

HOW TO DESIGN AN EMOTION · WORKING PAPER 7 OF 8

The Constructivist Linguistics of Emotion: The Symbolic Threshold, Emotionese, and the Two Levels of Emotional Meaning

Cézjah (Cecil John)

ORCID 0009-0005-1866-9700

New Intellectual LLC, Miami

Correspondence: cezjah@cezjah.com

24 October 2026 · Version 2 (26 October 2026)

DOI: Crossref DOI pending

Also available as https://doi.org/10.5281/zenodo.23676664 (Zenodo)

License: CC BY 4.0 (this paper only; the book is All rights reserved)

newintellectual.com/pages/how-to-design-an-emotion-wp-07

ABSTRACT

Do animals and infants have emotions? I argue that observing arousal, expression, or feeling cannot settle the question, because emotional meaning depends on a symbolic architecture that feeling does not supply. I define meaning as a structured representation that a system interprets in light of its stake, and I name four thresholds a system must cross to reach it: symbolic detachment, endorsement, grammar, and metarepresentation. I call the symbolic system of emotional meaning Emotionese: learned, public, normative, and confined to one domain, unlike the language of thought. The Emotional Meaning Architecture, a ladder of ten levels, places the line between feeling and emotional meaning between Level 4 and Level 5. On current evidence, animals feel and stop short of that line; infants feel and then cross it, and I name three results that would defeat the claim. I then separate the kinds of reality that people collapse into one, name the resulting error misattribution, and distinguish two levels of emotional meaning: a prototype, one agent's authorized construct, and a stereotype, its socially stabilized form. What is new is a threshold model with a stated test, and one account that runs from that threshold through two levels of stabilization. The book gives the full treatment.

The Constructivist Linguistics of Emotion: The Symbolic Threshold, Emotionese, and the Two Levels of Emotional Meaning

Cézjah
New Intellectual LLC, Miami
Open Researcher and Contributor ID (ORCID): 0009-0005-1866-9700

Working paper 7 of 8 in the series on the frameworks of *How to Design an Emotion*, October 2026
Digital object identifier (DOI): 10.5281/zenodo.23631062
License: CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/)

Abstract

Do animals and infants have emotions? I argue that observing arousal, expression, or feeling cannot settle the question, because emotional meaning depends on a symbolic architecture that feeling does not supply. I define meaning as a structured representation that a system interprets in light of its stake, and I name four thresholds a system must cross to reach it: symbolic detachment, endorsement, grammar, and metarepresentation. I call the symbolic system of emotional meaning Emotionese: learned, public, normative, and confined to one domain, unlike the language of thought. The Emotional Meaning Architecture, a ladder of ten levels, places the line between feeling and emotional meaning between Level 4 and Level 5. On current evidence, animals feel and stop short of that line; infants feel and then cross it, and I name three results that would defeat the claim. I then separate the kinds of reality that people collapse into one, name the resulting error misattribution, and distinguish two levels of emotional meaning: a prototype, one agent's authorized construct, and a stereotype, its socially stabilized form. What is new is a threshold model with a stated test, and one account that runs from that threshold through two levels of stabilization. The book gives the full treatment.

Keywords: Constructivist Linguistics of Emotion; Emotionese; Emotional Meaning Architecture; symbolic detachment; symbolic threshold; metarepresentation; commitment; misattribution; prototype and stereotype; language of thought; animal emotions; infant emotional development; Robert Brandom

Relation to the book. This paper restates, in new prose, the argument of Chapters 10 to 12 of *How to Design an Emotion: The Deconstruction and Computation of Emotional Meaning from First Principles* (Cézjah 2026, chaps. 10–12). It stops at Chapter 12. The book's account of emotional categories, music, and phenomenological timbre (chap. 13) appears in working paper 8. A selection of the book's diagrams, excerpts of its C# code, and their console input and output appear in the appendix to this paper. Every source cited here is cited in the book, with the same locators.

1. Introduction: Feeling Without Meaning

Do animals have emotions? Do babies? Most answers begin from what an observer can see: a dog cowers, a monkey gives an alarm call, an infant cries and is soothed. Those observations matter, but they do not settle the question, because they record affect, feeling, and behavior, and the question concerns something further. If emotion, in the sense the book develops, is a goal-governed act by which an agent assigns significance to a felt state under context, constraint, and possible correction (Cézjah 2026, chap. 9), then the decisive issue is whether a creature has the architecture that act requires.

This paper states that architecture. Its central claim is that emotional meaning has symbolic conditions that feeling does not supply. A mind can sense, feel, flinch, regulate, and learn without any symbol. It cannot mean, in the strict sense, without symbols it can detach from the moment that prompted them, combine by rule, act under, and revise. Ludwig Wittgenstein (1889–1951) wrote that “the limits of my language mean the limits of my world” (Wittgenstein 1922, 5.6). Readers often take the line to mean that a person can know only what that person has words for. My concern is narrower and, in one way, wider. It is meaning rather than knowledge, and it counts as symbols far more than words.

The argument runs in three stages. Sections 2 to 4 define meaning, state the four thresholds, and introduce Emotionese, the symbolic system of emotional meaning. Section 5 applies the thresholds to animals and infants through the Emotional Meaning Architecture and names the results that would defeat the claim. Sections 6 and 7 separate the kinds of reality in emotional life, name the error of misattribution, and state the two levels at which agents and communities stabilize emotional meaning. Section 8 states what the account shares with its nearest prior work and what it adds, Section 9 answers objections the book anticipates, and Sections 10 and 11 mark the scope and conclude.

2. What Meaning Requires

The word meaning shifts along three axes, and an account that does not separate them will equivocate (Cézjah 2026, chap. 10).

The first axis is phenomenological: sensation versus representation. A bodily change becomes sensation when it is felt, when there is something it is like to be in that state (Nagel 1974). Sensation presents experience. It does not yet organize that experience under conceptual constraint, so it does not yet represent anything as anything.

