moves either left or right; if he wants to make x_i true, then he moves to $\bar{X}_i$, otherwise to X_i. Let us assume that he moves right, towards $\bar{X}_i$. Then Teacher has exactly four moves to remove all the edges leading to the dead-end #. At this point, Teacher cannot remove any edge in some other place in the graph without losing. So, Learner reaches $\bar{X}_i$, and Teacher is forced to remove the edge that leads from B to X_i, because otherwise Teacher would allow Learner to reach the dead-end #. At this point Learner moves towards B, and in the next step exits the gadget. Moving back towards $\bar{X}_i$ would cause him to lose, because then Teacher could remove the edge between B and $\bar{X}_i$ and Learner would be forced to enter the goal.

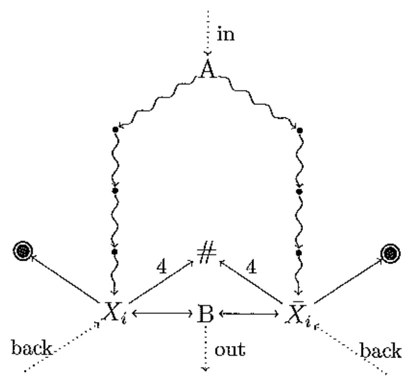

Figure 7.1: The ∃-gadget

The ∀-gadget. Figure 7.2 represents the situation in which Teacher chooses the assignment for x_i when i is even. This part corresponds to assigning the value to a universally quantified variable.

Let us assume that Teacher wants to make x_i false. Then she leads Learner towards X_i by successively removing the edges between C and $\bar{X}_i$. When Learner already is on the path to X_i, Teacher starts removing an edge going from X_i to the dead-end #. When Learner reaches X_i, he chooses to go towards B (because the other option is a goal). Then Teacher removes the edge that goes from B to $\bar{X}_i$, and Learner leaves the gadget.

Let us now assume that Teacher wants to make x_i true, and therefore wants Learner to reach $\bar{X}_i$. First she removes three of four edges from $\bar{X}_i$ to the dead-end #. Then Learner reaches C. Let us consider two cases:

1. Learner moves to $\bar{X}_i$. Teacher removes the last edge between $\bar{X}_i$ and the dead-end #, Learner moves to B, Teacher removes the edge to X_i and Learner leaves the gadget.

2. Learner moves to D. Teacher removes the four edges from X_i to the dead-end #, and then eliminates the remaining edge between $\bar{X}_i$ and the dead-end #. Learner leaves the gadget.

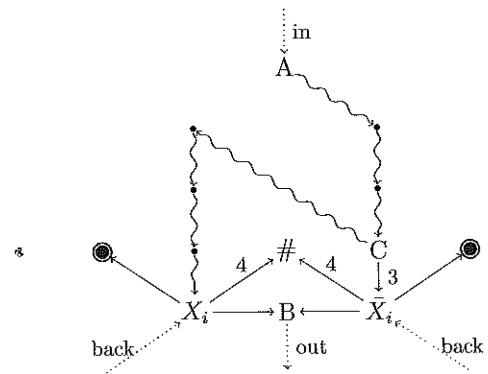

Figure 7.2: The ∀-gadget

The verification gadget. Figure 7.3 represents the situation in which Teacher chooses one clause from ψ. If Teacher chooses c, then Learner can choose one literal x_i from c. There are edges from x_i to an ∃-gadget if i is odd, and to a ∀-gadget otherwise, leading directly to X_i if x_i is positive in c, and to $\bar{X}_i$ otherwise. So, if the chosen assignment satisfies ψ, then for all clauses there is at least one literal which is true, and leads to the opposite truth value in the corresponding gadget, from which in turn Learner can get to the dead-end # (there are four edges left) and win the game.

For the converse, if the chosen assignment does not satisfy ψ, then Learner gets to a corresponding point in a proper gadget, Teacher removes the edge from this point to B, and the only option left for Learner is to enter the goal, which means that he loses the game.

[6]The proof uses the same strategy to the one of Theorem 7.3.6 of Löding & Rohde (2003b). We would like to thank Frank Radmacher for suggestions about this proof.

