

Theseus in the Epistemic Labyrinth: A Critical History of the Semantic Differential Method and the Apparent Weight of Color

Evan Donahue

Abstract

Critical scholars of AI have identified the text, images, and other media that constitute the data of AI systems as a promising site of encounter between AI and critical disciplines in the humanities and social sciences. However, questions remain concerning what role critical scholarship can play in guiding the development of technology towards more just futures. In this essay, I examine how AI researchers interpret the raw data at the center of their discipline, and how these interpretations inform the development of technology. I trace the history of the 'semantic differential', a dataset creation methodology in use in modern AI research, through four distinct epistemic moments, backwards from affective computing through Japanese industrial design informed by 18th-century German aesthetics, through structuralist investigations of the universal language of myth, and into early 20th-century psychological investigations of the connection in the brain between the perception of color and the sensation of weight. Attending to the role of the interpretation of data in guiding technical development in these four moments, I argue that critical scholarship has an important role to play in the creation of technology.

1. Introduction

'Why is a person with an empty stomach heavier than after a meal?' asks Odilo Schreger in 1755, before answering, 'because eating increases the quantity of the spirits, which owing to their airy and fiery nature lighten the human body... For the same reason a cheerful person is much lighter than a sad one' (Fleck, 2012: 200). Philosopher of science Ludwig Fleck uses Schreger's question to illustrate how part of the work of scholarship is to separate broad vernacular concepts such as 'weight' into distinct phenomena. Gravitation, glucose, and

serotonin all require distinct theories and practices of measurement, even if all contribute to a subjective experience of sluggishness. To confuse having one's spirits lifted with being lifted by spirits would invite disaster.

Viewed from the humanities and social sciences, the field of artificial intelligence (AI) often seems to overlook important nuances of human thought, language, emotion, and experience, as Schreger does with 'weight'. Scholars in these fields have engaged critically with AI on the intuition that more nuanced conceptions of the human have something to offer the development of AI technologies (Dreyfus, 1965; Forsythe, 2001; Suchman, 1987; Edwards, 1997; Collins, 1990; P. E. Agre, 1997). Such efforts have only become more pronounced as AI technologies become increasingly ubiquitous (Castelle, 2018; Seaver, 2019; Stark & Hoey, 2021; Selbst et al., 2019; Elish & Boyd, 2018; Campolo & Crawford, 2020; Markham, 2013; Schrock, 2017). As AI has become entangled with critical social and civic infrastructures, critics have been quick to identify how overly reductive treatments of race, ethnicity, gender, sexuality, disability and other essential dimensions of human experience lead to social harms.

This urgency has given way to a range of proposals for integrating the knowledge from the humanities and social sciences into AI education and practice (Jo & Gebru, 2020; Selbst et al., 2019; Karoff, 2019). These proposals have merit, but as anthropologist of AI Diana Forsythe reminds us, there is no substitute for the critical intuitions developed over the course of a career dedicated to the practice of the humanities and social sciences (2001). For this reason, I argue in this essay for the necessity of reimagining the work of AI in a manner that places the social sciences—and in particular in this essay the humanities—alongside computer science at the center of the discipline. I focus on the humanities in particular, recognizing their often porous boundaries, because they are well positioned to address a lacuna in critical engagements with AI. The social sciences, broadly construed, have found success in arguing for the necessity of studying the communities impacted by AI technologies developed and deployed by corporations and governments (Selbst et al., 2019). Yet they have thus far had less to say about the vast body of AI work that exists only as fragments of visions not yet realized or realizable, but which nevertheless creates the contexts from which those future applications emerge.¹ As I

will argue, AI research in its earliest stages revolves around questions of the interpretation of texts and media and what those texts tell us about the human experience. In this respect AI research shares much of its subject matter with scholarship in the humanities, making it possible to imagine a larger shared enterprise.

In this essay, I take the concept of 'emotion' in AI, particularly as it manifests in the subfield of Emotion-Based Textile Image Retrieval (EBTIR), as a case study to illustrate what I believe to be a more general dynamic in AI. EBTIR aims to produce algorithms to predict the emotional response a viewer will have to a visual textile pattern. This work illuminates how technical research takes shape around an interpretive choice such as that of reading textile design through the language of 'emotion'. More specifically, I focus on the semantic differential, a methodology at the heart of EBTIR work for eliciting judgments of the emotional qualities of textile patterns from human subjects in order to supply its algorithms with data.

I trace the semantic differential through four epistemic moments. Starting with the most recent, that of EBTIR work, I follow the method backwards through moments in which it was seen as a measure not of emotion but of aesthetics, of 'meaning,' and ultimately of psychosomatic reflexes. This history, like Fleck's example, helps to defamiliarize the present categories of AI research. It raises questions of whether 'emotion', however theorized, is the right vocabulary through which to grasp these technologies, and if not how we are to make sense of their meaning, their possibilities, and their risks.

In light of these radical shifts in interpretation, I argue that the categories through which AI grasps human experience, whether 'emotion', 'aesthetics', or otherwise, are only very loosely held and that critiquing them directly will have minimal impact on the direction of technology. Rather, these categories are the result of a continually unfolding struggle to make sense of how the text and media that form the data of AI capture or reflect the human essence the field wishes to bestow upon its machines. It is therefore by engaging with and interpreting not the rhetoric of AI per se but rather its data that the humanities can situate themselves centrally within the development of AI technologies.

In not recognizing the centrality of textual interpretation to its practice, the field of AI finds itself stumbling through an epistemic labyrinth. Every step risks being turned down some winding corridor, new epistemic regimes obscuring the meaning of its data and of the machines built upon it. Building on scholars who have called for critical work excavating the histories of the data of AI (Plasek, 2016; Acker, 2015), I argue that such histories not only of data but of the more general data *practices* such as that semantic differential through which that data is collected, transmitted, modified, interpreted, and reinterpreted can serve as a thread that allows us to unwind the labyrinth. Such histories allow us not only to critique AI's potential harms but also, more provocatively, to sketch new visions for technology and society (Stark & Hoey, 2021). If, however, we are to unravel the history of this textile technology, then like Theseus we must follow the thread.

