Lehr- und Arbeitsmaterial

MENSCH, Maschine!

Wer zeigt hier wem den Weg?

Eine Initiative des Bundesministeriums für Bildung und Forschung

Wissenschaftsjahr 2019

KÜNSTLICHE INTELLIGENZ

Nutzungsrechte

„Human vs Machine" – Who *shows* whom *the way?*

What is this about?

Unlocking your phone via face recognition, speaking with chatbots or navigating through the traffic: Artificial Intelligence (AI) has become an important part of our everyday lives. But how do AI and Machine Learning (ML) work? Machines are not intelligent themselves; they can just be trained through Machine Learning to perform some specific tasks quite well. The computer "learns" to do this by recognizing patterns in huge amounts of data and deriving rules for decisions.

That sounds complex! But in this activity, *"Human vs Machine"*, in the context of playing a game against a machine, school-aged young people experience how a machine is trained, and come to understand the basic principles behind "Artificial Intelligence" and "Machine Learning."

How does this work?

Martin Gardner introduced the game Hexapawn in 1962 as the "learning matchbox computer." In *"Human vs Machine"*, a human opponent plays Hexapawn against the machine, which is simulated by several other players. While playing the game, the players experience Machine Learning as an algorithmic process controlled by rules and get to know how the machine "learns" during the turns.

In the activity, students play Hexapawn multiple times. If the machine loses, the last move the machine made will be deleted—so that it will never make that mistake again. The better the human player plays against the machine, the faster the machine learns and improves its chances of winning. You can find out exactly how the game works and what you need to pay attention to as a teacher or supervisor in the printed instructions that come with the game materials and in the following description.

Dieses Modul stellt den Kern des Materials dar und sollte daher in jedem Fall mit den Kindern und Jugendlichen bearbeitet werden. Es bildet die Grundlage für eine tiefergehende pädagogische Auseinandersetzung mit KI und ML.

- ▶▶ Wer sich in die Idee der **Matchbox Computer** etwas vertiefen möchte, kann sich hier **MENACE** ansehen, einen wesentlich komplexeren Matchbox Computer für Tic-Tac-Toe (Englisch): **ybit.ly/2AMhxU3**
- ▶▶ Außerdem gibt es eine Simulation des **MENACE,** bei dem sein Lernprozess gut zu verfolgen ist: **bit.ly/30BUyYr**

Wie geht es danach weiter?

Auf Basis der Erfahrungen mit dem Lernspiel werden in [▷Modul 3] Unterschiede zwischen menschlicher und Künstlicher Intelligenz herausgearbeitet, in [▷Modul 5] werden ethische Fragen und Problemstellungen im Umgang mit und in der Nutzung von KI-Systemen behandelt und in [▷Modul 6] wird als Zukunftswerkstatt diskutiert, wie das zukünftige Zusammenleben zwischen Menschen und Maschinen (im Sinn von KI-Systemen) aussehen kann. Sie können alle Module nacheinander einsetzen oder auch nur die Themen entnehmen, die für Sie passen – sie bauen alle auf diesem Modul auf.

Before we talk about “Human vs Machine”, let’s review some of the theory behind it.

What are Artificial Intelligence (AI) and Machine Learning (ML)?

It is not easy to answer this question. Underlying terms such as “intelligence” and “learning” are not clearly defined. As a result, it is hard to construct simple definitions for Artificial Intelligence and Machine Learning. Many definitions look similar on the surface, but when you look at them in detail, they’re quite different. We must therefore communicate clearly which one we plan to use.

We’ll use this one for Artificial Intelligence, from Simmons and Chappell in 1988, throughout this teaching material:

▶▶ **Artificial intelligence denotes the behavior of a machine which, if a human behaves in the same way, is considered intelligent.**

We have still not defined “intelligent behavior,” but let us proceed assuming that we would recognize it if we saw it.

Zur Begriffsdefinition gibt es eine Menge Vorschläge. Wir orientieren uns in der Regel an den Aussagen, die in den folgenden Quellen zu finden sind:

▶▶ **Glossar der Plattform Lernende Systeme:** **bit.ly/2Le18xt**

▶▶ **BaFin-Studie 2018:** **bit.ly/2KXnBjx**

What is “Machine Learning”?