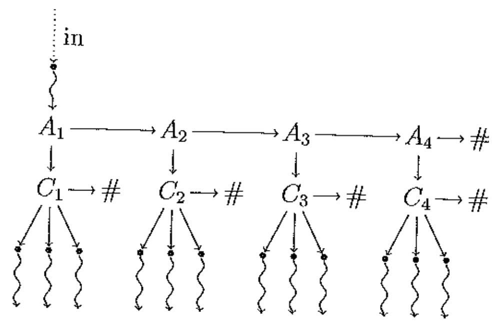

Figure 7.3: The verification gadget

With the above considerations, we can observe that Learner has a winning strategy in $\texttt{SLGhu}_\varphi$ iff φ is satisfiable. Moreover, the representation can clearly be done in polynomial time with respect to the size of φ. This finishes the proof. □

Helpful Teacher and eager Learner (SLGhe) Finally, let us have a look at SLGhe. This game is different from the two previous ones: L and T win or lose together. Then, a winning strategy for each of them need not take into account all possible moves of the other. This suggests that this version should be less complex than SLGue and SLGhu.

As mentioned in the proof of Theorem 7.3.5, L and T have a joint winning strategy if and only if the goal vertex is reachable from L's position. Thus, determining whether they can win SLGhe is equivalent to solving the REACHABILITY (ST-CONNECTIVITY) problem, which is known to be non-deterministic logarithmic space complete (NL-complete) (Papadimitriou, 1994).

Theorem 7.3.9. SLGhe *is NL-complete.*

Proof. Polynomial equivalence of SLGhe and REACHABILITY follows from the argument given in the proof of Theorem 7.3.5. □

Table 7.3 summarizes the complexity results for the different versions of SLG. The complexity of these problems can be interpreted as the difficulty of deciding whether certain aims of Teacher can be fulfilled. We can attribute the question of the existence of the winning strategy to Teacher, as she already has the global perspective on the situation anyway. Our results say how difficult it is for her to decide whether there is a chance of success. The results agree with our intuition, as coming up with a strategy to teach is easier if Learner and Teacher cooperate. Following our interpretation it is the easiest for Teacher to check whether the teaching will work out if Teacher assumes eagerness of Learner and she herself does her best to ensure that he succeeds. Moreover, the remaining two cases turn out to be equally difficult — it is as difficult to decide whether Teacher can force

an unwilling Learner to learn as it is to decide whether an eager Learner can learn in the presence of an unhelpful Teacher.

From the perspective of the standard sabotage games, our complexity result for SLGhu means that with an additional *safety*[7] winning condition, sabotage games are PSPACE-complete.

Game	Winning Condition	Complexity
SLGue	Learner wins iff he reaches the goal state, Teacher wins otherwise	PSPACE-complete
SLGhu	Teacher wins iff Learner reaches the goal state, Learner wins otherwise	PSPACE-complete
SLGhe	Both players win iff Learner reaches the goal state. Both lose otherwise	NL-complete

Table 7.3: Complexity Results for Sabotage Learning Games

7.4 Relaxing Strict Alternation

As mentioned above, in sabotage (learning) games the players' moves are asymmetric: Learner moves locally (moving to a vertex accessible *from the current one*) while Teacher moves globally (removing *any* edge from the graph, and thereby manipulating the space in which Learner is moving). Intuitively, both for a helpful and an unhelpful Teacher, it is not always necessary to react to a move of Learner by giving immediate feedback (here, by removing an edge). This leads us to a variation of a SLG in which Learner's move need not in principle be followed by Teacher's move, i.e., Teacher has the possibility of skipping a move.

Definition 7.4.1. *A* sabotage learning game without strict alternation *(for T) is a tuple* $\texttt{SLG}^* = \langle V, \mathcal{E}, v_0, v_g \rangle$. *Moves of Learner are as in* SLG *and, once he has chosen a vertex* v_1, *Teacher can choose between removing an edge, in which case the next game is given as in* SLG, *and doing nothing, in which case the next game is* $\langle V, \mathcal{E}, v_1, v_g \rangle$. *Again, there are three versions with different winning conditions, now called* SLG*ue, SLG*hu *and* SLG*he.

After defining the class of games SLG*, the natural question arises of how the winning abilities of the players change from SLG to SLG*, since in the latter

[7] Safety concerns those properties that we want to hold throughout the process, in this case it is L being away from the goal. Safety is usually contrasted with reachability, which requires the system to get into a certain configuration at some moment, here L reaching the goal (see, e.g., Radmacher & Thomas, 2008).

Teacher can choose between removing an edge or doing nothing. In the rest of this section, we show that for all three winning conditions (SLG*ue, SLG*hu, SLG*he), the winning abilities of the players remain the same as in the case in which players move in strict alternation. This is surprising in the SLGhu case. It is by no means obvious that if Teacher has a winning strategy in the game without strict alternation then she also has one in the regular version of the game, because we might expect that removing an edge instead of skipping the move could result in blocking the way to the goal.