2. Emotion

2.1. Emotion-Based Textile Image Recognition

The semantic differential has most recently found application in Emotion-Based Textile Image Retrieval (EBTIR) research, which emerged from a series of papers published by a group at Konkuk university in South Korea between 2005 and 2009 (E. Y. Kim *et al.*, 2005; S.-j. Kim *et al.*, 2006; N. Y. Kim, Shin & E. Y. Kim, 2007; N. Y. Kim, Shin, Y. Kim, *et al.*, 2008; Shin *et al.*, 2010) and subsequently spread to several groups in China and Taiwan (Li *et al.*, 2017; Wang *et al.*, 2011; Liu *et al.*, 2015; Su & Sun, 2017). EBTIR researchers study algorithms for predicting the emotional responses of a viewer to a textile pattern. Such a model would be useful, researchers imagine, for textile merchants who need to automatically annotate catalogs of textile images with their emotional qualities so that consumers can search for different patterns by their emotional mood (E. Y. Kim *et al.*, 2005).

Within EBTIR research, the semantic differential serves as a means of producing datasets from which algorithms can learn to predict the emotional qualities of textiles. Not all EBTIR researchers refer to the technique by its name or seem aware of its history or original intent. Nevertheless, whether aware of its history or merely adopting its procedures from prior work, they gather images of textile design patterns and ask human

annotators to rate them on a series of scales between opposing pairs of adjectives representing opposite emotional states. This process of image gathering and annotation constitutes the semantic differential procedure. The annotated images are then fed as data into a machine learning algorithm. The algorithm scans each image for patterns of color or texture that seem to consistently correlate with certain ratings, perhaps finding that brightly colored textiles are consistently rated as highly 'cheerful' by the human annotators. In theory, if the algorithm were able to identify such correlations, it would be able to act in place of the human subjects, labeling all brightly colored images as 'cheerful'.

Significantly however, although researchers often invoke the language of 'emotion', most of the adjectives with which annotators are presented are not emotions. In this section I will examine this discrepancy and use it to highlight the interpretive practices through which the field of EBTIR and of AI more broadly can arrive at interpretations of its algorithms that are at odds with its data. The semantic differential requires researchers to supply annotators with adjectives describing the emotions they are to look for in the textile patterns. This requirement assumes that visual patterns can make us *feel* specific emotions. However, what that means is unclear. Can a fabric pattern, like Solomon's ring, *make* one feel cheerful when one is not (Fitzgerald, 1887)? If it cannot, then in what sense *is* a pattern *cheerful*? Had Theseus taken the black sails from his ship, the sight of which caused his father to believe his son had perished in the labyrinth and hurl himself into the sea, and had them tailored into a little black dress for his wife Phaedra, would it have evoked sadness or romance, and to what could its affective tone be ascribed? EBTIR research trades on an intuition about the relationship between emotions and visuality. However, at no point do they make precise the hypothesized link among vision, textile, and affect.

Such ambiguities are common in AI, which often struggles to formalize human experience. However, they are particularly acute in EBTIR work in that, with the exception of a small number of terms such as 'cheerful', the adjectives presented to human annotators have little to do with 'emotion' per se. Terms such as 'cold' and 'warm', 'soft' and 'hard', 'dynamic' and 'static', 'stable' and 'unstable', 'luxuriant', 'elegant', 'vigorous', 'flowing', 'rich', 'light', and 'heavy' define the perimeter of a deeply polysemous descriptive terrain. A closer examination of this

data raises the question of whether a 'light' color and a 'heavy' emotion are weighted in the same way, or whether they express two different senses of the idea of *weight*?

This discrepancy stems from an interpretive strategy common in AI, although prefigured by the statistical tradition of psychological research from which EBTIR borrows.² Because AI research aims to produce technologies that 'work' rather than knowledge about the phenomena on which they work, it is not the researchers but rather the algorithms that must *know* what emotions are (Mitchell, 2018; Blackwell, 2019). Researchers can only speculate about what the algorithm has learned, which leads to a discourse of enchantment within which it becomes possible to imagine that a vague referent such as the emotional quality of a textile pattern has been detected by the algorithm—and indeed that the very working of the algorithm proves the coherence of the referent—without questioning its coherence (Campolo & Crawford, 2020; Elish & Boyd, 2018).

Researchers' speculations about what an algorithm *has* learned, however, must be rooted in assumptions about what *in principle* it was *possible* to learn. This possibility, in turn, is circumscribed by what knowledge they believed to be *contained* in the data, and it is here that the need for deeper critical engagement with the histories of data becomes apparent. Even brief consideration of the EBTIR data calls into question the claims of the field, yet I think it is worth resisting the temptation to dismiss it out of hand. Rather, I propose that EBTIR offers a way in to understanding just how such discrepancies can develop without notice as a means to think more deeply about how to integrate critical engagements with data into AI practice.

2.2. Affective Computing

EBTIR's conception of emotion is situated within the broader field of affective computing. As I will argue in this section, the language of 'emotion' in EBTIR does not, as it may first appear, directly inform most technical decisions. Rather, it helps to conjure an imagined future of machines that *behave* emotionally that EBTIR researchers attempt to bring about through their work. The question of what an emotion *is*, either neurologically or phenomenologically, is secondary. Therefore

to engage productively with AI, critical work must grapple with the field's imagined futures—the imagined textile search engines that justify the field's existence—and reconceptualize what futures the field is striving to realize.

Affective computing as a field owes at least its original vision of computational emotion to the work of AI researcher Rosalind Picard in the late 1990s (1997). Picard's book is divided into two sections, one describing scattered past AI work on emotion and one offering a more theoretical argument for why the study of 'emotion' should be central to AI research. In the theoretical section, Picard draws on neuroscientific research to argue that, because emotion is bound up with other cognitive processes in the human brain, it must also be integral to intelligence in machines. Whether the connection between emotion in the human and in the machine is fundamental or fleeting is, of course, a matter of interpretation. Picard offers parallels between common errors AI systems make and behavioral studies of neurological patients with conditions arguably connected to interruptions in emotion processing in the brain as possible evidence of the connection (1997). However, the more significant work the book accomplishes is that of sketching a vision of a future populated by emotional machines.

Affective computing researchers have long argued over how neuroscientific theories of emotion can guide the design of intelligent machines (Sloman, 1999; Stark & Joey, 2021). Likewise, scholars from beyond the field have on occasion attempted to contribute perspectives from affect theory and other disciplines to inform the work of affective computing (Wilson, 2011). However, these debates do not, at least in the case of EBTIR research, result in concrete blueprints for applying the lessons of neuroscience or affect theory to the design of AI systems. Rather, to the extent that they influence the direction of technical research, they serve primarily to bring researchers together around a commonly envisioned goal—that of designing emotionally intelligent machines.