In popular media about science, the terms *Artificial Intelligence* (AI) and *Machine Learning* (ML) are often used synonymously. But this oversimplifies the matter, because Machine Learning is a subarea of Artificial Intelligence. An easy and helpful definition of Machine Learning is found in the report of the BaFin-Studie 2018:

▶▶ **Machine Learning is, quite generally, the idea of giving computers the ability to learn from data and experiences through algorithms. Using Machine Learning, computers create their own models of the world and improve in solving their specific tasks.**

A Machine Learning algorithm often uses huge amounts of data to “train” these “models.” To simplify somewhat, these models consist of sets of rules which the computer uses to classify or rate a situation; the rules then let the algorithm make decisions. In our case, these decisions are game moves. Such a model, with all those rulesets, is called a “Machine Learner.”

While playing the simulation game *“Human vs Machine”*, students get to know a simple version of such a model. In this activity, the model is a sheet showing the set of all of the moves the machine might make in the game. By playing the game, students train the model, gradually reducing that set of available moves—and improving the machine’s performance.

With rapidly increasing computer power, researchers, governments, and companies in the private sector are using increasingly complex models, and training them with more and better data. You can find AI-systems that use Machine Learning in everyday applications such as apps in smartphones or various assistance systems. Computer programs like AlphaZero show how powerful such systems are already. They are able to beat even professionals in games like Go, Shogi or chess.

AlphaZero is an AI-System developed by DeepMind that learned how to play several board games by playing against itself. This strategy is called “reinforcement learning.” This computer system is described quite well *on Wikipedia.*

What is the difference between “classical” IT-systems and AI-systems with Machine Learning?

The difference is in the approach to problem solving.

In the “classical” approach, we analyze underlying problems, develop a problem-solving algorithm, and turn it into code.

▶ In our game, we would analyze the game thoroughly first, and then develop the algorithm. If we did it right, the algorithm would make the best possible move in every state of the game right away. For that to work, the algorithm would need to know every possible state the game could be in. Then it would track every possible move to the end of the game and determine whether that path resulted in a win or a loss. This is quite easy to implement for a simple game like Hexapawn, but in a more complex game with more states, this procedure is not feasible. Besides that, the program would not work if just a single state of the game or possible move were left out when the algorithm was developed.

Machine Learning, on the other hand, works by analyzing lots of data and using the data to “train” a model “on the fly,” assigning and updating values for states based on the outcomes of the games in which the states appear. The model therefore improves as it gets more data, solving the problem as well as possible.

▶ In our activity “Human vs Machine”, we start with an untrained model which contains every possible move the machine can make for every state of the game. All those moves are valued equally, thus they are not categorized as to whether they are good moves or not. The training happens while playing the game: If the machine loses, the particular move that made the machine lose is crossed out. Therefore, the model changes every time the machine loses, until all the moves that lead to a loss by the machine are deleted. After this procedure, the model is fully trained.

You can find the terms “weak AI” and “strong AI” quite often alongside all these definitions:

▶▶ **Weak AI** is like the AI in the present discussion: AI-systems that are optimized for a specific task and do not show behavior similar to human intelligence. Our AI system with Machine Learning is as dependent on its defined domain as the classical algorithms. It can develop an optimized procedure only for a specific task or at most a group of tasks.

▶▶ On the other hand, **Strong AI** means an actual, generalized intelligence at a human level. Often the passing of a Turing test is seen as evidence of a strong AI.

„MENSCH, Maschine!" im Einsatz

How to play "Human vs Machine"

The learning game "Human vs Machine" makes Machine Learning tangible for students.

We suggest that you set up the materials at tables before the start of the lesson. If possible, arrange students at the tables in groups of 5. However, the game can be played by 3 to 6 people at a time. For information on how to play the game with different numbers of people, and for other ways to vary the game, see the "variation of the game" below.

After explaining the rules of the game and distributing the roles to the students, consider playing the first game (of ten) step by step with the entire class, using the printed flowchart that accompanies the game.