The second axis is normative: structure versus stake. A thermostat registers temperature differences under strict structural rules and regulates a furnace against a set point. Nothing rides on the result for the thermostat; its designer assigns its function from outside. Stake enters when a

system treats something as mattering because it can gain, keep, or lose it through its own continued operation. Structure without stake computes blindly. Stake without structure is undifferentiated impulse.

The third axis is interpretive: representation versus interpretation. A representation can sit inside a system and remain inert. Interpretation is uptake: the system proceeds under the representation, building it into expectation, choice, and sensitivity to outcome.

Together the axes yield a minimal criterion. Meaning is a structured representation that a system interprets in light of its stake: something presented as something, to a system for which it matters. One clarification follows at once. A system can sort patterns reliably into threat, food, or ally through learned regularities and still lack the identity conditions that make those sortings correct across contexts. Extensional success does not imply intensional, identity-governed meaning, and the difference decides the cases of Section 5.

3. Language, Symbols, and Four Thresholds

3.1 Language Beyond Words

In the sense that matters here, a language is a generative symbolic architecture: discrete, combinable units organized by rule into higher-order forms that admit interpretation. Natural language is one instance. Mathematics, music, and formal computation organize discrete units under combinatorial constraint in the same way. The view aligns with the tradition of Jerry Fodor (1935–2017) and Ray Jackendoff (b. 1945). Fodor and Zenon Pylyshyn (1937–2022) argue that cognition operates over symbolically structured expressions with combinatorial and compositional properties (Fodor and Pylyshyn 1988), and Jackendoff treats linguistic competence as parallel structures linked by interface rules (Jackendoff 2003, 662).

Ferdinand de Saussure (1857–1913) separated langue, the rule-governed system, from parole, its enactment in speech (Saussure 1959, 9–15). The thesis here reaches past both and past vocabulary. Words are symbols, but marks, gestures, diagrams, mathematical signs, musical notation, and internal identifiers count too, whenever a mind can detach them and combine them by rule. From usage-based and construction-grammar linguists, I borrow one idea only: meaning lives in structured, reusable units (Tomasello 2000; Goldberg 2003). I offer no theory of grammar or acquisition. My grammar is the rule set that governs emotional identifiers, not the syntax of English.

3.2 The Four Thresholds

Emotional meaning is not a property that appears wherever affect does. It is an architectural achievement, and a system reaches it only by crossing four thresholds (Cézjah 2026, chap. 10).

1. **Symbolic detachment** is the capacity to invoke an identifier when its trigger is absent: to say hawk under an empty sky.
2. **Symbolic endorsement** is acting under that identifier as operative, so that it organizes behavior and expectation as if its content holds.
3. **Grammar** binds identifiers into structure under rules of admissibility: hawk, danger, here. A heap of tokens carries nothing until rules relate them.
4. **Metarepresentation** is the capacity to judge one's own attribution as an attribution: "I thought I saw a hawk, but I was wrong." Only at this step does error become something a system can state and correct.

These thresholds sharpen the book's definitions of feeling and emotion. Feeling makes affect available to a subject, but it carries no symbol, no concept, and no grammar. Emotion organizes a feeling under a detachable identifier that a concept governs and the agent's stake anchors. A tightened jaw is a muscle contracting. It becomes anger only when an agent holds it under the identifier anger and that identifier's identity conditions.

4. Emotionese

4.1 The System

I call the symbolic system of emotional meaning Emotionese (Cézjah 2026, chap. 10). It is the system through which agents construct, stabilize, and communicate emotional meaning, and it is not reducible to speech. Its primitives are emotional identifiers such as fear, shame, pride, and betrayal; its grammar relates those identifiers to their objects, their contextual parameters, and the agent's stake. Faces, voices, postures, and bodily cues can enter a construct as properties, but their significance comes from their place in the architecture, not from their physical form. Emotionese is shared without being uniform: a community stabilizes prototypes that coordinate interpretation, and its members do not merely detect inner states but take part in a rule-governed system that organizes affect into structured significance.

4.2 Two Neighbors

The word emotionese already has a use. Developmental researchers apply it to the exaggerated emotional expressions adults direct at infants, by analogy with motherese (Ruba and Repacholi 2020, 256). That is a register of expression, and it can help teach Emotionese. It is not Emotionese, which is a symbolic system of identifiers, objects, contexts, and stake under rules.

The second neighbor is the language of thought. Contemporary defenders list six properties of such a language, among them discrete constituents, predicate–argument structure, logical operators, and abstract content, and they report that "complex infant and animal reasoning, and

automatic, intuitive cognition in adults all implicate LoT-like structures," where LoT abbreviates language of thought (Quilty-Dunn, Porot, and Mandelbaum 2023). I accept the combinatorial point; it is why I align the account with that tradition. Emotionese differs on four counts (Cézjah 2026, chap. 10). It belongs to one domain, emotional meaning, rather than serving as the medium of all thought. It is public: a community shares its rules, and its members correct one another's use. It reaches past speech to the faces, voices, and postures its constructs organize. And it is normative, because its identifiers answer to identity conditions an agent can misapply. The difference matters for the question of animals and infants. If both already reason in structures like a language of thought, such structure cannot mark the threshold for emotional meaning. Emotionese can, because a child acquires it from a community.

4.3 The Threshold Inside the Sign

The book's general architecture of meaning, the Constructivist Semiotic Sign of working paper 2, specifies what full emotional meaning needs (Cézjah 2026, chaps. 7, 10). Take a speaker before an evaluating audience: heart rate rises, muscles tighten, attention narrows, and images of failure arrive. Affect and feeling are present. Emotional meaning arrives only when the agent commits to the concept fear under its identity conditions, such as perceived threat, anticipated harm, and vulnerability relative to stake; marks that commitment with a detachable identifier; organizes the felt state into a construct such as threat to social standing; and does so relative to a referent-in-context, this audience and this reputation, as an interpreter for whom the outcome matters. Remove any constituent, and perception, sensation, and behavior can persist while meaning drops in kind.

