This is the translation of: W. Balzer 2021: Was weiss ein Bot?, in: Brandl, J. L., Messelken, D., Wedman, S. (eds.), *Denken. Reden. Handeln – Thinking. Talking. Acting*, Open Access Publikationsserver der Universität Salzburg (ePLUS), ISBN 978-3-200-07526-9 (PDF), 133 -146.

What a Bot knows?

Wolfgang Balzer, LMU München, MCMP

At the moment, very few people own a bot.[1] But this will certainly change. People in mass societies are becoming lonelier. They increasingly live in the world of Internet. In some societies, more and more old people buy a bot – if they can afford it. Young people in such societies prefer at the moment to keep a pet. This, too, could change. Bots are becoming cheaper and easier to keep than animals.

In the original meaning, a bot was simply a 'working servant'. Today, there are many types of bots used in different applications (Baumunk, 2007). For example, there are bots for commodity production, for car control, for caring for the elderly, or for killing machines.

A bot is an artificial being that has many human characteristics and abilities, and can establish and maintain many relationships with people and other things. Here I would like to explore what a bot knows about living together with humans in a practical way.

Bot's Body

The bot has sensors, 'sense organs'. It can see, hear and touch similar to a human. So far, it can hardly smell and taste. The bot is often similar to a human from the outside. It has a body, two legs, two arms, two hands, a head, two eyes, a mouth and two ears. Producers can design the bot to look very much like a living person. The bot can walk upright and

[1] We use here the term 'bot' in the same meaning as the term 'robot'.

also run – though he still does both somewhat clumsily. He also does not (yet) move its hands and arms very elegantly.

Te pote can neither eat, digest, nor sleep in a natural way. In the bot, no food that is suitable for humans can be converted into movement. The processes of eating and digesting in humans work differently in the bot. The bot is 'fed' with electrical energy, which it charges into it's batteries. In its body, energy is converted into motion with electric motors. A bot does not need to sleep and without sleep it does not forget anything – unless it is supposed to forget something.

When one of its batteries is empty, certain movements can come to an abrupt halt. When all batteries are almost empty, she is usually still put into 'stand by mode'. She 'herself' has no possibility to become active again. While humans, at least in principle, can feed and maintain themselves on their own, the bot needs an 'outlet' and an owner to turn him on and off.

Some of the bot's movements are executed locally, some are controlled centrally. Many movements consist of multiple sub-movements. The more important a movement is to the bot, the more is it controlled centrally. The control center consists of a processor and a memory. In other words, the 'bot's brain' consists of a computing machine. Today, smaller processors and the memory also distributed in the bot in various places. Some smaller processores also run in parallel. This system is constructed in the same way that the human neural nervous system (brain included) and that the human body.

Movements create new processes – and thus events that did not exist before. Also the movements of a bot lead to many other events and some of them can be considered as effects of the bot. An effect can affect nature, or other people and/or bots. Often the effects remain mainly in the domain of speech. An utterance is heard and understood. Also talking has an effect. The bot does something, it is active, it changes the world, it acts.

Bot's basic functions

Through it's sensors the bot picks up stimuli – *patterns* –, which are represented in it's body by *terms*. In it's central computer, from the beginning, there is a store of terms that can be used in it's computations.

Some of these terms represent patterns, others are for other purposes, and still others are freely available as variables. Terms of the first kind I call *active* terms. By a computer program, by the active terms and by the body of the bot, movements are generated.

In the bot, *lines* (connections, paths) are established, which express abbreviations between terms. These lines lead from *nodes* to other nodes. Each active term is abstractly attached to exactly one node. Nodes and lines form a *net* and if the lines are transitive and not circular, such a net is a *Bayesian net*. In a bot, if a particular term is used more often than another term, the first term represents a pattern that is perceived with greater (subjective) probability than a pattern represented by the second term.

In the bot, visual, tactile, and sound patterns received by the sensors are filtered, compared, and classified. The bot picks up patterns through it's sensors, associates them with terms ('signs', 'representations'). Terms which already are bound to other patterns in other connections can also be untied from those.

The bot is already provided with some *elementary* terms which are incorporated already in it's construction. From elementary terms further *possible* terms can be constructed. All these terms can be represented in different levels. An elementary term can be represented as a bit or as a system of bits or by some other symbol. In Indo-European languages, for example, letters, words, and sentences are used as symbols.

