



Constructivist Approaches for Computational Emotions: A Systematic Survey

Alexander Viola¹, Vladimir Pavlovic², and Sejong Yoon¹(✉)

¹ The College of New Jersey, Ewing, NJ 08628, USA
{violaa1,yoons}@tcnj.edu

² Rutgers University, Piscataway, NJ 08854, USA
vladimir@cs.rutgers.edu

Abstract. Computational emotion, is naturally predicated on an operating theory of emotion. This paper seeks to explore the prevalence of three different approaches in the literature, namely basic emotion, dimensional emotion, and constructed emotion. Basic emotion maintains that there exists a discrete set of primitive emotions evolved as responses to certain stimuli; dimensional emotion sees different emotions as systematically related by two or more dimensions (typically valence and arousal); and constructed emotion describes emotional experience as a function of the brain's general predictive faculties applied to learned social concepts of different emotions. In order to see how these approaches are represented in affective computing literature, we conduct a systematic survey spanning the IEEE, ACM, ScienceDirect, and Engineering Village databases. Out of 204 selected papers, 151 apply basic emotion theory, 48 apply dimensional emotion, and 5 apply constructed emotion. We find promising representation of the constructed emotion theory in the affective computing literature and conclude that it provides a theoretical basis worth pursuing for affective engagement human computer interaction (HCI) applications.

Keywords: Constructed emotion · Affective computing · Systematic survey

1 Introduction

The very idea of affective computing, that is, the capacity for computers to perceive or express emotion, took off in Picard's seminal 1995 paper titled *Affective Computing* [43]. In it, she saw the technology of the time and imagined it would soon be capable of reckoning with human emotion in a robust way, imbuing it all with importance with an observation from the field of psychology: emotion is fundamental to the decision-making of all kinds, minor and major, frivolous and life-changing; it undergirds our values and impacts, literally, how we see the world; and at last, it is essential to communication. The idea is, if a computer could develop a sort of *empathy*, an awareness of the moods of its users, it could

become a more helpful tool for a great number of applications. Suppose learning software could detect the interest or frustration of a student and modify a lesson to suit that. Or consider a computer as a tool for those whose jobs are to play with emotion: computer-aided composing, visual art, clip selection. Even entertainment that can swell and recede and shift with its' viewers participation and emotion, or, perhaps, simply giving synthesized speech its proper tonality to convey subtle meaning. You can even find applications in health and safety – suppose a system could determine if a driver was angry and prone to aggressive driving or if a driver was inattentive and liable to cause an accident. These are all examples where computational emotion models can help human-machine interactions in various ways.

Applications of affective computing are so numerous by virtue that emotion is an undercurrent that influences nearly everything in our lives. The practice of affective computing is inherently multidisciplinary, drawing from psychology, neurology, mathematics, computer science, sociology, and linguistics [6] and so can be challenging – but potentially enriching – pursuit. Evidently, many see that potential. In the decades following Picard's paper, we see affective computing applied in as many ways she foresaw and more. We see papers pursuing emotion recognition in faces, speech, and gestures [28,38,58], or in brain scans, heart rates, or skin conductance [3,33,39], emotionality and other subjective attribute detection in music, movies, and visual art [40,54]. There are strides being made in artificial affective agents [34,60,61], and in sentiment analysis of forums, blogs, and social media posts [20,41,63]. The field is lush with a variety of diverse applications and holds promise in expanding the range of computers' usefulness and perhaps someday fundamentally changing the way we interact with them.

As affective computing grows in popularity and as machine learning has become ascendant, the ultimate aim of creating silicon systems that can effectively grasp at and reckon with human emotion seems ever more attainable. Amidst this promise and excitement, however, we argue that it is important to step back and examine the theoretical foundations of our very idea of emotion: how we think about these things informs how we develop affective systems, what we expect from them, how we conceptualize them, and ultimately, how we use them.

The predominant theory of emotion that largely guides current affective computing, **basic emotion**, holds that there are fundamental emotional experiences that a computer (or an observer) can correctly and objectively *detect* in a person. The underdog theory of **constructed emotion**, however, posits that emotion is inherently subjective, impossible to accurately *detect* in a person's face, behavior, or neural activity. It may seem, then, like the very concept of computational emotion prediction is wholly incompatible with this theory. Yet, we seek to find applications that reimagine the roles of these predictive systems in affective application design, creating programs that enhance a user's ability to examine the personal feelings only they are truly equipped to determine. As the main contribution, this paper puts forth the idea that affective computing informed by the *constructed theory of emotion* holds promise in creating systems that a user feels emotionally empowered by, rather than unsettlingly analyzed by, with a systematic survey of the theories.

2 Computational Models of Emotion

Before we discuss the survey methodology and results about various computational emotion models, it would be good to examine these theories more in-depth and grasp at the general form of affective computing papers that apply each of them. As reflected in the results above, both basic emotion and dimensional emotion share the lion’s portion of guiding thought in the field. Affective computing has historically accommodated both of these competing approaches and continues to do so. Picard’s paper over two decades ago mentions this capacity to pursue useful affective research in either vein of theory [43], and the significant presence of both basic and dimensional emotion papers to this day offers testament to this fact. Planting the seeds of constructed emotion in this fertile field very well may yield new, interesting, and applicable research.

2.1 Basic Emotion

The theory of basic emotion has enjoyed prevalence in the literature, introductory psychology courses, and the public’s general science consciousness. Its premise is intuitive and offers a digestible origin story to the sometimes primal-feeling emotions that color our lives in alternately beautiful, tragic, and frightening hues. One of its fundamental premises, universally understood emotion, is also a pleasing and hopeful conclusion to arrive at – it’s something exciting to communicate. In affective computing, its taxonomy of discrete emotions is also pleasantly well-suited for classification models of all stripes.

Summary. Most popular as Ekman’s theory of basic emotion, this approach maintains the existence of six emotions with distinct causal neurology and unique physical expression, developed in response to frequently-encountered situations in our evolutionary history [15]. Namely, the six emotions are anger, disgust, fear, surprise, happiness, sadness, and surprise. The classification of “basic” requires that these responses exhibit aforementioned causal circuits and, ideally, exist in other species as well; among other requirements, these rules differentiate these six from the myriad non-basic emotions that can be considered various modulations or alternations of these basic components. Given an evolutionary basis, this theory also goes hand-in-hand with the concept of universal emotion, i.e., that particular facial configurations and situations can be reliably and consistently classified as evoking one of these six emotions, especially across highly differing cultures. This theory evidently informs affective computing approaches that aim to classify “emotion signals” into corresponding discrete categories, often a subset of the above six emotions. A clear example would be a facial emotion classification model trained on emotion-labeled face images that considers success as an objective detection of emotion as it adheres to these labels (Fig. 1).