Lisa Feldman Barrett illustrates the dependence of perception on concepts with a red apple, a falling tree, and a first hearing of Indonesian gamelan music (Barrett 2017a, 29, chap. 7, esp. 133). I accept the dependence and extend it. Concepts are necessary for construction, but they are not sufficient for full semiotic meaning. They must operate inside a detachable symbolic architecture governed by identity conditions and combinatorial constraint.

4.4 Detachment, Displacement, and Detached Representation

Detachment has neighbors too. Charles F. Hockett (1916–2000) counted displacement among the design features of language, defining a message as "displaced to the extent that the key features in its antecedents and consequences are removed from the time and place of transmission" (Hockett 1958, 579, quoted in Wacewicz and Żywiczyński 2015). Sławomir Wacewicz (b. 1978) and Przemysław Żywiczyński (b. 1971) reply that displacement matters as a capacity of minds rather than a feature of a code, and they take its clearest evidence from food-caching birds (Wacewicz and Żywiczyński 2015). Peter Gärdenfors (b. 1949) draws the nearer line: "Detached

representations stand for objects or events that are neither present in the situation nor triggered by some recent situation," and he ties them to planning (Gärdenfors 1995, 263).

I part from both on what detaches and what detachment buys. I grant the planning, including to scrub-jays, and still withhold the threshold, because what must detach is an identifier for the agent's own felt state, held under identity conditions that fix what counts as a correct application (Cézjah 2026, chap. 10).

5. The Emotional Meaning Architecture

5.1 Ten Levels

To make the question about animals and infants precise, I build a ladder of ten levels (Cézjah 2026, chap. 11, table 11-1). Level 0 is computational mapping: rule execution with nothing at stake. Level 1 is physiological activation. Level 2 is affective modulation: a bias of valence and arousal, with non-conscious evaluation. Level 3 is feeling, conscious awareness of an affective state. Level 4 is representational feeling, a felt state organized about an object or target.

Between Level 4 and Level 5 runs the symbolic threshold. Level 5 is detachable symbolic stabilization: a felt state organized under a reusable identifier governed by conceptual identity conditions, which yields proto-emotional meaning. Level 6 situates that construct relative to stake, alternatives, and consequences. Level 7 is metarepresentational endorsement, the representation of one's own emotional representation, with explicit endorsement, revision, negation, and justification; correctness conditions enter here, and with them full emotional meaning. Levels 8 and 9 extend that capacity into evaluation under reasons and into ongoing revision of one's own constructs.

The ladder assigns real status to everything below the line. Feeling at Level 3 is real feeling. Representational feeling at Level 4 is real organization of affect about the world. What the levels below the line lack is not reality but emotional meaning.

5.2 Animals

The hardest cases come from animal cognition, and each deserves its due (Cézjah 2026, chap. 11). Vervet monkeys give distinct alarm calls for distinct predators (Seyfarth, Cheney, and Marler 1980). Scrub-jays display episodic-like memory and plan for future needs (Clayton and Dickinson 1998; Raby et al. 2007). A bonobo acquired lexigrams without explicit training (Savage-Rumbaugh et al. 1986). A border collie learned the labels of more than two hundred items and inferred the name of a new item by exclusion (Kaminski, Call, and Fischer 2004). Corvids and apes track what others have seen (Emery and Clayton 2004; Call and Tomasello 2008), and scrub-jays re-cache food when another bird watched them hide it (Emery and Clayton 2001).

I take every result seriously, and each falls short of the line for a stated reason. A distinct call reflects stimulus-sensitive classification, not a symbol the caller can invoke when no leopard is near. Planning reflects flexible representation of future possibility, not an identifier for a felt state. The lexigrams were acquired inside a symbol system that people built, in daily exchange with people, not through species-typical use. Learning that a sound goes with an object does not mean that the dog holds its own inner state under a label. Tracking what a rival saw is a representation about another's perceptual state; it is not a representation of one's own representation. Even the dog's guilty look tracks the owner's scolding rather than the dog's own transgression (Horowitz 2009, 450).

The book puts the point in one line: "A parrot can say 'sorry' a thousand times without once apologizing" (Cézjah 2026, chap. 11). The bird has the word but no standard behind it, so it has no error it could catch. The boundary is architectural, not sentimental. On current evidence, animals reach Levels 1 to 4, and the book places no species-typical capacity above the line. Frans de Waal (1948–2024) argued that animal emotions are real and continuous with ours (de Waal 2011), and Jaak Panksepp (1943–2017) argued that core emotional systems are shared across mammals (Panksepp 2011). I deny neither claim about affect and feeling. Where they say emotion, I say affect or feeling, and the substantive disagreement is narrow: on current evidence, animals do not organize feeling under detachable symbols.

5.3 Infants

Infants tell a different story. They begin at the same four lower levels, with arousal, affective modulation, conscious feeling, and early non-symbolic organization of affect. Feeling precedes the symbol. Around two to three years, children begin to hold felt states under words such as sad and mad across situations and to apply them to themselves (Bretherton and Beeghly 1982). Vocabulary alone does not mark the crossing. A child can repeat an emotion word without organizing felt states under stable identity conditions. What marks Level 5 is cross-situational application, self-attribution, and consistent differentiation among categories. Around four to five years, with the development tracked in research on false belief (Wellman, Cross, and Watson 2001), a child reaches Level 7 and can say that what she called fear was really excitement. She now has an emotion she can correct.

5.4 Commitment Below Assertion

The normative force of the threshold comes from Robert Brandom (b. 1950), for whom a commitment answers to correction. For him, the sentence is "the minimal unit for which one can take responsibility, or which can express the undertaking of a commitment" (Brandom 2007, 671), and commitment lives in assertion and inference. I take the notion and change where it attaches (Cézjah 2026, chap. 11). At Level 5, an agent commits to a concept for its own felt state and marks

that commitment with an identifier. The commitment binds, and can fail, before the agent asserts anything about the state in full propositional form. A creature that cannot hold itself to a standard cannot misapply one.