A probability is assigned to each active term of the bot. The bot uses special terms (numbers) that express – from a human point of view – probabilities. Such probabilities change over time. As the bot perceives new patterns, the associated probability of a term changes.

Today there are many different learning algorithms, learning methods, and learning programs, and many types of machine learning. Most of the learning programs are based on probabilistic Bayesian networks, keyword: '*Bayes network*' or '*Bayes statistics*'.[2]

In a learning program, a 'new' pattern is compared to an already existing set of patterns. More precisely, there is already a term in the bot that represents a certain set of patterns, and also 'some' elements from this set must already be connected with terms. The bot perceives the new pattern and assigns a term to it. This term is compared with

[2]See (Koch, 2000), (Riguzzi, 2018).

other terms that are already active.

In a first case, where the term which was just assigned, matches one of the already active terms, the bot inserts this term into it's system of active terms. All – and exactly these – active terms represent patterns and collections of patterns. In a second case, if the new term does not match other active terms, the term is detached again from the new pattern. It is again available for other calculations.

In the first case, the probability of 'the term' which represents the collection of patterns, will increase. Somewhat simply said, the frequency of occurrence of patterns of a certain kind is increased. The term for a given set of patterns 'becomes' more probable. In the second case, the probability of term under discussion decreases. In both cases the probabilities of all affected nodes in the network are 'updated' by the Bayes statistics.

In other words, the bot can identify and rank patterns. It binds patterns and collections of those to terms and classes of terms. It constantly expands it's system of terms and it changes the probabilities associated with it's terms. He learns.

For example, the bot compares a pattern of color pixels arriving at it's optical sensor with other patterns it has already received when the bot was built or that it learned over time. By touching, it can learn whether a real object it has 'at hand' is spherical or has corners. By hearing, he learns to distinguish a sound, such as 'oh' from other sounds. Whether the sound comes from a human or not is still difficult for him to decide.

Today, the bot has no problem forming pairs of terms. Therefore, it can also construct classes of pairs. However, it uses different algorithms for this than for the formation of pattern classes. In a further step, the bot can also create classes whose elements can come from different levels. For example, she can form a relationship between a pattern and a class of patterns. However, this requires that she has already perceived many patterns of a certain kind. She must have already learned something.

Through many repetitions, the bot forms a large system of active terms. This system represents a set of pattern classes. An active term, which represents a totality, is not simply stored as a list of terms. The bot also gives this totality a new 'name'. It activates a term that was not active at this moment. A class of patterns is represented by the term and the frequency of occurrence of patterns from this class by a probability.

Besides the three *basic patterns* (tactile, visual and phonetic patterns) there are also *complex patterns*. These consist of systems of different basic patterns. A complex pattern can contain all three basic patterns. The basic patterns are 'woven' into complex patterns. Complex patterns are not independent of each other. For example, a complex tactile pattern may also rely to visual constituents and/or phonetic constituents.

The Psyche of the Bot

A complex pattern can have multiple components that come from different dimensions. The first component of a pattern may be a part of the bot's external world, and the second may be, for example, a part of it's internal life. Patterns that are found inside the bot cannot be perceived by it's sensors at first. Nevertheless, he can infer some parts of such a pattern. For example, he can perceive external patterns that he himself has caused. For example, he hears that (and how) he talks, how he feels, how he touches his hand with his other hand. And this he can also perceive purely visually.

What is going on inside the human being is represented by scientific models. In such a model, processes are triggered, generated, caused, which are not directly perceived in the affected person. The inner world of the person remains opaque. In a bot, on the other hand, we know quite precisely how the inner processes take place.

When describing the inner world of the bot, some of the same terms are used as for humans. The memory of a bot is divided into three types, as in humans: sensory memory, working memory, and long-term memory. The first, stores raw patterns in milliseconds, which in the positive case are passed on to the working memory. There the raw pattern is recognized as a real pattern and compared with other patterns. If the bot considers the pattern important, it associates it with an already existing term. In long-term memory, this term is finally incorporated into a complex system of terms and connected to many other terms in the Bayesian network.

In general, terms representing patterns and classes are now classified into four dimensions. Beliefs, attitudes (intentions), desires, and emotions are distinguished. To express these different entities in a general way, the notion of *event* has become common. The world consists of events (Balzer & Brendel, 2019).