Picard's own path to work on emotion illuminates how the language of emotion helped to imagine the field's goal without necessarily offering insight into the path needed to achieve it. Picard's earliest work at the beginning of the 1990s involved designing algorithms for compressing and manipulating images (1992a). Over the course of the next few years, in conversation

with others in the field of image analysis and retrieval, she shifted from working with pure textures to analyzing the 'semantic' contents of an image (1992b), and from there to arguing for a science of 'subjectivity' concerned with identifying users' preferences to help recommend images they might like (1996). When she subsequently began to articulate her vision for affective computing, her most visible original contribution would be a proposal for a system rooted in her previous work that would allow users to search for images by 'emotional' mood very much in the spirit of subsequent EBTIR work (Sloman, 1999). The shift subtly transformed a technology of what might now also be called 'personalization' into a technology of emotion, despite the fact that the relationship of the technology to emotion as such remained unclear. There is no a priori reason to believe the images we want a system to show us are those that will elicit a measurable emotional reaction on sight. Yet in adopting this very assumption, EBTIR research is true in a very literal sense to the original vision of the technologically-aided practices affective computing aspired to enable even as the details of EBTIR systems defy easy categorization in terms of emotions.

The issue is not that EBTIR researchers do not notice the discrepancy between their theories of emotion and their data. Li et al., for instance, explicitly justify their divergence from more traditional emotions by arguing that textiles require their own fabric-oriented emotional vocabulary (2017). The problem is that 'emotion', like 'weight', does several different kinds of work in this instance. It encompasses both the joy sparked by a new article of clothing and the aesthetic characteristics of the types of clothing that spark that joy while also signaling membership in a field of research organized around the word as a technical term of art. These terms are emotions in the sense that the field understands its subject to be emotion while differing from emotions in the broader and less technical sense that forms the substance of the field's ultimate ambitions.

A critical engagement with AI must recognize the organizing function that broad discourses such as that of emotion perform in gathering researchers around a shared vision of a technical future. It is tempting to critique the apparently reductive categories of AI. However, such critiques will have little force if they do not suggest a different vision for emotional machines (P. Agre, 1997). The starting point for such engagements must be not an engagement with the theories, rhetorics, and

categories. Even among AI researchers such debates, I have argued, exist at a remove from the work of technology. Rather, it is necessary to turn such critical engagements towards the data of AI, because while the theoretical debates rage, it is the slight terminological slips that often go unnoticed and unremarked upon, such as Li *et al.*'s willingness to overlook the nuances of their adjectives in applying to them the broad term 'emotion', that have the most dramatic material consequences for the unfolding of a technical discipline. It is these slips of interpretation to which I contend the humanities are best positioned to attend.

3. Aesthetics

3.1. *Kansei* Engineering

Misclassifying technologies can have a profound impact on their future development, so closely intertwined are the broad visions and the technical minutiae in the act of invention (P. Agre, 1997). The history of EBTIR research highlights this dynamic by virtue of its similarity to the parallel tradition of *kansei* (感性) engineering that developed alongside it in Japan. *Kansei* is emphatically not translatable as 'emotion' (*kanjou* 感情), but rather means something closer to aesthetic sensibility or taste (Nagamachi, 2018). Yet their proximity has allowed two intimately related yet not quite identical traditions of technical research to develop around many of the same objects of study. In this section, I trace the evolution of *kansei* engineering, which inhabits an epistemic moment that overlaps with EBTIR research, and within which researchers understand the semantic differential to be a tool for measuring not 'emotion' but aesthetics. I examine how the same fundamental technologies, when viewed through a slightly different lens—*kansei* rather than emotions—can evolve in distinct technical directions. Building on the previous section, I argue that the language of *kansei* does not provide additional specific technical guidance for the design of AI systems, but by refocusing researchers' imaginaries from machines that deal with emotions to those that deal with aesthetics or taste, a seemingly minor terminological change nevertheless carries significant weight.

The semantic differential method first entered EBTIR research through the 1981 work of Shigenobu Kobayashi (1981), who

shared with EBTIR researchers a broad vision of creating a mathematical tool for use in industrial design (S.-j. Kim *et al.*, 2006). In that sense his work formed a natural foundation from which to further develop the statistical tools of textile analysis. However, unlike EBTIR researchers, for whom this task had already been pre-framed in the language of 'emotion', that term as such held little significance for Kobayashi. I will return later to the specifics of Kobayashi's work, but it suffices for the moment to note that it is only in the context of EBTIR research that researchers would come to view the semantic differential as a tool for measuring emotion. In Japan, the emerging discipline of *kansei* engineering would draw on Kobayashi and adopt the semantic differential as a tool for measuring *kansei* instead, which would lead to technical practices distinct from those of EBTIR research (Schütte *et al.*, 2004; Xue *et al.*, 2011; Tharangie *et al.*, 2010).

Kansei engineering developed in the mid-1980s out of research on ergonomics and industrial design (Nagamachi, 2018; Lévy *et al.*, 2007). With the subsequent emergence of image-based information retrieval in the 1990s—the discipline that gave rise to Picard's vision of affective computing—a tradition of *kansei*-based image retrieval emerged that applied the same methods EBTIR researchers inherited from Kobayashi to paintings, photographs, and even textiles, and bore many of the formal elements characteristic of the later EBTIR work (Yoshida *et al.*, 1998; Hayashi & Hagiwara, 1997; Sobue *et al.*, 2008). Yet, '*kansei*' is not 'emotion', and to the extent that these cognate fields have inevitably come into contact, there remains on the basis of that untranslatability a sense on both sides that they are related but not identical approaches to product design (Lévy *et al.*, 2007; Black Jr. *et al.*, 2004).

Despite their technical similarity, the very fact of their distinct names has tended to push each field in its own direction. Affective computing was from its inception firmly grounded in discourses of neuroscience and psychology. Much of the field remains dominated by a version of Basic Emotion Theory, which restricts the vocabulary through which researchers can frame their experiments to a handful of fixed categories authorized by scientific literature in the brain sciences (Stark & Hoey, 2021). It is for this reason that Li *et al.* feel the need to justify the divergence of EBTIR's emotional vocabulary from the norm in much of the rest of the field. Moreover, even having adopted this novel lexicon, they continue to treat it as

though it represented emotion in the traditional sense. They argue that even though they borrowed their adjective list in translation from a Chinese language work, this should not affect the results due to the universality of human emotional experience (Li *et al.*, 2017).³ Moreover, when annotators disagreed about which images to label as 'elegant', 'flowing', and 'romantic', they speculate that these 'emotions' may be more subjective and therefore require further study for algorithmic prediction (Li *et al.*, 2017). They do not, however, question whether or not the disagreement is due to the words' failure to meaningfully describe emotions. For Li *et al.*, the emotions are what they are because they are underwritten by human biology, and it is not within the scope of EBTIR research to re-taxonomize human emotion.