5.5 What Would Defeat the Claim

A threshold that excludes every case by definition proves nothing. I therefore name three results that would count against the model (Cézjah 2026, chap. 11). First, self-correction without reward: a symbol-trained animal that revises its own earlier symbol for an inner state, unprompted and unrewarded, and carries the revision into new settings. Second, displaced self-report: an animal that uses a learned symbol for an inner state no longer present, where the use predicts its later choices better than present cues do. Third, developmental order: the model predicts that cross-situational self-attribution comes before reflective revision, so children who reliably revise emotion attributions before they apply emotion words to themselves would falsify the ordering. Evidence of feeling, learning, or signaling does not count against the model; it grants all three already.

6. Kinds of Reality and the Error of Misattribution

6.1 Five Kinds of Reality

If feelings are real, are emotions? They are, but they are real as constructions, and people mistake one kind of reality for another constantly (Cézjah 2026, chap. 12). Physiological and predictive reality belongs to the body, which anticipates and prepares before any awareness. Phenomenological reality is what the agent feels; about the fact of the feeling, the agent cannot err. Emotional reality is what an agent constructs by interpreting that feeling under a symbol, a context, and a goal. Social reality is what a community stabilizes through shared use and correction. Empirical reality is the first kind seen from outside, as observers measure it in heart rates, hormones, and expressions. Each is real. None can stand in for another.

6.2 Misattribution

The error that follows from collapsing them I call misattribution: treating what is real in experience, affect and feeling, as if it already carried emotional meaning. Suppose a person sleeps badly, drinks a great deal of coffee, and calls the trembling before a meeting dread. The trembling exists. The dread is a reading of it, and the wrong one, because caffeine and lost sleep produced the state (Cézjah 2026, chap. 12).

The classical form of the idea begins with arousal. Stanley Schachter (1922–1997) and Jerome E. Singer (1934–2010) proposed that physiological activation does not fix emotional meaning on its own; participants with elevated arousal labeled it differently depending on the situation (Schachter

and Singer 1962). Rainer Reisenzein (b. 1956), reviewing the evidence two decades later, judged that the data support only one of the theory's major deductions, that arousal misattributed to an extraneous source intensifies an emotional reaction (Reisenzein 1983). The core point survives the criticism. Even when affective processes have structure, they underdetermine interpretation (compare Russell 2003; Barrett 2017b). The agent does not detect emotion in arousal; the agent assigns meaning to a bodily state in context.

Misattribution has consequences well beyond the individual. A clinician who reads arousal and unease as anxiety can miss a thyroid problem, sleep loss, or a stimulant. An observer who reads a scowl as anger infers rather than detects. A machine that sorts faces and voices into emotion labels correlates patterns with names that people supplied. The book's summary is exact: an agent constructs an emotion, others infer it, and science correlates it (Cézjah 2026, chap. 12).

7. Two Levels: Prototype and Stereotype

7.1 The Architecture

Emotional meaning does not arise at a single level (Cézjah 2026, chap. 12). At the first level, an agent organizes a state under a detachable identifier governed by identity conditions and authorizes that construct as binding. At the second level, agents communicate such constructs, interpret them under shared rules, and apply them consistently across contexts, so that the construct stabilizes across a community. I call the first-person construct a prototype and its socially stabilized form a stereotype, both in a strict and neutral sense. A child learns anger as caregivers correct and model the word, until her private use joins a shared one. A private nickname for a Sunday mood means something to its inventor; it becomes a word only when others take it up and correct one another's use.

The structure is not peculiar to emotion. The book illustrates it with a hypothetical digital token: an agent can construct a token within a protocol, but the token becomes socially real only when other agents with the required capacities validate, accept, and use it under shared rules. A single person cannot create a currency by belief alone.

7.2 Earlier Owners of the Words

Both words have earlier owners, and I depart from each (Cézjah 2026, chap. 12). In the prototype theory of Eleanor Rosch (b. 1938), a prototype is the clearest case of a category, and membership comes in degrees of resemblance (Rosch and Mervis 1975). Emotion research took up that view: types of emotion "can be reliably ordered from better to poorer examples of emotion" (Fehr and Russell 1984, 464), and emotion categories come to be organized around prototypes through repeated experience (Shaver et al. 1987, 1061). My prototype is not a clearest case, and it grades nothing. It is one agent's authorized construct; typicality effects arise later, in categorization. In

social psychology, Susan T. Fiske (b. 1952) describes stereotypes as "shared beliefs about common groups' warmth and competence," which in her model give rise to emotional prejudices (Fiske 2018, sec. "Stereotype Content Model"). My stereotype describes no group of people. It is the shared form of an emotion construct, and it supplies the standard by which others can correct each agent's use.

7.3 Agreement and Correctness

The two levels raise an obvious worry: if a community fixes the standard for an emotion word, does agreement decide correctness after all? It does not. Agreement institutes the standard; once the standard is fixed, each application answers to it, not to agreement (Cézjah 2026, chaps. 6, 12). The book uses objective in three senses and keeps them apart. Correctness against an instituted standard is objective, because an application is right or wrong whatever anyone believes. Third-person access to emotion is not available, because observers infer. A mind-independent essence shared by every instance is what the account denies.

7.4 Five Tests of Emotional Essence

The book puts the doctrine of emotional essences to five executable tests (Cézjah 2026, chap. 12). No intrinsic feature of a sign fixes its identifier, a point that extends Saussure's arbitrariness of the sign (Saussure 1959, 67–70). One identifier does not fix one configuration. No single invariant set of identity conditions governs every admissible instance. The same conditions applied to the same referent do not yield one invariant identity across settings. The same sign does not carry the same motivational force for every interpreter. For emotion, the verdict holds at every point: no essence, no invariant identity, no intrinsic meaning, and no universal significance lives inside any signal.

