

Computational Modeling of Emotion: Towards Improving the Inter- and Intradisciplinary Exchange

Rainer Reisenzein, Eva Hudlicka, Mehdi Dastani, Jonathan Gratch, Koen Hindriks,
Emiliano Lorini, and John-Jules Meyer

Abstract—The past years have seen increasing cooperation between psychology and computer science in the field of computational modeling of emotion. However, to realize its potential, the exchange between the two disciplines, as well as the intradisciplinary coordination, should be further improved. We make three proposals for how this could be achieved. The proposals refer to: (1) systematizing and classifying the assumptions of psychological emotion theories; (2) formalizing emotion theories in implementation-independent formal languages (set theory; agent logics); and (3) modeling emotions using general cognitive architectures (such as Soar and ACT-R), general agent architectures (such as the BDI architecture) or general-purpose affective agent architectures. These proposals share two overarching themes. The first is a proposal for modularization: deconstruct emotion theories into basic assumptions; modularize architectures. The second is a proposal for unification and standardization: Translate different emotion theories into a common informal conceptual system or a formal language, or implement them in a common architecture.

Index Terms—Computational emotion modeling, affective computing, emotion theories, theory formalization, agent logics, cognitive architectures

1 COMPUTATIONAL MODELING OF EMOTION AS AN INTERDISCIPLINARY PROJECT

1.1 Definition and Aims of Computational Modeling of Emotion

As used in this article, *computational modeling of emotion* refers to attempts to develop and validate computational models of human emotion mechanisms. Computational modeling of emotion is an interdisciplinary endeavor between, in particular, psychology and computer science [135]. The goals of computational modeling of emotion largely correspond to the general goals of AI (e.g., [134]), when these are restricted to the domain of emotions: To achieve a better theoretical understanding of emotions in natural (specifically human) and artificial agents by creating computational models of them; and to enrich the architecture of artificial agents with emotion mechanisms similar to those of humans, and thus with

the capacity to “have” emotions.¹ Computational emotion modeling projects that have the first goal can be called *theoretical* and those that have the second goal *applied* [31]. Psychologists are typically more interested in theoretical and computer scientists in applied computational emotion modeling; however, some researchers in both disciplines pursue both goals. In fact, the two goals are closely linked: Endowing artificial agents with truly human-like emotion mechanisms presupposes reasonably faithful computational models of these mechanisms (e.g., [68]); conversely, one (some would say the best) way to attain a deep theoretical understanding of mental processes, including emotions, is to try to synthesize them in artificial agents (e.g., [25]).²

- Rainer Reisenzein, Institute of Psychology, University of Greifswald, Franz-Mehring-Str. 47, 17487 Greifswald, Germany. E-mail: rainer.reisenzein@uni-greifswald.de
- Eva Hudlicka, Psychometrix Associates Inc., PO Box 9409, North Amherst, MA 01059, US. E-mail: hudlicka@ieee.org
- Mehdi Dastani, Department of Information and Computing Sciences, University of Utrecht, PO Box 80.089, 33507 TB Utrecht, The Netherlands. E-mail: M.M.Dastani@uu.nl
- Jonathan Gratch, Institute of Creative Technologies, University of Southern California, 12015 Waterfront Way/Playa Vista, CA 9009, US. E-mail: gratch@itc.usc.edu
- Koen Hindriks, Delft University of Technology, Department of Intelligent Systems, PO Box 5031, 2600 GA Delft, The Netherlands. E-mail: K.V.Hindriks@tudelft.nl
- Emiliano Lorini, Centre National de la Recherche Scientifique Toulouse, Université Paul Sabatier, 118 route de Narbonne, 31062 Toulouse, Cedex 9, France. E-mail: Emiliano.Lorini@irit.fr
- John-Jules Meyer, Department of Information and Computing Sciences, University of Utrecht; and Alan Turing Institute Almere, PO Box 80.089, 33507 TB Utrecht, The Netherlands. Email: J.J.C.Meyer@uu.nl

¹Based on a functionalist view of mental states (e.g., [24]), we propose to understand this capacity as the capacity of artificial agents to have internal states that are functionally equivalent or at least similar to emotions in humans, i.e. that play causal roles in the agent architecture that mimic those played by emotions in humans. The causal effects of emotions that can be modelled in artificial agents include self-awareness of emotions when they occur (e.g., [20], [135]); however, no claim is made that artificial agents, at least those that currently exist, are conscious of their emotions in the sense of having qualitative phenomenal experiences (see [154]).

²An additional goal of affective computing is to provide artificial agents with a theory of mind module that allows them to recognize, explain, predict, and conceive of ways to influence, the emotions of other agents ([135]; see also [152]). Computational models of emotion can be helpful for this purpose, too: Declarative versions of these models (possibly in a simplified form) can be used as components of the theory-of-mind module of artificial agents; alternatively or in addition, agents can be programmed to use their emotion mechanisms in “simulation mode” to predict the emotions of others (as well as their own) (as realized, for example, in Pynadath and Marsella’s [139] *PsychSim* agent).

The close connection between theoretical and applied computational emotion modeling is one reason why the past years have seen increasing cooperation between psychology and computer science in research on emotions (see e.g., [11], [78], [168]). Some see in this development a revival of the cognitive revolution of the 1960ies (e.g., [116], [119]) in the cloak of affect [67]. While this may seem to be a grand claim, there is at least one sense in which it seems justified: Because emotions are thought to be primarily useful for resource-bounded autonomous agents acting in dynamic and only partially known environments (e.g., [118], [135], [100], [179]), the computational modeling of emotion requires returning to the original goal of artificial intelligence to create “complete” artificial agents (see e.g., [89], [164]); and likewise, a refocusing on the original goal of cognitive science, to create a unified theory of the mind [121]. The cooperation between psychology and computer science in the field of emotion therefore has the potential to contribute to unification, not only in emotion psychology, but in psychology at large.

1.2 Realizing the Potential of Computational Emotion Modeling

However, to realize its potential, the exchange between as well as the coordination within the two disciplines should be further improved. As to the *interdisciplinary* exchange, the export of theories has so far been mainly from psychology to affective computing, rarely in the reverse direction (see [31]). For example, as Broekens [31] notes, the *Handbook of Emotions* [94], a standard psychological reference work on emotions, contains virtually no citations to the affective computing literature. This is so despite the fact (a) that psychologists frequently emphasize the need to better understand the computational mechanisms underlying emotions (e.g., [167]) and (b) that at least some of the emotion models developed by computer scientists can claim to be serious candidates for computational theories of human emotion, or at least to elucidate and concretize the psychological theories on which they are based (for supportive evidence see parts 3 and 4, and [31], [108]). Concerning the *intradisciplinary* situation, both disciplines could profit from more comparative efforts, coordination, systematization, and focus. Within affective computing research, there are now numerous computational models of emotions (for overviews, see e.g., [108], [78]), but this work is rarely cumulative or even comparative; rather, as noted by Marsella et al. [108], different researchers have tended to start anew from the same or different psychological sources (for notable exceptions, see [27], [53]). Partly as a consequence, the theoretical landscape in affective computing increasingly resembles that in the psychology of emotion, where numerous psychological theories of emotion coexist side by side (see e.g., [167], [155]), with few attempts to integrate or even to systematically compare them (for an exception, see [148]). From the psychologist’s perspective, affective computing researchers may even seem to further contribute to theoretical fragmentation by developing different computational models of the same psychological theory (as documented, for

example, by the many implementations of the OCC theory of emotion; [129]).

1.3 Aims and Overview of this Article

Our aim in this article is to make several proposals for further improving the interdisciplinary exchange between, as well as the intradisciplinary coordination within, psychology and artificial intelligence research in the domain of emotion modeling. A complementary perspective—in the form of a set of guidelines for prospective computational emotion modelers—is provided by Hudlicka [80]. In part 2, we make proposals for systematizing the psychological theories of emotion that typically serve as blueprints for creating computational emotion models. These proposals can also be helpful for systematizing and comparing existing computational emotion models. In part 3, we discuss possibilities of formalizing psychological (and philosophical) emotion theories in implementation-independent formal languages. Again, these languages can also be useful for formalizing (at least part of) the assumptions of existing computational models. In part 4, we consider options for modeling emotions in existing, comparatively general-purpose cognitive and agent architectures, including general-purpose affective agent architectures. Due mainly to the composition of the group of authors, our discussion focuses on *symbolic* models of *psychological* (specifically cognitive) emotion theories.³ Hence, we ignore the computational modeling of neurophysiological emotion theories (e.g., [8]) as well as attempts to implement psychological emotion theories in subsymbolic neural networks (e.g., [12]). Note, however, that the symbolic modeling of psychological emotion theories is by far the dominant approach in the field of computational modeling of emotion [78], [108].

2. SYSTEMATIZING PSYCHOLOGICAL EMOTION THEORIES

2.1 The Situation in Emotion Psychology

From the perspective of an applied affective computing researcher, psychology would ideally be at an advanced stage where *the ultimate correct theory of emotion* (UCTE) or even better, *the ultimate unified theory of the mind* (UUTM) had been attained. Furthermore, that theory would ideally be formulated as a computational model, or would at least be available in a format that lends itself readily to implementation as a computer program. Unfortunately, psychology has not yet arrived at this stage. First, there is not one emotion theory; there are many theories: Strongman [190] lists no less than 150 psychological and philosophical emotion theories that have been proposed during history, and these are still not all. To be sure, some of these “theories” are no more than vague hypotheses, there is much redundancy, and many theories are of historical interest only. However, even if one restricts atten-

³Although the views expressed in this article mostly reflect a majority consensus, agreement among all authors on all points is not implied. However, considering the many issues discussed in this article and the different backgrounds and theoretical predilections of the authors, it was encouraging to see how much agreement could be achieved.

tion to the currently discussed emotion theories, many do remain (e.g., [80], [155], [167]). Making matters worse, the exact relations between the different emotion theories are sometimes difficult to judge even by experts. Second, the existing psychological emotion theories are not formulated as computational models and most are not described in a sufficiently precise and complete way to be directly implemented as computer programs.⁴

Theoretical fragmentation and a comparative lack of precision are not peculiar to emotion psychology but are by and large characteristic of psychology in general. Psychologists often excuse these problems of their field by pointing to the complexity of psychology's subject matter and the comparatively young age of psychology as a science. There is truth to these claims. However, there are also other reasons for the present state of psychology: Partly as an after-effect of the behavioristic phase in psychology (about 1920-1960) with its associated positivist philosophy of science, a premium is still placed in today's psychology on empirical as compared to theoretical research [177].⁵ Furthermore, novelty and surprisingness of claims and findings are honored more than are attempts to integrate and unify. The recent interaction between psychology and computer science in the field of emotion has created a need to provide concise and precise summaries of psychological theories to nonspecialists who are at the same time accustomed to a high degree of precision. This need could provide an impetus to psychologists to try to clean up and organize their theoretical landscape. In this, they can be helped by computer scientists and logicians (see [41], [80], and part 3).

2.2 Towards Systematization and Unification in Emotion Psychology

There are several, not necessarily exclusive ways how this aim could be achieved. One possibility is to systematically reconstruct the existing emotion theories; possibly using a formal language (e.g., [204]; see part 3 for more detail). However, this approach, although appropriate for selected theories of particular importance, is probably too cumbersome as a general strategy. Not every theory needs to be precisely reconstructed; in particular, it is not necessary to formalize theories that have turned out to be (or at least, most likely are) false. One should focus on the promising alternatives.

An alternative, quicker and more viable strategy may be to (1) break up existing emotion theories into their component assumptions and (2) reformulate these as-

sumptions in a common conceptual framework (e.g., [80], [108], [147]). This common framework could be a formal language, possibly enriched with a set of special concepts (e.g., *belief* and *desire*) or it could be a general-purpose cognitive architecture (see parts 3 and 4); however, already semi-formal or informal theory reconstructions can significantly contribute to clarification and unification if concepts are used in a clear and consistent way [41]. The ultimate aim of this second strategy would be to reduce the many existing emotion theories to a smaller (hopefully *much* smaller) set of basic assumptions. These building blocks of emotion theories can then be separately discussed and refined. Furthermore, similar assumptions of different emotion theories can be more easily compared, vague assumptions can be more easily clarified, and conflicting assumptions more easily identified, when they are recast in a common conceptual framework (see also, parts 3 and 4).

In the remainder of part 2, we make proposals for breaking up and organizing the assumptions of emotion theories in a principled way. These proposals derive from a reflection on the basic task of emotion psychology. The mainstream position in today's psychology, at least its cognitive branch—which is not only compatible with, but in fact rests on, a computational view of the mind (see Footnote 4)—conceives of humans as autonomous cognitive systems, or intelligent agents, whose basic architecture emerged in evolution, but is significantly elaborated and possibly even transformed, by individual and cultural learning. This theoretical perspective implies that, to explain human behavior and experience, it is necessary to decode the structure and functioning of the human information processing system. For emotion psychology, this means that its central task is to reverse-engineer the structure and functioning of the human emotion system, including its connections to other subsystems of the mind (e.g., [155]; as mentioned in section 1.1, to understand emotions, it may in fact be necessary to reconstruct significant parts of the whole cognitive architecture). As already emphasized by McDougall [110], the task of reverse-engineering the emotion system comprises identifying the inherited versus acquired components of the emotion system.

2.2.1 On the Definition of Emotion

Before proceeding, a note on the definition of emotion is in order. Just as there is at present no generally accepted *theory* of emotion, there is currently no generally accepted *definition* of emotion—at least if by that one means a specification of the singularly necessary and jointly sufficient features of emotion. Causal-role functionalism about mental states ([24], Footnote 1) suggests that the latter is a direct consequence of the former [149], [157]. According to functionalism, mental states are defined (at least in part) in terms of their causal-functional role in the mental economy. Because this role is specified by a theory of emotion, a precise definition of emotion can only be given in the context of an emotion theory that is presupposed as correct. Hence, a precise definition of emotion is necessarily a *theoretical* (theory-based) *definition*. In addition, be-

⁴This is actually surprising if one considers that, since the cognitive revolution of the 1960ies, the psychological mainstream (at least in cognitive psychology) subscribes officially to the “information processing paradigm”, that takes information processing in computers as the guiding metaphor for understanding mental processes in humans. According to the information-processing paradigm, the task of (cognitive) psychology is to reconstruct the “rules and representations” that characterize mental processes [131]. In practice, however, the development of theories of mental processes that fully meet this description (i.e., computational models) in psychology is still restricted to a comparatively small community of experts (see e.g., [191]) for an overview) and only very few of them are working on computational models of emotion.

⁵One indicator of this “bias for the empirical” is that, different from most other sciences, psychology has no official theoretical branch: There is no official subdiscipline called *Theoretical Psychology* [177].

cause a theoretical definition of emotion stands or falls with the emotion theory that it presupposes, it is not a linguistic convention or stipulation about the use of the term “emotion”, but an empirical hypothesis about the nature of emotions [149], [157].

However, a precise definition of emotion is not needed to conduct meaningful research on emotions. All that is needed for this purpose is a rough, preliminary demarcation of the object domain “emotion”—what can be called a *working definition of emotion* [149], [157]. Minimally, such a working definition is already provided by a list of typical examples of emotion. We believe that, despite their differences of opinion regarding the true nature or correct theoretical definition of emotions, most emotion researchers are by and large in agreement on such a list. That is, most emotion researchers would agree among themselves, as well as with lay persons, that the objects of their inquiry comprise, centrally, those mental (or perhaps mental-behavioral) states of humans that in common language are denoted by words such as “joy”, “sadness”, “fear”, “anger”, “pity”, “pride” and so on. Beyond this, a working definition of emotion can include those features of emotion that are largely uncontroversial and therefore do not imply a decision for or against a particular emotion theory. Although there is perhaps no feature of emotion that has not been questioned by one or the other theorist, it is widely accepted (a) that emotional states (joy, fear etc.) are normally reactions to the perception, imagination or thought of certain things (typically events or states of affairs); and (b) that emotions have characteristic subjective as well as objective aspects. Subjectively, emotions manifest themselves, centrally, in characteristic phenomenal experiences or feelings, that appear to be directed at their eliciting objects: For example, one feels happy about the arrival of a friend, or experiences surprise about the election victory of a political party. Objectively, emotions manifest themselves at least occasionally in certain characteristic actions (e.g., flight or avoidance in the case of fear), expressive reactions (e.g., smiling in joy) and physiological changes (e.g., a rise of blood pressure in anger).