Example. *Image based Static Facial Expression Recognition with Multiple Deep Network Learning* is a paper published in 2015 by Yu and Zhang [58] for the Emotion Recognition in the Wild Challenge of that same year. They propose a model

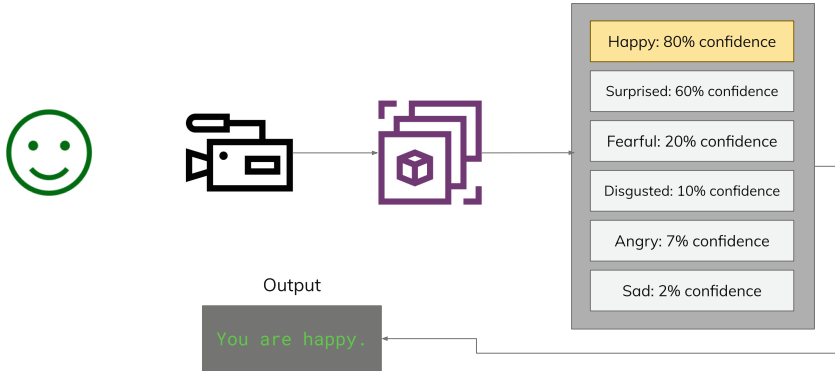


Fig. 1. An illustration of a sample basic emotion approach applied to a facial expression recognition task. A face being examined with a camera attached to an FER model, with outputs showing confidence levels for a variety of emotion classes. The label with the highest confidence is taken as the answer.

to perform an emotion categorization task on the Static Facial Expression in the Wild (SFEW) dataset, placing movie frames of human faces into seven categories, namely *Angry*, *Disgust*, *Fear*, *Happy*, *Neutral*, *Sad*, and *Surprise*. This model is first built on a robust, multi-level facial detection system, with the largest detected area across all levels being used as input for prediction. The highest level is a joint cascade detection and alignment detector, as it is reasonably robust to image perturbations and offers better face localization, the second level a deep CNN detector that offers more robustness in the case of occluded or sharply angled faces, and the last a mixture of trees detector. The prediction model itself is formed by five convolutional layers with three stochastic pooling layers interspersed between to reduce overfitting, three final densely connected layers, and a softmax layer followed with negative log-likelihood loss. For robustness, the paper also generates randomly perturbed images as a part of the input. It considers both the original and perturbed images in prediction and outputs the average voting response of all forms of the image. To further improve performance, multiple differently initialized copies of the model are ensemble, with learned ensemble weights using either optimal ensemble log-likelihood loss or optimal ensemble hinge loss. The network pre-trains on the FER dataset and is fine-tuned on the SFEW training set to the tune of 61.29% accuracy on the challenge’s SFEW test set. This significantly surpasses the challenge baseline accuracy of 39.13% and so proves to be an effective basic emotion classification model that improves on its predecessors via a variety of smart changes.

2.2 Dimensional Emotion

A dimensional representation of emotion aims largely to address perceived shortcomings in a discrete basic emotion approach, primarily issues of applicability to actual emotion experience due to a lack of nuance [24]. Proponents believe that

breaking emotion down into two (or more) dimensions provides such nuance and creates room to render systematic relationships between emotions in the space. Papers applying dimensional theory are free to predict continuous values for various emotion dimensions and leave that as is or may use those values to place a reading within discrete emotion regions in the emotional dimension space [43].

Summary. A dimensional emotion approach relies on Russell’s circumplex model of affect [47], which is based on the hypothesis that emotions may be represented by particular combinations of various dimensions. Russell’s model focuses particularly on the dimensions of valence and arousal (or activity). For example, a state assessed as highly negative (i.e., low valence) with low arousal might be classed as a depressive state; a state assessed as more or less neutral (i.e., moderate valence) with high arousal might be classed as a state of surprise. These states are not entirely independent as in basic emotion, instead of exhibiting systematic relationships to one another – in comparing, say, fear (negative, high arousal) and contentment (positive, lower arousal), they can be considered opposites. Optionally, a dimensional model in this vein may include additional dimensions such as dominance (a Pleasure-Arousal-Dominance (PAD) model), expectation, or intensity depending on desired complexity and nuance (Fig. 2).

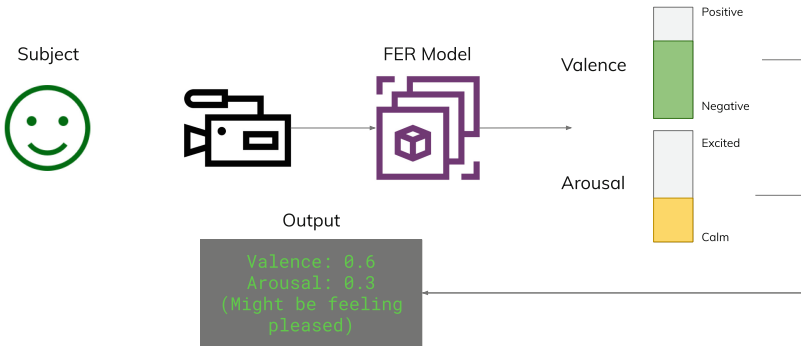


Fig. 2. An illustration of a sample dimensional emotion approach applied to an FER task. A face being examined with a camera attached to an FER model, with outputs showing meters that display valence and arousal levels. This is connected to a terminal “reading” these results and inferring an emotion label.

Example. *Continuous Prediction of Spontaneous Affect from Multiple Cues and Modalities in Valence-Arousal Space* is a 2011 paper written by Nicolaou, et al. that “presents the first approach in the literature towards automatic, *dimensional* and continuous affect prediction in terms of arousal and valence based on facial expressions, shoulder gesture, and audio cues” [38]. The model operates on the *Sensitive Artificial Listener Database* (SAL-DB), which contains spontaneously-elicited emotion data in the form of audio/video samples with continuous human-generated annotations. Based on these annotations, the data has been normalized

to account for positive emotion bias in the dataset and segmented into roughly equal quantities of positive and negative emotion clips. The authors designed features for this data in three separate modalities: for audio, Mel-frequency cepstral (MFC) coefficients over time, and prosody features like energy and pitch; for the face, a mapping of 20 facial feature points represented by video frame-based vectors of the 2D coordinates of these points; and for the shoulders, there are similar sets of points, two on each shoulder and one on a stable central point. Comparing the performance of SVMs for regression and Bidirectional LSTMs (BLSTMs), the authors find better affect prediction performance from the BLSTMs on all input modalities (audio, video) and for all emotion dimensions (valence, arousal), suggesting the importance of the proper representation of temporal data in continuous prediction. Also comparing feature fusion (feature concatenation as input into a single model), model-level fusion (fusion of individual predictions of a particular emotion dimension from facial expression cues and audio cues into another LSTM for final prediction of the same dimension), and output-associative fusion (the combination of both valence/arousal predictions for all cues into another model to yield a single prediction for valence or arousal), they find the best performance out of output-associative fusion. This output-associative fusion appropriately represents observed systematic relationships between valence and arousal values, i.e., the model changes its final arousal prediction based on its prior valence predictions. Improved performance, in this case, suggests the importance of representing this relationship in effective dimensional emotion prediction. Overall, the paper finds promise in the temporal representation of affect via LSTMs and in the representation of these systematic relationships between valence and arousal.