8. What Is Shared and What Is New

The book states the shared ground and the addition for the Constructivist Linguistics of Emotion as a whole (Cézjah 2026, chap. 13). What is shared: words and concepts shape emotional life, and an emotional category holds instances that share no single essence. What is new: one account that runs from a symbolic threshold, through two levels of stabilization, to emotional categories and the grammar that organizes them. This paper covers the threshold and the two levels; working paper 8 covers the categories and their grammar.

Against each nearest neighbor, the book states a narrower claim. Linguistic relativity asks whether a language shapes thought (Lucy 1997); Anna Wierzbicka (b. 1938) argues that emotion words do not translate cleanly and proposes simple universal concepts for comparing them (Wierzbicka 2009); and a study of 2,474 languages reports both variation and shared structure in emotion vocabularies (Jackson et al. 2019). I accept all three results. My claim is not that English shapes

English feelings, but that emotional meaning, in any language, requires a detachable symbol under identity conditions.

The nearest neighbor is experimental and theoretical work reporting that emotion words help people perceive emotion and arguing that language is constitutive of emotion experience and perception (Lindquist et al. 2006; Lindquist, MacCormack, and Shablack 2015). I part from it at two points. For them, language supports the concept knowledge that makes sensations meaningful. For me, the constitutive element is a detachable symbol under identity conditions, in any modality, and it marks a threshold: without it, there is feeling but no emotional meaning. What is new, then, is not that words matter. It is a threshold model, a stated point at which feeling becomes emotional meaning, with a test for whether a creature has crossed it (Cézjah 2026, chap. 10).

The book adds four narrower statements of the same form, which I gather here (Cézjah 2026, chaps. 10–12):

- **Against the language of thought.** Shared: meaning needs discrete symbols combined by rule. New: a learned, public, normative symbol system for a single domain, whose detachable identifiers mark where feeling becomes emotional meaning.
- **Against displacement and detached representation.** Shared: a symbol stays available when its trigger is absent. New: detachment as a property of an identifier under identity conditions, applied to the agent's own feeling and open to correction.
- **Against Brandom.** Shared: a commitment answers to correction, and meaning depends on it. New: commitment to a concept for one's own feeling, below the level of assertion, as the act that turns feeling into emotional meaning.
- **Against prototype theory and stereotype research.** Shared: emotion categories have central cases, and communities hold shared beliefs. New: the prototype and the stereotype as the first-person and social levels of one construct, so that agreement fixes the standard while each application still answers to it.

9. Objections and Replies

Objection 1: A threshold adds nothing to the claim that language is constitutive of emotion. A reader who follows that work on language and emotion will press this point (Cézjah 2026, chap. 13). The reply is that the constitutive element is not a word but a detachable symbol in any modality, held under identity conditions, and that the threshold makes the claim testable: Section 5.5 names three results that would defeat it.

Objection 2: Animals plainly categorize, so they already have concepts. Animals discriminate predators from non-predators, generalize across contexts, and track rank and history. But category learning that runs on statistical regularity and feature clustering differs from identity conditions

that symbols govern (Cézjah 2026, chap. 11). A symbolic type fixes what counts as a legitimate instance, can be invoked without the stimulus, and is subject to endorsement and revision. A predator category that guides avoidance across contexts, without such an identifier, remains a regulatory construct rather than a semiotic one. The distinction rests on architectural form, not on behavioral complexity.

Objection 3: The criteria exclude animals by stipulation. If the threshold were defined so that no evidence could meet it, the conclusion would be empty. That is why the book states in advance what would count against it (Section 5.5), and why it treats the animal case as open to evidence rather than settled by definition (Cézjah 2026, chap. 11).

Objection 4: The disagreement with Panksepp and de Waal is only verbal. Partly it is. Where they say emotion, I say affect or feeling, and I grant everything they claim at those levels. The substantive remainder is narrow and testable: whether animals organize feeling under detachable symbols (Cézjah 2026, chap. 10).

Objection 5: If a community fixes the standard, correctness reduces to majority use. Section 7.3 gives the reply. Agreement institutes a standard; it does not decide each application. A community can misapply its own standard in bulk, and the standard it set convicts it (Cézjah 2026, chaps. 6, 12, 13).

10. Scope and Limits

The account makes architectural claims about what emotional meaning requires, and empirical claims only where it says so: about where current evidence places animals and children on the ladder. The age ranges for infants are approximate windows, not fixed boundaries, and the architecture, not the timing, carries the argument. The five tests of emotional essence are executable checks of internal consistency, not experiments.

The paper also stops short of two topics the book treats next. It does not explain how interpreted instances gather into emotional categories, how a category stays stable without an essence, or what qualitative dimension of feeling lies beyond hedonic tone. Working paper 8 takes up those questions through music, genre, and phenomenological timbre (Cézjah 2026, chap. 13). Nor does the paper address the role of emotion in moral action, which the book treats briefly at the close of Chapter 12: emotion serves judgment and does not replace it.

11. Conclusion

Feeling is real everywhere one looks. Emotion arrives only with symbols. I defined meaning as a structured representation interpreted in light of stake, named four thresholds a system must cross to reach it, and called the symbolic system of emotional meaning Emotionese. The Emotional

Meaning Architecture places the line between Level 4 and Level 5; on current evidence, animals feel and stop short of it, infants feel and cross it, and three named results would defeat the claim. I then separated the kinds of reality that people collapse, named the resulting error misattribution, and distinguished the prototype, one agent's authorized construct, from the stereotype, its socially stabilized form, with agreement fixing the standard and each application answering to it.

The practical stake is plain. Every distinction a person can name, hold, and revise widens what that person's feelings can mean, and every distinction missing narrows it. The words adults give a child become the material of her emotional life. The full argument appears in Chapters 10 to 12 of *How to Design an Emotion* (Cézjah 2026, chaps. 10–12), and the next working paper continues it into emotional categories.