Although the proposed working definition of emotion is necessarily fuzzy, this is no objection to it. Just as providing a precise definition of emotions is the result rather than the precondition of theorizing about emotions, so too is providing a sharp demarcation of the domain of emotions, as well as a principled taxonomy of emotions (see also, [129], [150]).

In this article, we will however make one additional assumption about emotions: We will assume that whatever emotions are more precisely, they are *mental states* which are caused by, and in turn cause, other mental states and behaviors. Harmless and even self-evident as this assumption may seem to some readers, it is (at first sight at least) in conflict with those emotion theories that view emotions as *response syndromes* including both mental and behavioral components (e.g., [92], [166]; see [149], for a discussion). However, the syndrome theories, too, assume a generating process that maps (representations of) objects or events into one or several emotion-specific

mental states (e.g., appraisals, bodily feelings, action tendencies), which then, singularly or in combination, influence further cognitive processes and behavior. With respect to the syndrome theories, our term “emotion” refers to the emotion-specific internal states assumed in these theories.

2.2.2 Reverse-Engineering the Emotion System: “Horizontal” and “Vertical” Task Divisions

The task of emotion psychology—the reverse-engineering of the emotion system—can be broken down both “horizontally” and “vertically”: Horizontally, by following the temporal and causal flow of information from the perception (or an internally generated representation) of an eliciting object to emotion, and from there to the cognitive and behavioral effects of emotion; vertically, by studying the emotion system at different levels of analysis. In particular, one can focus on either the “intentional level” of system analysis or on the “design level” (Dennett, [51]; similar distinctions have been drawn by Marr [106] between the computational and algorithmic level, and by Newell [122] between the knowledge level and the symbol level).

The proposed *horizontal division* suggests three main questions for emotion psychology: (1) How are emotions generated? (2) What are their effects on subsequent cognitive processes and behaviors? and (3) related to both 1 and 2 (see below), what is the nature of the emotion itself? ([80], [108], [155], cf. 2.2.1).

The *vertical division*, between the intentional and the design level of system analysis, separates (1) theoretical descriptions of emotion processes in terms of what are ultimately common-sense psychological concepts (*perception, belief, desire, intention, fear, hope* etc.) from (2) theories that (also) refer to the internal representations (sentences in a language of thought, images, subsymbolic signals etc.) and computations on these representations, that according to the computational view of the mind underlie the intentional-level phenomena [51]. The term “intentional” in Dennett’s distinction expresses the idea that intentional-level theories explain the behavior of a system by attributing to it *intentional* (meaning here: *representational*) *mental states* familiar from common-sense psychology, such as beliefs and desires (e.g., [70], [171]): Internal system states that refer to or represent aspects of the world, including the own system, in particular ways (e.g., as being true in the case of belief, or as to-be-made true if possible in the case of desire). In adopting the “intentional stance” towards a system [51], we assume—either seriously or only as an “as-if” explanatory strategy—that the behavior of the system is controlled by representational mental states of this kind ([171], [70]).⁶

Simple as it is, the proposed classification scheme

⁶Note that the distinction between the intentional and the design level is orthogonal to the distinction between deliberate/conscious and automatic/unconscious mental processes referred to later in our discussion of cognitive versus noncognitive emotion elicitation. Both deliberate/conscious and automatic/unconscious emotion generation processes can in principle involve folk-psychological mental states such as beliefs and desires (e.g., [33]); and both processes can also be described in terms of representations and computations.

brings several important facts about psychological emotion theories into focus. Concentrating on the vertical division between the intentional and the design level of system analysis brings out that, as already mentioned, the majority of the existing psychological emotion theories are formulated on the intentional level (e.g., “fear is experienced if the person believes that an undesired event might occur”). Nevertheless, it should be noted that some psychological emotion theories (e.g., [62], [181], [166], [151]) provide at least sketches of the computational processes assumed to underlie the intentional-level phenomena. That is, whereas “pure” intentional-level theories of emotion only describe inputs and outputs of presumed mental computations, these “process theories” also make some assumptions about the computational processes that transform inputs into outputs (e.g., the nature of the underlying representations, the identity and sequence of the involved subprocesses; see [147]). As such, these process theories can play the role of intermediaries between intentional-level theories of emotion and full-fledged computational theories ([32], [100]).

Other important facts about emotion theories become transparent if one focuses on the proposed horizontal division, the flow of processing. First, quite a few historically important and contemporary emotion theories are primarily concerned with the question of the *nature* of emotion and say comparatively little (at least little specific) about how the different emotions are produced, and what useful or other effects they have (cf. [154]). These theories include James’s [82] theory of emotion and neo-Jamesian emotion theories (e.g., [86]); cognition-arousal theory ([16], [101], [165]) and “dimensional” emotion theories such as the pleasure-arousal model ([162], [163], [144]). The question that these theories mainly seek to answer is what kind of mental states emotions are, which is interpreted as meaning: From which building blocks, if any, emotions are composed, and how they fit into a general taxonomy of mental states. Specifically, are emotions mental states *sui generis* (e.g., centrally generated feelings of pleasure and displeasure) or can they be reduced, at least in part, to known mental states (e.g., feelings of bodily changes, a combination of cognitions and bodily feelings, patterns of appraisal or of beliefs and desires, or felt action tendencies; see [154])? Theorists interested in these questions have typically proceeded on the assumption that they can be answered without *detailed* theories of the generation and the effects of emotion (although *some* causal assumptions seem unavoidable). However, because of their restricted focus, the resulting theories of the nature of emotion are alone insufficient as psychological foundations of computational emotion models; they must be combined with theories of the elicitation and the effects of emotions. This becomes evident, at the latest, when one tries to use these theories to construct emotion mechanisms for artificial agents. For example, Becker-Asano and Wachsmuth [20] have implemented a version of the pleasure-arousal theory of emotion in an artificial agent; however, to account for the generation of emotions, they found it necessary to enrich pleasure-arousal theory with the OCC theory [129].

Second, even theories that are quite specific about emotion generation (in particular, some cognitive emotion theories) are sometimes relatively silent about the *effects* of emotions (cf. [80]). Again, if these theories are to be used as starting points for creating emotion mechanisms for artificial agents—which are usually desired because one hopes that they will enhance the functioning of the agents in some way—they must be enriched with assumptions about the effects of emotions (see e.g., [20], [68], [53]).

Interestingly, most theories of emotion generation seem to be compatible with rather different assumptions about the nature of emotions and their effects. For example, different cognitive emotion theorists have proposed that emotions are bodily feelings, feelings of pleasure-displeasure, mixtures of pleasure-displeasure and activation-deactivation, felt action tendencies, evaluative beliefs, patterns of belief and desire, various combinations of these states, and even syndromes including both mental and bodily reactions (cf. 2.1.1). Likewise, most theories of emotion generation seem to be compatible, in principle, with a variety of hypotheses about how emotions influence subsequent cognitive processes and actions: by creating hedonistic desires, by creating nonhedonistic action tendencies, by focusing attention or shifting goal priorities etc. (see 2.2.4). Furthermore, most cognitive theories of emotion generation seem equally compatible with a highly modular, “discrete emotions” view of the emotion system (e.g., [92], [161]) and with a fairly domain-general view (e.g., [150], [151], [163]). The compatibility of the assumptions of different emotion theories about the generation, nature, and effects of emotions may seem desirable from an engineering perspective, but it can be argued that it actually reveals a weakness of the respective emotion theories: Apparently, they are unable to impose strong constraints on what goes with what.⁷

2.2.3 Classifying Theories of Emotion Elicitation

2.2.3.1 Cognitive versus Noncognitive Theories of Emotion

The assumptions of the different existing psychological (and philosophical) theories about the process of emotion elicitation can be classified into two main groups, “cognitive” and “noncognitive”. Although the usefulness of this distinction has been questioned because of the vagueness and ambiguity of the term “cognitive”, it seems that a meaningful dividing line can still be drawn as follows: Cognitive theories assume that emotions (or at least a core subset of emotions) require certain “higher-order” mental representations, in particular, they require what are commonsensically called *beliefs* and *desires* [111], [70].⁸

⁷Closer investigation of the issue suggests, however, that there are at least tensions between the assumptions of different emotion theories about the causes, nature and effects of emotions. For example, it can be argued that, at closer look, a cognitive view of emotion elicitation does not fit comfortably with a discrete emotions perspective [148].

⁸Alternative terms used for “desire” in the emotion literature are “motivational” (e.g., [92]), “goal” (e.g., [129], [166]) and “concern” (e.g., [62]). Although all of these terms are meant to denote similar *motivational representations*, there are some differences between theorists about how, exactly, they are conceptualized. Here, we understand desires as motivational representations with propositional contents, i.e. as motivational representations of states of affairs [70], [171]. Desires include both performative desires (desires to do something) and nonperformative desires

In their more extreme forms, cognitive theories of emotion even assume that emotions *are* a form of cognition (e.g., evaluative beliefs, or belief-desire compounds), or contain beliefs or belief-desire compounds as essential components (see [37], [153], [154]). In contrast, “noncognitive” emotion theories assume that (a) certain kinds of emotions in a broad sense (e.g., sensory affects such as the hedonic pleasure evoked by the smell of a rose), or even (b) some prototypical emotions such as fear, are always (sensory affects) or at least sometimes (fear elicited by stimuli via a more direct route, that circumvents the mentioned higher cognitive processes. According to the most interesting interpretation of this hypothesis (e.g., [82], [128]), this “noncognitive” pathway involves only *nonpropositional* and possibly even *nonconceptual* mental representations, such as simple sensations, or images.

2.2.3.2 Noncognitive Emotion Elicitation in Computational Models

There is a longstanding and still ongoing discussion in psychology about the importance, and even existence, of “noncognitive” emotion generation (e.g., [91], [93], [189], [209]). This discussion has suffered, among other things, from the failure to distinguish clearly between the two described versions of the hypothesis of noncognitive emotion elicitation. Hypothesis [a], that certain kinds of affect (such as sensory pleasures and displeasures) are noncognitively generated in the described sense—i.e., presuppose only nonpropositional or even just nonconceptual representations—is intuitively plausible [148]. In contrast, hypothesis [b], that prototypical emotions such as fear, anger or joy can be (and perhaps even, often are) noncognitively caused (e.g., that fear can be elicited by the visual appearance of an abyss or by loud sounds, without any mediating thoughts), is more controversial. On closer inspection, the data that have been adduced to support this hypothesis turn out to be less convincing than is often claimed (see e.g., [151], [175], [173]). For example, it has been argued that noncognitive emotion generation is demonstrated by the finding that certain emotional stimuli elicit physiological reactions even if they are presented below the threshold of awareness (e.g., [128]; see [189], for a review). However, recent studies point to the possibility that the observed physiological reactions were shown by a subset of the participants who were still able to recognize the emotional stimuli at the presumed “subliminal” exposure durations (see e.g., [132]). Furthermore, even if it is a genuine phenomenon, emotion elicitation by subliminal stimuli could be mediated by automatized and unconscious cognitive appraisal processes (e.g., [174], [181]).

It should be noted, however, that regardless of how this issue is decided, the computational modeling of “reactive” emotion elicitation by specific objects or patterns of sensory stimulation is a comparatively easy task (e.g., [3], [53]). The real computational challenge is posed by the cognitively mediated emotions [68].

2.2.3.3 Cognitive Emotion Elicitation: Two Theories

Philosophers of emotion (e.g., [70]) have drawn a distinc-

tion between two basic kinds of cognitive emotion theories: *cognitive-evaluative* (or evaluative-beliefs) theories, and *cognitive-motivational* (or belief-desire) theories. The crucial difference between the two concerns the proximate mental causes (or, depending on viewpoint, the constituents; [154], [153]) of emotions: Are they only beliefs (including evaluative beliefs); or are they beliefs and desires? To illustrate the difference with a simple example, according to the cognitive-evaluative theory (which arguably comprises most psychological appraisal theories), John is happy about a state of affairs p if he believes that p is the case and believes that p is good for him (e.g., [9]). In contrast, according to the belief-desire theory, John is happy about p if he believes p and desires p (e.g., [37], [70], [126], [151]). Hence, according to the belief-desire theory, the evaluative beliefs (appraisals) featured in the cognitive-evaluative theory of emotion are in fact not required for emotions; all that is needed are factual beliefs and desires.

The distinction between the two theories rests on the assumption, which is widely shared in contemporary philosophy, that beliefs and desires are fundamentally distinct kinds of mental states that cannot be reduced to one another (e.g., [70], [171]). Beliefs are *informational* mental states: their function is to provide the agent with information about the world and the self. By contrast, desires are *motivational* states: their function is to provide the agent with goals, states of affairs whose realization is attempted under suitable circumstances (e.g., [70]; see also [143]).

Proponents of the belief-desire theory of emotion have argued that this theory avoids a number of objections to the cognitive-evaluative theory (see [70], [151]), including the objection that evaluative and factual beliefs are, as a matter of fact, neither necessary nor sufficient for emotions. A discussion of these arguments, which would require an in-depth comparison of the foundational assumptions of the two kinds of cognitive emotion theory, is beyond the scope of the present article. However, we would like to point out that it is possible to a large degree to “map” the assumptions of each theory into the conceptual framework provided by the other (although with a corresponding shift in meaning!). This remapping is possible because most appraisal theories assume that appraisals (evaluations of perceived events as good or bad) result from the comparison of perceived or believed events with goals [148]. Psychological appraisal theory can therefore be regarded as a special version of the belief-desire theory of emotion which assumes that the causal link between beliefs and desires on the one hand (e.g., the belief that p and the desire for p), and emotions on the other hand (e.g., the emotion of joy about p) is mediated by the appraisal of p [148]. Conversely, the belief-desire theory can be recovered from appraisal theory by linking belief and desire directly to the emotion. This can be achieved by either *dropping* the mediating appraisal from the theory, or by *reinterpreting* the appraisal (the outcome of the comparisons of perceived events with goals) as *being the emotion* (e.g., as a preverbal form of evaluation that is experienced as a feeling of

(desires that something should be the case) [195].

pleasure or displeasure; see [151]). Several recent formalizations of appraisal theories have in fact implicitly used this approach and thus have effectively reinterpreted appraisal theory as belief-desire theory (e.g., [1], [187]; see section 3.2). Analogously, some computational models of emotion (e.g., [20]) treat emotions as direct effects of goals and beliefs, without a separate, mediating appraisal state (see section 4.2).

2.2.4 Classifying Hypotheses about Emotion Effects

Compared to the broad consensus about how emotions are elicited, there is considerably less agreement in the psychological (and philosophical) literature about how, exactly, emotions influence subsequent cognitive processes and behavior. Many different hypotheses about the effects of emotions have been proposed (see e.g., [80], for an overview). To get a systematic grip on them, it is useful to distinguish, first, between *functional* and *non-functional* effects of emotions. The concept “functional” in this distinction is best explicated in evolutionary terms (e.g., [117]): Functional effects of emotions are those of their effects that were responsible for the selection of the corresponding emotion mechanisms in evolution, whereas non-functional effects of emotions are neutral, or even negative effects of emotions that are tolerated to reap the benefits of emotions (see also, [120]).⁹ To illustrate, one proposed functional effect of emotions is the focusing of attention on the emotion-evoking events (see below). A nonfunctional and potentially even harmful side-effect of this is a reduced capacity for engaging in other, simultaneous tasks. Ideally, only functional effects of emotions would have to be *explicitly* considered in computational emotion models, because non-functional effects would arise naturally as side-effects of the activity of the mechanisms that produce the functional effects (e.g., [160]).