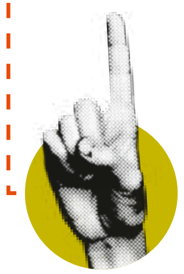

Bitte planen Sie ausreichend Zeit ein
Von der Erklärung bis zum Ende der Durchführung zwischen 45 und 60 Minuten, für die anschließende Reflexion weitere 15 bis 30 Minuten.

The games are usually different in each group. By discussing the varied results, you can help students see what they are supposed to do when they win or lose against the machine. In addition, exchanging results among the groups can provide various insights.

Afterwards, the groups should complete at least 10 games and write down their results. If they play fewer games, the machine's improvement – i.e., learning – might not be visible.

If a group finishes faster, they can play a few more games. Their model can be interesting when you debrief the activity, because their machine will have had more games to train and – assuming a good human player – may well be better-trained, with more eliminated moves.

▶▶ IMPORTANT NOTICE
The machine determines which move to make by comparing the input data (the current state) with the model. After each defeat, the algorithm improves the model by deleting the machine's last move. Should the machine win, the model remains unchanged.

Please make sure that the human player always tries to win the game, as this is the only way to improve the model!

When debriefing the lesson and reflecting on the game, you can discuss how the strength of the human player influences the model.

After the game: debriefing

After playing "Human vs Machine", the most important phase of the module follows with reflection about the process of playing the game and updating the model. The groups gain different experiences while playing and come to different conclusions. It is therefore very important to facilitate a discussion of this subject in order to gain further insights. You can use [Worksheet 1] to structure this process.

Why?

Removing moves is changing the probability of a move being randomly selected. Adjusting - increasing or decreasing - these probabilities by eliminating the wrong moves will cause the machine to learn a "wrong strategy." Recognizing this is a crucial step towards understanding how an AI system learns.

Pitfalls and other learning opportunities

It can happen that the students accidentally remove one of the machine's winning moves, so that the machine always loses. This is usually noticed during the debriefing and can lead to additional learning outcomes: The machine learns by removing losing moves, but if you remove winning moves, it impairs its operation. It is "broken."