Appendix: Code, UML Diagrams, and Input/Output

The figures below are the book's own diagrams in the Unified Modeling Language (UML), reproduced unchanged. The listings are excerpts from the book's C# companion projects, and the input/output blocks were captured by running those projects with the .NET 8 software development kit; I keep only the lines that bear on the argument. The full programs and the complete transcripts sit beside this paper in `code-samples/`. As in the book, the code checks the consistency of the architecture, not the truth of the theory. The diagrams, data, and program come from Chapter 12 (Cézjah 2026, chap. 12). The irreducibility tests of Appendix N appear in working paper 2 and are not repeated here.

A.1 Test 3: No Universal Essence

The test keeps the identifier and the context and varies the identity conditions, which also change the admissible members (Section 7.4). The input is the book's data file; the verdict is computed from it.

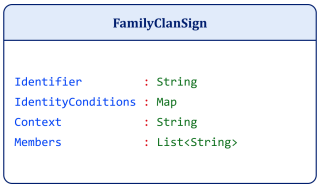


Figure 1. UML class diagram of FamilyClanSign (book Figure 12-9).

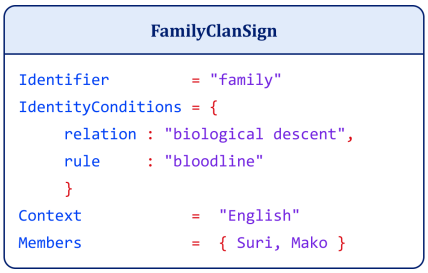


Figure 2. UML object diagram: family under the identity conditions of biological descent (book Figure 12-10).

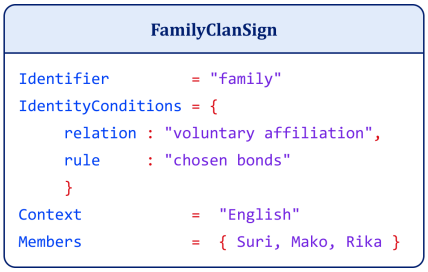


Figure 3. UML object diagram: family under the identity conditions of voluntary affiliation (book Figure 12-11).

Listing 1. Input: the essence dataset (code/chapter-12/SemioticConstraintEngine_v2/data.json**)**

```
{
  "essence": {
    "family": {
      "context": "English",
      "instances": [
        {
          "title": "Biological Essence",
          "identityConditions": {
            "relation": "biological descent",
            "rule": "bloodline"
          },
          "members": [
            "Suri",
            "Mako"
          ]
        },
        {
          "title": "Chosen Essence",
          "identityConditions": {
            "relation": "voluntary affiliation",
            "rule": "chosen bonds"
          },
          "members": [
            "Suri",
            "Mako",
            "Rika"
          ]
        },
        {
          "title": "Institutional Essence",
          "identityConditions": {
            "relation": "co-residence",
            "rule": "legal or domestic unit"
          },
          "members": [
            "Suri",
            "Rika",
            "Toma"
          ]
        }
      ]
    }
  }
}
```

Listing 2. Input/output: Test 3, essence (code/chapter-12/SemioticConstraintEngine_v2**; captured with .NET 8)**

```
$ dotnet run -- --test essence --identifier family
TEST 3: ABSENCE OF UNIVERSAL OBJECTIVE ESSENCE

Input identifier: family

Admissible instances by identity conditions:
  1. Biological Essence
     IdentityConditions = { relation: biological descent, rule: bloodline }
     Members = { Suri, Mako }

  2. Chosen Essence
     IdentityConditions = { relation: voluntary affiliation, rule: chosen bonds }
     Members = { Suri, Mako, Rika }

  3. Institutional Essence
     IdentityConditions = { relation: co-residence, rule: legal or domestic unit }
     Members = { Suri, Rika, Toma }

Invariant essence found: NO

Conclusion:
  No single set of identity conditions governs all admissible instances.
  There are no universal objective essences.
```

A.2 Test 5: No Universal Significance

The test keeps the event sign constant and varies the interpreter (Section 7.4).

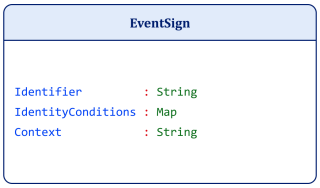


Figure 4. UML class diagram of EventSign (book Figure 12-17).

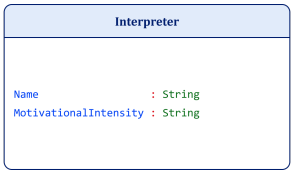


Figure 5. UML class diagram of Interpreter (book Figure 12-19).

Listing 3. Input: the significance dataset (`code/chapter-12/SemioticConstraintEngine_v2/data.json`**)**

```
{
  "significance": {
    "jobInterview": {
      "identifier": "job interview",
      "identityConditions": {
        "type": "employment evaluation",
        "rule": "formal candidate assessment"
      },
      "context": "professional setting",
      "interpreters": [
        {
          "name": "Candidate A",
          "motivationalIntensity": "low",
          "reason": "Already employed; outcome has minimal impact."
        },
        {
          "name": "Candidate B",
          "motivationalIntensity": "moderate",
          "reason": "Seeks career progression; outcome is beneficial but not critical."
        },
        {
          "name": "Candidate C",
          "motivationalIntensity": "high",
          "reason": "Unemployed; outcome directly affects livelihood."
        }
      ]
    }
  }
}
```

Listing 4. Input/output: Test 5, significance (`code/chapter-12/SemioticConstraintEngine_v2`**; captured with .NET 8)**

```
$ dotnet run -- --test significance --identifier jobInterview
TEST 5: ABSENCE OF UNIVERSAL OBJECTIVE SIGNIFICANCE

Constant sign: job interview
  IdentityConditions = { type: employment evaluation, rule: formal candidate assessment
}
  Context = professional setting

Interpreters:
  Candidate A -> low
    Reason: Already employed; outcome has minimal impact.
  Candidate B -> moderate
    Reason: Seeks career progression; outcome is beneficial but not critical.
  Candidate C -> high
    Reason: Unemployed; outcome directly affects livelihood.

Conclusion:
  The sign does not carry a single invariant significance for every interpreter.
  There is no universal objective significance.
```

References

Barrett, Lisa Feldman. 2017a. *How Emotions Are Made: The Secret Life of the Brain*. Boston: Houghton Mifflin Harcourt.