Although many specific functions of emotions have been proposed (e.g., [63], [70], [85]), it appears that most of them can be classified into one of three categories (e.g., [130]):¹⁰

1. *Informational or epistemic functions of emotions.* According to this hypothesis, a main function of emotions is to provide adaptively useful information to other cogni-

tive (sub-) systems, including other agents (via nonverbal or verbal communication). The information in question is held to be in particular information about (1a) the occurrence of appraisals (e.g., [170]) or of changes in the person’s belief-desire system ([151], [157]) and (1b) related to this, information about the value of objects and events, including action consequences (e.g., [43], [111], [180]).

2. *Attentional or resource-allocation functions of emotions.* According to this hypothesis, emotions serve to shift the focus of attention to (or computationally speaking, allocate central processing resources to) their eliciting events, or to give priority to them in processing (e.g., [176], [151], [178]).

3. *Motivational functions of emotions.* This hypothesis claims that a main adaptive function of emotion is to reprioritize existing or to generate new goals or intentions (e.g., [62], [127]). With respect to the generation of new goals by emotions, the two main proposals are: (3a) emotions create hedonistic desires (desires to reduce or avoid negative, or to maintain or produce positive feelings; e.g., [18], [113]); (3b) emotions induce emotion-specific goals or action tendencies (e.g., fear causes the desire to flee, anger to aggress, pity to help) *directly*, that is, without the mediation of hedonistic desires (e.g., [62], [109], [161], [203]). A possible mediating mechanism for the latter effect could be that emotions draw attention to, or increase the salience or persistence of goals currently at stake (e.g., [3], [77], [118]; see also [145]). Hence, the motivational effects of emotions could be partly based on their attentional and informational effects [151].

2.2.5 A Note on the Relation between Emotion and Motivation

We end this part with a few comments on the relation between the psychology of emotion and motivation (including decision making).¹¹ Although there is general agreement that emotion and motivation are related and several early theorists have strongly emphasized their connection (e.g., [22], [109]), during the 20th century the fields of motivation and emotion psychology have to a large degree developed independently. As a consequence, some of the best-known psychological theories of motivation and decision making, such as the theory of planned behavior [2] or prospect theory [84], do not consider emotions, at least not explicitly. Although it can be argued that even in these formally affect-free theories, emotions are implicitly considered as a possible source of utility (i.e., as hedonistic desires, as explained in 2.2.4; see [10], [42]), the explicit recognition of emotions as separate variables in modern motivation and decision theories is of relatively recent origin (e.g., [113], [180], [203]). Likewise, 20th century emotion theorists have typically paid little attention to theories of motivation and decision making. That is, in developing their theories of the generation and

⁹ The assumption that emotions are overall adaptive is widely accepted in today’s psychology. The best general argument for this assumption is that it is hard to see how complex mechanisms like emotions could have emerged in evolution other than by natural selection, and hence because of their useful effects [see 49].

¹⁰ It has also been proposed that the main evolutionary function of emotions is to coordinate the subsystems of the organism into an adaptive whole-organism response to evolutionarily significant challenges and opportunities (e.g., [110], [127], [166], [196]). However, the available empirical evidence suggests that the components of emotional syndromes are in fact only loosely associated (e.g., [17]) and the observed clustering can probably be explained as the result of the three described functional effects.

Other documented effects of (some) emotions comprise the priming of emotion-congruent material in memory and the biasing of information processing in emotion-congruent ways (e.g., [7], [23], [170]); as well as physiological activation. Whereas physiological activation can be given a reasonably plausible functional explanation (it served to prepare fight/flight reactions in evolutionary predicaments; [35], [166]), the functionality of the biasing effects of emotion on memory and information processing is debated (for one possible functional explanation see [158]).

¹¹ According to Weiner [202], motivation psychology is concerned with explaining why humans behave as they do. If one accepts that human behaviors are for the most part goal-directed, this means that motivation psychology is centrally concerned with explicating the processes that generate goal-directed behavior [42]. Thus, a theory of motivation is essentially a theory of the mechanisms of action and goal generation, including a theory of basic goals or fundamental desires.

nature, and even the effects of emotions, these emotion theorists have typically not tried to explain (at least not in detail) how emotions fit into the mechanisms of goal and action selection described by existing theories of motivation [145].

To further clarify the relation between emotion and motivation, we briefly consider the question from the perspective of motivation theory. It can be argued that nearly all current theories of motivation and action generation are variants of a singly basic theory, the belief-desire theory of action (for a selective sample of the relevant literature see [29], [41], [59], [60], [66], [112], [172], [183], [206], [202]). The basic assumption of the belief-desire theory of action is that actions—or more precisely, intentions to act, which then cause actions under favorable circumstances—are products of beliefs about one's current state and means-ends relations, and desires for the ends.¹² Philosophical proponents of the theory have argued that, to produce adaptive actions, exactly these two kinds of representational states, the one informational and the other motivational, are needed [e.g., 183]: Without beliefs (about the current state of the world and about ways to achieve goals), desires would be blind; without desires, beliefs would be inert. These considerations suggest that—contrary to the claims of some emotion theorists (e.g., [22]; and possibly [43])—emotions are *not indispensable* for the generation of adaptive actions; even though affect-free actions may well be overall less adaptive than actions that are also informed by emotions (2.2.4). Support for this conclusion might be seen in the fact that it is possible to develop workable theories of motivation and decision making without explicitly considering emotions (e.g., [2], [84]). Support might also be seen in the fact that it has been possible to create artificial agents that show reasonably adaptive goal-directed behavior even though their actions are not explicitly influenced by internal states corresponding to emotions (e.g., BDI agents; [142]; see section 4.2).¹³ By contrast, the generation of intentional actions is simply unthinkable without motivational mechanisms—mechanisms of goal selection and action generation that invoke desires and beliefs (see also, [199]). In line with this, it appears that to the degree that the effects of emotions on intentional action have been discussed in some detail by emotion theorists, they presuppose for the most part, at least implicitly, a traditional motivation or decision theory (e.g., [9], [43], [62], [92]).

Perhaps these considerations may be taken to suggest the following: First, the central evolutionary function of emotions is to improve the generation of adaptive action

¹²Decision theories (including subjective expected utility theory, prospect theory, and diverse psychological expectancy-value theories such as the theory of planned behavior) can be regarded as different attempts to quantify a generalized version of the belief-desire theory of action [145]. However, the relation between *specific versions* of belief-desire theory, such as BDI theory, and decision theory is not straightforward (e.g., [47]).

¹³As mentioned, the belief-desire theory of action and its quantitative version, decision theory, do not exclude an influence of emotions (in the form of hedonistic desires). Our argument here is that these theories do not *require* emotions; in particular, nothing in these theories enforces a hedonistic interpretation of desires. Actions require desires for goals, but as already stressed by McDougall [109], the goals need not be hedonistic.

over and above what is possible with beliefs and desires alone. Second, emotions achieve this by influencing, in one way or other, the motivational (belief-desire) machinery that proximately controls actions. This may be what unites the above-described informational, attentional and motivational effects of emotions.

2.2.6 Formal Tools for Representing Emotion Theories

Although informal analyses of emotion theories can go a long way towards clarifying and systematizing them, even the reconstructed theories will usually not be sufficiently precise and complete to be directly implemented as computational models. Of course, increasing the precision and completeness of a theory is precisely one goal of creating a computational model of that theory (e.g., [61], [192]). However, at this point in the history of computational modeling, the *direct* implementation of an emotion theory as a computer program is usually not the best way to go—at least if this means to program everything from scratch. For one reason, this is a cumbersome way to proceed, as it requires to invent or at least program much anew; for another reason, too many potentially irrelevant implementation choices have to be made along the way. In this and the following part, we consider two alternatives: (1) Formalizing emotion theories in a precise but implementation-independent formal language; (2) implementing emotion theories in a (comparatively) general-purpose cognitive or agent architecture.

3. THEORY FORMALIZATIONS: SET THEORY AND AGENT LOGICS

The first proposal consists of translating a psychological emotion theory (or selected assumptions of one or more emotion theories, as discussed in 2.2), as an intermediate step, into a precise, but implementation-independent language. Such a formalization has two main benefits (e.g., [74]). The first coincides with the general benefit of theory formalization. This benefit derives to a large extent from the fact that the formalization of a theory is usually much more than a simple *translation* of the theory into a formal language; it is simultaneously a process of *explicitating* the theory: resolving ambiguities and possible inconsistencies, uncovering hidden assumptions, adding missing assumptions, and the like (e.g., [205]). The result, ideally, is an improved version of the original theory: a version that is precise, complete, consistent, and concise. In addition, formalizations provide improved opportunities for comparing and evaluating the assumptions of different theories described in the same format, they facilitate the derivation of predictions, and they can raise new questions that may even suggest additional empirical research.

However, it can be argued that the general benefits of formalization are shared by computational models (which can be regarded as formalizations of substantive theories in a programming language; e.g., [108]).¹⁴ What are the

¹⁴More precisely, one should say that *certain parts* of computational models (those corresponding to the theory) algorithmically *implement*, or *realize*, substantive theories (e.g., [102]).

special benefits of first creating a formal specification of a psychological emotion theory, that is then implemented as a computational model? Perhaps the main benefit is that such a formalization yields a precise model of the psychological theory but avoids a commitment to potentially irrelevant implementation details (that are a necessary part of a computational model). That is, only assumptions regarded as essential for the emotion theory in question need to be considered in the formalization, whereas computational details—even important ones such as the structure of the agent architecture—can be ignored if they are considered irrelevant [74]. As a consequence, less effort is needed to verify that a formal model is a correct representation of an emotion theory, than would be required to verify a computational model. In addition, a precise and unambiguous formal version of an emotion theory facilitates its subsequent implementation as a computational model (see also, [45]). Finally, like the source theories on which they are based, formalized theories still allow to deduce implications (make predictions) *directly*, whereas computational models need to be *run* to generate predictions. Because of their special virtues, theory formalizations in implementation-independent languages can also be useful for comparing existing computational models of emotions.

3.1 Set-theoretical Formalization of Emotion Theories

Broekens, DeGroot and Kusters [32] propose to use *set theory* as a language for formalizing psychological emotion theories (see also, [147]). Set theory recommends itself as a formalism for the precise reconstruction of theories because of its expressiveness; in particular, (nearly) all mathematical concepts can be defined in terms of set-theoretical concepts [28].¹⁵ Choosing set theory as the formalization language therefore makes all of classical mathematics available for theory reconstruction and essentially amounts to the mathematization of a theory (by which we understand the construction of a mathematical, although not necessarily a quantitative model). A set-theoretical formalism is also well-suited to represent computational theories (as well as psychological theories interpreted as information-processing theories), because a common way of describing computational processes is as mathematical functions that map input to output information. This formal representation translates fairly directly into familiar graphical descriptions of systems, such as the block diagrams used in cybernetics and computer science and associated methods such as functional decomposition, and is also closely related to causal modeling formalisms used in psychology (e.g., [26]). These parallels can be exploited to depict computational functions in a more accessible graphical format (e.g., [32]) and conversely, to translate graphical theory proposals into set-theoretical language.

For example, Broekens et al. [32] break up the emotion generation process assumed by cognitive appraisal theo-

ries into three linked functions (see Figure 1): *Perceptual processes* map external objects into percepts (or mental objects); *appraisal processes* map mental objects into different appraisal dimension values; and *mediating processes* map appraisals into components of emotional reactions (e.g., feelings, expressions, actions). Using this basic functional (or process) decomposition, Broekens et al. then propose set-theoretical formalizations of two specific appraisal theories, Scherer's [166] stimulus evaluation check theory, and Smith and Kirby's [181] appraisal detector theory.




Figure 1. Basic framework for a set-theoretical formalization of appraisal theories [32].

The set-theoretical formalization of emotion theories used by Broekens et al. [32] ties in naturally with a still more rigorous set-theoretical formalism for theory reconstruction, developed as part of the so-called *structuralist program* in the philosophy of science (e.g., [14], [15], [204]). Essentially, this formalism consists of set theory (including whatever parts of mathematics are needed) enriched with a set of precisely defined special concepts that describe common, recurring components of empirical theories. According to structuralism, empirical theories are best viewed, from the systematic perspective, as set-theoretic structures composed of so-called *theory-elements*. At minimum, a structuralist reconstruction of a theory consists of a single theory element, but frequently it consists of several hierarchically connected theory-elements, called *theory-nets*. In the typical case, theory-nets comprise a basic element that represents the fundamental assumptions of a theory, plus a number of specialized elements that result from the basic element through the more precise specification of existing or the addition of further assumptions. In this way, theory nets allow to represent both refinements and different elaborations of a common core idea. In addition, theory elements belonging to different theory nets can be connected through *intertheoretical links*. Links are used to formally express the fact that empirical theories typically presuppose other theories (e.g., theories of measurement of their basic quantities). The basic building blocks of theories, the theory elements, have a precisely defined internal structure: They are defined as tuples $\langle \mathbf{K}(\mathbf{T}), \mathbf{I}(\mathbf{T}) \rangle$, where $\mathbf{K}(\mathbf{T})$ is the mathematical core of the theory, and $\mathbf{I}(\mathbf{T})$ is a set of intended applications of the theory core. $\mathbf{K}(\mathbf{T})$ consists in particular (but not exclusively) of the classes of potential models (\mathbf{M}_p) and the class of models (\mathbf{M}) of the theory. Potential models are all entities (set-theoretical structures, or systems) that can be described in terms of the theory's basic concepts, and for which it is therefore meaningful to ask whether they fulfill the substantive laws of the theory; whereas models are potential models that also fulfill these laws.

To use an elementary example, consider a structuralist

¹⁵ Note that the Bourbaki approach uses informal rather than formal (in the sense of axiomatic) set theory; hence, this approach can be described as "informal formalization" [184].

reconstruction of a miniature quantitative belief-desire theory of hope and fear (e.g., [151]). The class of potential models of the basic element of this theory could be defined as the set of systems $x = \langle PROP, TIMES, b, d, hope, fear \rangle$ consisting of a set of propositions $PROP = \{p_1, p_2, \dots, p_n\}$, a set of time points $TIMES = \{t_1, t_2, \dots, t_m\}$, a belief function $b: PROP \times TIMES \rightarrow [0,1]$, a desire function $d: PROP \times TIMES \rightarrow \mathbb{R}$ (the set of real numbers), and functions for *hope* and *fear*, both of which are also defined on $PROP \times TIMES \rightarrow \mathbb{R}$. A system x belongs to the class of models of this mini-theory of hope and fear if x is a potential model of the theory and in addition fulfils the following two laws (which are possible specializations of the more general functions proposed in [151]; see also, [95]):

$$\begin{aligned} hope(p, t) &= b(p, t) \times d(p, t) \quad \text{if } 0 < b(p, t) < 1 \text{ and } d(p, t) > 0; \text{ else } 0. \\ fear(p, t) &= |b(p, t) \times d(p, t)| \quad \text{if } 0 < b(p, t) < 1 \text{ and } d(p, t) < 0; \text{ else } 0. \end{aligned}$$

Because the definitions of set-theoretic predicates never yield a sufficient characterization of an empirical theory, the theory's range of application is determined by describing a set of intended applications $I(T)$. This is done by (a) explicitly listing a few real systems, called *paradigms*, and declaring them to be intended systems, or applications of the theory core; and by (b) declaring all other systems to be intended applications if they are sufficiently similar to the paradigmatic cases. For example, one paradigmatic application of hope-fear theory would be the experiences of hope and fear of a person awaiting a possible gain (hope) or loss (fear) in a two-outcome lottery where the other outcome is zero [156]; whereas another intended application might be the hope and fear experiences of gamers at certain stages in a computer game [69].