The term '*kansei*', by contrast, has no such association with a small and specific set of scientifically vetted affective states. As a result, *kansei* engineers are at liberty to incorporate new adjectives at the discretion of the researcher, perhaps gathering hundreds of candidate adjectives from relevant magazines, manuals, or experts without the need to justify them as rooted in neuroscience (Schütte, 2005). As a result, if annotators disagree, it is not assumed to be because of the subjectivity of the emotion the adjective describes but rather because of its lack of meaning and therefore descriptive utility in the domain. Such adjectives need not warrant further study, as they do for Li *et al.*, but may simply be discarded (Yanagisawa, 2011). Whereas EBTIR research takes the adjectives as given, *kansei* engineering begins with a search for the right words.

The adjectives ultimately included or omitted in this type of research define the limits of what the resulting technologies will be able to do. If images are not annotated with 'emotions', then no algorithm will be able to extract emotion from the dataset. Researchers' intuitions about emotions or *kansei* that inflect how they prepare their data already determine what it is possible for the algorithm to discover before it has even been run.

The situation of *kansei* engineering makes clearer than that of affective computing how it is the *word itself* rather than the scientific reality conjectured to underlie it that structures work in the field. *Kansei* engineers, lacking a neuroscientific discourse of *kansei* in which to ground the discipline, have developed an etymological tradition instead, plumbing the

history of the word *kansei* in search of spiritual guidance for the field. Researchers have traced the term's modern usage to a translation by early 20th century Japanese philosopher Teiyu Amano of the late 18th century German philosopher Immanuel Kant's *Critique of Pure Reason*, and in particular to a passage in which he takes issue with the *Aesthetica*, published in Latin, by his contemporary, the philosopher Alexander Baumgarten (Levy, 2013). The term at issue in Kant is the German word '*sinnlichkeit*', by which the philosophers mean the pre-conceptual sensory impression left upon the human sensorium by the stimuli of the external world (1998), and with which Baumgarten argues for a universal science of aesthetic form (Gregor, 1983). It is therefore in Baumgarten, some *kansei* scholars claim, that one can find the meaning of *kansei* (Lee et al., 2002).

Baumgarten, whose project was a science of beauty (Gregor, 1983), may well have approved of the discipline that has become his legacy. Yet, as I argued in the case of affective computing, it would be difficult to trace a direct arc from the *Aesthetica* to contemporary *kansei* textile retrieval practice in the same way that it would be difficult to see many of the neurological connections Picard draws in *Affective Computing* in contemporary EBTIR work. These words create research collectives, the collectives create shared visions, and it is within the epistemic fabric of these visions that subtle changes to technical practice begin to alter the warp and weft of the technologies produced. Even a slight difference in terminology can contribute to the divergence of two distinct technical divisions on the basis of the realities they make imaginable. As psychologist Gregory Kimble remarks of psychology, although I suggest with strong relevance to the present discussion, 'the problems of psychology are the same as those that frustrate public understanding, and for the same reason: The language of psychology is also that of common sense... If there is a word for it, there must be a corresponding item of reality. If there are two words, there must be two realities and they must be different' (1995: 137).

The repeated emergence of shared visions anchored by the language of emotion or *kansei* speaks to the necessity of such visions in organizing AI research. To engage critically with AI cannot simply be to find philosophical fault with how these terms are operationalized. Rather, it must be to understand more precisely what the technologies produced within these

fields actually do, and to balance the sometimes reductive demands of the broader visions against the specificities of human expression that resist reduction to uniform data. The language of emotion allows a broad range of heterogeneous technologies to come together into a shared vision of a technical future, yet without attending to the specifics of how these technologies inflect the concept of emotion, one risks becoming lost in a labyrinth in which every turn is identical, the most valuable insights into the nature of human expression our machines might offer us overlooked in our hurry to reach that future.

4. Meaning

4.1. The Measurement of Meaning

It would be easy to imagine that, whatever one called it, EBTIR and *kansei* engineering both circled around an as yet poorly understood but nevertheless stable scientific reality of mood, feeling, or aesthetics that underwrote the eventual success of such programs regardless of the minor confusions caused by unstable terminology. However, the long history of the semantic differential suggests that technologies do not always converge on a scientific real but may sometimes remain radically and indefinitely open to reinterpretation. It suggests that it is fully possible to wander forever amidst the corridors of the labyrinth without ever finding the exit. Admittedly, the small technical changes from moment to epistemic moment discussed in the previous sections mean that, as the decades pass, the boundaries around what constitutes a seemingly coherent and mathematically precise technology in the present become blurred, and the crisp form of the semantic differential undergoes modification to the point at which it is arguably no longer the same technology, complicating attempts to trace its history. Moving backwards into this history, high resolution textile images are replaced with colored paper, and the list of adjectives undergoes slight, almost unconscious modification to better reflect the intuitions of the moment. However, like Theseus's ship in Plutarch's riddle, returned from his heroic journey into the labyrinth, docked in the great harbor of Athens years after the hero's death, board after rotting board replaced until not a single piece of the original remained, so too does the familiar structure of the semantic differential remain recognizable by its telltale features even as each piece is

modified and ultimately replaced. Following the semantic differential back into its early history, therefore, like Fleck's excursion into the history of the scientific concept of 'weight', helps to reveal how the discourses of aesthetics and affect that now seem so natural to so many researchers in AI guarantee nothing in and of themselves about the teleology of technology over the long course of technological development.

In both EBTIR and *kansei* engineering research, Kobayashi's work participates in discourses that would have been unfamiliar to Kobayashi himself. Working before the advent of modern machine learning techniques, he used simple colors in place of complex textile patterns, but the methods by which he elicited annotations using scales of opposing adjectives remained the same.⁴ He understood the semantic differential to be a tool for measuring neither 'emotion' nor '*kansei*' per se, but rather what he referred to as 'color images'. He posited that all mental concepts—or 'images'—had an independent mental existence to which adjectives and colors could both refer. To call a red color 'warm' was to suggest that there was some mental concept that the color red and the word 'warm' both referred to as if two separate dialects of the same language. Even though he aspired to create a scientific tool for industrial design, this notion of color images was colored by an earlier moment surrounding the semantic differential that reflected no such aspiration and from which he in turn had borrowed. Rather, in this moment, the semantic differential was not a tool for industrial design, but one for probing the mysteries of 'meaning' in the human psyche.