We start with the case of an unhelpful Teacher and an eager Learner, SLG*ue. Note that although in this new setting matches can be infinite, e.g., in the game with unwilling L, if T skips her moves indefinitely, L will just keep moving, and hence the game will not be finished in finite time. However, in fact if Learner can win the game, he can do so in a finite number of rounds. We start with a lemma stating that if Learner can win some SLGue in some number of rounds, then he can do so also if the underlying multi-graph has additional edges.

Definition 7.4.2. *Let $\Sigma = \{a_1, \dots a_n\}$ be a finite set of labels. For directed labeled multi-graphs, $G^\Sigma = (V, \mathcal{E})$ and $G'^\Sigma = (V', \mathcal{E}')$, we say that G^Σ is a* subgraph *of G'^Σ if $V \subseteq V'$ and $\mathcal{E}_{a_i} \subseteq \mathcal{E}'_{a_i}$ for all labels $a_i \in \Sigma$.*

Lemma 7.4.3. *If Learner has a strategy for winning the* SLGue *$\langle V, \mathcal{E}, v_0, v_g \rangle$ in at most n rounds, then he can also win any* SLGue *$\langle V, \mathcal{E}', v_0, v_g \rangle$ in at most n rounds, where $(V, \mathcal{E})$ is a subgraph of $(V, \mathcal{E}')$.*

Proof. The proof is by induction on n. In the inductive step, for the case that T removes an edge which was not in the original multi-graph, note that the resulting graph is a supergraph of the original one. Then we can use the inductive hypothesis. □

Theorem 7.4.4. *Let us consider the* SLG *$\langle V, \mathcal{E}, v_0, v_g \rangle$ with $(V, \mathcal{E})$ being a directed labeled multi-graph and $v, v_g \in V$. Learner has a winning strategy in the corresponding* SLGue *iff he has a wining strategy in the corresponding* SLG*ue.

Proof. From left to right, we show by induction on n that if L can win the SLGue in at most n rounds, then he can also win the SLG*ue in at most n rounds. In the inductive step, in the case that T responds by not removing any edge, we first use Lemma 7.4.3 and then can apply the inductive hypothesis.

The direction from right to left is immediate: if L has a winning strategy for SLG*ue, then he can also win the corresponding SLGue by using the same strategy. □

The case of helpful Teacher and unwilling Learner is more interesting. One might expect that the additional possibility of skipping a move gives more power to Teacher, since she could avoid removals that would have made the goal unreachable from the current vertex. However, we can show that this is not the case. First, we state the following lemmas.

Lemma 7.4.5. *Consider the* SLG*hu *$\langle V, \mathcal{E}, v_0, v_g \rangle$. If there is a path from v_0 to v_g and there is no path from v_0 to a state from where v_g is not reachable, then Teacher has a winning strategy.*

Proof. Let us assume that all states reachable from v_0 are on paths to v_g. Then, even if T refrains from removing any edge, L will be on a path to the goal. Now, either the path to the goal does not include a loop or it does. If it does not then T can simply wait until L arrives at the goal. If it does, T can remove the edges that lead to the loops in such a way that v_g is still reachable from any vertex. □

Lemma 7.4.6. *For all* SLG*hu *$\langle V, \mathcal{E}, v_0, v_g \rangle$, if Teacher has a winning strategy and there is an edge $(v, v') \in \mathcal{E}_a$ for some $a \in \Sigma$ such that no path from v_0 to v_g uses (v, v'), then Teacher also has a winning strategy in $\langle V, \mathcal{E}', v_0, v_g \rangle$, where $\mathcal{E}'$ results from removing (v, v') from $\mathcal{E}_a$.*

Proof. If v is not reachable from v_0, it is easy to see that the claim holds. Assume that v is reachable from v_0. T's winning strategy should prevent L from moving from v to v' (otherwise L wins). Hence, T can also win if (v, v') is not there. □

Theorem 7.4.7. *If Teacher has a winning strategy in the* SLG*hu *$\langle V, \mathcal{E}, v_0, v_g \rangle$, then she also has a winning strategy in which she removes an edge in each round.*

Proof. The proof is by induction on the number of edges $n = \sum_{a \in \Sigma} |\mathcal{E}_a|$.

The base case
Straightforward: there is no round because L cannot move.

The inductive case
Assume that T has a winning strategy in SLG*hu $\langle V, \mathcal{E}, v_0, v_g \rangle$ with $\sum_{a \in \Sigma} |\mathcal{E}_a| = n + 1$.