2.3 Constructed Emotion

Constructed emotion, compared to basic emotion and dimensional emotion represents something of a paradigm shift. It aims to bring emotion theory up-to-date with modern neurology research, dispelling outdated ideas of ‘regions’ of emotions and fully dissolving the arbitrary philosophical barrier between “thought” and “emotion” [9]. Emotion becomes a complex but almost romantic process of social construction, with sophisticated neural predictive processes opening up potentially infinite varieties of affective experience. It remains a minority theory, especially in affective computing where it has scarcely penetrated, but it has its growing, enthusiastic supporters [7–9, 19].

Summary. In simplest terms, the theory of constructed emotion holds that emotion is in the eye of the beholder and in the heart of the feeler. Emotion is held to be an experience created within and between human beings through complex predictive processes, and so is something sheerly subjective. The theory suggests, then, that it is impossible to objectively detect emotion as a predictable, well-formed response to certain stimuli.

This approach refutes the idea of basic emotions with distinct mechanisms or expressive “fingerprints,” instead maintaining that emotions, in the confluence of

context, verbal emotion conceptualization, interoception, social agreement, and personal history, are constructed by the brain into a unique experience. These influences feed into the brain’s default mode of prediction, where input is constantly presaged (and corrected, if it varies from what’s predicted), and appropriate responses occur based on these predictions. This general mechanism may be taken as the evolutionary development of a highly efficient, highly flexible response system to an infinite variety of situations. Like experiences of emotion, perceptions of others’ emotional displays are based upon prediction and thus are not infallible and rely extensively on context. In another sense, emotions do not exist objectively to be reliably “detected,” rather, they are powerful instances of human-created social reality. In this vein of logic, the constructivist approach calls emotion universality into doubt, often citing flaws in the methodology of universality research.

Example. *Mirror Ritual: An Affective Interface for Emotional Self-Reflection*, a 2020 paper written by Rajcic and McCormack [46], describes work done on an affective interface that integrates existing emotion perception and text generation technologies to create emotionally meaningful experiences for users. The system takes on the external appearance of a smart mirror with a concealed camera and a reflective display. A user looks at the mirror, and the system uses OpenCV’s Haar cascade classifier to detect their face. The affective mirror then performs real-time facial emotion detection based on a CNN trained on the FER-2013 dataset and generates an emotional seed-word based on perceived emotion and intensity. A mild grimace and furrowed brow, for example, might generate the seed-word “irritated,” and a beaming grin might generate the seed-word “ecstatic.” After the seed-word is generated, it is then fed into a fine-tuned GPT-2-345M text generation model from OpenAI to generate brief, user-engaging poetry based on their perceived emotion. This text generation model

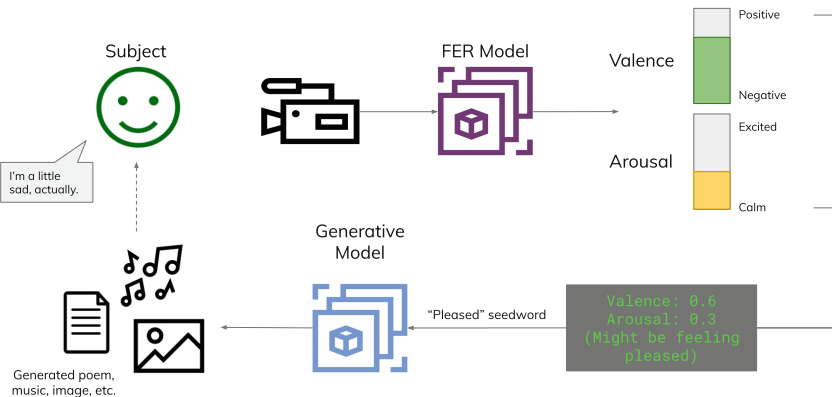


Fig. 3. An illustration of a sample constructivist approach that uses a FER model. A face being examined with a camera attached to a FER model, with outputs showing valence and arousal levels. These levels are used to generate an appropriate emotion seed word for another model that will generate affective content for the subject. The subject reflects on this content and arrives at their own assessment of their emotion.

is trained on a variety of sources, including postmodern poetry, in order to yield poetry that’s accessible but still open to interpretation in order to best facilitate a sort of affective relationship between the mirror and a user. User assessment of the mirror described moments of uncanny appropriateness and great relevance to personal events, though on occasion, users reported a dip in their affective engagement when poems did not seem relevant (Fig. 3).

This affective mirror paper describes an imperfect but still quite promising HCI application that successfully integrates Barrett’s theory of constructed emotion with existing AI and affective computing technologies, like FER and text generation. Importantly, it reconciles the apparent conflict between the constructed emotion theory and the prescriptive nature of most emotion assessment systems. Simply put, Rajcic and McCormack relegate the emotion perception and subsequent poetry generation to a position of non-authority in the overall design of the mirror. Ultimately, the mirror’s capabilities are tools for humans to make sense of their own emotions and relationships – the agency and interpretive work is given to the users. The emotion prediction aspect refrains from acting as an authoritative, correct *recognition* of human emotions as is common in other applications like surveillance. Given a poem instead, a user is free to reject or accept its implications. The tool combines constructed emotion with affective computing in a truly inspiring way.

3 Systematic Survey

The aim of a systematic survey is to provide a reproducible, rigorous, and accountable process for creating questions and finding answers in related literature. The purpose of these questions may be to inquire about the effectiveness of relevant technologies, to provide a valuable introductory summary to the surveyed field, or to suggest an area worthy of additional research. To achieve reproducibility and accountability, a systematic survey publishes its database search queries and

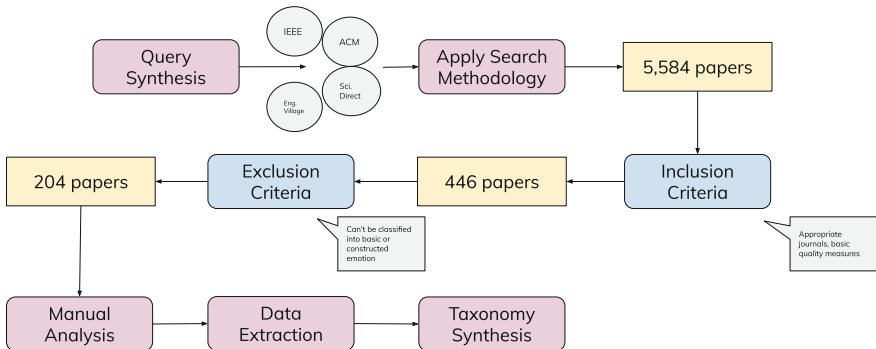


Fig. 4. An illustration of the sequence of paper gathering and selection. After search string generation and database querying, a series of selections reduces the number of papers to an amount tractable for manual analysis.

maintains consistent and documented criteria by which papers are deemed pertinent or impertinent. Following these overarching steps of search and then selection, finally qualitative and/or quantitative analysis in service of the survey’s purpose is performed on the remaining papers (Fig. 4).