Like humans, the bot can also form beliefs. Through many similar experiences that are *quasi* repeated, he binds a pattern class to a 'new' term and to a probability. He relates this term to other terms that have been bound to other pattern classes, thus changing a set of probabilities. When the probability of the new term has become large enough, the associated term acquires a new status, which can also be expressed linguistically. All these beliefs (terms) form a large system, a network. Until now, however, the beliefs of a bot and it's net are structured in a rather simple way.

Intentiones, desires and emotions are only partially implemented so far and are not formed so often in the bot. The bot can have, for example, the intention to reach a certain goal. In the car industry, a bot's 'intention' to fix a special part of a car is already technically present. In other applications, the development of a bot's intention becomes more difficult if complex plans and goals come into play. The bot can also have wishes, but at the moment only wishes which are simply structured. For example, she says that her battery is almost empty, or that she always prefers products from company XYZ. She can also have simple feelings. For example, she can feel temperature differences and material pressure.

As long as the bot is not switched off, it constantly discovers new patterns. She represents them by terms and relates them to the already existing terms. As the central memory of a bot today is as large, or even larger, than the human brain, the bot always has at hand a 'free' term with which he can expand it's belief system.

In this way the world of possibilities is opened for a bot. The bot begins to form 'it's' inner world, it's psyche.

The Language of the Bot

A human being has an unlimited supply of signs (words, phrases, sentences, pictures, etc.), which he generates and uses especially for reading and writing. In the bot, this works in a similar way.

The patterns discussed are also used in learning a language. More specific types of patterns come into play in the learning of language. The bot learns to distinguish whether a pattern that is 'new' for him comes from a human (or a bot) or from 'nature'. Further, he learns whether an incoming pattern is used mainly for speech and communication, or for other purposes. Third, he learns to distinguish whether an utterance (which is regarded as a pattern) is addressed directly to the bot or not (Meggle, 2010). Complex patterns usually consist of a mixture of tactile, visual, and phonetic patterns. For example, the bot sees 'its master' standing directly in front of her, she hears a sound sequence that probably comes from the mouth of her owner, and 'her' hand is grasped by the hand of the master.

All these patterns, which are constantly repeated, are represented by associated terms and probabilities. The patterns that emerge during listening and speaking lead to a separate domain: the world of the bot's language. Some words, phrases (in the technical sense), and commands, as well as some grammar rules, already exist in his language as terms. Such a term is used both for something perceived and at the same time (or at another time) for a message to the bot. In (Balzer, Kurzawe, Manhart, 2014), the term *inel* was introduced for this purpose, which comes into play in several ways. First, the term is learned, second, a connection is drawn between the term and the object. And third, the term is used as a medium for messages in information transfer and communication. In the case of Indo-European languages, the term is recognized as a noun or a name. In the case of verbs, such a recognition is not so easy. For verbs, classes of phonetic sequences and visual and tactile sequences must be identified and learned.[3] In principle, the bot can also learn a verb via class formation.

In a natural language there are different rules to combine terms from a language, to more complex terms. When using such a rule, a relation of two terms is perceived, where this relation concerns both external reality and speech. Between two non-linguistic 'things' already associated with two terms, a new relation is discovered, perceived, and assigned. A new term and probability is developed by many repetitions. Such a relationship is then also represented by a term in the bot's language domain.

More rules are added in the computer over time. Some of these rules only refer to relationships of linguistic entities. Terms representing relationships can also be learned with Bayesian programs. Such terms can be used to form sentences and similar constructs that are known in the

[3]It is difficult to assess how far these constructions have progressed, since the best bots are constructed in the armor sector and are thus subject to the greatest secrecy.

natural languages. For example, for the English language, there are several approaches about the formulation of rules by which sentences and similar constructs can be composed or be demarcated (Chomsky, 2002).

In the reduced 'language' of the bot, terms can be translated into words, sentences and commands. In this way, the bot can express its beliefs linguistically. Through rules that apply to a language, and through further formal rules that come from logic and computer science, the bot can form systems of beliefs. In computer science, this involves the adaption of the rules of logic, which are infinitely repeatable, to finite systems.

In a computer language, in addition to sets or lists of terms, another kind of elementary term plays a role that does not exist in natural language, namely the *variable*. In computer languages – and in logic as well – the function and deconstruction of 'new' possible terms, sentences, and situations is represented by variables. In natural languages, on the other hand, there are no words that indicate a particular blank space in a sentence. Among the elementary terms, which are already given to the bot during construction, special terms play a central role, which are used in bot's language as 'possible names' or names for variables.