Kobayashi borrowed the semantic differential—including the list of adjectives—from the 1962 work of Oyama *et al.* (1962). In this work, Oyama *et al.* describe an experiment in which a group of American and a group of Japanese subjects were shown colored cards and asked to rate them on a familiar set of scales described by opposing adjectives. The goal of this work was not to discover consumer preference, but rather to map the structure of the human mind itself.

Oyama *et al.* were colleagues of American psychologist Charles Osgood, who originally introduced the 'semantic differential' method in his book, *The Measurement of Meaning* (1964).⁵ Osgood locates inspiration for its development in a collection of ethnographic field reports documenting commonalities in the mythic traditions of widely separated

cultures that became an object of fascination of his during work on his undergraduate thesis at Dartmouth College (1964: 23). He reports being struck by the seemingly common motif the world over of the original human beings climbing, 'from the dark, cold, wet, sad world below the ground up to the light, warm, dry, happy world on the surface', and suspected that this seeming commonality pointed to a universal mental structure (Osgood *et al.*, 1964: 23). The semantic differential emerged out of his efforts to measure this mythical space he imagined existed within the human mind. Where Kobayashi saw a language of images, Oyama *et al.*'s experiment incorporated both English and Japanese speakers, partly at Osgood's urging, in an effort to locate a Levi Straussian universal language of myth (Lévi-Strauss, 1963).

Insofar as Oyama *et al.*'s understanding of the meaning of the data produced by the semantic differential differed from later EBTIR researchers, their methods of analysis and conclusions drawn did as well. The autoencoders characteristic of the latest generation of EBTIR research are designed to predict the best adjectives to describe a textile. The theory of EBTIR underlying them holds that in the responses of the human annotators lies an ineffable recognition of a particular visual feature with a particular emotional adjective. Oyama *et al.*'s work, by contrast, analyzes the results with a statistical technique known as factor analysis.

In a factor analysis, the annotator responses are analyzed not for how the adjectives correspond to the visual pattern, but for how adjectives correlate with one another in terms of which patterns they apply to. It is significant whether colors score similarly on the 'happy/sad' scale and on the 'light/heavy' scale, as a strong correlation would suggest that happiness and lightness are, as Schreger proposed, two surface manifestations of the same deeper structure of human meaning. In the original semantic differential, researchers used visual patterns only as a way to test how adjectives correlated with one another, and their goal was not a system for prediction but a systematization of the hypothesized natural categories of the mind. Predicting the adjectives that a specific color would evoke was, for the earliest users of the semantic differential, a meaningless task. Yet, Osgood shared with the later work a familiar sense of the enchantment of statistical tools, imagining that the factor analysis was producing a model of an as yet undiscovered neurological mechanism in the human brain that in his moment

could only be understood through interpretation of the statistical models capable of teasing out its effects (1962). This hypothesized neurological mechanism would take on an even more specific, literal character in the sources from which Oyama *et al.* in turn drew the raw materials for their semantic differential analysis.

5. Psychosomatic Reflex

5.1. The Apparent Weight of Color

In a strict sense, the 'semantic differential' cannot precede Osgood and his colleagues who would first name it. The adjectives, the opposing scales, and the factor analyses disappear once one moves past that point. Yet, in selecting the lists of adjectives that would endure for decades afterwards, Oyama *et al.* drew on an early 20th-century psychological literature in which the rudiments of the later semantic differential method are still visible.

The psychologists on whom Oyama, Tanaka, Chiba, and Osgood drew were part of an epistemic moment that took as its central hypothesis the contention that linguistic metaphor was in fact an expression of a deeper neurophysiological literalism. Researchers in this moment grappled with the individual adjectives that Osgood and his colleagues would later assemble into the more familiar form of the semantic differential. Warm colors did not make one feel emotionally warm, but literally, cutaneously warm (Tinker, 1938). Heavy colors were experienced not affectively, but kinesthetically. To see something painted in a heavy color was, literally, to perceive it pre-conceptually as weighing more—to feel the 'apparent weight of color' (Payne, 1958; Oyama *et al.*, 1962).

Schreger viewed weight as a physical sensation and as an emotional experience, both products of the action of fiery spirits. In the century long history of the semantic differential, unseen brain structures took the place of fiery spirits, but the boundary between kinesthetic and affective weight remained blurry. At no point in this history were discoveries made or new theories proposed that better explained the empirical phenomenon of human annotators ascribing adjectives to color patterns. There were no clear moments of scientific revolution (Kuhn, 1970). Rather, the same phenomenon was found to fit

naturally into successive generations of explanations without ever offering resistance, seemingly little noticed or remarked upon by the researchers themselves.

If in the present moment it seems self evident that color patterns do not mimic the physical experience of warmth or weight, but that they do have some form of emotional, affective, or aesthetic effect, what empirical phenomena could we offer to Schreger to convince him that cheerfulness and ponderousness were two distinct phenomena? Would he immediately agree to the rightness of our divisions, asks Fleck, or would they seem as alien to him as do his to those acquainted with a modern conception of weight (2012)?

The semantic differential, and indeed many of the methods of modern AI, are radically open to interpretation. Because of the subtleties of concepts and the coarseness of the language through which we grapple with them, it is easy to slip without realizing it from one notion to another, and to potentially great technical and social consequence. It is for this reason that I have argued for a deeper critical engagement with the data of AI, because interpretations of that data give meaning to the system as a whole. When annotators appear as shadowy figures in the discussion of the research methodology, taking the stage for only a single sentence, it is easy to imagine that the act of annotating was simple and unproblematic. However, on further consideration it is unclear, when annotators are asked to render emotional responses in a strange language of fabrics, what interpretive process they use to produce their judgments. All of the work in EBTIR, as in much of modern AI, rests on the assumption that annotators possess an inherent expertise that can guide algorithms to a humanlike sensibility. However, once the act of annotation is complete, the precise contexts that led to their decisions are often rendered invisible by a literature focused on algorithms rather than on the material conditions of the production of its data and it is this choice of algorithms over data that the history of the semantic differential calls into question more than a century before current work in EBTIR.