To substantiate our claims on the status quo of affective computing and the promise of constructivist-inspired program design, we have conducted a systematic survey of the field and found a crucial representation of constructivist approaches in recent papers. The primary impetus for conducting this survey is to get a grasp on the field of affective computing as a whole, especially as it applies emotion theory to various applications. This is a crucial part of our research that’s been conducted so far because our aim is to reconsider existing practices and offer a constructivist-based approach that has the potential to create novel experiences of affective engagement in Human-Computer Interaction (HCI).

3.1 Description

The following section includes a breakdown of our key systematic literature review steps as they appear in Silva and Neiva’s guide to the practice [50]. Grouping minor and similar tasks for the sake of organization, these include: formulating the research question(s); generating, testing, and refining search strings, conducting the searches and storing data, and finally parsing through the data to select and then analyze relevant papers. In each of these, we will briefly introduce the task, discuss methods, and offer an evaluation on the process and results.

3.2 Methodology

Problem Formulation. In some ways, the questions we posit reflect the suspicions we have about the topic. Our paper primarily seeks to examine the efficacy of existing emotion inference methods, ponder the potential effectiveness of constructivist methods, and question whether emotion inference technologies will provide lasting value in in-the-wild settings. These topics and rationale for asking them will be discussed in greater detail below. Some of them arise in part due to conclusions drawn in Barrett’s *How Emotions Are Made* [9].

Our first overarching question: *How effective are existing emotion inference methods based on basic emotion theory, and how well will they generalize to real-world, in-the-wild applications?* Though, say, facial expression classification may be growing increasingly robust, it is reasonable to question whether or not these discrete classification models will be able to classify less well-formed facial input well. In addition, generalizability gets called into question if models are trained on acted, stereotypical expressions of emotion—these are clear signals, but in actual scenarios, you are unlikely to find these perfect matches. When systems like these are integrated into aspects of HCI (robot or apps), will the user find the classifica-

tion of their feelings into six firm categories robust or reductive? If overly reductive, an application integrating such technology may seem either toy-like or at worst presumptive, and in either case, will fail to be useful. This question plays the role of acting as a primary impetus for our research. It represents one of the key questions that we are overall seeking to prove or disprove.

The above question presumes some level of widespread adoption of basic emotion-based inference techniques, however, and so we are also responsible for confirming this presumption. We therefore have a few more key questions on our plate. *What does the field of affective computing look like? Are approaches either explicitly or implicitly based on basic emotion theory very prevalent, to begin with? Are there other, more widespread approaches that we should instead ask questions of? What are typical applications for these affective computing technologies?* Seeking an answer to these questions acts as a key grounding element that ensures we have an accurate and less-skewed perspective of the field. If basic emotion approaches turn out to be relatively uncommon, or applications largely shy away from actually predicting emotions, then perhaps there is less of a need for our question to be asked in the first place. Perhaps others have had the same hypothesis and arrived at the same conclusion already. Essentially, this question helps ensure that our research is relevant, representative, and fair.

Our second big question: *Would a constructivist (or some other) approach be more effective than the dominating approaches? Would this approach capture more nuance in an emotion prediction system?* Of course, we must also examine whether or not a system guided by the constructivist approach would be better to begin with—regardless of our hypotheses, we can't in good faith assume so. This question essentially asks us to justify the inclinations we may have towards the approach and asks us to provide a basis for arguing for the pursuit of constructivist-based affective computing. If we can find no compelling reasons or promises, then there would be no point in encouraging computing research based on this approach.

At last, we must ask: *What does affective computing informed by a constructivist approach even look like?* This is a key question for two reasons: (a) we may lack examples because systems following the constructivist approach are relatively few; and (b), Barrett's theory posits ideas that may fundamentally conflict with the idea of computational emotion prediction. In simplest terms, the theory of constructed emotion holds that emotion is in the eye of the beholder and in the heart of the feeler. Emotion is held to be **an experience created within and between human beings through complex predictive processes**, and so is something sheerly subjective. The theory suggests, then, that it is impossible to objectively detect emotion as a predictable, well-formed response to certain stimuli. Barring completely abandoning the premise of affective emotion prediction, then, how do we reconcile the practice to this theory? Could a predictive agent act

like another subjective observer of others’ emotions, with biases based on training data instead of human experience? There seems to be an added complexity to designing a constructivist-based emotion perceiver or seems to require some re-conceptualization. These questions serve to explore what practical implementation might look like, as well as to consider how a “paradigm shift” might be necessary to attain the benefits of a constructivist-based approach.

Search Methodology. With the above questions in mind, the next task is to create the search string that will be used to query various published-paper databases, and we focused on the computer science literature.

The first step is to consider our research questions and create a preliminary search string that may lead us to papers that can answer these questions. We then take this string and query three databases, recommended by Silva and Neiva’s guide [50] for their prevalence in computer science and overall comprehensiveness: *IEEE Explore* [26], *ACM Library* [1], and *Elsevier ScienceDirect* [16]. Examining the quality, quantity, and relevance of results each round, the string is iteratively revised to yield a set of more promising results. With each revision, we also take care to ensure that the string is properly adapted to the syntax of each database we query, so it retains the same search semantics. For reference, the aim was to retrieve approximately 3,000 to 5,000 papers on the topic of various approaches (basic emotion, constructivist, dimensional) in the field of affective computing. In particular, we wanted to ensure that any constructivist approaches are represented and so take additional care to modify our search accordingly.

Between each iteration of the string was a process of experimenting with syntax, search parameters, and sample searches to get ideas of how different keywords were represented in the databases. For example, searches of just “*affect*” and “*affect NOT affective*” were compared to get an idea of how many papers might be captured by the homonym verb “affect” but not be related to emotion. This assumes that a paper containing “affect” but not “affective” is less likely to be about emotion and more likely to include the word as an incidental verb. Respectively, “*affect*” alone returned 57k results in the IEEE database, and “*affect NOT affective*” returned 56k, suggesting that the majority of papers included by the term “affect” was probably not related to emotion or affective computing. This informed the change from querying for “affect” to “affective.” Similarly, searches of the names representing various emotion theories (i.e., Ekman for basic emotion, Barrett for constructed emotion) returned very few results and so informed additional changes. We arrive at the following string and have used it to conduct our search: (*“affective” OR emotion OR mood*) AND (*prediction OR inference*) AND (*“basic emotion” OR “theory of constructed emotion” OR constructivist OR Plutchik*).

With the search strings finalized and the searches complete, we must proceed with passing eyes over our results to begin collecting information and start answering the questions we posed in earlier steps of the survey process. This proves to be an intensive process that examines papers in rounds with increasing levels of detail. This and other ancillary tasks are as follows.