With each kind of bot another variant of language is used. But in each of these 'languages' further terms are formed with different rules from existing terms. These rules can be from different kinds. Some of these rules come from logic, for example, the construction of conjunction, adjunction, etc. The handling of negation and sentences beginning with 'for all', on the other hand, is handled differently in computer languages than in logic – this also depends on the particular programming language used.

Every bot has a wide range of mathematical construction rules, which were already given to her when she was produced. In humans, these are learned – or not. The bot, on the other hand, can immediately distinguish natural, integer, and real numbers – as long as those numbers do not get too big. The bot can add natural numbers, for example, but also 'simple' real numbers. Most people do not have an algorithm 'at hand' to calculate real exponentials, for example; the bot can. Meanwhile, the bot also uses statistical rules that come from Bayesian statistics. Certain types of probabilities are updated in different ways. In humans, these rules are developed in their evolution. How these rules work, most people do not know and this is similarly true for bots.

Every 'real' bot uses language rules in addition to logical and mathematical rules. These rules stem from logical, mathematical and statistical rules, but they usually also depend on language constructs of the particular natural language in which the bot has to learn to speak. For example there are constructions with 'if' and 'not' in the Chinese language that are not present in the US language – and vice versa.

The Bot and it's Group

The bot does not 'live' alone; it belongs to at least one group of people. In the following, for simplicity, we also refer to bots as persons. So the group contains at least the bot and a natural person – who is also the owner of the bot. Often, also other persons belong to the group. All of these speak the same language that the bot uses in a reduced way. The bot communicates with others. He receives instructions, but he may also express an instruction to others.

In the learning phase of the bot, she must first learn the language of the group. By talking and listening, she increases her vocabulary. If she is intelligent enough, she can also learn language rules. In this environment, there is a lot of talking, listening and, in a sense, understanding. She understands a sentence, which she heard, if the appertaining pattern enters her working memory and, in short, fits her inner world and is stored there. This does not mean that the bot can classify the sentence as right or wrong. One reason, why this does not function, is that many sentences have an ambiguous truth value. For example, is the sentence just uttered 'That was a murder' correct? This leads into the[4] ethics.

In each group, some practices exist that the bot must also learn. For example, in the evening he will turn off the light in a certain room at the right time, because that is just how it is done in the group. Or he will greet another person by a practice common in the group. These practices do not have to be explicitly described in the group. Nevertheless, they are valid in this group.

In addition to practices, the group also has rules that are formulated

[4] I use the word 'ethics' in the singular, but only because this is so common among philosophers. In the real world, there are various approaches, 'theories', about ethics: practical, normative, utilitarian ethics, more recently now 'machine ethics' (Misselhorn, 2018)

explicitly. These mainly have the purpose of keeping the members in the group. The members have something in common that other people do not have. If a person does not observe the group rules, he is – at least in the long run – excluded. She no longer belongs to the group. To 'practice' a rule does not mean that all members must also have read or understood this rule.

The rules lead into the world of norms. A norm is a rule that is constitutive for a group. Norms form a necessary condition to hold the group together. Normative rules do not usually adhere to the distinction between the natural world and the social world.

Rules that the bot must observe in his group would be, for example, the following. He must not cross a street when he is in front of a red light. He must swerve if a bicyclist is coming directly at him. He must stop his action of throwing a large iron bar downward because 'down below' a human has suddenly appeared. If the bot lives in a group of creationists, he must be convinced that the world was created only 6000 years ago – also this 'fact' has a very strong normative background.

Whether the bot itself can also join with other groups remains unclear at the moment. He will simply be accepted. Whether he can also leave the group depends very much on the rules of the group and on his beliefs formed up to now.

All of this leads to moral questions that must also be asked of the bot. However, this paper is not about humanity or human rights in general. We are talking here only about rights and duties, which regulate the norms *of a certain group*. In other words, the bot is not a supermoralist who must consider for every possible action whether the action contradicts a philosophical approach about ethics.

The norms discussed here are formulated rather asymmetrically so far. For example, the bot may not harm or even kill any member of 'it's' group. This, of course, *not* applies to people from other groups, for example, to 'the enemy' in war. The bot should also pay attention to the owner's beliefs. The bot has to report if he 'lacks something'. He has the right to observe a person who does not belong to his group. The owner has other rights. He can violate and shut down ('kill') his bot. He does not have to discuss his beliefs with the bot. He does not have to tell the bot how he feels all the time and he does not have – at least in many existing groups – the right to watch every stranger all the time.