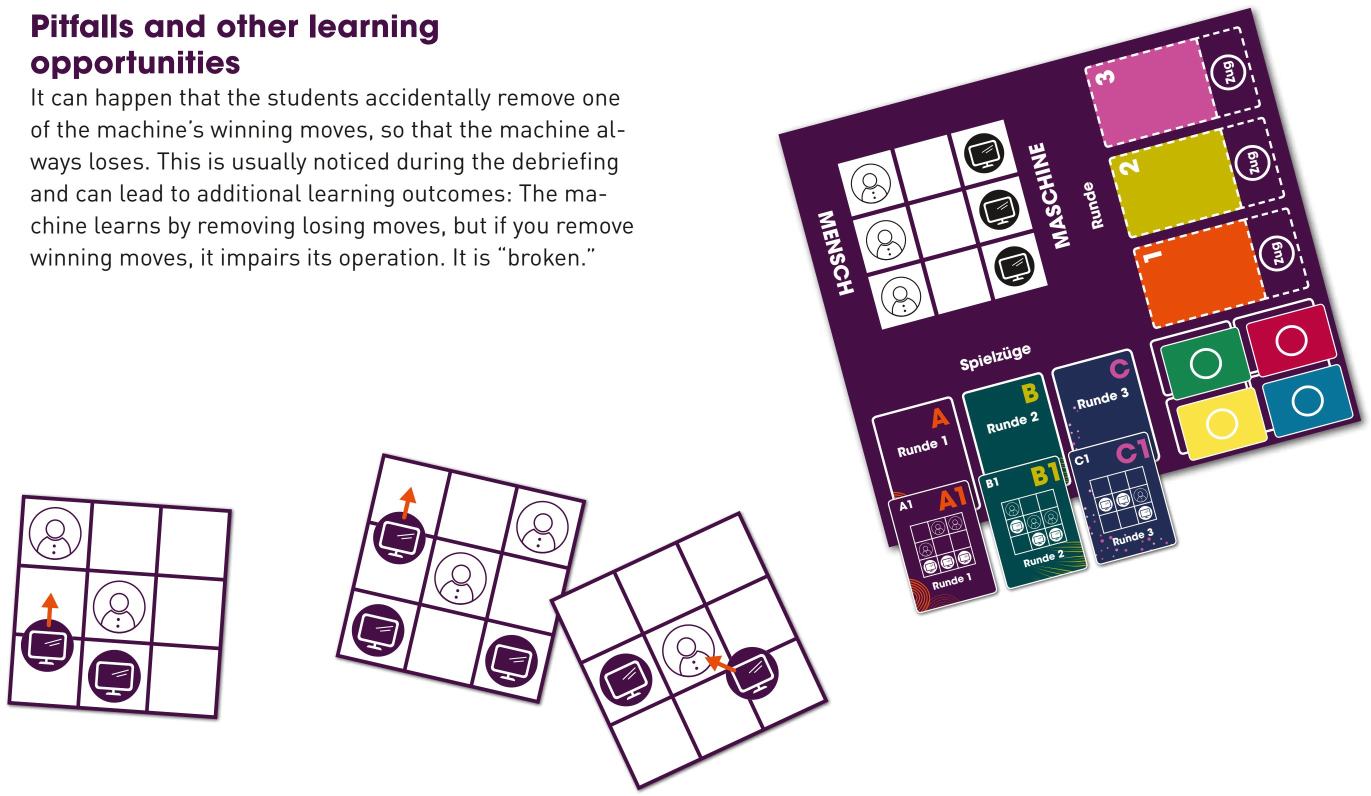

Worksheet 1

Findings from the game "Human vs Machine"

You have now played at least ten games of "Human vs Machine" between a human and a simulated machine, and noted down the results. Now it's time to find out together what this game has to do with Artificial Intelligence and Machine Learning.

1. My experience with the game "Human vs Machine"

Did the machine play better as it gained experience?

◯ Ja ◯ Nein

Explain your answer. ______________________________

____

____

____

2. "Human vs Machine" learns

Describe how the machine "learned". ______________________________

____

Can you use "Human vs Machine" to play tic-tac-toe or chess? Please explain your answer.

____

____

____

____

3. "Human vs Machine" and Artificial Intelligence

You now know how "Human vs Machine" works and how it can learn. Try to explain what Machine Learning means in your own words. ______________________________

____

So, what is Artificial Intelligence? ______________________________

____

____

____

Notes on worksheet 1: Findings from the game "Human vs Machine"

Towards the end of the game, students can begin to fill in the worksheet individually. It is important to discuss their responses together with all group members afterwards.

The question whether the machine plays better as the module progresses will usually be answered and justified differently, depending on the course of play in the various groups.

In the first ten rounds, human and machine often win similarly often, so from a statistical perspective, there may be no obvious change in winning behavior. Here you can ask the groups who answered "No" (the machine did not show improvement) whether it became more difficult for the human player to win against the machine during the activity. In our experience, these players often confirm that it became more difficult because individual moves that quickly led to a defeat of the machine are no longer in play.

Even those groups who state that their machine played better as the games progressed often say that the machine has to get better, because the bad moves are eliminated. Students can use this insight to investigate how the machine has learned.

Classification of the learning process.

Learning through reinforcement
The machine in "Human vs Machine" is a so-called Machine Learner. This type of Machine Learning process is technically called "Reinforcement Learning." This procedure is commonly used for learning strategies.

What does "Reinforcement Learning" mean?
In general, this means that actions that lead to desired results are strengthened, while actions that lead to undesired results are weakened. The "Human vs Machine" machine uses a very simple procedure: Unsuccessful moves are strongly weakened by taking them out. The probability that they will be drawn again in the future is then 0. This increases the probability of choosing one of the moves remaining on the state card. On the other hand, the algorithm does not explicitly reinforce successful actions. Thus the higher probability of good moves comes only from the elimination of the bad choices. In this way the machine will play "perfectly" quite quickly, as all unsuccessful moves will be quickly eliminated.

▸▸ Note:
If you want to focus on changing the model by deleting the unsuccessful moves, i.e. the process of "learning," or more advanced aspects of machine learning, you should consistently encourage the groups to delete eliminated moves from the overview of moves as well as from the state cards.