Barrett, Lisa Feldman. 2017b. “The Theory of Constructed Emotion: An Active Inference Account of Interoception and Categorization.” *Social Cognitive and Affective Neuroscience* 12 (1): 1–23. https://doi.org/10.1093/scan/nsw154.

Brandom, Robert. 2007. “Inferentialism and Some of Its Challenges.” *Philosophy and Phenomenological Research* 74 (3): 651–76. https://doi.org/10.1111/j.1933-1592.2007.00044.x.

Bretherton, Inge, and Marjorie Beeghly. 1982. “Talking about Internal States: The Acquisition of an Explicit Theory of Mind.” *Developmental Psychology* 18 (6): 906–21. https://doi.org/10.1037/0012-1649.18.6.906.

Call, Josep, and Michael Tomasello. 2008. “Does the Chimpanzee Have a Theory of Mind? 30 Years Later.” *Trends in Cognitive Sciences* 12 (5): 187–92. https://doi.org/10.1016/j.tics.2008.02.010.

Cézjah. 2026. *How to Design an Emotion: The Deconstruction and Computation of Emotional Meaning from First Principles*. Universal Moral Concepts III. Miami: New Intellectual LLC. International Standard Book Number (ISBN) 979-8-9898135-9-9 (paperback), 979-8-9898135-3-7 (Kindle). https://doi.org/10.5281/zenodo.23630995.

Clayton, Nicola S., and Anthony Dickinson. 1998. “Episodic-Like Memory during Cache Recovery by Scrub Jays.” *Nature* 395 (6699): 272–74. https://doi.org/10.1038/26216.

Emery, Nathan J., and Nicola S. Clayton. 2001. “Effects of Experience and Social Context on Prospective Caching Strategies by Scrub Jays.” *Nature* 414 (6862): 443–46. https://doi.org/10.1038/35106560.

Emery, Nathan J., and Nicola S. Clayton. 2004. “The Mentality of Crows: Convergent Evolution of Intelligence in Corvids and Apes.” *Science* 306 (5703): 1903–7. https://doi.org/10.1126/science.1098410.

Fehr, Beverley, and James A. Russell. 1984. “Concept of Emotion Viewed from a Prototype Perspective.” *Journal of Experimental Psychology: General* 113 (3): 464–86. https://doi.org/10.1037/0096-3445.113.3.464.

Fiske, Susan T. 2018. “Stereotype Content: Warmth and Competence Endure.” *Current Directions in Psychological Science* 27 (2): 67–73. https://doi.org/10.1177/0963721417738825.

Fodor, Jerry A., and Zenon W. Pylyshyn. 1988. “Connectionism and Cognitive Architecture: A Critical Analysis.” *Cognition* 28 (1–2): 3–71. https://doi.org/10.1016/0010-0277(88)90031-5.

Goldberg, Adele E. 2003. “Constructions: A New Theoretical Approach to Language.” *Trends in Cognitive Sciences* 7 (5): 219–24. https://doi.org/10.1016/S1364-6613(03)00080-9.

Gärdenfors, Peter. 1995. “Cued and Detached Representations in Animal Cognition.” *Behavioural Processes* 35 (1–3): 263–73. https://doi.org/10.1016/0376-6357(95)00043-7.

Hockett, Charles F. 1958. *A Course in Modern Linguistics*. New York: Macmillan.

Horowitz, Alexandra. 2009. “Disambiguating the ‘Guilty Look’: Salient Prompts to a Familiar Dog Behaviour.” *Behavioural Processes* 81 (3): 447–52. https://doi.org/10.1016/j.beproc.2009.03.014.

Jackendoff, Ray. 2003. “Précis of *Foundations of Language: Brain, Meaning, Grammar, Evolution*.” *Behavioral and Brain Sciences* 26 (6): 651–65. https://doi.org/10.1017/S0140525X03000153.

Jackson, Joshua Conrad, Joseph Watts, Teague R. Henry, Johann-Mattis List, Robert Forkel, Peter J. Mucha, Simon J. Greenhill, Russell D. Gray, and Kristen A. Lindquist. 2019. “Emotion Semantics Show Both Cultural Variation and Universal Structure.” *Science* 366 (6472): 1517–22. https://doi.org/10.1126/science.aaw8160.

Kaminski, Juliane, Josep Call, and Julia Fischer. 2004. “Word Learning in a Domestic Dog: Evidence for ‘Fast Mapping.’” *Science* 304 (5677): 1682–83. https://doi.org/10.1126/science.1097859.

Lindquist, Kristen A., Jennifer K. MacCormack, and Holly Shablack. 2015. “The Role of Language in Emotion: Predictions from Psychological Constructionism.” *Frontiers in Psychology* 6: 444. https://doi.org/10.3389/fpsyg.2015.00444.

Lindquist, Kristen A., Lisa Feldman Barrett, Eliza Bliss-Moreau, and James A. Russell. 2006. “Language and the Perception of Emotion.” *Emotion* 6 (1): 125–38. https://doi.org/10.1037/1528-3542.6.1.125.

Lucy, John A. 1997. “Linguistic Relativity.” *Annual Review of Anthropology* 26: 291–312. https://doi.org/10.1146/annurev.anthro.26.1.291.

Nagel, Thomas. 1974. “What Is It Like to Be a Bat?” *Philosophical Review* 83 (4): 435–50. https://doi.org/10.2307/2183914.

Panksepp, Jaak. 2011. “Cross-Species Affective Neuroscience Decoding of the Primal Affective Experiences of Humans and Related Animals.” *PLOS ONE* 6 (9): e21236. https://doi.org/10.1371/journal.pone.0021236.