These formulations are currently at the bot's limit of comprehension.

How far sentences (terms) concerning rights and obligations and containing operators are comprehensible for the bot can only be discussed for the individual case today. It's understanding depends on the quality of the communication and the time the bot has been active so far. If he has generated many language elements of his own and formed many beliefs, he may also be able to distinguish and 'understand' terms that contain deontic operators. For example, can he distinguish whether a person is dead or alive? Or, if he finds the verbs 'hurt' and 'eat' in his long-term memory, can he distinguish whether he hurt or fed a member of the group? If he finds a norm that says that the bot must not hurt a group member, this usually leads to further activities that were already build in to the bot at it's construction.

The Knowledge of a Bot

When the probability of a set of patterns has become 'large enough', as said before, a new, not yet used term is assigned to this set, which is also included in the bot's belief system. When a belief is shared across the group, we speak of *knowledge*. A belief is shared in the group if every member has that belief, if every other member has that belief, if everyone is convinced that this belief is also present in the other members, and so on (Balzer & Tuomela, 1999).

The bot shares this knowledge. Whether she also critically examines this knowledge is not implicated. After the first period, in which the probability of a particular term has become large enough for the bot to insert it into it's language, it can also use filtering and checking methods. Whether she does so depends heavily on the bot's constructors.

Just as a human pushes aside many sensory patterns, after a while, the bot will also block more and more patterns she receives from the outside and not pay further attention to them. This happens for two reasons. First, she has already perceived and also incorporated so many patterns of this type that the probability for this pattern class can only marginally increase. It is also no longer worthwhile for the bot to activate the elementary update mechanism. The second reason would be that the pattern does not match any of the patterns she has perceived in her life so far and that the pattern is not important for her in the given situation.

However, this does not mean that a bot gives up knowledge that he has acquired. In order to give up knowledge, there must be structures and rules that make this possible at all.[5] The bot does weigh which incoming patterns it can set aside. Terms that he has already stored as knowledge are not abandoned in this way, however. The revision of knowledge (Rott, 1991) digs into the depth of the overall system of beliefs. To deactivate a single known term, the beliefs of the rest of the members of the group may also need to be changed.

Once the bot has formed a system of terms, relations,[6] and paths between terms, it can also, in some sense, give reasons for using a particular term.[7] Such reasoning is done by a separate program module, which leads to different results depending on the constructor, which also depends on whether probabilities are used.

A 'good' reason for a term from the bot's knowledge system entails two components: a set of active terms and the trust of members. Between the terms used in a reasoning, there are network relations; paths that lead from one term to other terms. An important condition assumed in Bayesian networks involves circularity. When the bot needs to reason about a term, she looks for paths that all start from the term. If there is a circular path, the reasoning may not find an end. In Bayes nets this is excluded.

Trust in a group is difficult to describe so far and even more difficult to explain. How far can a person trust a member? This leads to many psychological, political, economic questions and the question of power (Balzer, 1993). In the bot, this leads to further modules that represent parts of the inner human world, which I can only briefly outline here, see e.g. (Balzer, Kurzawe, Manhart, 2014).

When the bot did pass it's first learning phase, he can point to something and utter the associated word 'this'. We now look at the group, at the owner M and the bot R. Let us assume that both R and M are looking at the same picture. M says 'this is my house' and R is said to utter in a similar way. In this case, the visual patterns are very similar for both individuals. If both subjects come into many similar situations where similar images are seen, R can use a Bayesian learning algorithm. If the house is not only in the picture but also real, M will say 'this is

[5]Such rules are difficult to enforce for the Internet, for example

[6]A relationship is also represented by a term.

[7]In declarative computer languages, such as *Prolog*, this is easily accomplished by 'backtracking'.

correct'. In this case, M refers to it's corresponding gesture. In the class of these situations, R will most likely classify the sentence 'this is my house' as 'correct' in his language.

But what happens if the house in the picture is not similar to what is seen in real life. If M sees an entity that is not similar to his house, he will say 'this is not my house'. In this case, the bot must be able to distinguish whether he sees two houses or one house and, say, a garage. I assume that he has already learned such distinctions.