Because although an improvement in the game is usually not obvious, the strategy for selecting the machine's moves changes with every move that is cancelled. You can see this state of the decision model of the machine easily in the overview of moves and compare between the groups.

Use the second part of the worksheet for whole-class exploration of the learning procedure.

Use the questions to prompt discussion and summary in the debriefing. The answer to the question of how the machine has learned is usually derived from the class discussion. When the students write their individual responses, the complexity of the writing will depend on the learning group; in any case, important insights for the students should be:

What prompts the machine to learn?

When the machine loses, it changes its strategy; it learns from its mistakes.

Possible addition for strong or older learning groups: It is about Reinforcement Learning. Playing represents the training of the model, with the games providing the data needed for training. In order not to encourage student misconceptions, you should point out during the discussion that the learning in "Human vs Machine" is an exception in that it uses only negative reinforcement. Most machine learners also use positive reinforcement.

How does the machine learn?

It learns by deleting unsuccessful moves.

Possible addition for strong or older learning groups: By deleting the moves, the probabilities of making these moves again are set to zero.

At this point you can also discuss how you could optimize the training process (see "Variation by adapting the game system" in the booklet on the website). This can lead to exciting discussions in some groups—especially if they already understand probability.

After students have understood how "Human vs Machine" learns, it is important to make it clear that the machine can only play Hexapawn. It is trained and optimized for this specific task and the learning process is adapted to this task. However, this question offers a starting point for interesting discussions, for example about the universality of different AI applications that can already be found in everyday life. Because even though "machine learners" outperform humans in some areas (e.g. AlphaGo for the Go game or OpenAI for the game Dota) due to today's ever-increasing amounts of data and computing power, these learners are usually specifically designed for their task. In contrast, the development of a general AI for arbitrary fields of activity, a so-called strong AI, is not yet foreseeable.

The third section of the worksheet, which asks students to define Artificial Intelligence and Machine Learning, serves to generalize and consolidate the knowledge acquired so far specifically for "Human vs Machine" and to correct any misconceptions at an early stage.

Especially if they have room for discussion, students are often able to formulate initial definitions that essentially describe the modification of the model by data or the removal of moves.

If you decide not to work on [Worksheet 1], you should discuss what the term "intelligent" in the human sense means with the students, especially in the individual roles of the players and what appears intelligent in the machine. In this way, students can fundamentally recognize what distinguishes human intelligence from artificial intelligence.

An important realization is that the machine does not need to understand the problem, in this case the game, in order to win. In fact, it cannot develop an understanding—of anything. This means that the algorithm does not learn the game in a direct sense, but only analyzes the inputs—the states and available moves—and reacts to them using the model it has constructed.

Notes on worksheet 2
Deeper understanding of "Human vs Machine"

This worksheet can follow the debriefing in worksheet 1.

The aim of these exercises is to recognize that although the design of the machine shows intelligence, the roles of the machine are not intelligent—and the machine is not intelligent either.

Nevertheless, the machine seems to show intelligent behavior because of the way it is designed: it "learns from its mistakes," a behavior that seems quite human even though the behavior arises not from insight or understanding, but from following simple rules. This can be compared with the thought experiment of the "Chinese Room" by John Searle.
You can implement this worksheet in your class using the jigsaw technique, where players of each role get together in a homogeneous, "expert" group. The last task on the worksheet, where students return to their original "game" groups for discussion, is quite openly phrased, in order to give you flexibility in bringing closure to the activity. For example, you could initiate discussions within the game groups, discuss the findings in the whole learning group, or let the groups create posters about their work on worksheet 2.

Gedankenexperiment des chinesischen Zimmers

Eine gut verständliche Erklärung findet sich auf **Wikipedia**.

Eine kurze, prägnante Erklärung samt kurzem Film (Englisch): **bit.ly/2ZsbaAj**

Eine kurze deutschsprachige Erklärung, die von der Uni Frankfurt veröffentlicht wurde und zudem zur Diskussion über die Möglichkeit, harte KI zu entwickeln, anregt: **bit.ly/2Lagbsb**

Worksheet 2

„MENSCH, Maschine!“ genauer auf den Zahn gefühlt

We would like to know a little more about what exactly defines the machine.