The viability of the structuralist approach to theory reconstruction is documented by the fact that at least 50 reconstructions of different theories from diverse disciplines have been published (see e.g., [15], [204]), including several theories from psychology. Westmeyer [205] therefore recommended the structuralist approach as a general methodology for theoretical psychology; while Balzer and Moulines [13] recommended the structuralist formalism of theory representation (specifically to computer scientists) as a good tool for representing scientific knowledge in a general context.¹⁶ Furthermore, the structuralist view of theories has inspired a simulation language for social science applications (MIMOSE; now obsolete) that facilitates the transformation of a structuralist theory reconstruction into an executable program ([197]; see [198], for other simulation languages that show a similar correspondence to set-theoretic formalizations¹⁷). Although so far only one serious structuralist reconstruction of an emotion theory (the three-dimensional theory of emotion proposed by Wundt [207]) seems to have been

¹⁶According to the authors, "there seems to be a fact of the matter that requires a set-theoretic representation of scientific knowledge, all weaker (and logically nicer) formal systems [being] too poor in expressive power or otherwise too restrictive for the working scientist" ([13], p. 9).

¹⁷Thanks are due to Klaus Troitzsch for providing this information.

attempted [146], the sketch of the belief-desire model of hope and fear given above and the set-theoretic formalizations of the two emotion theories proposed by Broekens et al. [32] suggest that rigorous structuralist reconstructions of other emotion theories are certainly feasible. Importantly, as mentioned above, structuralist reconstructions of psychological theories are not restricted to intentional-level theories [51], but extend to design-level (computational) theories. Hence, the structuralist idiom could also be used for an abstract description of computational emotion models realized in cognitive and agent architectures such as Soar (e.g., [68], [104]; 4.1), the belief-desire-intention architecture (e.g., [20]; 4.2), or affective agent architectures (e.g., [77]; 4.3). In fact, a structuralist reconstruction of a precursor of ACT-R (ACT*, [4]) has been presented by Heise [72].

3.2 Formalizing Emotion Theories in Agent Logics

Even though Balzer and Moulines [13] recommended the set-theoretic formalization of theories for representing scientific knowledge in a general context, they acknowledged that more restrictive formalisms may be more useful for specific applications. In particular, applied computer scientists who seek to develop a computational model of an emotion theory may prefer a formal specification language that, while less expressive than set theory, provides them with more structure, to facilitate the transformation from specification to implementation. One particularly interesting alternative possibility of formalizing emotion theories—at least intentional-level, cognitive emotion theories (see part 2)—is to represent them in a *logic language*. During the last 20 years, several logic languages have been developed for modeling cognitive autonomous agents that are suited for this purpose. Most of these so-called *agent logics* belong to the class of (broadly conceived) *BDI (Belief-Desire-Intention) logics*, that describe autonomous agents on the intentional level [51] in terms of beliefs, desires (goals), intentions (roughly, the goals that the agent has committed to) and possibly other related attitudes, such as capabilities (see e.g., [54], [199], [206]).

More precisely, BDI logics are formal logic languages that arise from the combination of several modal logics: a temporal or a dynamic logic used to capture the dynamic nature of agents, and logics for the mental states of belief, desire and intention. Each of the modal operators is given a precise syntactical definition in terms of a set of axioms, and a precise semantics in terms of possible worlds models (Kripke structures; see e.g., [81]). The axioms for belief, desire and intention and their interaction encapsulate principles of thinking (belief-formation) and deciding for an idealized rational agent that reasons qualitatively [206]. Put differently, these axioms encapsulate a qualitative version of rational decision theory [47], whose basic decision principle is simply that rational agents intend to do what they believe will lead to what they desire ([21], [138]). To this, BDI theory adds the important idea that rational decision makers need to show some commitment to their intentions [29].

As mentioned in section 2.2.3, most (at least most "higher") emotions are thought to depend on beliefs and—directly or indirectly—on desires. Therefore, BDI

logics are natural candidates for formalizing theories of (these) emotions. A second reason for formalizing emotion theories in BDI logics is that these logics already contain axioms for decision making and hence describe a workable, autonomous goal-directed agent. Formalizing an emotion theory in a BDI logic thus amounts to introducing emotion concepts into a formal theory of rational agency, or a formal qualitative decision theory. This approach facilitates the modeling of the effects of emotion on action (e.g., [46]) and the systematic investigation of the benefits of emotion mechanisms for goal-directed rational agents (e.g., [45], [97], [99]).

Although semi-formal belief-desire analyses of some emotions have been proposed by philosophers at least since the 1970ies (e.g., [48], [50], [70], [171]), rigorous formalizations of belief-desire analyses of emotions in agent logics are of relatively recent origin. Most of these formalizations focus on the cognitive and motivational preconditions (or, depending on viewpoint, the components; see [70]) of emotions (e.g., [1], [36], [114], [133], [186], [187]). For example, Castelfranchi and Lorini ([36]; see [96] for an update) formalized the belief-desire preconditions of a set of emotions related to expectations (hope, fear, disappointment, and relief) using one of the first BDI logics, proposed by Cohen and Levesque [40]. Meyer [114] used KARO, an agent logic based on a dynamic logic augmented with BDI-type modal operators [200] to formalize the belief-desire preconditions of four basic emotions (happiness, sadness, anger, and fear) proposed in Oatley and Johnson-Laird's [127] theory of emotion, another belief-desire theory (see also [126]). More recently, Steunebrink et al. [187] used KARO to formalize the cognitive-motivational preconditions of the 22 emotions considered in the OCC theory [129]. Another formalization of the OCC theory, using an extended version of the Cohen-Levesque logic, was proposed by Adam et al. [1]. As mentioned in part 2, these formalizations involve a reinterpretation of the OCC theory as a belief-desire theory of emotion.

In addition to formalizing hypotheses about emotion generation, recent work has formalized some of the proposed effects of emotion (see section 2.2.4); specifically their effects of actions ([187], [46]), including expressive speech acts [71].

The main motivation behind these logic formalizations of emotion theories has been to provide rigorous formal specifications for the implementation of emotions in autonomous agents. However, as mentioned, these formalizations can also be useful for psychologists and philosophers of emotion: Apart from providing a concise summary of a psychological or philosophical source theory, they can help to clarify and disambiguate assumptions of the source theory, reveal problematic implications that may otherwise remain hidden or go unnoticed, facilitate the derivation of consequences, and raise new questions [1].¹⁸ Several examples can be cited to support these

claims. (1) Steunebrink, Dastani, and Meyer [186] as well as Adam et al. [1] found that their logical reconstructions of the OCC theory revealed a number of ambiguities in that theory, such as whether event-based emotions are caused by perceived events or the believed consequences of events. (2) Adam et al.'s [1] reconstruction of the OCC theory led them to propose, for example, that an agent is disappointed about a state of affairs $\text{not-}p$ if it believes $\text{not-}p$ to be the case, desired p to occur, and believes that he previously expected p . Hence, agents need to form a metacognition (a belief about an own belief) to experience disappointment. Although this requirement may be unproblematic for artificial agents, at least in principle (though not for some concrete realizations of BDI agents; e.g., [73]), it could be problematic for humans [151]. Adam et al.'s formalization makes this assumption transparent and thereby presents it for discussion. In fact, Castelfranchi and Lorini [36] and Steunebrink et al. [187] have proposed alternative analyses of disappointment that avoid postulating metacognitions. For example, Castelfranchi and Lorini assume that an agent is disappointed about $\text{not-}p$ if it previously hoped that p will occur and then comes to believe $\text{not-}p$ (see also [151]). (3) The two formalizations of the OCC theory by Adam et al. [1] and Steunebrink et al. [187] also differ in several other respects (see [187], for a discussion) that could stimulate discussion, and possibly even empirical research. (4) The logical formalizations bring into sharp focus the question of the nature of emotions in belief-desire theory. In the majority of the existing formalizations (the exception is [187]), emotions are *defined* in terms of beliefs and desires—e.g., $\text{joy}(p) =^{\text{df}} \text{bel}(p) \ \& \ \text{des}(p)$ —which might suggest that emotions are to be *equated* with particular belief-desire structures. This has indeed been proposed by some belief-desire theorists of emotion (e.g., [105]; see also [70]). However, this is neither the only nor necessarily the most plausible version of the belief-desire theory of emotion (see [153], for a detailed discussion). In fact, most authors make clear that they do not intend their formalizations as implying a reduction of emotions to beliefs and desires, or at least do not wish to commit to this theory of the nature of emotion. Rather, the formalization is meant to express, or at least to be also compatible with, the idea that beliefs and desires are necessary, and under normal circumstances also sufficient, *causes* of emotions.

Let us now consider two possible, general objections to formalizations of emotion theories in BDI logics. The first is that BDI logics *are not expressive enough* to accommodate all assumptions of psychological emotion theories. Three concrete forms of this objection are the following: (1) The above-mentioned formalizations are purely qualitative; however, emotions differ not only in quality but also in intensity, and intensity is important for explaining certain effects of emotions [64]. To account for emotion intensity, quantitative concepts of belief and desire are needed. (2) Even on the qualitative level, BDI logics may not be rich enough to formalize the cognitive-motivational preconditions of all emotions. For example, it has been argued that

¹⁸The fact that a variety of different BDI-type logics are available to choose from should probably be regarded as an asset rather than as a disadvantage. First, it means that one need not commit to a particular logic; different routes can be tried and compared. Second, the existence of

different logics suggests that basic conceptual issues need to be, and can be, explored while formalizing emotion theories.

feeling guilty about p presupposes not only the belief that p is the case and the desire that p should not have occurred, but also the belief that one is responsible for p (e.g., [129], [203]). (3) Although BDI logics go beyond traditional “static” decision and action theories in considering the dynamics of mental states and actions, they are still restricted to the intentional level of system analysis [51]; they do not address the underlying representations and computational processes.

These objections can be answered as follows. As to the first objection (BDI emotion models are qualitative), it can be argued that for one thing, precise qualitative formulations are no less important than quantitative ones (e.g., [129]). For another thing, it is possible to extend BDI logics to account for emotion intensity. For example, Lorini [96], [46] proposed to combine a dynamic logic of knowledge with a logic for graded beliefs and desires for this purpose. Similarly, Steunebrink, Dastani and Meyer [185] have incorporated quantitative aspects of emotion into their KARO model.

The second objection can be answered similarly. For example, to explicate the concept of responsibility and related concepts such as preventing and causing, Lorini and Schwarzenhuber [98] have proposed to combine BDI logic with STIT logic (a logic of “seeing to it that”).

As to the third objection (logical emotion models are restricted to the intentional level of system analysis), this restriction is deliberate and does not preclude combining the logical models with design-level assumptions, as are in fact provided by BDI architectures (section 4.2). On the contrary, as mentioned, because of their greater precision, logical BDI models can *assist* the creation of computational models. Nevertheless, at this juncture, the interests of psychologists and computer scientists may partly diverge. From the perspective of computer scientists, at least those pursuing applied goals, a specification of an emotion theory that abstracts completely from questions concerning representational formats and computational mechanisms may be regarded as a benefit. In contrast, as mentioned in part 2, psychologists adhering to the computational paradigm are equally interested in certain design-level questions, although not in all (e.g., [68]; see also, [151], [178]).

A second general objection that could be raised against formalizing emotion theories in BDI logics is that the rationality assumptions incorporated in the usual axioms for belief, desire and intention imply that the logical BDI models can only be regarded as strong idealizations of human belief-formation and decision making. For example, standard BDI logics assume that agents believe and desire all logical implications of their beliefs and desires, respectively, and are also perfectly aware of the beliefs and desires that they have (e.g., [58]). These are not plausible assumptions to make for real agents with limited memory, limited processing time and (most likely) limited metacognitive abilities. Because these rationality assumptions carry over to beliefs and desires as presuppositions of emotion, the logical models of “emotion-enhanced” BDI agents can likewise be regarded only as strong idealizations of human emotion mechanisms.

To some degree this criticism can be met by formulating at least the emotion axioms in ways that do not demand more from agents than is necessary and psychologically plausible (cf. [36], [187]). Beyond that, attempts have been made to develop logics that describe “less than ideal” agents (e.g., resource-bounded epistemic logics, [115]). However, this project has its own problems; for example, Duc [56] argues that by axiomatically restricting the rationality of agents, nothing is left that can account for their rationality any more. From the perspective of psychology, two additional considerations may be more important: First, it should be remembered that all theories are idealizations and that even strong idealizations can be quite useful (consider, for example, the case of game theory; e.g., [34]). In particular, the strong rationality assumptions of BDI logic can be mentally set aside when the logic is used as a tool to clarify hypotheses (cf. [41]); for this purpose, what mainly counts is the increased precision provided by the formal language. Second, *computer implementations* of BDI agents (as opposed to their logical descriptions) have to operate under time and memory constraints like humans. As a consequence, they become *boundedly rational agents*, agents that behave as well as possible given their limited resources and constraints (e.g., [169]). To obtain more realistic models of motivational and emotional processes, psychologists can therefore focus on the BDI architectures (section 4.2).

4. MODELING EMOTIONS USING COGNITIVE AND AGENT ARCHITECTURES

Another proposal for bridging the gap between psychological emotion theories and computational emotion modeling, as well as for increasing the comparability of different computational emotion models, is to encourage the use of (suitably extended) *general-purpose cognitive and agent architectures* to model emotion theories. Cognitive architectures and agent architectures are theories of the *basic design-level structure* [51] of cognitive systems or agents, respectively: theories that describe their domain-independent, basic computational structures (e.g., memories, representational systems) and processes (e.g., a repeated sense-think/decide-act cycle) (e.g., [6], [88], [164], [192]). From the perspective of the user (modeler), an architecture appears as the “hardwired” part of an agent that cannot be changed (at least not easily).

Architectures are typically used to facilitate the creation of computational models of specific psychological phenomena of interest, such as the performance on a particular task (e.g., in the case of emotion research, the performance on the Iowa gambling task; [188]). The main advantage of using an architecture for this purpose, compared to creating a computational model from scratch, is that the architecture takes care of numerous implementation choices that otherwise would have to be made. Thereby, it relieves the modeler of the need to invent everything anew and helps—similar to the formal languages discussed in part 3, but in a different way—to solve the problem of irrelevant design specifications: Whereas in the formal models, design assumptions con-

sidered irrelevant are *omitted*, in an architecture, they are *assigned plausible default values*. By letting the architecture take care of design decisions, the modeler implicitly accepts that the respective assumptions of the architecture are correct, or at least not too far off the mark in ways that matter to the phenomenon of interest. If one is willing to make this assumption, the use of an architecture can greatly simplify the creation of a computational emotion model [108]. Furthermore, to the degree that architectures are becoming reasonably easy to use even for non-professional programmers, they could significantly lower the threshold for more psychologists becoming interested in, and even engaging themselves in the computational modeling of emotions. In addition, the research efforts of affective computing scientists could be focused, and competing proposals more easily compared, by modeling emotions in otherwise well-understood architectures.

However, the role of architectures in computational modeling extends beyond these practical benefits: Architectures can also play important *theoretical roles*. In particular, to the degree that an architecture is well-anchored in theory and research and hence contains many at least approximately correct psychological assumptions, the computational distinctions incorporated in the architecture (e.g., between deliberative and automatic processing) can be exploited to clarify similar distinctions proposed in a psychological process theory of emotion (see e.g., [68], and 4.1). Sloman [178] even argues that an understanding of emotions can *only* be attained by considering the role of emotions in cognitive architectures. As Sloman puts it: “A proper analysis of the concept of an ‘affective’ state or process must be based on a more general theory of the coarse-grained architecture of mind” ([178], p. 233). This said, it should be acknowledged that the fixed computational structure provided by a given architecture can also be a disadvantage: namely, if it constrains what one wants to model in crucial respects. This consideration is particularly important in the computational modeling of emotion, because emotions may be implicated in basic psychological processes at the architectural level (section 2.2.4; see also [151]).