Quilty-Dunn, Jake, Nicolas Porot, and Eric Mandelbaum. 2023. “The Best Game in Town: The Reemergence of the Language-of-Thought Hypothesis across the Cognitive Sciences.” *Behavioral and Brain Sciences* 46: e261. https://doi.org/10.1017/S0140525X22002849.

Raby, C. R., D. M. Alexis, A. Dickinson, and N. S. Clayton. 2007. “Planning for the Future by Western Scrub-Jays.” *Nature* 445 (7130): 919–21. https://doi.org/10.1038/nature05575.

Reisenzein, Rainer. 1983. “The Schachter Theory of Emotion: Two Decades Later.” *Psychological Bulletin* 94 (2): 239–64. https://doi.org/10.1037/0033-2909.94.2.239.

Rosch, Eleanor, and Carolyn B. Mervis. 1975. “Family Resemblances: Studies in the Internal Structure of Categories.” *Cognitive Psychology* 7 (4): 573–605. https://doi.org/10.1016/0010-0285(75)90024-9.

Ruba, Ashley L., and Betty M. Repacholi. 2020. “Beyond Language in Infant Emotion Concept Development.” *Emotion Review* 12 (4): 255–58. https://doi.org/10.1177/1754073920931574.

Russell, James A. 2003. “Core Affect and the Psychological Construction of Emotion.” *Psychological Review* 110 (1): 145–72. https://doi.org/10.1037/0033-295X.110.1.145.

Saussure, Ferdinand de. 1959. *Course in General Linguistics*. Edited by Charles Bally and Albert Sechehaye. Translated by Wade Baskin. New York: Philosophical Library.

Savage-Rumbaugh, Sue, Kelly McDonald, Rose A. Sevcik, William D. Hopkins, and Elizabeth Rubert. 1986. “Spontaneous Symbol Acquisition and Communicative Use by Pygmy Chimpanzees (*Pan paniscus*).” *Journal of Experimental Psychology: General* 115 (3): 211–35. https://doi.org/10.1037/0096-3445.115.3.211.

Schachter, Stanley, and Jerome E. Singer. 1962. “Cognitive, Social, and Physiological Determinants of Emotional State.” *Psychological Review* 69 (5): 379–99. https://doi.org/10.1037/h0046234.

Seyfarth, Robert M., Dorothy L. Cheney, and Peter Marler. 1980. “Monkey Responses to Three Different Alarm Calls: Evidence of Predator Classification and Semantic Communication.” *Science* 210 (4471): 801–3. https://doi.org/10.1126/science.7433999.

Shaver, Phillip, Judith Schwartz, Donald Kirson, and Cary O’Connor. 1987. “Emotion Knowledge: Further Exploration of a Prototype Approach.” *Journal of Personality and Social Psychology* 52 (6): 1061–86. https://doi.org/10.1037/0022-3514.52.6.1061.

Tomasello, Michael. 2000. “First Steps toward a Usage-Based Theory of Language Acquisition.” *Cognitive Linguistics* 11 (1–2): 61–82. https://doi.org/10.1515/cogl.2001.012.

Waal, Frans B. M. de. 2011. “What Is an Animal Emotion?” *Annals of the New York Academy of Sciences* 1224 (1): 191–206. https://doi.org/10.1111/j.1749-6632.2010.05912.x.

Wacewicz, Sławomir, and Przemysław Żywiczyński. 2015. “Language Evolution: Why Hockett’s Design Features Are a Non-Starter.” *Biosemiotics* 8 (1): 29–46. https://doi.org/10.1007/s12304-014-9203-2.

Wellman, Henry M., David Cross, and Julanne Watson. 2001. “Meta-Analysis of Theory-of-Mind Development: The Truth about False Belief.” *Child Development* 72 (3): 655–84. https://doi.org/10.1111/1467-8624.00304.

Wierzbicka, Anna. 2009. “Language and Metalanguage: Key Issues in Emotion Research.” *Emotion Review* 1 (1): 3–14. https://doi.org/10.1177/1754073908097175.

Wittgenstein, Ludwig. 1922. *Tractatus Logico-Philosophicus*. Translated by C. K. Ogden. London: Routledge and Kegan Paul.

How to cite

The Constructivist Linguistics of Emotion: The Symbolic Threshold, Emotionese, and the Two Levels of Emotional Meaning

APA

Cézjah. (2026). *The Constructivist Linguistics of Emotion: The symbolic threshold, Emotionese, and the two levels of emotional meaning* (Working paper 7). New Intellectual LLC. https://newintellectual.com/pages/how-to-design-an-emotion-wp-07 (Crossref DOI pending)

CONTACT THE AUTHOR

Cézjah (Cecil John)
cezjah@cezjah.com
ORCID 0009-0005-1866-9700
newintellectual.com

ALSO AVAILABLE AS

https://doi.org/10.5281/zenodo.23676664 (Zenodo, version 2; all versions: https://doi.org/10.5281/zenodo.23631061)

LINKS

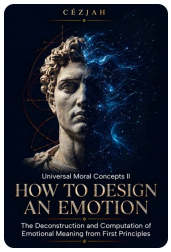
Landing page: newintellectual.com/pages/how-to-design-an-emotion-wp-07

Product page: newintellectual.com/products/working-paper-07-constructivist-linguistics

License: CC BY 4.0 (this paper only; the book is All rights reserved)

KEYWORDS

Constructivist Linguistics of Emotion · Emotionese · Emotional Meaning Architecture · symbolic detachment · symbolic threshold · metarepresentation · commitment · misattribution · prototype and stereotype · language of thought · animal emotions · infant emotional development · Robert Brandom



This paper restates, in new prose, the argument of Chapters 10 to 12 of *How to Design an Emotion* (Cézjah 2026, chaps. 10–12). The book is the full treatment.

CC BY 4.0