1. **Form groups with people who had the same role as you during the games**
2. **If you have been responsible for more than one role, choose the role you think is the most important.**
3. **Answer the following questions together with your group.**

"What am I?"

Which task is performed by your role? Describe the task in your own words.

Do any other names for your role come to mind? What are these names?

Intelligent or not?

What aspects of your role would you call "intelligent"?

What aspects of the machine would you call "Intelligent"?

▶▶ Return to the group you played the game with and discuss your responses to this worksheet.

This worksheet consists of two parts:

The first part "What am I?" encourages your "expert" group to explore the specific task of their role. The explicit description of their role provides clarification about the different roles in the game and prepares the second question. The discussion about other names of their role should help them understand the whole mechanism of the game more deeply – a deeper understanding than the simple phrase "Unsuccessful moves are deleted."

The following overview shows some examples from a field test of this module with school-aged students at the Paderborn University. If some of the students in class are struggling to find appropriate alternative names, because none of them cover all aspects of their role, you can use this struggle as a new occasion for a discussion.

„Who am I?"

Human
teacher
▶ Teach the machine
trainer
▶ Train the machine
guinea pig
▶ (Unknowingly) help the machine
sparring partner
▶ Use the machine as a training partner / train the enemy.

Situation Evaluator
move verifier
▶ Check: Did the human already win?
state card searcher
▶ Find the card that matches the game board
referee
▶ Decide, after the check, if the human already won

Move Finder
editor
▶ Cross out unsuccessful moves, provide options
card shuffler
▶ Present cards for random pick

Move Picker
random selector
▶ Select a color at random and move the piece

Win Checker
final analyzer
▶ Check: Did the machine already win?
referee
▶ Decide, after the check, if the machine already won

The second part "intelligent or not" supports learners in recognizing that the different roles within the machine are not intelligent. The exception is the human player, who plays the game using the rules. The human builds up strategies, thinking ahead to what the machine might do, and adapting to the machine's moves, all of which seems to be intelligent.

When you ask the students what aspects of the machines' moves have been intelligent, some will recognize that the individual instructions or tasks the machine performs are not intelligent, whereas the function of the machine as a whole looks intelligent.

But the machine's behavior doesn´t mean the machine actually understands what it is doing or how it does it. The machine just performs this specific task – winning Hexapawn. It is a carefully defined "Machine Learner."

Some students confuse the appearance of intelligence with a genuine, underlying intelligence. We humans appear intelligent because our inner intentions, motivations, planning, understanding, and experience combine to produce what we recognize as intelligent from years of interacting with other humans. But the "intelligence" the machine shows – on the outside – is based only on algorithms and rules, even though it resembles intelligence to an observer. It is quite important in this case, that you work with your class to understand this distinction between appearance and the underlying mechanism.

It may help to compare the human`s winning strategy with the strategy the machine uses: "What is the difference between the human and the machine when they are making a move?" The human acts intentionally and analytically, while the machine simply reacts on instructions in each role, without understanding what to do. Therefore, the machine is not really intelligent.

▶▶ Möchten Sie sich genauer mit den Unterschieden von menschlicher und Künstlicher Intelligenz beschäftigen, sollten Sie im Weiteren [▶Modul 3] bearbeiten.

Erweiterung der Spielanleitung für Lehrkräfte und Gruppenleitungen

„Mensch, Maschine!" im Einsatz – Game variations

This sheet describes various options for modifying "Human vs Machine" to make it as useful and effective as possible in your learning group.

Game variations for different numbers of players

The activity is designed for five players. Every player takes one of five roles. One role represents the human player, while the other four roles take over the tasks of the machine. The following roles are usually filled:

- Human
- Machine
  - Situation Evaluator
  - Move Finder
  - Move Picker
  - Win Checker

However, these roles can be split or merged to make the game work for different numbers of players.

Variation for three players

With only three players combine these machines roles:

- Merge the Situation Evaluator and the Move Finder
- Merge the Move Picker and Win Checker

Note that it is important that the Move Finder and Move Picker are in two different players, so that the randomness for picking moves for the machine becomes apparent.