Table 1. Papers by Category. *Theoretical* papers are those that discuss applying a given theory of emotion to affective computing. *Implementation* papers refer to those that explicitly or implicitly use a theory of emotion in the creation of an affective computing application. *Datasets/Other* refer to training data created for model prediction in a particular vein of emotion theory. Irrelevant papers and those whose theory is not apparent have been omitted for clarity.

Category	Number of papers
Basic Emotion	–
Theoretical	38
Implementation	106
Datasets/Other	6
Subtotal	150
Dimensional Emotion	–
Theoretical	18
Implementation	30
Datasets/Other	1
Subtotal	49
Constructed Emotion	–
Theoretical	1
Implementation	2
Datasets/Other	2
Subtotal	5
Total	204

The first step to this larger task was exporting all of the 5500+ results from our databases—often requiring page-by-page exporting—and saving them to a local archive. A reference management software [29] was used extensively for this purpose, as we were able to easily import paper metadata and abstracts in the bibtex format into it. Once imported, then began the task of broadly classifying all of the papers as irrelevant or relevant. If relevant, a paper was also organized by the apparent theory of emotion the paper’s method ascribes to, based on the title and abstract, and whether a paper appears to be implementation-based or theoretical. If a paper was decidedly relevant but didn’t ascribe to either basic or constructed emotion theory, it was placed in the Relevant/Other category. When classified, a paper was marked as ‘skimmed’ to indicate completion and facilitate useful grouping and sorting functionality in JabRef. Table 1 summarizes this step.

To narrow down 5500+ papers manually tractable, some heuristics were applied to classify papers as irrelevant. If a paper is: a) older than 2004, b) not in English, c) lacking title or abstract, d) is an inaccessible book, or e) published in a most likely irrelevant journal, it is classified as out-of-scope for this survey. Note that we post-processed the resulting list to include some key papers published before 2004. To illustrate the last criterion, an article published in *Poultry*

Science or *Poetics*, for example, is most likely not relevant to our survey. These heuristics a) and e) mostly culled results in pure psychology or neurobiology, as well as other miscellany venues. Roughly four thousand results were culled from our pool of 5,583 via these heuristics.

After irrelevant papers were sorted away and relevant papers coarsely classified into emotion theory groups, the relevant papers were passed over once more to gather additional useful information. To gauge the relative popularity and importance of a paper in its field, we used citation counts. To accomplish this, paper titles were used as queries into Google Scholar, and the citation count was gathered into our JabRef archive as additional metadata.

Beyond coarse classifications, the second pass over relevant results involved scanning titles and abstracts once more, with an eye on two particular aspects, namely, the affective computing method used and its application, if one is apparent (e.g., for gauging student interest in a virtual classroom setting). These two aspects were concatenated and appended as additional metadata to relevant results, in the form of the string, e.g., “*artificial affective agents for human-robot interaction*,” for example. The purpose of this step was to get an idea of where and how affective computing is frequently applied and what technologies are frequently pursued.

Search Results. Final searches also included results from *Engineering Village* [17], rounding out results with an additional 42 papers and completing the list of databases that were recommended by Silva and Neiva [50] and were accessible through our institutional resources. The final tally of results are as follows: 4,846 papers from *ScienceDirect*, 92 from *IEEE*, 604 from *ACM Library*, and 42 from *Engineering Village*, for a total of 5,584 papers. Trimming the irrelevant papers using the method explained in the previous section, we ended up with **204** papers as shown in Table 1 and Table 2.

As a qualitative overview, a couple of applications saw considerable representation in this survey, particularly facial emotion recognition (FER) and textual emotion recognition (TER), the latter primarily for sentiment analysis applications. Interestingly, a non-negligible amount of papers discussed the application of affect modeling for the sake of artificial affective agents, like game AI or human-robot interaction. Another common application was multimedia sentiment analysis, mostly of videos and images, but occasionally of music, as shown in Table 2.

Outcomes. General classifications of papers into emotion theory groups followed most of the original hypotheses. A significant portion of the relevant papers fell under the basic emotion category (151 of 204 papers, nearly three quarters). However, a significant amount fell under the “Other” category. A good amount of these fell under a **dimensional emotion** approach, which assessed emotions based on several dimensions – typically, but not always, these were of valence (positive/negative) and arousal (high energy/low energy). Despite not explicitly addressing dimensional approaches in our search string, this is a surprising turnout that suggests that dimensional approaches are another popular con-

Table 2. Papers by Application. A breakdown of collected papers by applied field. Papers in the “other” category frequently discuss theory of applying a given emotion theory to affective computing, as well as includes miscellaneous singleton applications.

Category	Number of papers
Facial expression recognition	64
Textual emotion recognition	56
Speech-based emotion recognition	27
Biometric emotion recognition	22
Multimedia emotion classification	22
Multimodal emotion recognition	14
Other	8
Total	204

tender. The majority of the “Other”-categorized papers fell under “unspecified other,” however, mostly because many papers made no implicit or explicit mention of their approach for their emotion models. Many of them had ambiguous or brief abstracts and titles that made categorization difficult from this short pass-over and so have been dropped from the results to preserve a list of papers with definitely known emotion theories. A closer reading of these papers yield mostly basic emotion and dimensional emotion categorizations, and constructivist papers represented only a little over 2 percent of all relevant papers.

Yet, finding even a few papers that fall under this non-prescriptivist constructed emotion heading is an important result that suggests interest in a constructivist-informed approach to affective computing, especially in HCI. Below we will summarize this particularly relevant paper as well as prominent and illustrative examples applying the other theories of emotion for future references. Table 3 shows representative samples from the resulting survey database.

4 Discussion

After surveying the affective computing literature and examining a few notable papers in-depth, we now revisit a few of our initial questions and draw conclusions.

Broadly, what does the field of affective computing look like in terms of the theory of emotion? As initially expected, there seems to be a very significant representation of basic emotion theory at work in the field, informing many papers on a variety of tasks, particularly emotion classification. Dimensional emotion represents a significant second theory alive in the literature with a moderate showing in the survey, though it is important to consider that the final query string did not explicitly search for dimensional approaches. Having so many dimensional papers turn up without “dimensional” literally within the search string may suggest that dimensional papers represent a much greater portion than represented in this survey. Another look, next time not focusing primarily *basic emotion vs. constructed emotion* may yield an answer to this open question and provide a more

Table 3. A subset list of collected papers grouped by emotion theory category and sorted by year published. The acronym *EP* refers to emotion prediction. *BASIC* refers to Basic Emotion, *DIM.* refers to Dimensional Emotion, and *CON.* refers to Constructed Emotion. For detailed discussions on the definitions of these, please refer to the main text.