Of all the works belonging to the extended history of the semantic differential that I have traced throughout this essay, only one engages critically with the process of annotation at the heart of the method. In 1907, the psychologist E. Bullough, partly out of a suspicion concerning the very types of statistical methods that would later characterize modern AI, engaged the

subjects in his study in conversation about their interpretive decisions. His method was unsystematic, but what he discovered can only be described as hermeneutic chaos.

Much of Bullough's experiment involved showing subjects pairs of geometric shapes, such as a pair of triangles, each divided into two halves horizontally (1907). The first triangle was colored with one pair of colors for each of its two halves, such as light and dark green, and the other triangle was colored in the reverse fashion, with the darker and then the lighter green. Bullough then asked subjects which coloration they preferred.

To explain this experimental design, he invokes an unsourced 'decorative cannon' that the eye prefers a wall to be painted with darker colors at the bottom and lighter colors at the top (Bullough, 1907: 113). Perhaps, he speculates, the truth of this apparent fact lies in an instinctive human sense that darker colors communicate an unconscious 'moreness', as of a darker claret communicating a more concentrated wine, and give the viewer a greater sense of stability and weight (Bullough, 1907: 113). Even in reasoning through his experiment, he blends categories of heaviness, darkness, and moreness across different media, unsure of how to draw the ontological lines.

Bullough's subjects, lacking access to his rationale, experience the resulting questions concerning colored geometric patterns as deeply ambiguous, and cast out desperately for any interpretive device they can think of to make sense of these unusual questions. Some subjects perceived such triangles as abstract shapes, and attempted to determine a preference on aesthetic grounds. Others viewed them as depictions of landscapes (such as a distant green field behind a sunlight grassy meadow), and so chose the triangle that could most easily be made sense of in those terms. Others viewed them simply as quantities of color, and chose the figure in which the larger half contained the color they preferred.

Bullough mused that perhaps had he been able to paint an actual wall, per his decorative cannon, he would have been able to reduce the interpretive ambiguity in his study. Perhaps he could have, but for that matter even a slight change in the representational medium might have had profound consequences for the experimental results of any of the works that would build on his results. Would Li et al.'s annotators,

over a century later, have had the same disagreement about which images depicted 'flowing' or 'elegant' fabrics had they been annotating pictures of evening dresses rather than swatches of abstract colors? There is as much a difference among an image of a pattern, a photograph of a textile product, and the textile product itself as there is between a little black dress and the black sails of Theseus. Entire research fields live or die on these distinctions without ever seeming to consider them, and it is for this reason that I have argued that AI requires a critical, medium-specific analysis of its data (Hayles, 2004).

6. Conclusion

In this essay, I have followed the thread of the semantic differential through a century of reinterpretations. The semantic differential, like the ship of Theseus, remained unchanged as each piece was gradually replaced. What changed was not the ship, but rather the generations of Athenians who would come to see the ship in the harbor of Athens, and who in turn would retell to one another the myth of the man as, bit by bit, he became legend. If we are to imagine new futures for these technologies, much hinges on our ability to understand the pasts that have been obscured by these myths.

I have argued that critical histories of data processes can offer a thread that can guide us through the labyrinth, but they can only take us so far. If AI were to be infused with the intuitions regarding the human that it seems to lack, the task would still remain to decide what to build. The epistemic labyrinth is a labyrinth without a door—without a future that can be reached purely by unwinding its twisting passages. Only Theseus can ever truly leave the epistemic labyrinth this way, guided by Ariadne's magic thread, because Theseus is not a man but a legend—like a hero of some Levi Straussian myth.

What such histories can do is help us remember what it was once possible to imagine. They can teach us to understand the labyrinth so that one day we might leave the labyrinth as did Daedalus—through the spark of invention. Imprisoned in his own creation, Daedalus invented wings, for the labyrinth has no roof.

What, after all, are warm colors and cool colors, heavy colors and light colors? Where do these distinctions intersect with language and culture and how through them can we better understand ourselves? If we set aside for the moment the lenses of emotion, *kansei*, color images, mythemes, and kinesthetic reflexes, what if anything does a century of shifting intuitions about color, form, aesthetics, design, consumption, language, metaphor, and myth tell us about ourselves? What would we find if we more carefully attended to how the data of AI was created—how forms of human experience became crystallized as data—and what futures might it enable us to imagine? How can the texts authored by our machines help us understand what it means to be human? What winding paths might such knowledge lead us down? What AI needs from criticism more than method is purpose.

Daedalus built a magnificent labyrinth, but it brought only despair for the Athenians and death for his son, Icarus. Is this not the situation AI finds itself in and with which its critics are principally concerned—that of building technical marvels that may harm rather than heal? Affective computing has produced technologies to scan crowds for 'agitated' individuals and detect 'deceitfulness' in courtroom testimony (Andalibi & Buss, 2020; Wu *et al.*, 2017). If these technologies worked as they claimed, we might not even desire the futures they would help bring about (Andalibi & Buss, 2020; Reynolds & R. Picard, 2004). If, however, claims about such technologies rested on data that, like the semantic differential, was deeply ambiguous, it is chilling to think what injustices could be wrought by their misapplication. However, all the social power of these algorithms rests crucially on interpretation. If they measure not emotion but aesthetics, these powers may vanish. If they measure meaning, perhaps they lead in new directions. It is ultimately through textual criticism, carried out often unconsciously within the field of AI, through which algorithms achieve their social functions and through which researchers imagine new possibilities—new futures—towards which to build.

Only once Daedalus had experienced and come to understand the horror of the labyrinth he created could he begin to imagine building not a labyrinth but a temple. Criticism must not satisfy itself with stopping the construction of labyrinths. Rather, it should ask how can the same craftsmanship be directed to the building of temples? After the labyrinth, Daedalus built a

temple to Apollo to atone for his sins. On its gilded doors, he carved the story of the labyrinth up to the point where Icarus fell into the sea, his wings unable to bear his weight. At that point Daedalus's own hands fell, his spirits unable to raise them, and he understood at last just how much can turn on a simple concept like 'weight'.