General overviews of cognitive and agent architectures are provided by Chong, Tan, and Ng [38]; Langley, Laird, and Rogers [90]; Vernon, Metta, and Sandini [201]; and Zacharias, Macmillan, and Van Hemel [208]; for overviews of emotional agent architectures, see Hudlicka [78] and Marsella et al. [108]. For the purpose of our review, we have grouped the (comparatively) general-purpose cognitive and agent architectures that can be and in fact have been used for emotion modeling into (a) by now “classical” cognitive architectures such as Soar or ACT-R (section 4.1); (b) autonomous agent architectures developed since the 1990ies in artificial intelligence, the most widely used probably being the BDI (belief-desire-intention) architecture (4.2); and (c) (comparatively) general-purpose architectures specifically designed for the creation of affective agents, such as FATiMA Modular and MAMID (4.3). Despite important differences between these architectures, they share some basic similarities. In particular, all of them can be taken to describe goal-based

(or goal-directed) agents [57] rather than (merely) reactive agents (although reactive behaviors can also be modeled); and all use symbolic internal representations to represent the agent’s cognitions. Furthermore, none of these architectures does anything on its own; it is only when they are supplied with data (knowledge/beliefs, goals) and placed in a suitable task environment that they “come to life” and generate cognitive processes and behavior.

The main difference between the affective agent architectures and the more traditional cognitive and agent architectures is that the latter do not, in their original form, consider emotions. Nevertheless, these architectures have been used to model both processes of emotion elicitation ([68]; using Soar) and the effects of emotion on action and task performance (see again [68] for Soar, and e.g., [160], for ACT-R). This has been possible because both Soar and ACT-R contain, at least implicitly and in an approximate fashion, central cognitive and motivational constructs needed to account for emotions (see part 2)—constructs relating to the beliefs (knowledge) and the motivation (desires, goals) of agents (for more detailed discussions, see [57], [65], [107]).

4.1 Modeling Emotions in Cognitive Architectures

Classical cognitive architectures such as Soar and ACT-R are “echos of the original goal [of AI] of creating an intelligent machine faithful to human intelligence” (Taatgen & Anderson, [193], p. 170). At the same time, they are attempts at theoretical unification in the field of cognitive psychology (again [193], p. 170).

4.1.1 Modeling Emotions in Soar

One popular cognitive architecture, that has also been used to concretize emotion theories, is Soar ([88], [123]). Soar is a general cognitive architecture for human cognition realized as a production system, the basic operation of which consists of matching the contents of working memory to the precondition of rules (see [164]). Following the classic work of Newell and Simon [124] on the *General Problem Solver*, the basic assumption of Soar is that intelligent behavior is a form of problem solving. In Soar, tasks are represented as problem spaces and problem solving consist of a search (involving both automatic and sequential processes) for sequences of operators that will reach a goal state in problem space. At each step in this search, Soar rapidly and in parallel recruits knowledge from long-term memory that is relevant to the current situation. After this knowledge recruitment processes reaches quiescence, the decision procedure executes a single most-preferred operator, records its consequences to working memory or the outside environment, and then recruits knowledge relevant to this new state. If a problem arises due to the fact that certain knowledge is lacking, resolving this problem becomes a subgoal of the original goal. Soar is also capable of learning about how subgoals were reached, and what the consequences of operators are. It may be noted that considered from the intentional stance [51], the operation of Soar is essentially a form of practical reasoning ([29]; see also [65]): Goals are desired states to be reached by actions; operators cor-

respond to possible actions, and knowledge about the effects of operators correspond to means-ends beliefs.

Soar was not designed to model emotions; nevertheless, two different computational emotion models have been built upon Soar ([68], [104]). These two models represent the two principal alternative paths available to model emotions in Soar and similar cognitive architectures, and they also illustrate how the theoretical assumptions embedded in an architecture can influence modeling choices.

EMA, proposed by Gratch and Marsella ([68]; see also, [107]), implements the cognitive process theory of emotion proposed by Smith and Lazarus [182]. In EMA, the basic architecture of Soar is left largely untouched, by implementing appraisal processes as production rules encoded in Soar's long-term memory. In doing this, EMA exploits Soar's theoretical distinction between deliberative and automatic processes to concretize a similar distinction proposed by Smith and Lazarus between deliberative (conscious) and automatic appraisal processes: Appraisals are realized in Soar as "elaboration" productions that fire rapidly and in parallel, and that bias sequential operator decisions. Thus, *by following Soar's own theoretical assumptions*, EMA achieves a tight and cyclical coupling between automatic and deliberative appraisal processes in a way that can be regarded as a computational specification of Smith and Lazarus's process theory of emotion.

In contrast to EMA, Marinier et al.'s [104] PEACTIDM model introduces emotions into Soar by changing Soar's core architecture: Inspired by Scherer's [166] theory of sequential appraisal checks, PEACTIDM attempts to capture the sequential nature of appraisals posited by in that theory by adding sub-stages to Soar's decision procedure. Specific appraisals are encoded into the architecture and are associated with these processing stages. As a result, PEACTIDM, like EMA, computationally implements the assumption of a tight coupling between automatic and deliberative appraisal processes; however, it takes a very different path to realize this idea.

4.1.2 Modeling Emotions in ACT-R

ACT-R ([5], [6]) is another influential cognitive architecture that shares many similarities with Soar; in particular, like Soar, ACT-R is a production system with goals that operates on symbolic representations. Also like Soar, the basic operation of ACT-R is to achieve goals by applying suitable procedures (actions). The differences between Soar and ACT-R concern mainly a variety of design features that together are intended to make ACT-R more psychologically plausible than Soar in its original form ([6]; for a discussion of some differences see also [83]). These design features include an explicit distinction between a procedural and a declarative memory and multiple modality-specific memory buffers.¹⁹ A particular strength of ACT-R is its well-developed model of long-

term memory, in which information is represented as networks and is activated by a subsymbolic process of spreading activation (in fact, ACT-R developed out of a theory of human associative memory). In this sense, ACT-R can be said to represent an integration of symbolic and subsymbolic mechanisms [160].

Although ACT-R has so far only been used to a limited extent to model emotions, it deserves particular attention because it is probably the best-known cognitive architecture in psychology and has been used to create precise computational models of many psychological theories, including several models of decision-making (e.g., [103]). Ritter et al. [160] used ACT-R to demonstrate two general ways of how some of the effects of emotions can be modeled in cognitive architectures, namely (a) by varying the content of the (procedural or declarative) knowledge of the agent; and (b) by varying certain architectural parameters (as proposed by Hudlicka [76]; see also [77], [79]). The first approach was used to model the effects of worry on performance, by including a task-irrelevant production in ACT-R that could fire at any time, simulating the effects of distracting thoughts on working memory. The second approach was used to model the effects of threat versus challenge appraisals of a stressful task on performance, by varying the level of randomness in the conflict resolution process (which determines which production rule will fire when more than one rule matches the goals of the system). Other noteworthy applications of ACT-R to emotion modeling are Cochran, Lee and Chown's [39] model of the impact of arousal on memory, and Stocco, Fum and Zalla's [188] model of the Iowa gambling task, the main experimental paradigm used to test Damasio's [43] somatic marker hypothesis of the effects of emotion on decision making. Although *emotion elicitation processes* do not seem to have been modeled in ACT-R, the similarities of this architecture to Soar suggest that these processes could be considered in one of two ways in ACT-R: (a) by modeling appraisals as production rules (analogous to the approach taken by [68] in Soar); or by extending the basic ACT-R architecture (analogous to [104]).

4.2 Modeling Emotions in BDI Agent Architectures

In part 3.2, we discussed logical models of BDI (Belief-Desire-Intention) agents and noted that, among other uses, such models can aid the creation of corresponding computational models. Here, we discuss the associated computational BDI architecture (e.g., [30], [142], [143]). To recall, logical BDI models are formalized theories of qualitative rational decision making [29]. However, although these models are precise renderings of informal intentional-level [51] theories, they remain silent about the design level: They make no assumptions about how beliefs, desires and intentions, and the mental operations performed with them, are computationally implemented. BDI architectures, in contrast, (also) address these questions: They are pieces of software that implement the abstract principles of BDI agents.

The BDI architecture, which exists in several variants, has become one of the most frequently used software

¹⁹Recent versions of SOAR (see [87]) incorporate many of ACT-R's features, including the distinction between procedural, declarative and episodic memory, and sub-symbolic processes such as reinforcement learning and mental imagery.

architectures for autonomous intelligent agents. The central assumption incorporated in BDI agents is that actions are generated by a process of practical reasoning [29] that comprises two computational steps. In the first step, a set of desires (goals) is selected; in the second step, the agent determines how these goals can be achieved by means of the available actions or plans [206]. Characteristic for BDI agents is the separation of the process of *deliberation* from the process of *execution* of currently active plans. This allows BDI agents to balance the time spent on goal and plan selection (choosing what to do) and executing plans (acting), which seems to be essential for adaptive action in resource-bounded agents [29].

Arguments in favor of BDI architectures as blueprints for autonomous agents are their foundation in philosophical (and psychological) theory; the existence of logical models that allow to define, and reason about, BDI agents in a precise manner (see 3.2); and the existence of several software systems for programming BDI agents that are close to these logical specifications (e.g., JAM [75]; JADEX [136]; 2APL [44]; GOAL: [73]). Compared to cognitive architectures such as Soar and ACT-R, BDI architectures are more abstract; for example, they do not require a commitment to specific memory types. Depending on one's modeling goals, this can be an advantage or a disadvantage (e.g., it is a disadvantage if one's interest is to model the effects of emotion on memory). Most BDI systems execute a fixed perception-deliberation-action (sense-think-act) execution cycle comprising the following steps: (1) perception of events; (2) updating of the belief and desire (goal) base (adding new beliefs and desires, dropping false beliefs and satisfied or impossible desires); (3) generating options (plans) (deliberating); (4) choosing the most suitable plan (generating an intention) and (5) executing intentions (acting). The basic BDI architecture has been extended in several ways, either by modifying the agent's internal processes or by extending the BDI model to the social environment (communication, multi-agent BDI systems). „Internal“ extensions of the BDI architecture comprise the inclusion of learning mechanisms, the consideration of personality differences (e.g., [125]), and the inclusion of emotions.

BDI architectures are in principle well-suited to create computational models of emotions given that most (at least most “higher”) emotions are thought to depend on beliefs and—directly or indirectly—on desires (section 2.2.3; see also, [74]). Furthermore, the formalization of emotion theories in BDI logics discussed in section 3.2 should facilitate their implementation in BDI software systems. Indeed, several attempts have already been made to incorporate (aspects of) psychological emotion theories into BDI systems (e.g., [140], [141], [20]). For example, Rank [140], [141], has extended JAM agents [75] by including an appraisal mechanism that compares perceived facts concerning own actions and those of other agents (success, failure, attempt) with the agent's goals and computes positive or negative “relevance values” on the basis of the utility of the affected top-level goal and the match or mismatch of the action-relevant facts to that goal. Successful actions are additionally matched to stan-

dards [129] that are associated with virtual goals to uphold them. If the appraisal of an action-relevant fact exceeds a threshold of intensity, up to three effects occur: The creation of an “impulse goal” to display a fitting emotional expression, the creation of a coping goal that may lead to a corresponding intention (e.g., retribution, trying to hinder), and a change of the evaluation of the agent who showed the action. These “emotion-enhanced” JAM agents have been used to steer the characters of a simple interactive drama featuring a hero, a victim, a mentor and a villain.

A somewhat different approach to enhancing BDI-agents with emotions has been taken by Becker-Asano and Wachsmuth ([20]; see also, [19]). In the emotional BDI agent *Max* constructed by these authors, an emotion system (a model of the pleasure-arousal theory of emotion; [162]) is added to the BDI architecture as a separate module that receives signals from, and sends signals to, several other modules. For example, on the input side, the emotion module receives a positive signal if a goal is achieved, which results in an increase in pleasure; whereas on the output side, emotions (pleasure-arousal states) among other effects generate beliefs about having them, amounting to self-awareness of the agent of its own emotions. In yet another “emotional” extension of the BDI architecture, Lorini and Piunti [97] have modified the processing cycle of the JADEX agent by including a surprise-based filter for belief revision ([97]; see also [100]).

To allow the modeling of different emotion theories or specific assumptions of different theories (see part 2), it would be useful if prospective emotion modelers could more easily customize and extend BDI systems than is currently possible. One way of achieving this would be to break up the original BDI interpreter cycle into a small set of processes that can be combined as needed [137], as well as extended with additional processes. For example, it would be useful to be able to extend the BDI architecture with a set of basic desires from which concrete desires can be autonomously derived ([11], [99]), or with a more detailed model of memory similar to ACT-R (e.g., to model effects of emotions on memory). Also, the BDI architecture could be parameterized to allow the modeling of certain personality differences (e.g., [125], [76]). In fact, some of the proposed extensions are already being realized; for example, Ritter [159] has enriched a BDI agent architecture with elements from ACT-R in CoJACK.

4.3 Modeling Emotions in Affective Agent Architectures

Several specialized “affective” agent architectures have been developed by affective computing researchers. Some of these are (relatively) general-purpose architectures and thus suited to model different emotion theories. To illustrate this approach, we describe FATiMA Modular ([52]; [11]) and MAMID ([77], [79]).

4.3.1 Modeling Emotions in Fatima Modular.

FATiMA Modular is an autonomous agent architecture with reactive and deliberative (BDI-based) behavior designed to allow emotions and personality characteristics

to influence the agent's behavior. It was created by refactoring a previous version of the architecture [53] into a modular version, consisting of a core that defines the general processing cycle of the agent, and several specialized components that are specific realizations of the generic agent functions (e.g., specific appraisal computation and action selection functions). The basic processing cycle of FAtiMA agents, as well as several details of the architecture, have been influenced by the EMA model [68]. The basic processing cycle comprises the following steps: (1) perception of events; (2) on this basis, memory is updated and (3) appraisal processes are triggered; (4) an affective state is generated; (5) actions are executed based on goal-based planning (a BDI-style reasoning process); and (6) actions are executed. As concerns emotion generation, FAtiMA combines a reactive layer that generates emotions on the basis of a set of domain-specific rules with a deliberative layer that generates emotions via an appraisal of the agents' current plans (cf. 2.3.3). The core of the deliberative module is a continuous planner that was extended to include probability information about actions and to perform emotion-focused coping strategies. When an event is received, active plans are updated and prospect-based emotions (hope, fear, satisfaction) based on the agent's plans and goals are generated.

Different cognitive emotion theories (or components of such theories) can be implemented in FAtiMA Modular by adding specific affect generation functions. Likewise, specific action selection functions can be implemented by adding specific behavior components. This philosophy is in agreement with the "modular" philosophy of theory systematization endorsed in part 2. Recently, FAtiMA Modular has been further extended to include, among others, an autobiographical memory for remembering emotionally significant events, a module for basic motives (adopted from Dörner's, [55], PSI theory; see also [12]) from which concrete goals are derived, and provisions for modeling cultural differences [11].

4.3.2 Modeling Emotions in MAMID

Similar to FAtiMA, MAMID [77], [79] implements an extended sense-think/feel-act processing cycle. In this case, the involved processes are: (1) sensory pre-processing (translating incoming data into task-relevant cues); (2) attention (filtering incoming cues and selecting a subset for processing); (3) situation assessment (integrating individual cues into an overall situation assessment); (4) expectation generation (projecting the current situation onto possible future states); (5) emotion generation (deriving emotions on the basis of a variety of cues); (6) goal selection and (7) action selection (selecting the best action for goal achievement). The computations of each module are supported by long-term memories represented as belief nets or rules. The memories of MAMID agents can vary in type and structure to reflect differences in individual histories, sensitivities, and responsiveness.