It becomes more difficult when M is to distinguish the house and the house seen in the picture *only* by the word 'my'. In the picture, M's house (*his* house) can be seen. However, the real house he sees directly does not belong to him. There are also such cases in sufficient number. The owner of R must look more closely at the surroundings of the house. Even this is not always sufficient to distinguish 'mine' from 'not mine'. For example, if the group (M, R and others) lives communistically, M may draw the distinction between 'mine' and 'yours' differently.

In the social world in which the group lives, M can turn a belief into a truth through his experience. R can do this as well. Through his experience, he forms a class of terms represented by the word 'knowledge'. R knows that a term (for example a sentence) is correct for him, if for R the (subjective) probability for this class of experiences has reached a certain size and if this probability is similarly high for the other members of the group.

But what does R do with sentences that are correct in a first group but not in another group? For example, M utters the sentence 'I should not eat pork', whereas M and R live in a Christian group. The problem obviously lies in the normative operator 'should'. In R's first learning phase, he cannot do much with such operators. R has to form new classes of different levels. R has to learn, for example, that M treats R with care, that R does not kill M's dog, or that M does not eat toller cherries from his tree. These sentences refer to concrete events that are happening here and now. Examples in which at least one event class is used would include 'bots should be treated with care', 'Christians should not kill animals', or 'apples from the paradise tree should not be eaten'. The word 'paradise tree' can only be learned by R through many different text examples.

Other sentence examples that transcend the group of M and R can be found in large numbers. In order to incorporate a norm-laden sentence into R's knowledge system, R must first ask itself which norm system, which moral system, is used in a person's utterance. Depending on the norm system, the sentence may be true or false. The meaning of the sentence often depends on the surrounding text. For example, R as an observer will consider the sentence 'I should not eat pork' to be wrong in a Bavarian restaurant, but right in Mecca. **A Philosophical Outlook**

The issues discussed here are mainly discussed from the point of view of an owner or user of a bot. From this point of view, it is 'only' about the knowledge of the bot, which is the property of M. In this case, it will probably be said that the knowledge of the bot depends very much on M and on M's group in which M lives. M may well influence the beliefs of his bot. The converse is only true to a small extent to date. However, M is also willing to give up a certain belief if his bot tells him new information from a reliable source that M did not know yet and that makes a belief of his quite unlikely.

What *knowledge for all* could mean was not discussed here. In particular, there was no discussion of what it might mean for an owner of the bot to be an anonymous or legal person. Nor was there any discussion of what it might mean, property-wise and ethically, for a pure Internet actor to be the owner and the guilty party. There was no discussion of when a bot can, may, or must do something in common with other bots. Finally, military bots could not be discussed; they remain top secret.

All of these questions are relevant today. As it looks now, in the future, people will probably be influenced by bots as well. Some first steps can already be seen. First big 'influencers' and 'changers' are well known. The question of power could soon also be asked for beliefs in general. In the most extreme case, this could end with the question: Are humans perhaps only kept in a reservation for scientific reasons, which the bots maintain?

Bibliography

Balzer, W., Moulines, C. U., Sneed, J. D. 1987: *An Architectonic for Science*, Reidel, Dordrecht.
Balzer, W. 1993: *Soziale Institutionen*, de Gruyter, Berlin/New York.

Balzer, W. & Tuomela, R. 1999: Eine Theorie des Geneinschaftlichen, *Facta Philosophica* 1, 55 - 76.
Balzer, W., Kurzawe, D., Manhart, K. 2014: *Künstliche Gesellschaften - mit Prolog*, V & R unipress, Göttingen.
Balzer, W. & Brendel, K.R. 2019: *Theorie der Wissenschaften*, Springer VS, Wiesbaden.
Baumunk, B.-M. 2007: *Die Roboter kommen. Mensch, Maschine, Kommunikation*, Wachter Verlag, Heidelberg.
Chomsky, N. 2002: *Syntactic Structures*, Mouton de Gruyter, Berlin/New York.
Koch, K.-R. 2000: *Einführung in die Bayes-Statistik*, Berlin etc., Springer.
Meggle, G. 2010: *Handlungstheoretische Semantik*, de Gruyter, Berlin/ New York.
Misselhorn, C. 2018: *Grundfragen der Maschinenethik* (3. Aufl.), Reclam, Stuttgart.
Riguzzi, F. 2018: *Foundations of Probabilistic Logic Programming. Languages, Semantics, Inference and Learning*, River Publishers Series in Software Engineering, http://ml.unife.it/plp-book/
Rott, H. 1991: *Reduktion und Revision*, Frankfurt/Main.