Notes

1. This is not of course to overlook the significant ethnographic work that has documented the construction of scientific knowledge within AI laboratories (Forsythe, 2001; Collins, 1990; Hoffman, 2015). Rather, it is to argue that the work of AI takes place within larger cultural imaginaries shared among researchers across time and space and that the boundaries of these imaginaries only become visible through a critical engagement with the documents that register their formation and evolution.
2. See Yarkoni (2020) for further discussion of this issue in psychology, although ironically in this instance he considers the example of machine learning as a possible remedy.
3. Naoki Kawamoto and Toshiichi Soen, on whose work EBTIR researchers would also draw, speculated that advances in image processing might allow psychological work on color to address more complicated fabric patterns, prefiguring the link between this work and that of later of EBTIR researchers (1993).
4. Contra Li *et al.*'s claims, the list of adjectives registers acutely the effects of translation across multiple languages. The word 'clear,' for instance, appears variously throughout this history paired with 'indistinct' (Sobue *et al.*, 2008), 'greyish' (Kobayashi, 1981), and 'muddy' (Oyama *et al.*, 1962), in the last of which it appears as a translation of the pair '*sunda-nigotta*' (澄んだ-濁った) and captures the sense of clear or muddied water rather than grayscale or resolution as suggested by the other pairs. Moreover, the word 'romantic', which appears in Li *et al.*'s own work, was not originally part of the list, but was rather borrowed by Kim *et al.* from a section of Kobayashi's paper where he proposes ex nihilo a list of 'fashion terms' and analyzes them *in terms of* the list of color-images (N. Y. Kim, Shin & E. Y. Kim, 2007; Kobayashi, 1981).

5. It was Osgood himself who proposed adding the cross-cultural dimension to the work the better to explore universal characteristics of meaning (Oyama *et al.*, 1962: 78).

References

Acker, A. (2015) “Toward a Hermeneutics of Data”, *IEEE Annals of the History of Computing* 37. No. 3: 70–75.

Agre, P. (1997) “Lessons Learned in Trying to Reform AI”, in *Social Science, Technical Systems, and Cooperative Works: Beyond the Great Divide*, (eds.) Bowker Geof et al. New Jersey: Erlbaum Associates.

Agre, P. (1997) *Computation and Human Experience*. Cambridge: Cambridge University Press.

Andalibi, N. & Buss, J. (2020) “The Human in Emotion Recognition on Social Media: Attitudes, Outcomes, Risks”, *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. (April 25th – 30th): 1-16.

Black Jr., J. et al. (2004) “Indexing Natural Images for Retrieval Based on Kansei Factors”, *Proceedings volume 5292, Human Vision and Electronic Imaging IX*. (June 7th): 363–375.

Blackwell, A. (2019) “Objective Functions: (In) Humanity and Inequity in Artificial Intelligence”, *HAU: Journal of Ethnographic Theory* 9. No. 1: 137–146.

Bullough, E. (1907) “On the Apparent Heaviness of Colours”, *British Journal of Psychology* 2. No. 2: 111–152.

Campolo, A. & Crawford, K. (2020) “Enchanted Determinism: Power Without Responsibility in Artificial Intelligence”, *Engaging Science, Technology, and Society* 6. 1–19.

Castelle, M. (2018) “The Linguistic Ideologies of Deep Abusive Language Classification”, in *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*. (October): 160–170.

Collins, H. (1990) *Artificial Experts: Social Knowledge and Intelligent Machines*. Cambridge: MIT press.

Dreyfus, H. (1965) *Alchemy and AI*. Santa Monica: RAND Corporation.

Edwards, P. (1997) *The Closed World: Computers and the Politics of Discourse in Cold War America*. Cambridge: MIT Press.

Elish, M. & Boyd, D. (2018) “Situating Methods in the Magic of Big Data and AI”, *Communication Monographs* 85. No. 1: 57–80.

Fitzgerald, E. (1887) *Works of Edward Fitzgerald*. New York: Houghton, Mifflin & Co.

Fleck, L. (2012) *Genesis and Development of a Scientific Fact*. Chicago: University of Chicago Press.

Forsythe, D. (2001) *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford: Stanford University Press.

Gregor, M. (1983) “Baumgarten's ‘Aesthetica’”, *The Review of Metaphysics* 37. No. 2: 357–385.

Hayashi, T. & Hagiwara, M. (1997) “An Image Retrieval System to Estimate Impression Words from Images Using a Neural Network”, in *1997 IEEE International Conference on Systems, Man, and Cybernetics. Computational Cybernetics and Simulation*. (October 12th – 15th): 150–155.

Hayles, N. (2004) “Print is Flat, Code is Deep: The Importance of Media-Specific Analysis”, *Poetics Today* 25. No. 1: 67–90.

Hoffman, S. (2015) “Thinking Science with Thinking Machines: The Multiple Realities of Basic and Applied Knowledge in A Research Border Zone”, *Social Studies of Science* 45. No. 2: 242–269.

Jo, E. & Gebru, T. (2020) “Lessons from Archives: strategies for collecting sociocultural data in machine learning”, *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. (January 27th – 30th): 306–316.

- Kant, I. (1998) *Critique of Pure Reason*. Cambridge: Cambridge University Press.
- Karoff, P. (2019) “Embedding Ethics in Computer Science Curriculum”, *The Harvard Gazette* (January 25th): <https://news.harvard.edu/gazette/story/2019/01/harvard-works-to-embed-ethics-in-computer-science-curriculum/>.
- Kawamoto, N. & Soen, T. (1993) “Objective Evaluation of Color Design II”, *Color Research & Application* 18. No. 4: 260–266.
- Kim, E. et al. (2005) “Emotion-Based Textile Indexing Using Colors and Texture”, in *International Conference on Fuzzy Systems and Knowledge Discovery*. (September 24th – 28th): 1077–1080.
- Kim, N. et al. (2007) “Emotion-based Textile Indexing using Neural Networks”, in *2007 IEEE International Symposium on Consumer Electronics*. (June 20th – 23rd): 1-6.
- Kim, N. et al. (2008) “Emotion recognition using color and pattern in textile images”, in *2008 IEEE Conference on Cybernetics and Intelligent Systems*. (September 21st – 24th): 1100 – 1105.
- Kim, S. et al. (2006) “Emotion-Based Textile Indexing Using Colors, Texture and Patterns”, in *2nd International Symposium Advances in Visual Computing*. (November 6th – 8th): 9–18.
- Kimble, G. (1995) *Psychology: The Hope of a Science*. Cambridge: MIT Press.
- Kobayashi, S. (1981) “The Aim and Method of the Color Image Scale”, *Color Research & Application* 6. No. 2: 93–107.
- Kuhn, T. (1970) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Lee, S. et al. (2002) “Design Based on Kanse”, in *Pleasure with Products: Beyond Usability*. (eds.), W. Green & P. Jordan. New York: Taylor and Francis.