Variation for four players

With four players one player will get two roles: There are three reasonable combinations:

- Merge the Situation Evaluator and Win Checker
- Merge the Situation Evaluator and the Move Finder
- Merge the Move Picker and Win Checker

Variation for six players

With six players, introduce the additional role of the observer, who takes over the recording tasks of the Situation Evaluator and Win Checker. That is, when the machine wins or loses, the observer writes down the results and crosses out the bad moves in the game cards and move overview as necessary.

Variation by adapting the game system

Reinforcement Learning

In general, this means that actions that lead to desired results are strengthened, while actions that lead to undesired results are weakened. The machine in "Human vs Machine" uses a very simple procedure of Reinforcement Learning to learn its strategy. That procedure requires data, in this case, plays of the game. The amount of data (the number of games played) and the quality of the data (in this case: playing strength of the human player) are both important for the machine's learning process. Unsuccessful moves are strongly weakened by taking them out. The probability that they will be drawn again in the future is then 0. "Human vs Machine" does not use positive reinforcement, that is, it does not reward successful moves by increasing of the probability for drawing them. It generally takes at least 10 rounds for the machine to get through a learning progress. An optimal model is reached after the machine loses ten to 15 times.

Game variation by the integration of positive reinforcement

In this game variation the positive actions of the machine are strengthened by doubling the probability for successful moves – thus the machine is getting a "reward" for successful moves alongside the "punishment" for unsuccessful moves. To make this work, the Move Finder circles the color of the last move on the situation card and the situation overview, when the machine has won. Circled moves identify the successful moves. The Move Finder provides two color cards for each circled move and one color card for other moves indicated on the situation card. However, the machine´s learning process needs more rounds than in the originally variation, because this machine wins more often.

Game variation: fast learning by using symmetric game situations

As in the game instructions stated, moving the left piece in the first move of the human is equal to moving the right piece. In this game variation, you will use the symmetry of the game situations even more now. If students explore the possible game situations more deeply, they will find there are five symmetric (mirrored) situations and therefore there are five pairs of symmetric situation cards (e.g. C5 and C9). You can remove one of each pair, for example, cards C8, C9, B11, C10 and C11 so that you only play with 19 situation cards.

This accelerates the computer´s learning progress because duplicates of unsuccessful moves are deleted as well. This does make it harder to find the right situation card while playing the game, however, so you need to weigh the advantages and disadvantages to decide whether use this variation or not.

The "treat" version

Replace the color cards in the game with different snacks (e.g. fruits or sweets)! The treats should have the same size and they should feel the same. Every treat is placed on one of the colored boxes on the gaming board. To allow the Move Picker to draw the "colors" blind, the Move Finder places corresponding treats in a small bag or suitable container. Once a treat has been drawn, the color of the treat is checked on the game board and the corresponding move is made. The trick: If the game is played with treats, the colors of moves where the machine loses are simply eaten up. However, the move should still be crossed out on the move overview so that the learning progress of the machine is clear.

Impressum

Herausgeber
Bundesministerium für Bildung und Forschung
Referat Wissenschaftskommunikation;
Wissenschaftsjahre
10117 Berlin

Idee, Redaktion, Gestaltung
DLR Projektträger; familie redlich AG Agentur für Marken und Kommunikation / KOMPAKTMEDIEN Agentur für Kommunikation GmbH

Konzeption und Texterstellung
Universität Paderborn, Fachgruppe Didaktik der Informatik mit Unterstützung der Fachgruppe Didaktik der Mathematik

Druck
Druck- und Verlagshaus Zarbock GmbH & Co. KG

Stand
September 2019

Die Wissenschaftsjahre sind eine Initiative des Bundesministeriums für Bildung und Forschung (BMBF) gemeinsam mit Wissenschaft im Dialog (WiD). Sie tragen als zentrales Instrument der Wissenschaftskommunikation Forschung in die Öffentlichkeit und unterstützen den Dialog zwischen Forschung und Gesellschaft.

„Mensch, Maschine!“ ist ein gemeinsames Projekt von:

Bundesministerium für Bildung und Forschung

Deutsche Telekom Stiftung

UNIVERSITÄT PADERBORN