If $v_0 = v_g$, it is obvious. Otherwise, since T can win, there is some $v_1 \in V$ such that $(v_0, v_1) \in \mathcal{E}_a$ for some $a \in \Sigma$ and for all such v_1 we have:

1. There is a path from v_1 to v_g, and
2. (a) T can win $\langle V, \mathcal{E}, v_1, v_g \rangle$, or
   (b) there is a $((v, v'), a) \in (V \times V) \times \Sigma$ such that $(v, v') \in \mathcal{E}_a$ and T can win $\langle V, \mathcal{E}', v_1, v_g \rangle$ where $\mathcal{E}'$ is the result from removing (v, v') from $\mathcal{E}_a$.

If 2b holds, since $\sum_{a \in \Sigma} |\mathcal{E}'_a| = n$, we are done — we use the inductive hypothesis to conclude that T has a winning strategy in which she removes an edge in each round (in particular, her first choice is $((v, v'), a)$). Let us show that 2b holds.

If there is some $(v, v') \in V \times V$ such that $(v, v') \in \mathcal{E}_a$ for some $a \in \Sigma$ and this edge is not part of any path from v_1 to v_g then by Lemma 7.4.6, T can remove this edge and 2*b* holds, so we are done.

If all edges in $(V, \mathcal{E})$ belong to a path from v_1 to v_g, from 1, there are two cases: either there is only one, or there is more than one path from v_1 to v_g.

In the first case (only one path) (v_0, v_1) can be chosen since it cannot be part of the *unique* path from v_1 to v_g. Assume now that there is more than one path from v_1 to v_g. Let $p = (v_1, v_2, \ldots, v_g)$ be the/a shortest path from v_1 to v_g. This path cannot contain any loops. Then, from this path take v_i such that i is the smallest index for which it holds that from v_i there is a path $(v_i, v'_{i+1}, \ldots v_g)$ to v_g that is at least as long as the path following p from v_i (i.e., $(v_i, v_{i+1}, \ldots, v_g)$). Intuitively, when following path p from v_1 to v_g, v_i is the first point at which one can deviate from p in order to take another path to v_g (recall that we consider the case where every vertex in the graph is part of a path from v_1 to v_g). Now it is possible for T to choose $(v_i, v'_{i+1}) \in \mathcal{E}_a$. Let $\mathcal{E}'$ be the resulting set of edges after removing (v_i, v'_{i+1}) from $\mathcal{E}_a$. Then we are in the game $\langle V, \mathcal{E}', v_1, v_g\rangle$. Note that due to the way we chose the edge to be removed, in the new graph it still holds that from v_0 there is no path to a vertex from which v_g is not reachable (this holds because from v_i the goal v_g is still reachable). Then by Lemma 7.4.5, T can win $\langle V, \mathcal{E}', v_1, v_g\rangle$, which then implies 2*b*.

Hence, we conclude that 2*b* is the case and thus using the inductive hypothesis, T can win $\langle V, \mathcal{E}, v_0, v_g\rangle$ also by removing an edge in every round. □

Corollary 7.4.8. *Teacher has a* SLG*hu*-winning strategy in* $\langle V, \mathcal{E}, v_0, v_g\rangle$ *iff she has a* SLGhu*-winning strategy.*

As the reader might have noticed, the result that if Teacher can win a SLG*hu, then she can also win the corresponding SLGhu, relies on the fact that Learner is the first to move. For instance, in a graph with two vertices and and one edge — leading from the initial vertex to the goal vertex — if Teacher was to move first, she can win the SLG*hu only by skipping the move.

Finally, let us consider the case of helpful Teacher and eager Learner.

Theorem 7.4.9. *If Learner and Teacher have a joint* SLG*he*-winning strategy in* $\langle V, \mathcal{E}, v_0, v_g\rangle$, *then they have a joint* SLGhe*-winning strategy*

Proof. If the players have a joint SLG*he-winning strategy, then there is an acyclic path from v_0 to v_g, which L can follow. At each round, T has to remove the edge that has just been used by L. □

Let us briefly conclude this section. We have shown that allowing Teacher to skip moves does not change the winning abilities of the players. Using these results, both the complexity and definability results from the previous section also apply to their versions without strict alternation, in which Teacher can skip a move.