Lévi-Strauss, C. (1963) *Structural Anthropology*. New York: Basic Books.

Levy, P. (2013) “Beyond Kansei Engineering: The Emancipation of Kansei Design”, *International Journal of Design* 7. No. 2: 83-94.

Lévy, P. et al. (2007) “On Kansei and Kansei Design: A Description of Japanese Design Approach”, *Proceedings of the Second Congress of International Association of Societies of Design Research*. (November 12th – 15th).

Li, Z. et al. (2017) “Emotional Textile Image Classification Based on Cross-Domain Convolutional Sparse Autoencoders with Feature Selection”, *Journal of Electronic Imaging* 26. No. 1: 1–16.

Liu, J. et al. (2015) “Could Linear Model Bridge the Gap Between Low-Level Statistical Features and Aesthetic Emotions of Visual Textures?”, *Neurocomputing* 168. 947–960.

Markham, A. (2013) “Undermining ‘data’: A critical examination of a core term in scientific inquiry”, *First Monday* 18. No. 10.

Mitchell, C. (2018) “Whether Something Works: Finding the Human in the Margins of Machine Translation”, *Amodern* 8: *Translation-Machination* 8.

Nagamachi, M. (2018) “History of Kansei Engineering and Application of Artificial Intelligence”, in *Advances in Affective and Pleasurable Design*, (eds.) W. Chung & C. Shin. Cham: Springer International Publishing.

Osgood, C. (1962) “Studies on the Generality of Affective Meaning Systems”, *American Psychologist* 17. No. 1: 10.

Osgood, C. et al. (1964) *The Measurement of Meaning*. Urbana: University of Illinois Press.

Oyama, T. et al. (1962) “Affective Dimensions of Colors”, *Japanese Psychological Research* 4. No. 2: 78–91.

Payne, M. C. (1958) “Apparent Weight as a Function of Color”, *The American Journal of Psychology* 71. No. 4: 725–730.

Picard, R. (1992a) “Gibbs random fields: temperature and parameter analysis”, in *IEEE International Conference on Acoustics, Speech, and Signal Processing*. (March 23rd – 26th): 45–48.

Picard, R. (1992b) “Random Field Texture Coding”, in *Society for Information Display International Symposium Digest*, (ed.) Society for Information Display. Boston: SID.

Picard, R. (1997) *Affective Computing*. Cambridge: MIT Press.

Picard, R. et al. (1996) “Modeling User Subjectivity in Image Libraries”, in *Proceedings of 3rd IEEE International Conference on Image Processing*. (September 19th): 777–780.

Plasek, A. (2016) “On the Cruelty of Really Writing a History of Machine Learning”, *IEEE Annals of the History of Computing* 38. No. 4: 6–8.

Reynolds, C. & Picard, R. (2004) “Affective Sensors, Privacy, and Ethical Contracts”, in *CHI'04 Extended Abstracts on Human Factors in Computing Systems*. (April 24th – 29th): 1103–1106.

Schrock, A. (2017) “What Communication Can Contribute to Data Studies: Three Lenses on Communication and Data”, *International Journal of Communication* 11. 701-709.

Schütte, S. (2005) “Engineering Emotional Values in Product Design”. PhD thesis. Linköping: Linköping University.

Schütte, S. et al. (2004) “Concepts, Methods and Tools in Kansei Engineering”, *Theoretical Issues in Ergonomics Science* 5. No. 3: 214–231.

Seaver, N. (2019) “Knowing Algorithms”, in *DigitalSTS: A Field Guide for Science & Technology Studies* (eds), J. Vertesi & D. Ribes. Princeton: Princeton University Press.

Selbst, A. et al. (2019) “Fairness and Abstraction in Sociotechnical Systems”, in *Proceedings of the Conference on*

Fairness, Accountability, and Transparency. (January 29th – 31st): 59–68.

Shin, Y. et al. (2010) “Automatic Textile Image Annotation by Predicting Emotional Concepts from Visual Features”, *Image and Vision Computing* 28. No. 3: 526–537.

Slooman, A. (1999) “Review of Affective Computing”, *AI Magazine* 20. No. 1: 127.

Sobue, S. et al. (2008) “Mapping Functions Between Image Features and KANSEI and Its Application to KANSEI Based Clothing Fabric Image Retrieval”, in *ITC-CSCC: International Technical Conference on Circuits Systems, Computers and Communications*. (July 6th – 8th): 705–708.

Stark, L. & Joey, J. (2021) “The Ethics of Emotion in Artificial Intelligence System”, *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. (March): 782-793.

Su, Y. & Sun, H. (2017) “Emotion-Based Classification and Indexing for Wallpaper and Textile”, *Applied Sciences* 7. No. 7: 691.

Suchman, L. (1987) *Plans and Situated Actions: The Problem of Human-Machine Communication*. Cambridge: Cambridge University Press.

Tharangie, K. et al. (2010) “Kansei Colour Concepts to Improve Effective Colour Selection in Designing Human Computer Interfaces”, *International Journal of Computer Science Issues (IJCSI)* 7. No. 3: 21.

Tinker, M. (1938) “Effect of Stimulus-Texture Upon Apparent Warmth and Affective Value of Colors”, *The American Journal of Psychology* 51. No. 3: 532–535.

Wang, X. et al. (2011) “Modeling the Relationship Between Texture Semantics and Textile Images”, *Research Journal of Applied Sciences, Engineering and Technology* 3. No. 9: 977–985.

Wilson, E. (2011) *Affect and Artificial Intelligence*. Seattle: University of Washington Press.

Wu, Z. et al. (2017) “Deception Detection in Videos”, *CoRR* abs/1712.04415. arXiv: 1712.04415.

Xue, Y. et al. (2011) “An Analysis of Emotion Space of Bra by Kansei Engineering Methodology”, *Journal of Fiber Bioengineering and Informatics* 4. No. 1: 96–103.

Yanagisawa, H. (2011) “Kansei Quality in Product Design”, in *Emotional Engineering*, (ed.) S. Fukuda. London: Springer.

Yarkoni, T. (2020) “The Generalizability Crisis”, *Behavioral and Brain Sciences*: 1–37.

Yoshida, K. et al. (1998) “Image Retrieval System Based on Subjective Interpretation”, *International Journal of Biomedical Soft Computing and Human Sciences: the Official Journal of The Biomedical Fuzzy Systems Association* 4. No. 1: 65–74.