In MAMID, emotional states influence processing in three basic ways, the first two of which correspond to those featured in section 2.2.4: First, emotions influence the allocation of processing resources and thus determine indirectly whether a specific cue or situation is processed

or a specific goal is selected. Second, emotions directly influence which goals and actions are selected. Third, emotions can change the speed and capacity of the different modules of the architecture. This effect is meant to model system-wide subsymbolic (neuromodulatory) effects of emotions. In addition, both sensory processing and attention are affected by global architectural parameters that determine the speed and capacity of each processing module, as well as the ordering of the constructs within each module's processing queue. Together, these parameters allow to model biases for each of the basic processes (attention, situation assessment, expectation generation, emotion generation, and goal selection)

A key objective in developing MAMID has been to provide flexibility in the modeling of assumptions about the generation and particularly the effects of emotions (in fact, the primary focus of MAMID is on the modeling of *emotion effects* on attention, perception and cognition, as the modelling of such effects has been neglected in previous computational models). The methodology used to achieve this flexibility combines two ideas: (1) a high degree of parameterization of the processes and structures that comprise the different modules of the architecture; and (2) the modeler's interactive access to these parameters and their weights. For example, the valence (pleasure-displeasure quality) and intensity of emotions are computed in MAMID as linear combinations of several eliciting factors. The weights controlling the contributions of the individual eliciting factors can be modified interactively by the model developer. This provides a simple and quick way of implementing alternative hypotheses regarding the eliciting conditions of different emotions, as well as of tuning models to reflect emerging empirical data. Similarly, emotion-relevant individual differences (e.g., differences in threat sensitivity) can be modeled by interactively varying parameters. For example, in a stressful encounter, a trait-anxious agent is likely to consider its competence and degree of control, whereas a non-trait anxious agent would not consider these factors. This methodology for modeling the effects of emotions as well other behavior moderators (e.g., personality traits, see [158]) has several advantages [77], including the possibility to rapidly model a broad range of distinct individual profiles and to integrate the (possibly conflicting) effects of multiple interacting emotions and traits.

5. CONCLUSION

The proposals made in this article for increasing the interdisciplinary exchange between psychology and computer science, as well as the intradisciplinary coordination, in the field of computational emotion modeling share two overarching themes. The first is a proposal for *modularization*: deconstruct emotion theories into basic assumptions; modularize architectures. The second is a proposal for *unification and standardization*: Translate different emotion theories into a common informal conceptual system or a formal language, or implement them in a common architecture. In making these proposals, we find ourselves in agreement with developments already un-

derway in the domain of cognitive architectures. Specifically, Taatgen and Anderson [194] predict that, given that different cognitive architectures already share many identical or similar mechanisms, they will converge even more in the future, as a consequence of which the choice of a particular architecture will no longer imply strong theoretical commitments. Furthermore, by increasing the degree of modularization of architectures, it will be possible to build models out of components and thus avoid the constraints imposed by fixed architectures. Analogously, one may envision the development of a theoretical toolbox of basic theory-elements, formulated in a common language, from which theories of emotional agents (or of emotion modules for agents) can be constructed.

REFERENCES

- [1] Adam, C., Herzog, A., & Longin, D. (2009). A logical formalization of the OCC theory of emotions. *Synthese*, 168, 201-248.
- [2] Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50, 179-211.
- [3] Allen, S. R. (2001). *Concern processing in autonomous agents*. PhD thesis, Faculty of Science of the University of Birmingham, School of Computer Science.
- [4] Anderson, A. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- [5] Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111, 1036-1060.
- [6] Anderson, J. R., & Lebiere, C. (1998). (Eds.). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- [7] Angie, A. D., Connelly, S., Waples, E. O., & Kligyte, V. (2011). The influence of discrete emotions on judgement and decision-making: A meta-analytic review. *Cognition and Emotion*, 25, 1393-1422.
- [8] Armony, J. L., Servan-Schreiber, D., Cohen, J. D., & LeDoux, J. E. (1997) Computational modeling of emotion: Explorations through the anatomy and physiology of fear conditioning. *Trends in Cognitive Sciences*, 1, 28-34.
- [9] Arnold, M. B. (1960). *Emotion and personality*. New York: Columbia University Press.
- [10] Atkinson, J. W. (1964). *An introduction to motivation*. Princeton, NJ: Van Nostrand.
- [11] Aylett, R., & Paiva, A. (2012). Computational modelling of culture and affect. *Emotion Review*, 4, 253-263.
- [12] Bach, J. (2009). *Principles of synthetic intelligence. An architecture for motivated cognition*. New York: Oxford University Press.
- [13] Balzer, W., & Moulines, C. U. (2000). Introduction. In W. Balzer, J. D. Sneed, & C. U. Moulines (2000). *Structuralist knowledge representations: Paradigmatic examples (Poznan studies in the philosophy of the sciences and the humanities, Vol. 75, pp. 3-18)*. Amsterdam: Rodopi.
- [14] Balzer, W., Moulines, C. U., & Sneed, J. D. (1987). *An architectonic for science: The structuralist program*. Dordrecht: Reidel.
- [15] Balzer, W., Sneed, J. D., & Moulines, C. U. (2000). Structuralist knowledge representations: Paradigmatic examples (*Poznan studies in the philosophy of the sciences and the humanities, Vol. 75*). Amsterdam: Rodopi.
- [16] Barrett, L. F. (2006a). Solving the emotion paradox: Categorization and the experience of emotion. *Personality and Social Psychology Review*, 10, 20-46.
- [17] Barrett, L. F. (2006b). Are emotions natural kinds? *Perspectives on Psychological Science*, 1, 28-58.
- [18] Baumeister, R. F., Vohs, K. D., DeWall, C. N., & Zhang, L. (2007). How emotion shapes behavior: Feedback, anticipation, and reflection, rather than direct causation. *Personality and Social Psychology Review*, 11, 167-203.
- [19] Becker-Asano, C. (2008). *WASABI: Affect simulation for agents with believable interactivity*. US: IOS Press.
- [20] Becker-Asano, C., & Wachsmuth, I. (2008). Affect simulation with primary and secondary emotions. *Intelligent Virtual Agents*, 8, 15-28.
- [21] Bennett, J. (1991). Analysis without noise. In R. J. Bogdan (Ed.), *Mind and common sense* (pp. 15-36). Cambridge: Cambridge University Press.
- [22] Bentham, J. (1789/1970). *An introduction to the principles of morals and legislation*. London: Athlone Press.
- [23] Blanchette, I., & Richards, A. (2010). The influence of affect on higher level cognition: A review of research on interpretation, judgement, decision making and reasoning. *Cognition and Emotion*, 24, 561-595.
- [24] Block, N. (1980). Introduction: What is functionalism? In N. Block (Ed.), *Readings in philosophy of psychology* (Vol. 1, pp. 171-184). Cambridge, MA: Harvard University Press.
- [25] Boden, M. (1988). *Computer models of mind: Computational approaches in theoretical psychology*. Cambridge: Cambridge University Press
- [26] Bollen, K. (1989). *Structural equations with latent variables*. NY: Wiley.
- [27] Bosse, T., Gratch, J., Hoorn, J. F., Pontier, M., & Siddiqui, G. F. (2010). Comparing three computational models of affect. In Y. Demazeau, F. Dignum, J. Corchado, & J. Prez (Eds.), *Advances in practical applications of agents and multiagent systems* (Advances in soft computing, Vol. 70, pp. 175-184). Heidelberg: Springer.
- [28] Bourbaki, N. (1968). *Elements of mathematics: Theory of sets*. Reading, Mass.: Addison-Wesley.
- [29] Bratman, M. E. (1987). *Intentions, plans, and practical reason*. Mass: Harvard University Press.
- [30] Bratman, M., Israel, D. J., & Pollack, M. E. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4, 349-355.
- [31] Broekens, J. (2010). Modeling the experience of emotion. *International Journal of Synthetic Emotions*, 1, 1-17.
- [32] Broekens, J., DeGroot, D., & Kusters, W. A. (2008). Formal models of appraisal: Theory, specification, and computational model. *Cognitive Systems Research*, 9, 173-197.
- [33] Brook, J. A. (1992). *Psychoanalysis and commonsense psychology. Annual of Psychoanalysis*. Hillsdale, NJ: The Analytic Press.
- [34] Camerer, C. F. (2003). *Behavioral game theory. Experiments in strategic interaction*. Princeton, NJ: Princeton University Press.
- [35] Cannon, W. B. (1927). The James-Lange theory of emotions: A critical examination and an alternative theory. *American Journal of Psychology*, 39, 106-124.
- [36] Castelfranchi, C., & Lorini, E. (2003). Cognitive anatomy and functions of expectations. *Proceedings of IJCAI'03 workshop on cognitive modeling of agents and multi-agent interactions*.
- [37] Castelfranchi, C., & Miceli, M. (2009). The cognitive-motivational compound of emotional experience. *Emotion Review*, 1, 223-231.
- [38] Chong, H-Q, Tan, A-H, & Ng, G-W (2007). Integrated cognitive architectures: a survey. *Artificial Intelligence Review*, 28, 103-130.
- [39] Cochran, R. E., Lee, F. J., & Chown, E. (2006). Modeling emotion: Arousal's impact on memory. *Proceedings of the 28th Annual Conference of the Cognitive Science Society* (pp. 1133-1138). Vancouver, British Columbia, Canada
- [40] Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42, 213-261.
- [41] Conte, R., & Castelfranchi, C. (1995). *Cognitive and social action*. London: University College London Press.
- [42] Cox, W. M. & Klinger, E. (2004). Motivation and the theory of current concerns. In W. M. Cox & E. Klinger (Eds.), *Handbook of motivational*

- counseling* (pp. 1-28). Chichester: Wiley.
- [43] Damasio, A. R. (1994). *Descartes' error*. New York: Avon.
- [44] Dastani, M. (2008). 2APL: a practical agent programming language. *International Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 16, 214-248.
- [45] Dastani, M. & Meyer, J.-J. Ch. (2010). Agents with emotions. *International Journal of Intelligent Systems*, 25, 636-654.
- [46] Dastani, M., & Lorini, E. (2012). A logic of emotions: from appraisal to coping. In W. van der Hoek, L. Padgham, V. Conitzer, & M. Winikoff (Eds.), *Eleventh International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 1133-1140.
- [47] Dastani, M., Hulstijn, J., & van der Torre, L. (2005). How to decide what to do? *European Journal of Operational Research*, 160, 762-784.
- [48] Davis, W. (1981). A theory of happiness. *Philosophical studies*, 39, 305-317.
- [49] Dawkins, R. (1983). Universal Darwinism. In D. S. Bendall (Ed.), *Evolution from molecules to man* (pp. 403-425). Cambridge: Cambridge University Press.
- [50] Day, J. P. (1970). The anatomy of hope and fear. *Mind*, 79, 369-384.
- [51] Dennett, D. C. (1971). Intentional systems. *Journal of Philosophy*, 68, 87-106.
- [52] Dias, J., Mascaranhas, S., & Paiva, A. (2011). Fatima modular: Towards an agent architecture with a generic appraisal framework. *Proceedings of the International Workshop on Standards for Emotion Modeling*.
- [53] Dias, J., & Paiva, A. (2005). Feeling and reasoning: A computational model for emotional agents. *Proceedings of the 12th Portuguese Conference on Artificial Intelligence (EPIA 2005)* (pp. 127-140). Portugal: Springer.
- [54] Dix, J., & Fisher, M. (2011). Where logic and agents meet. *Annals of Mathematics and Artificial Intelligence*, 61, 15-28.
- [55] Dörner, D. (2001). *Bauplan für eine Seele [Blueprint for a mind]*. Reinbek: Rowohlt.
- [56] Duc, H. N. (2001). Resource-bounded reasoning about knowledge. *Ph.D. thesis, Faculty of Mathematics and Informatics, University of Leipzig*.
- [57] Elio, R. (2002). Belief-desire-intention agency in a general cognitive architecture. *Cognitive Science Quarterly*, 2, 321-340.
- [58] Fagin, R., Halpern, J. Y., Moses, Y., & Vardi, M. Y. (1995). *Reasoning about knowledge*. Cambridge, MA: MIT Press.
- [59] Feather, N. T. (Ed.). (1982). *Expectations and actions: Expectancy-value models in psychology*. Hillsdale, NJ: Erlbaum.
- [60] Fox, C. R., & Poldrack, R. A. (2009). Prospect theory and the brain. In Glimcher, P. W., Camerer, C. F., Fehr, E., & Poldrack, R. A. (Eds.), *Neuroeconomics: Decision Making and the Brain* (pp. 145-174). London: Elsevier/Academic Press.
- [61] Frijda, N. H. (1967). Problems of computer simulation. *Behavioral Science*, 12, 59-67.
- [62] Frijda, N. H. (1986). *The emotions*. Cambridge, England: Cambridge University Press.
- [63] Frijda, N. H. (1994). Emotions are functional, most of the time. In Ekman, P., & Davidson, R. J. (Eds.), *The nature of emotion* (pp. 112-136). NY: Oxford University Press.
- [64] Frijda, N. H., Ortony, A., Sonnemans, J., & Clore, G. L. (1992). The complexity of intensity: Issues concerning the structure of emotion intensity. In M. S. Clark (Ed.), *Review of Personality and Social Psychology* (Vol. 13, pp. 60-89). Newbury Park: Sage.
- [65] Georgeff, M., Pell, B., Pollack, M., Tambe, M. & Wooldridge, M. (1998). The Belief-Desire-Intention model of agency. In J. P. Müller, M. P. Singh, & A. S. Rao (1999). *Intelligent Agents V: Agent Theories, Architectures, and Languages. 5th International Workshop, ATAL'98* (pp. 1-10). Berlin: Springer.
- [66] Goldthorpe, J. H. (1998). Rational action theory for sociology. *British Journal of Sociology*, 49, 167-192.
- [67] Gratch, J. (2012). The sciences of the artificial emotions: Comment on Aylett and Paiva. *Emotion Review*, 4, 266-268.
- [68] Gratch, J., & Marsella, S. (2004). A domain independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5, 269-306.
- [69] Gratch, J., Marsella, S., Wang, N. & Stankovic, B. (2009). Assessing the validity of appraisal-based models of emotion. *Proceedings of the International Conference on Affective Computing and Intelligent Interaction*.
- [70] Green, O. H. (1992). *The emotions: A philosophical theory*. Dordrecht: Kluwer.
- [71] Guiraud, N., Longin, D., Lorini, E., Pesty, S., & Rivière, J. (2011). The face of emotions: a logical formalization of expressive speech acts. *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 1031-1038. ACM Press.
- [72] Heise, E. (1992). Anderson's ACT*-theory applied to skill acquisition: From the theoretical framework to empirical content. In H. Westmeyer (Ed.), *The structuralist program in psychology: Foundations and applications* (pp. 87-103). Toronto: Hogrefe & Huber.
- [73] Hindriks, K. V. (2009). Programming rational agents in GOAL. In: R. H. Bordini, M. Dastani, J. Dix, & A. El Fallah Seghrouchni (Eds.), *Multi-agent Programming: Languages, tools and applications* (pp.119-157). Berlin: Springer.
- [74] Hindriks, K. V., & Broekens, J. (2011). *Comparing formal cognitive emotion theories*. Manuscript, Delft University of Technology.
- [75] Huber, M. J. (1999). JAM: A BDI-theoretic mobile agent architecture. *Proceedings of the Third International Conference on Autonomous Agents (Agents'99)*, pp. 236-243.
- [76] Hudlicka, E. (1998). Modeling emotion in symbolic cognitive architectures. *AAAI Fall Symposium*, TR FS-98-03. Menlo Park, CA: AAAI Press.
- [77] Hudlicka, E. (2004). Two sides of appraisal: Implementing appraisal and its consequences within a cognitive architecture. *Proceedings of the AAAI Spring Symposium, Architectures for modeling emotion*. Technical Report SS-04-02. Menlo Park, CA: AAAI Press.
- [78] Hudlicka, E. (2008a). Review of cognitive-affective architectures. In G. Zacharias, J. McMillan, and S. Van Hemel (Eds.), *Organizational modeling: From individuals to societies* (pp. 167-184). Washington, DC: National Academies Press.
- [79] Hudlicka, E. (2008b). Modeling the mechanisms of emotion effects on cognition. *AAAI Fall Symposium: Biologically inspired cognitive architectures* (TR FS-08-04, pp. 82-86). Menlo Park, CA: AAAI Press.
- [80] Hudlicka, E. (2011). Guidelines for developing computational models of emotions. *International Journal of Synthetic Emotions*, 2, 26-79.
- [81] Hughes, G. E., & Cresswell, M. J. (1996). *A new introduction to modal logic*. London: Routledge.
- [82] James, W. (1890/1960). *Principles of psychology* (Vol. 1 & 2). New York: Dover.
- [83] Jones, R. M., Lebiere, C., & Crossman, J. A. (2007). Comparing modeling idioms in ACT-R and Soar. In: *Proceedings of ICCM 2007 (Eighth International Conference on Cognitive Modeling)*, pp. 49-54. Oxford, UK: Taylor & Francis/Psychology Press.
- [84] Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263-291.
- [85] Keltner, D., & Kring, A. M. (1998). Emotion, social function, and psychopathology. *Review of General Psychology*, 3, 320-342.
- [86] Laird, J. D. (2007). *Feelings: The perception of self*. New York: Oxford University Press.
- [87] Laird, J. E. (2012). *The Soar cognitive architecture*. Cambridge, MA: MIT Press.
- [88] Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33, 1-64.
- [89] Langley, P. (2006). Cognitive architectures and general intelligent