7.5 Conclusions and Perspectives

We have provided a game theoretical approach to learning that takes into account different levels of cooperativeness between Learner and Teacher in a game of

perfect information based on sabotage games. Because of our new interpretation we were able to define sabotage learning games with three different winning conditions. Then, following the strategy of Löding & Rohde (2003b), we have shown how sabotage modal logic can be used to reason about these games and, in particular, we have identified formulae of the language that characterize the existence of winning strategies in each of the two remaining cases. We also provided complexity results for the model-checking problem of these formulae. Our complexity results support the intuitive claim that cooperation of agents facilitates learning. Moreover, in our framework it turns out to be as difficult to decide whether a Teacher can force an unwilling Learner to learn as it is to decide whether an eager Learner can learn in the presence of an unhelpful Teacher.

Viewed from the perspective of the standard sabotage games, our complexity result for the game between a helpful Teacher and an unwilling Learner means that also with a *safety* winning condition, sabotage games are PSPACE-complete.

From the game-theoretical perspective, sabotage learning games can be extended to more general scenarios by relaxing the strict alternation. The moves of the players are of a different nature. Learner's moves can be seen as internal ones, moving to a state that is reachable *from the current one*, while Teacher's moves can be interpreted externally, removing *any* edge of the underlying graph. Once this asymmetry is observed, it becomes natural to ask what happens if from time to time Learner's move is not followed by Teacher's move (e.g., Learner can perform several changes of his information state before Teacher makes a restriction). Our results of Section 7.4 show that if we allow Teacher to skip a move, the winning abilities of the players do not change with respect to the original versions of the games. In the case of helpful Teacher and unwilling Learner, the result is quite surprising since it says that if Teacher can force Learner to learn in the game with non-strict alternation, then she can also do it when she is forced to remove edges in each round. This result crucially depends on the fact that Learner is the first to move, and does not hold in case Teacher starts the game.

In this chapter, we have described the learning process purely as *changes in information states*, without going further into their epistemic and/or doxastic interpretation.

We understand successful learning as the ability to reach an appropriate information state, not taking into account what happens afterwards. Formal learning theory that treats the inductive inference type of learning situates our work close to the concept of *finite identification* (Mukouchi, 1992) treated extensively in previous chapters. In particular, we are not concerned with the stability of the resulting state. *Identification in the limit* (Gold, 1967) extends finite identification by looking beyond reachability in order to describe 'ongoing behavior'. Fixed-point logics, such as the modal μ-calculus (Kozen, 1983; Scott & Bakker, 1969), can provide us with tools to express this notion of learnability. Further work involves investigating how fixed points can enrich sabotage-based learning analysis.

Moreover, in natural learning scenarios, e.g., language learning, the goal of the learning process is concealed from Learner. An extension of the framework of randomized sabotage games (Klein, Radmacher, & Thomas, 2009) could then be used to model the interaction between Learner and Teacher.

Chapter 8

The Muddy Scientists

Imagine you are one of ten prisoners locked up for extensive use of logic. To make you even more miserable, the guard comes up with a puzzle. He gathers all ten of you and says: 'Each of you will be assigned a random hat, either black or white. You will be lined up single file where each can see the hats in front of him but not behind. Starting with the prisoner in the back of the line and moving forward, you must each, in turn, say only one word which must be 'black' or 'white'. If the word you uttered matches your hat color you are released, if not, you are killed on the spot. You have half an hour to pray for your life.' Then he leaves. One of the prisoners says: 'I have a plan! If you agree on it, 9 of us 10 will definitely survive, and the remaining one has a 50/50 chance of survival.' What does he have in mind?

Most probably the strategy that he wants to implement is as follows. First, the prisoners have to agree on the following meaning of the utterance of the one who is the last in the line. If he says 'white', it means that he sees an even number of white hats in front of him. If he says 'black' it means that he sees an odd number of white hats in front of him. Hence, his utterance has nothing to do with what he thinks his own hat is. There is a 50/50 chance of the total number black hats being odd or even, and a 50/50 chance of his hat being black or white, so his chance of survival is the same. However, after this utterance the prisoner that stands in front of him knows for sure the color of his hat—he compares the utterance of his predecessor with the number of white hats he sees in front of him. If the parity is the same, he concludes that his hat is black, otherwise it is white. He makes his guess aloud. Now the person in front of her takes into account the first announcement and the second utterance, sees the number of white hats in front of her, and now she is also certain about her hat's color, etc.

This epistemic scenario shows the power of multi-agent information exchange. A very simple *quantitative* public announcement carries powerful *qualitative* information. Agents can easily deduce nontrivial facts from implicit and indirect information. Obviously, the information has to be relevant to make certain deductions possible. For example, in the above scenario announcing 'At least 5