Citation	Author	Year	Category	Mode and Methodology
[14]	Domínguez-Jiménez et al.	2020	BASIC	physiological signals EP; comparison of multiple methods
[2]	Ahmad et al.	2020	BASIC	text; English to Hindi emotion embedding transfer learning, CNN/BLSTM
[54]	Yadav and Vishwakarma	2020	BASIC	movie trailer EP via ILDNet
[25]	Hameed et al.	2019	BASIC	respiration-based EP; FFT analysis
[18]	Feng	2019	BASIC	text; sentiment analysis of social media[...]
[12]	Chatterjee et al.	2019	BASIC	text; sentiment analysis using deep learning
[48]	Sajjad et al.	2019	BASIC	FER; Oriented FAST and Rotated BRIEF features supporting an SVM
[36]	Löffler et al.	2018	BASIC	affective agent; multimodal expression
[59]	Zeng et al.	2018	BASIC	FER; high-dimensional facial appearance features as input to DSAE
[5]	Arnau-González et al.	2017	BASIC	EEG EP; EEG feature combination
[65]	Zhou et al.	2015	BASIC	FER; emotion distribution learning
[58]	Yu and Zhang	2015	BASIC	FER; ensemble face detection, CNN
[30]	Khezri et al.	2015	BASIC	multimodal physiological signals, SVM/KNN
[57]	Yu and Wang	2015	BASIC	text; Twitter sentiment analysis
[40]	Orellana-Rodriguez et al.	2015	BASIC	multimedia affect contextualization
[61]	Zhang et al.	2015	BASIC	FER for AAAs, robust facial point detection
[34]	Lin et al.	2015	BASIC	AAAs for composite emotion study
[37]	Majumder et al.	2014	BASIC	FER using KSOM
[33]	Kukolja et al.	2014	BASIC	physio. EP method comparison
[60]	Zhang et al.	2013	BASIC	FER and topic analysis for affective agent
[44]	Purver and Battersby	2012	BASIC	text; automatic labelling for EP models
[27]	Ilbeygi and Shah-Hosseini	2012	BASIC	FER using fuzzy inference
[13]	Chen et al.	2012	BASIC	SER; multilevel models w/ SVMs
[52]	Wu et al.	2011	BASIC	SER using modulation spectral features
[31]	Kim et al.	2010	BASIC	text; comparison of unsupervised ER models
[28]	Iliev et al.	2010	BASIC	SER; glottal features on OPF model
[45]	Quan and Ren	2009	BASIC	text; creation of Chinese emotion corpus
[21]	Gill et al.	2008	BASIC	text; sentiment analysis via LIWC and LSA
[4]	Alm et al.	2005	BASIC	text; sentiment analysis via SNoW ML
[22]	Goldman and Sripada	2005	BASIC	FER via simulationist models
[35]	Liu et al.	2003	BASIC	text; 'common sense' affect detection
[11]	Calder et al.	2001	BASIC	FER; PCA for facial features
[42]	Pantic and Rothkrantz	2000	BASIC	FER; facial action-based EP
[49]	Scheirer et al.	1999	BASIC	wearable FER for expression detection
[51]	Wang et al.	2020	DIM	text; sentiment analysis, regional CNN-LSTM
[64]	Zhou et al.	2020	DIM	FER via bilinear CNN
[53]	Xiaohua et al.	2019	DIM	FER; two-level attention with Bi-RNN
[62]	Zhang et al.	2018	DIM	multimodal smartphone-based EP
[3]	Al Zoubi et al.	2018	DIM	EEG-based EP via liquid state machine
[55]	Yin et al.	2017	DIM	multimodal physio. EP with SAE ensembles
[20]	Giatsoglou et al.	2017	DIM	text; sentiment analysis comparison
[63]	Zhao et al.	2016	DIM	user-unique image EP
[32]	Koelstra and Patras	2013	DIM	FER and EEG fusion for affect tagging
[39]	Nogueira et al.	2013	DIM	DIM. regression to BASIC physio EP
[24]	Gunes and Schuller	2013	DIM	DIM. vs. BASIC comparison survey
[56]	Yoon and Chung	2013	DIM	EEG w/ ML classifier
[41]	Ortígosa-Hernández et al.	2012	DIM	text; sentiment analysis w/ semi-supervised models
[10]	Cai and Lin	2011	DIM	EP for driving safety analysis
[38]	Nicolau et al.	2011	DIM	multimodal EP using BLSTMS
[23]	Grimm et al.	2007	DIM	SER emotion primitive analysis
[46]	Rajcic and McCormack	2020	CON	FER for affective poem generation

accurate view of the affective computing field. As for constructed emotion, this survey found that this theory has not quite taken a significant foothold in the literature yet, though the presence of the promising *Mirror Ritual* paper [46] may be a sign of breakthrough and future growth of the theory in the literature.

What does a constructed emotion approach look like in affective computing? *Mirror Ritual* [46] provides one possible answer to this question. We see that this paper doesn't necessarily reject the existing methods of basic emotion classification and dimensional emotion prediction, but rather it leverages them to achieve a slightly different goal than the others. Instead of aiming to directly classify a user as experiencing a particular emotion (or as in some combination of valence and arousal), the idea is to use whatever credence existing prediction methods have to incorporate some form of generated art with the emotion the model perceives. The model may or very well may not be correct, but its direct assessment of the user is downplayed in favor of providing a tool for emotional reflection. This way, a given user retains agency and self-definition of their own internal state, choosing to integrate an emotionally relevant generated poem into their own understanding of their feelings or reject an irrelevant one. In this formulation, more accurate emotion prediction would be helpful, but if the capacity for a computer to *perceive* emotion is fundamentally limited by stipulations posed by constructed emotion theory, that is still okay. The ultimate goal is to create something evocative and emotionally salient for users, in some ways more in the wheelhouse of art than anything else. Furthering of constructivism in affective computing may very well resemble pursuits of AI art creation. This assessment provides some valuable insight into our next question.

Would a constructed emotion approach be more effective than approaches based on other theories? Given the above assessment, this question may very well have been a flawed one to ask. Ultimately, the methods are not necessarily competing, to begin with, as their goals are fundamentally different. It doesn't do much good to try and compare how accurately a basic-emotion predictive model classifies faces into emotion categories and how well a constructed-emotion approach creates opportunities for valuable emotional reflection. One may ask "*Which will ultimately prove more useful to society and helpful to human emotion modeling?*", but it stands outside of the scope of this survey.

5 Conclusion

This survey has systematically examined over 200 papers in the field of affective computing, and in doing so, has arrived at the following conclusions: (a) Basic emotion classification and analysis tasks are presently the most popular, representing a majority of papers. (b) Facial, speech, and text-based emotion recognition tasks, regardless of emotion theory, are the most popular tasks in the field. (c) Constructed emotion in affective computing does not compete with emotion prediction methods of other stripes but instead utilizes them to achieve an entirely different goal. (d) Constructed emotion approaches represent a tiny minority of papers, but sample papers nonetheless represent potential for a new class of 'affective engagement' HCI applications. Future directions include further exploration into the potential of constructivist-based affective computing applications, the creation of a constructed emotion HCI device prototype, and the pursuit of generative art models inspired by users' emotions, as in the *Mirror Ritual* paper [46].