- systems. *AI Magazine*, 27, 33-44.
- [90] Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10, 141-160.
- [91] Lazarus, R. S. (1982). Thoughts on the relations between emotion and cognition. *American Psychologist*, 37, 1019-1024.
- [92] Lazarus, R. S. (1991). *Emotion and adaptation*. New York: Oxford University Press.
- [93] Leventhal, H., & Scherer, K. R. (1987). The relationship of emotion and cognition: A functional approach to a semantic controversy. *Cognition and Emotion*, 1, 3-28.
- [94] Lewis, M., Haviland-Jones, J. M., & Barrett, L. F. (2008). *Handbook of Emotions (3rd Ed.)*. New York: Guilford Press.
- [95] Lorini, E. (2011). A dynamic logic of knowledge, graded beliefs and graded goals and its application to emotion modelling. In H. van Ditmarsch, J. Lang, & S. Ju (Eds.), *LORI-III workshop on logic, rationality and interaction* (Vol. 6953 of LNAI, pp. 165-178). Berlin: Springer.
- [96] Lorini, E. (2012). The cognitive anatomy and functions of expectations revisited. In F. Paglieri, L. Tummolini, R. Falcone, & M. Miceli (Eds.), *The goals of cognition. Essays in honor of Cristiano Castelfranchi* (pp. 273-288). London: College Publications.
- [97] Lorini, E., & Piunti, M. (2007). The benefits of surprise in dynamic environments: From theory to practice. *Affective Computing and Intelligent Interaction* (Lecture Notes in Computer Science 4738), 362-373.
- [98] Lorini, E., & Schwarzenhuber, F. (2011). A logic for reasoning about counterfactual emotions. *Artificial Intelligence*, 175, 814-847.
- [99] Macedo, L., & Cardoso, A. (2012). The exploration of unknown environments populated with entities by a surprise-curiosity-based agent. *Cognitive Systems Research*, 19-20, 62-87.
- [100] Macedo, L., Cardoso, A., Reizenzein, R., Lorini, L., & Castelfranchi, C. (2009). Artificial surprise. In J. Vallverdú, & D. Casacuberta, (Eds.) *Handbook of research on synthetic emotions and sociable robotics: New applications in affective computing and artificial intelligence* (pp. 267-291). Hershey, PA: IGI Global.
- [101] Mandler, G. (1984). *Mind and body*. New York: Norton.
- [102] Manhart, K. (1989). Können AI-Programme als Theorien betrachtet werden? [Can AI programs be regarded as theories?]. *Fifth Austrian Meeting for Artificial Intelligence. Informatik-Fachberichte*, 208, 346-358.
- [103] Marewski, J. N., & Mehlhorn, K. (2011). Using the ACT-R architecture to specify 39 quantitative process models of decision making. *Judgment and Decision Making*, 6, 439-519.
- [104] Marinier, R., Laird, J. & Lewis, R. (2009). A computational unification of cognitive behavior and emotion. *Journal of Cognitive Systems Research*, 10, 48-69.
- [105] Marks, J. (1982). A theory of emotion. *Philosophical Studies*, 42, 227-242.
- [106] Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Freeman.
- [107] Marsella, S., & Gratch, J. (2009). EMA: A process model of appraisal dynamics. *Journal of Cognitive Systems Research*, 10, 70-90
- [108] Marsella, S., Gratch, J., & Petta, P. (2010). Computational models of emotion. In K. R. Scherer, T. Bänziger, & E. Roesch (Eds.), *Blueprint for affective computing: A sourcebook* (pp. 21-46). Oxford: Oxford University Press.
- [109] McDougall, W. (1908/1960). *An introduction to social psychology*. London: Methuen.
- [110] McDougall, W. (1923). *An outline of psychology*. London: Methuen.
- [111] Meinong, A. (1894). *Psychologisch-ethische Untersuchungen zur Werttheorie* [Psychological-ethical investigations in value theory]. Graz: Leuschner & Lubensky.
- [112] Mele, A. R. (1992). *Springs of action: Understanding intentional behavior*. NY: Oxford University Press.
- [113] Mellers, B. A. (2000). Choice and the relative pleasure of consequences. *Psychological Bulletin*, 126, 910-924.
- [114] Meyer, J.-J. Ch. (2006). Reasoning about emotional agents. *International Journal of Intelligent Systems*, 21, 601-619.
- [115] Meyer, J.-J. Ch., & van der Hoek, W. (1995). Epistemic logic for AI and computer science. *Cambridge Tracts in Theoretical Computer Science 41*. Cambridge: Cambridge University Press.
- [116] Miller, G., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York: Holt.
- [117] Mitchell, S. (1995). Function, fitness and disposition. *Biology and Philosophy*, 10, 39-54.
- [118] Moffat, D., Frijda, N., & Phaf, H. (1993). Analysis of a model of emotions. In A. Sloman, D. Hogg, G. Humphreys, A. Ramsay, & D. Partridge (Eds.), *Prospects for Artificial Intelligence* (pp. 219-228). Amsterdam: IOS Press.
- [119] Neisser, U. (1967). *Cognitive psychology*. NY: Appleton-Century-Crofts.
- [120] Nesse, R. M. (2001). The smoke detector principle: Natural selection and the regulation of defenses. *Annals of the New York Academy of Sciences*, 935, 75-85.
- [121] Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283-308). New York: Academic Press.
- [122] Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18, 87-127.
- [123] Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard.
- [124] Newell, A. & Simon, H. A. (1963). GPS, a program that simulates human thought. In E. A. Feigenbaum & J. Feldman (eds.), *Computers and thought* (pp. 279-293). New York: McGraw-Hill.
- [125] Norling, E., & Ritter, E. (2004). Towards supporting psychologically plausible variability in agent-based human modelling. *AAMAS 2004*, 758-765.
- [126] Oatley, K. (2009). Communications to self and others: Emotional experience and its skills. *Emotion Review*, 1, 206-213.
- [127] Oatley, K., & Johnson-Laird, P. N. (1987). Towards a cognitive theory of emotions. *Cognition and Emotion*, 1, 29-50.
- [128] Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: towards an evolved module of fear and fear learning. *Psychological Review*, 108, 483-522.
- [129] Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions*. New York: Cambridge University Press.
- [130] Ortony, A., Norman, D., & Revelle, W. (2005). Affect and proto-affect in effective functioning. In J. M. Fellous & M. A. Arbib (Eds.), *Who needs emotions: The brain meets the machine*. New York: Oxford University Press.
- [131] Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 259-302). Hillsdale, NJ: Erlbaum.
- [132] Peira, N., Golkar, A., Öhman, A., Anders, S., & Wiens, S. (2012). Emotional responses in spider fear are closely related to picture awareness. *Cognition & Emotion*, 26, 252-260
- [133] Pereira, D., Oliveira, E., & Moreira, N. (2008). Formal modelling of emotions in BDI agents. In F. Sadri, & K. Satoh (Eds.), *Eighth Workshop on Computational logic in Multi-Agent Systems (Lecture Notes in Computer Science 5056)*, 62-81. New York: Springer.
- [134] Pfeifer R. (1996). Symbols, patterns, and behavior: towards a new understanding of intelligence. In *Proceedings of the Japanese Conference on Artificial Intelligence*, 1-15.
- [135] Picard, R. W. (1997). *Affective computing*. Cambridge: MIT Press.
- [136] Pokahr, A., Braubach, L., & Lamersdorf, W. (2005a). JADEX: a BDI reasoning engine. In R. H. Bordini, M. Dastani, J. Dix & A. El Fallah-

- Seghrouchni (Eds.), *Multi-agent programming: Languages, platforms and applications* (pp. 149-174). New York: Springer.
- [137] Pokahr, A., Braubach, W., & Lamersdorf, W. (2005b). A flexible BDI architecture supporting extensibility. In A. Skowron, J.-P. Barthès, L. C. Jain, R. Sun, P. Morizet-Mahoudeaux, J. Liu, & N. Zhong (Eds.), *Proceedings of the 2005 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, 379-385, IEEE Computer Society.
- [138] Pollock, J. L. (1989). OSCAR: A general theory of rationality. *Journal of Experimental and Theoretical Artificial Intelligence*, 1, 209-226.
- [139] Pynadath, D. V., & Marsella, S. (2005). PsychSim: Modeling theory of mind with decision-theoretic agents. *International Joint Conference on Artificial Intelligence 2005*, 1181-1186.
- [140] Rank S. (2004). *Affective acting: An appraisal-based architecture for agents as actors*. Master's thesis, Institut für Medizinische Kybernetik und Artificial Intelligence, University of Vienna, Austria. <http://www.ofai.at/~stefan.rank/StefanRank-AAAThesis.pdf>
- [141] Rank, S. (2005). Towards reusable roleplayers using an appraisal-based architecture. *Applied Artificial Intelligence, Special issue: Educational Agents and (e-)Learning*, 19, 313-340.
- [142] Rao, A. S., & Georgeff, M. P. (1991). Modeling rational agents within a BDI-architecture. In J. Allen, R. Fikes, & E. Sandewall (Eds.), *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, (pp. 473-484). Morgan Kaufmann.
- [143] Rao, A. S., & Georgeff, M. P. (1995). BDI Agents: From theory to practice. In Lesser, V. R., & Gasser L. (Eds.): *Proceedings of the First International Conference on Multiagent Systems (ICMAS 1995)*, pp. 312-319. San Francisco: MIT Press.
- [144] Reizenzein, R. (1994). Pleasure-arousal theory and the intensity of emotions. *Journal of Personality and Social Psychology*, 67, 525-539.
- [145] Reizenzein, R. (1996). Emotional action generation. In W. Battmann & S. Dutke (Eds.), *Processes of the molar regulation of behavior* (pp. 151-165). Lengerich: Pabst Science Publishers.
- [146] Reizenzein, R. (2000). Wundt's three-dimensional theory of emotion. In W. Balzer, & J. D. Sneed, & C. U. Moulines (Eds.), *Structuralist knowledge representations: Paradigmatic examples (Poznan studies in the philosophy of the sciences and the humanities, Vol. 75)*, pp. 219-250). Amsterdam: Rodopi.
- [147] Reizenzein, R. (2001). Appraisal processes conceptualized from a schema-theoretic perspective: Contributions to a process analysis of emotions. In K. R. Scherer, A. Schorr & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, methods, research* (pp. 187-201). NY: Oxford University Press.
- [148] Reizenzein, R. (2006). Arnold's theory of emotion in historical perspective. *Cognition and Emotion*, 20, 920-951.
- [149] Reizenzein, R. (2007). What is a definition of emotion? And are emotions mental-behavioral processes? *Social Science Information*, 46, 424-428.
- [150] Reizenzein, R. (2009a). Emotional experience in the computational belief-desire theory of emotion. *Emotion Review*, 1, 214-222.
- [151] Reizenzein, R. (2009b). Emotions as metarepresentational states of mind: Naturalizing the belief-desire theory of emotion. *Cognitive Systems Research*, 10, 6-20.
- [152] Reizenzein, R. (2010). Broadening the scope of affect detection research. *IEEE Transactions on Affective Computing*, 1, 42-45.
- [153] Reizenzein, R. (2012). What is an emotion in the Belief-Desire theory of emotion? In F. Paglieri, L. Tummolini, R. Falcone, & M. Miceli (Eds.), *The goals of cognition. Essays in honor of Cristiano Castelfranchi* (pp. 181-211). London: College Publications.
- [154] Reizenzein, R., & Döring, S. (2009). Ten perspectives on emotional experience: Introduction to the special issue. *Emotion Review*, 1, 195-205.
- [155] Reizenzein, R., & Horstmann, G. (2006). Emotion. In H. Spada (Ed.), *Lehrbuch Allgemeine Psychologie* (3rd Ed., pp. 435-500). Bern: Huber.
- [156] Reizenzein, R., & Junge, M. (2006). Überraschung, Enttäuschung und Erleichterung: Emotionsintensität als Funktion von subjektiver Wahrscheinlichkeit und Erwünschtheit [Surprise, disappointment and relief: Emotion intensity as function of subjective probability and desirability]. *Talk presented at the 45th Congress of the German Psychological Association in Nuremberg*.
- [157] Reizenzein, R., & Junge, M. (2012). Language and emotion from the perspective of the computational belief-desire theory of emotion. In P. A. Wilson (Ed.) *Dynamicity in emotion concepts (Lodz Studies in Language*, 27, 37-59). Frankfurt am Main: Peter Lang.
- [158] Reizenzein, R., & Weber, H. (2009). Personality and emotion. In P. J. Corr & G. Matthews (Eds.), *Cambridge handbook of personality psychology* (pp. 54-71). Oxford: Oxford University Press.
- [159] Ritter, F. E. (2012). CoJACK: A high-level cognitive architecture with demonstrations of moderators, variability, and implications for situation awareness. *Biologically Inspired Cognitive Architectures*, 1, 2-13.
- [160] Ritter, F. E., Avraamides, M., & Councill, I. G. (2002). Validating changes to a cognitive architecture to more accurately model the effects of two example behavior moderators. *Proceedings of the 11th Computer-Generated Forces and Behavior Representation Conference* (pp. 29-40), 02-CGF-002. Orlando, FL: University of Central Florida.
- [161] Roseman, I. J. (2001). A model of appraisal in the emotion system. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, Methods, Research* (pp. 68-91). Oxford: Oxford University Press.
- [162] Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39, 1161-1178.
- [163] Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, 110, 145-172.
- [164] Russell, S. J., & Norvig, P. (2010). *Artificial Intelligence: A modern approach (3rd Ed.)*. Prentice Hall.
- [165] Schachter, S. (1964). The interaction of cognitive and physiological determinants of emotional state. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 1, pp. 49-80). New York: Academic Press.
- [166] Scherer, K. R. (2001). Appraisal considered as a process of multilevel sequential checking. In K. R. Scherer, A. Schorr, & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, Methods, Research* (pp. 92-129). Oxford: Oxford University Press.
- [167] Scherer, K. R. (2010). Emotion and emotional competence: conceptual and theoretical issues for modelling agents. In K. R. Scherer, T. Bänziger, T., & E. Roesch (Eds.), *Blueprint for affective computing: A sourcebook* (pp. 3-20). Oxford: Oxford University Press.
- [168] Scherer, K. R., Bänziger, T., & Roesch, E. (2010). *Blueprint for affective computing: A sourcebook*. Oxford: Oxford University Press.
- [169] Schut, M., & Wooldridge, M. (2001). The control of reasoning in resource-bounded agents. *Knowledge Engineering Review*, 16, 215-240.
- [170] Schwarz, N., & Clore, G. L. (2007). Feelings and phenomenal experiences. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (2nd ed., pp. 385-407). New York: Guilford.
- [171] Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- [172] Shah, J. Y., & Gardner, W. L. (2008). *Handbook of motivation science*. New York: Guilford Press.
- [173] Siemer, M., & Reizenzein, R. (2007a). The process of emotion inference. *Emotion*, 7, 1-20.
- [174] Siemer, M., & Reizenzein, R. (2007b). Appraisals and emotions: Can you have one without the other? *Emotion*, 7, 26-29.
- [175] Siemer, M., Mauss, I., & Gross, J. J. (2007). Same situation-different emotions: how appraisals shape our emotions. *Emotion*, 7, 592-600.