Acknowledgement. This work was supported in part by NSF Awards IIS 1955404 and 1955365.

References

1. ACM Digital Library (2020). <https://dl.acm.org/>. Accessed 19 Sept 2020
2. Ahmad, Z., Jindal, R., Ekbal, A., Bhattacharyya, P.: Borrow from rich cousin: transfer learning for emotion detection using cross lingual embedding. *Expert Syst. Appl.* **139**, 112851 (2020)
3. Al Zoubi, O., Awad, M., Kasabov, N.K.: Anytime multipurpose emotion recognition from EEG data using a liquid state machine based framework. *Artif. Intell. Med.* **86**, 1–8 (2018)
4. Alm, C.O., Roth, D., Sproat, R.: Emotions from text: machine learning for text-based emotion prediction. In: *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, HLT 2005*, pp. 579–586. Association for Computational Linguistics, USA (2005)
5. Arnau-González, P., Arevalillo-Herráez, M., Ramzan, N.: Fusing highly dimensional energy and connectivity features to identify affective states from EEG signals. *Neurocomputing* **244**, 81–89 (2017)
6. Arya, R., Singh, J., Kumar, A.: A survey of multidisciplinary domains contributing to affective computing. *Comput. Sci. Rev.* **40**, 100399 (2021)
7. Aviezer, H., et al.: Angry, disgusted, or afraid?: studies on the malleability of emotion perception. *Psychol. Sci.* **19**(7), 724–732 (2008)
8. Bar, M.: The proactive brain: using analogies and associations to generate predictions. *Trends Cogn. Sci.* **11**(7), 280–289 (2007)
9. Barrett, L.F.: *How Emotions are Made*. PAN Books (2018)
10. Cai, H., Lin, Y.: Modeling of operators' emotion and task performance in a virtual driving environment. *Int. J. Hum. Comput. Stud.* **69**(9), 571–586 (2011)
11. Calder, A.J., Burton, A., Miller, P., Young, A.W., Akamatsu, S.: A principal component analysis of facial expressions. *Vis. Res.* **41**(9), 1179–1208 (2001)
12. Chatterjee, A., Gupta, U., Chinnakotla, M.K., Srikanth, R., Galley, M., Agrawal, P.: Understanding emotions in text using deep learning and big data. *Comput. Hum. Behav.* **93**, 309–317 (2019)

13. Chen, L., Mao, X., Xue, Y., Cheng, L.L.: Speech emotion recognition: features and classification models. *Digit. Signal Process.* **22**(6), 1154–1160 (2012)
14. Domínguez-Jiménez, J., Campo-Landines, K., Martínez-Santos, J., Delahoz, E., Contreras-Ortiz, S.: A machine learning model for emotion recognition from physiological signals. *Biomed. Signal Process. Control* **55**, 101646 (2020)
15. Ekman, P.: An argument for basic emotions. *Cogn. Emot.* **6**(3–4), 169–200 (1992)
16. Elsevier ScienceDirect (2020). <https://www.sciencedirect.com/>. Accessed 19 Sept 2020
17. Engineering Village (2020). <https://www.engineeringvillage.com/home.url?redir=t>. Accessed 19 Sept 2020
18. Feng, Z.: Hot news mining and public opinion guidance analysis based on sentiment computing in network social media. *Pers. Ubiquit. Comput.* **23**(3–4), 373–381 (2019)
19. Gendron, M., Roberson, D., Vyver, J., Barrett, L.: Perceptions of emotion from facial expressions are not culturally universal: evidence from a remote culture. *Emotion (Washington, D.C.)* **14**, 251–62 (2014)
20. Giatsoglou, M., Vozalis, M.G., Diamantaras, K., Vakali, A., Sarigiannidis, G., Chatzivasvas, K.C.: Sentiment analysis leveraging emotions and word embeddings. *Expert Syst. Appl.* **69**, 214–224 (2017)
21. Gill, A.J., French, R.M., Gergle, D., Oberlander, J.: The language of emotion in short blog texts. In: *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work, CSCW 2008*, pp. 299–302. Association for Computing Machinery, New York (2008)
22. Goldman, A.I., Sripada, C.S.: Simulationist models of face-based emotion recognition. *Cognition* **94**(3), 193–213 (2005)
23. Grimm, M., Kroschel, K., Mower, E., Narayanan, S.: Primitives-based evaluation and estimation of emotions in speech. *Speech Commun.* **49**(10), 787–800 (2007)
24. Gunes, H., Schuller, B.: Categorical and dimensional affect analysis in continuous input: current trends and future directions. *Image Vis. Comput.* **31**(2), 120–136 (2013)
25. Hameed, R.A., Sabir, M.K., Fadhel, M.A., Al-Shamma, O., Alzubaidi, L.: Human emotion classification based on respiration signal. In: *Proceedings of the International Conference on Information and Communication Technology, ICICT 2019*, pp. 239–245. Association for Computing Machinery, New York (2019)
26. IEEE Explore Database (2020). <https://ieeexplore.ieee.org/Xplore/home.jsp>. Accessed 9 Sept 2020
27. Ilbeygi, M., Shah-Hosseini, H.: A novel fuzzy facial expression recognition system based on facial feature extraction from color face images. *Eng. Appl. Artif. Intell.* **25**(1), 130–146 (2012)
28. Iliev, A.I., Scordilis, M.S., Papa, J.P., Falcão, A.X.: Spoken emotion recognition through optimum-path forest classification using glottal features. *Comput. Speech Lang.* **24**(3), 445–460 (2010)
29. JabRef - Free Reference Manager (2020). <https://www.jabref.org/>. Accessed 8 Sept 2020
30. Khezri, M., Firoozabadi, M., Sharafat, A.R.: Reliable emotion recognition system based on dynamic adaptive fusion of forehead biopotentials and physiological signals. *Comput. Methods Programs Biomed.* **122**(2), 149–164 (2015)
31. Kim, S.M., Valitutti, A., Calvo, R.A.: Evaluation of unsupervised emotion models to textual affect recognition. In: *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*,

- CAAGET 2010, pp. 62–70. Association for Computational Linguistics, USA (2010). 125, textual affect recognition; survey
32. Koelstra, S., Patras, I.: Fusion of facial expressions and EEG for implicit affective tagging. *Image Vis. Comput.* **31**(2), 164–174 (2013)
 33. Kukulja, D., Popović, S., Horvat, M., Kovač, B., Čosić, K.: Comparative analysis of emotion estimation methods based on physiological measurements for real-time applications. *Int. J. Hum. Comput. Stud.* **72**(10), 717–727 (2014)
 34. Lin, J., Yu, H., Miao, C., Shen, Z.: An affective agent for studying composite emotions. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Richland, SC*, pp. 1947–1948. International Foundation for Autonomous Agents and Multiagent Systems (2015)
 35. Liu, H., Lieberman, H., Selker, T.: A model of textual affect sensing using real-world knowledge. In: *Proceedings of the 8th International Conference on Intelligent User Interfaces, IUI 2003*, pp. 125–132. Association for Computing Machinery, New York (2003)
 36. Löffler, D., Schmidt, N., Tscharn, R.: Multimodal expression of artificial emotion in social robots using color, motion and sound. In: *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction, HRI 2018*, pp. 334–343. Association for Computing Machinery, New York (2018)
 37. Majumder, A., Behera, L., Subramanian, V.K.: Emotion recognition from geometric facial features using self-organizing map. *Pattern Recogn.* **47**(3), 1282–1293 (2014)
 38. Nicolaou, M.A., Gunes, H., Pantic, M.: Continuous prediction of spontaneous affect from multiple cues and modalities in valence-arousal space. *IEEE Trans. Affect. Comput.* **2**(2), 92–105 (2011)
 39. Nogueira, P.A., Rodrigues, R., Oliveira, E., Nacke, L.E.: A hybrid approach at emotional state detection: merging theoretical models of emotion with data-driven statistical classifiers. In: *Proceedings of the 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), WI-IAT 2013*, vol. 02, pp. 253–260. IEEE Computer Society, USA (2013)
 40. Orellana-Rodriguez, C., Diaz-Aviles, E., Nejd, W.: Mining affective context in short films for emotion-aware recommendation. In: *Proceedings of the 26th ACM Conference on Hypertext & Social Media, HT 2015*, pp. 185–194. Association for Computing Machinery, New York (2015)
 41. Ortigosa-Hernández, J., Rodríguez, J.D., Alzate, L., Lucania, M., Inza, I., Lozano, J.A.: Approaching sentiment analysis by using semi-supervised learning of multi-dimensional classifiers. *Neurocomputing* **92**, 98–115 (2012)
 42. Pantic, M., Rothkrantz, L.: Expert system for automatic analysis of facial expressions. *Image Vis. Comput.* **18**(11), 881–905 (2000)
 43. Picard, R.W.: *Affective computing* (1995)
 44. Purver, M., Battersby, S.: Experimenting with distant supervision for emotion classification. In: *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2012*, pp. 482–491. Association for Computational Linguistics, USA (2012)

45. Quan, C., Ren, F.: Construction of a blog emotion corpus for Chinese emotional expression analysis. In: Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing, EMNLP 2009, vol. 3, pp. 1446–1454. Association for Computational Linguistics, USA (2009). 123, textual emotion detection for analysis
46. Rajcic, N., McCormack, J.: Mirror ritual: an affective interface for emotional self-reflection. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI 2020, pp. 1–13. Association for Computing Machinery, New York (2020)
47. Russell, J.: A circumplex model of affect. *J. Pers. Soc. Psychol.* **39**, 1161–1178 (1980). <https://doi.org/10.1037/h0077714>
48. Sajjad, M., Nasir, M., Ullah, F.U.M., Muhammad, K., Sangaiah, A.K., Baik, S.W.: Raspberry Pi assisted facial expression recognition framework for smart security in law-enforcement services. *Inf. Sci.* **479**, 416–431 (2019)
49. Scheirer, J., Fernandez, R., Picard, R.W.: Expression glasses: a wearable device for facial expression recognition. In: Extended Abstracts on Human Factors in Computing Systems, CHI 1999, CHI EA 1999, pp. 262–263. Association for Computing Machinery, New York (1999)
50. Silva, R., Neiva, F.: Systematic literature review in computer science - a practical guide (2016)
51. Wang, J., Yu, L.C., Lai, K.R., Zhang, X.: Tree-structured regional CNN-LSTM model for dimensional sentiment analysis. *IEEE/ACM Trans. Audio Speech Lang. Proc.* **28**, 581–591 (2020)
52. Wu, S., Falk, T.H., Chan, W.Y.: Automatic speech emotion recognition using modulation spectral features. *Speech Commun.* **53**(5), 768–785 (2011)
53. Xiaohua, W., Muzi, P., Lijuan, P., Min, H., Chunhua, J., Fuji, R.: Two-level attention with two-stage multi-task learning for facial emotion recognition. *J. Vis. Commun. Image Represent.* **62**, 217–225 (2019)
54. Yadav, A., Vishwakarma, D.K.: A unified framework of deep networks for genre classification using movie trailer. *Appl. Soft Comput.* **96**, 106624 (2020)
55. Yin, Z., Zhao, M., Wang, Y., Yang, J., Zhang, J.: Recognition of emotions using multimodal physiological signals and an ensemble deep learning model. *Comput. Methods Programs Biomed.* **140**, 93–110 (2017)
56. Yoon, H.J., Chung, S.Y.: EEG-based emotion estimation using Bayesian weighted-log-posterior function and perceptron convergence algorithm. *Comput. Biol. Med.* **43**(12), 2230–2237 (2013)
57. Yu, Y., Wang, X.: World cup 2014 in the Twitter world: a big data analysis of sentiments in U.S. sports fans' tweets. *Comput. Hum. Behav.* **48**, 392–400 (2015)
58. Yu, Z., Zhang, C.: Image based static facial expression recognition with multiple deep network learning. In: Proceedings of the 2015 ACM on International Conference on Multimodal Interaction, ICMI 2015, pp. 435–442. Association for Computing Machinery, New York (2015)
59. Zeng, N., Zhang, H., Song, B., Liu, W., Li, Y., Dobaie, A.M.: Facial expression recognition via learning deep sparse autoencoders. *Neurocomputing* **273**, 643–649 (2018)
60. Zhang, L., Jiang, M., Farid, D., Hossain, M.: Intelligent facial emotion recognition and semantic-based topic detection for a humanoid robot. *Expert Syst. Appl.* **40**(13), 5160–5168 (2013)
61. Zhang, L., Mistry, K., Jiang, M., Chin Neoh, S., Hossain, M.A.: Adaptive facial point detection and emotion recognition for a humanoid robot. *Comput. Vis. Image Underst.* **140**, 93–114 (2015)

62. Zhang, X., Li, W., Chen, X., Lu, S.: MoodExplorer: towards compound emotion detection via smartphone sensing. *Proc. ACM Interact. Mob. Wearable Ubiquit. Technol.* **1**(4) (2018)
63. Zhao, S., et al.: Predicting personalized emotion perceptions of social images. In: *Proceedings of the 24th ACM International Conference on Multimedia, MM 2016*, pp. 1385–1394. Association for Computing Machinery, New York (2016)
64. Zhou, F., Kong, S., Fowlkes, C.C., Chen, T., Lei, B.: Fine-grained facial expression analysis using dimensional emotion model. *Neurocomputing* **392**, 38–49 (2020)
65. Zhou, Y., Xue, H., Geng, X.: Emotion distribution recognition from facial expressions. In: *Proceedings of the 23rd ACM International Conference on Multimedia, MM 2015*, pp. 1247–1250. Association for Computing Machinery, New York (2015)