- [176] Simon, H. A. (1967). Motivational and emotional controls of cognition. *Psychological Review*, 74, 29-39.
- [177] Slife, B. D., & Williams, R. N. (1997). Toward a theoretical psychology: Should a subdiscipline be formally recognized? *American Psychologist*, 52, 117-129.
- [178] Sloman, A. (1992). Prolegomena to a theory of communication and affect. In A. Ortony, J. Slack, & O. Stock (Eds.), *Communication from an artificial intelligence perspective: Theoretical and applied issues* (pp. 229-260). Heidelberg: Springer.
- [179] Sloman, A., & Croucher, M. (1981). Why robots will have emotions. In P. J. Hayes (Ed.), *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI-81)*, pp. 197-202. Vancouver, CA: William Kaufmann.
- [180] Slovic, P., Finucane, M., Peters, E., & MacGregor, D. (2004). Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk, and rationality. *Risk Analysis*, 24, 311-322.
- [181] Smith, C. A., & Kirby, L. D. (2000). Consequents require antecedents: Toward a process model of emotion elicitation. In J. P. Forgas (Ed.), *Feeling and thinking: The role of affect in social cognition* (pp. 83-105). Cambridge: Cambridge University Press.
- [182] Smith, C. A., & Lazarus, R. S. (1990). Emotion and adaptation. In L. Pervin (Ed.), *Handbook of Personality: Theory and Research* (pp. 609-637). NY: Guilford.
- [183] Smith, M. (1997). *The moral problem*. Oxford: Blackwell.
- [184] Stegmüller, W. (1986). *Theorie und Erfahrung. Dritter Teilband: Die Entwicklung des neueren Strukturalismus seit 1973. [Theory and experience, Vol 3: The development of recent structuralism since 1973]*. Berlin: Springer.
- [185] Steunebrink, B. R., Dastani, M., & Meyer, J.-J. Ch. (2008). A formal model of emotions: integrating qualitative and quantitative aspects. In G. Mali, C. D. Spyropoulos, N. Fakotakis & N. Avouris (Eds.), *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI08)* (pp. 256-260). Amsterdam: IOS Press.
- [186] Steunebrink, B. R., Dastani, M. M. & Meyer, J.-J. Ch. (2009). The OCC model revisited. In D. Reichardt (Ed.), *Proceedings of the 4th Workshop on Emotion and Computing - Current Research and Future Impact* (pp. 40-47). Paderborn, Germany.
- [187] Steunebrink, B. R., Dastani, M., & Meyer, J.-J. Ch. (2012). A formal model of emotion triggers: an approach for BDI agents. *Synthese*, 185, 83-129.
- [188] Stocco, A., Fum, D., & Zalla, T. (2005). Revising the role of somatic markers in the Gambling Task: A computational account for neuropsychological impairments. In B. Bara, L. Barsalou & M. Bucciarrelli (Eds.), *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates.
- [189] Storbeck, J., & Clore, G. L. (2007). Emotional controls on cognitive processes. *Cognition and Emotion*, 21, 1212-1237.
- [190] Strongman, K. T. (2003). *The psychology of emotion: From everyday life to theory (5th ed.)*. New York: Wiley.
- [191] Sun, R. (2008a). *Cambridge handbook of computational psychology*. Cambridge: Cambridge University Press.
- [192] Sun, R. (2008b). Introduction to computational cognitive modeling. In R. Sun (Ed.), *Cambridge handbook of computational psychology* (pp. 3-19). Cambridge: Cambridge University Press.
- [193] Taatgen, N. A., & Anderson, J. R. (2008). Constraints in cognitive architectures. In R. Sun (Ed.), *Cambridge Handbook of Computational Psychology* (pp. 170-185). Cambridge: Cambridge University Press.
- [194] Taatgen, N. A., & Anderson, J. R. (2010). The past, present, and future of cognitive architectures. *Topics in Cognitive Science*, 4, 693-704.
- [195] Taylor, C. C. W. (1986). Emotions and wants. In J. Marks (Ed.), *The ways of desire* (pp. 217-231). Chicago, IL: Precedent.
- [196] Tooby, J., & Cosmides, L. (1990). The past explains the present: Emotional adaptations and the structure of ancestral environments. *Ethology and Sociobiology*, 11, 375-424.
- [197] Troitzsch, K. G. (1996). Simulation and structuralism. In R. Hegselmann, U. Mueller, & K. G. Troitzsch (Eds.), *Modelling and simulation in the social sciences from a philosophy of science point of view* (pp. 183-208). Kluwer: Dordrecht.
- [198] Troitzsch, K. G. (2012). Theory reconstruction of several versions of modern organisation theories. In A. Tolk (Ed.): *Ontology, epistemology, and teleology for modeling and simulation. Philosophical foundations for intelligent M&S applications* (pp. 121-140). Heidelberg: Springer.
- [199] van der Hoek, W., & Wooldridge, M. (2003). Towards a logic of rational agency. *Logic Journal of the IGPL*, 11, 133-157..
- [200] van Linder, B., van der Hoek, W., & Meyer, J.-J. Ch. (1998). Formalizing abilities and opportunities of agents. *Fundamenta Informaticae*, 34, 53-101.
- [201] Vernon, D., Metta, G., & Sandini, G. (2007). A survey of artificial cognitive systems: implications for the autonomous development of mental capabilities in computational agents. *IEEE Transactions on Evolutionary Computation*, 11, 151-180.
- [202] Weiner, B. (1992). *Human motivation: Metaphors, theory and research*. New York: Wiley.
- [203] Weiner, B. (1995). *Judgments of responsibility. A foundation for a theory of social conduct* New York: Guilford.
- [204] Westmeyer, H. (1992a). (Ed.). *The structuralist program in psychology: Foundations and applications*. Toronto: Hogrefe & Huber.
- [205] Westmeyer, H. (1992b). *Introduction to the structuralist program in psychology*. In H. Westmeyer (Ed.), *The structuralist program in psychology: Foundations and applications* (pp. 1-12). Toronto: Hogrefe & Huber.
- [206] Wooldridge, M. J. (2000). *Reasoning about rational agents*. Cambridge, MA: MIT Press.
- [207] Wundt, W. (1896). *Grundriss der Psychologie [Outlines of Psychology]*. Leipzig: Engelmann.
- [208] Zacharias, G. L., Macmillan, J., & Van Hemel, S. B. (2008). *Behavioral modeling and simulation: From individuals to societies*. Washington DC: National Academies Press.
- [209] Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35, 151-175.



Rainer Reisenzein received the PhD degree in psychology from the University of Salzburg, Austria. After spending a postdoctoral research year at the University of California Los Angeles, he was an assistant professor of psychology at the Free University Berlin and then at the University of Bielefeld. Since 2002, he has been a full professor of General Psychology at the University of Greifswald, Germany. His research focuses on theoretical and empirical questions related to emotion and motivation and is marked by an interdisciplinary orientation toward, in particular, philosophy (philosophy of science and philosophy of mind) and cognitive science. Rainer Reisenzein has published more than 100 theoretical, review, and empirical articles and book chapters in English and German on diverse topics, mostly related to emotion, including emotional experience, emotion concepts, the facial expression of emotion, basic emotions, cognitive emotion theory, and the history of emotion research. A special research focus is the phenomenon of surprise in humans, considered from the perspective of emotion theory. Rainer Reisenzein is also a co-author (with Wulf-Uwe Meyer und Achim Schützwohl) of a German three-volume introduction to the psychology of emotions. In 2009, he edited (together with philosopher Sabine Döring) a special issue of *Emotion*

Review on the topic of emotional experience. He is a member of the HUMAINE Association (the international society for research on emotion and human-computer interaction) and of the International Society for Research on Emotion (ISRE), and is currently on the editorial board of the journals *Cognition and Emotion* and *Experimental Psychology*. Homepage: <http://www.mnf.uni-greifswald.de/institute/institut-fuer-psychologie/lehrstuehle/allgemeine-psychologie-ii/personal/homepage-rainer-reisenzein.html>



Eva Hudlicka was born in Prague, Czech Republic. She received her BS in Biochemistry from Virginia Tech, her MS in Computer Science from The Ohio State University, and her PhD in Computer Science from the University of Massachusetts-Amherst. Subsequently she was a Senior Scientist at Bolt Beranek & Newman, Cambridge, MA. In 1995, she founded Psychometrix Associates, Inc. to pursue research in affective computing. Her primary research focus is the development of computational models

of emotion; both the cognitive processes involved in appraisal, and the effects of emotions on cognition. The aim of this work is an improved understanding of the mechanisms of emotion generation and emotion-cognition interactions, including the nature of affective biases on cognition. The findings have practical implications for the development of affective user models and affective virtual agents across a range of applications, including serious games for training and therapy, and for cognitive-affective engineering in general. Specific applications include: training and education, decision-support systems, assistive technologies, assessment and treatment of cognitive-affective disorders, and affective gaming. Hudlicka's research is conducted within the context of a computational cognitive-affective architecture developed by Psychometrix, the MAMID architecture (Methodology for Modeling and Analysis of Individual Differences), which implements a generic methodology for modeling the interacting effects of multiple affective factors and individual differences on decision-making. In 2008, Eva Hudlicka was a member of a National Research Council committee on "Organizational Models: From Individuals to Societies", contributing her expertise on modeling individual behavior with a focus on affective processes. She is an associate editor of the *International Journal of Synthetic Emotions* and a member of the editorial board of the Oxford University Press series on *Cognitive Models and Architectures*, the *International Journal of Machine Consciousness* and the *International Journal of Biologically Inspired Cognitive Architectures*. She has taught courses and tutorials in "Affective Computing", "Affective Computing for Game Design" and "Computational Emotion Modeling", and has authored numerous book chapters, articles and reports in this research area. Homepage: <http://psychometrixassociates.com/>



Mehdi Dastani is associate professor in computer science at Utrecht University in the Netherlands. He received his Ph.D. degree in humanities from the University of Amsterdam for his dissertation on "Languages of Perception". He holds a MS degree in computer science and a MS degree in philosophy from the University of Amsterdam. His research focuses on intelligent systems, in particular logical and computational models of social and cognitive phenomena such as emotion, norms, decision-making, interaction, and perception. Currently, Mehdi Dastani is working in the area of multi-agent systems, multi-agent logics and multi-agent programming. He has developed various multi-agent (programming) frameworks such as BOID (Beliefs, Obligations, Intentions and Desires), 2APL (A Practical Agent Programming Language) and 2OPL (Organisation Oriented Programming Language) which support the development and implementation of multi-agent systems. These frameworks are being used by many universities for teaching and training purposes.

An important line of research in Dastani's work concerns emotion modelling. In this domain, Dastani has developed several logical models for emotion processes such as appraisal, emotion experience and coping, which are formalizations of the psychological emotion models proposed by Ortony, Clore and Collins, Frijda, and Lazarus. In addition to being the author of numerous technical articles, Mehdi Dastani has co-organized numerous international events related to multi-agent systems, including various editions of the International conference of Autonomous Agents and Multi-Agent Systems (AAMAS), the International workshop on Multi-Agent Programming (ProMAS), the Lorentz seminar on Multi-Agent Organizations and the Dagstuhl seminar on Multi-Agent Programming. He has also been on the chair of the European Agent System Summer Schools (EASSS) between 2007 and 2010, the Nineteenth Belgian-Dutch Conference on Artificial Intelligence (BNAIC 2007), and the Fifth European workshop on Multi-Agent Systems (EUMAS 2007). In addition, Mehdi Dastani has been a member of the senior programme committee of many international conferences in artificial intelligence (e.g., IJCAI, ECAI, AAI, AAMAS). Homepage: <http://www.cs.uu.nl/staff/mehdi.html>



Jonathan Gratch is an Associate Director for Virtual Human Research at the University of Southern California's (USC) Institute for Creative Technologies, Research Associate Professor of Computer Science and Psychology at USC and co-director of USC's Computational Emotion Group. He completed his Ph.D. in Computer Science at the University of Illinois in Urban-Champaign in 1995. Dr. Gratch's research focuses on computational models of human cognitive and social processes, especially emotion, and explores these models' role in shaping human-computer interactions in virtual environments. He studies the relationship between cognition and emotion, the cognitive processes underlying emotional responses, and the influence of emotion on decision making and physical behavior. He is the founding and current Editor-in-Chief of IEEE's *Transactions on Affective Computing*, Associate Editor of *Emotion Review* and the *Journal of Autonomous Agents and Multi-agent Systems*, former President of the HUMAINE Association, the international society for research on emotion and human-computer interaction, and is a member of IEEE, the Association for the Advancement of Artificial Intelligence (AAAI) and the International Society for Research on Emotion (ISRE). Dr. Gratch is the author of over 200 technical articles. Homepage: <http://www.ict.usc.edu/~gratch>

Koen Hindriks is Assistant Professor at the Interactive Intelligence group at the Faculty of Electrical Engineering, Mathematics and



Computer Science of the Delft University of Technology in the Netherlands. He studied computing science and completed his PhD at Utrecht University on agent programming languages. He has published more than 100 articles on agent technology and organized various events in this area including the 2012 Dagstuhl Seminar Engineering Multi-Agent Systems. His main research interests are cognitive agent technology and coordination models for effective multi-agent interaction. His research focuses on the analysis, modeling, and development of agent technology that integrates different aspects of intelligence such as reasoning, decision-making, planning, learning and interaction, as well as aspects of emotional intelligence. This multi-agent technology has been applied, among others, in micro-simulations of domains such as traffic, logistics and supply chain management, serious gaming, negotiation, socio-cognitive robotics, and user modeling. He has designed and developed the agent programming languages 3APL and GOAL and has worked on the verification and specification of agent programs. Homepage: <http://mmi.tudelft.nl/~koen/>



Emiliano Lorini is a research scientist at the Institut de Recherche en Informatique de Toulouse (IRIT) of the French National Research Council (CNRS). His research topics include epistemic logic and formal epistemology, game theory, logics of intention, logics of agency, deontic logic, the theory of emotion, and the theory of trust and reputation. He has authored or co-authored more than 80 articles in journals and international conferences in the fields of AI, logic, cognitive science and game theory.

Emiliano Lorini has been a reviewer for major journals and conferences in AI, cognitive science and economics, including *AI Journal*, *Mathematical Social Sciences*, *Journal of Autonomous Agents and Multi-Agent Systems* (JAAMAS), *Synthese*, *Cognitive Systems Research*, *Journal of Philosophical Logic*, AAI, IJCAI, KR, AAMAS, and ECAI. He has been a member of the program committee of several major conferences in the field of AI and cognitive science, including AAMAS, ECAI, IJCAI 2011, and CogSci, and has co-organized several conferences and workshops in the fields of AI, game theory, and logic. He has also been editor of special issues of, respectively, the *Journal of Philosophical Logic* and *Synthese*. Homepage: <http://www.irit.fr/~Emiliano.Lorini/>



John-Jules Ch. Meyer studied Mathematics with Computer Science and Digital Signal Processing at Leyden University. He obtained his Ph.D. from the VU University in Amsterdam on a subject in theoretical computer science. From 1988 to 1993 he was a full professor both at the VU University in Amsterdam and the University of Nijmegen. Since 1993, he has been a full professor of computer science at Utrecht University (UU). Currently he is heading the Artificial Intelligence Section within the Department of

Information and Computing Sciences and is programme leader of the AI Master of the UU. He is also the CSO of the Alan Turing Institute Almere, which specializes in multi-disciplinary research for the health sciences. He has well over 400 peer-reviewed publications in international journals and conference proceedings, and an h-index of 42 (Google scholar). He is a member of the editorial boards of several international journals including the *Journal of Intelligent Agents & Multi-Agent Systems*. His current research interests include artificial intelligence, and intelligent agents in particular, as to both theory (e.g. agent logics) and applications such as virtual characters in serious games. In 2005 he was appointed as a Fellow of the European Coordinating Committee for Artificial Intelligence (ECCAI). He has been a honorary member of the national research school SIKS (School of Information and Knowledge Systems) since 2007. Homepage: <http://www.alanturinginstitutealmere.nl/john-jules-meyer>