

JOURNAL > SITE 1: LOGIC GATE, THE POLITICS OF THE ARTIFACTUAL  
MIND | 2017

## A Theory of Vibe

Peli Grietzer

*Across the foliated space of the twenty-seven equivalents, Faustroll conjured up into the third dimension: From Baudelaire, E. A. Poe's Silence, taking care to retranslate Baudelaire's translation into Greek. From Bergerac, the precious tree into which the nightingale king and his subjects were metamorphosed, in the land of the sun. From Luke, the Calumniator who carried Christ on to a high place. From Bloy, the black pigs of Death, retinue of the Betrothed. From Coleridge, the ancient mariner's crossbow, and the ship's floating skeleton, which, when placed in the skiff, was sieve upon sieve.*

—Alfred Jarry, *Exploits & opinions of Doctor Faustroll, pataphysician: a neo-scientific novel*, 1929

I. An *autoencoder*<sup>1</sup> is a neural network process tasked with learning from scratch, through a kind of trial and error, how to make facsimiles of worldly things. Let us call a hypothetical, exemplary autoencoder 'Hal.' We call the set of all the inputs we give Hal for reconstruction— let us say many, many image files of human faces, or many, many audio files of jungle sounds, or many, many scans of city maps—Hal's 'training set.' Whenever Hal receives an input media file  $x$ , Hal's *feature function* outputs a short list of short numbers, and Hal's *decoder function* tries to recreate media file  $x$  based on the feature function's 'summary' of  $x$ . Of course, since the variety of possible media files is much wider than the variety of possible short lists of short numbers, something must necessarily get lost in the translation from media file to feature values and back: many possible media files translate into the same short list of short numbers, and yet each

short list of short numbers can only translate back into one media file. Trying to minimize the damage, though, induces Hal to learn—through trial and error—an effective schema or ‘mental vocabulary’ for its training set, exploiting rich holistic patterns in the data in its summary-and-reconstruction process. Hal’s ‘summaries’ become, in effect, cognitive mapping of its training set, a kind of gestalt fluency that ambiently models it like a niche or a lifeworld.

2. What an autoencoder algorithm learns, instead of making perfect reconstructions, is a system of features that can generate *approximate* reconstruction of the objects of the training set. In fact, the difference between an object in the training set and its reconstruction—mathematically, the trained autoencoder’s *reconstruction error* on the object—demonstrates what we might think of, rather literally, as the excess of material reality over the gestalt-systemic logic of autoencoding. We will call the set of all possible inputs for which a given trained autoencoder *S* has zero reconstruction error, in this spirit, *S*’s ‘*canon*.’ The canon, then, is the set of all the objects that a given trained autoencoder—its imaginative powers bounded as they are to the span of just a handful of ‘*respects of variation*,’ the dimensions of the features vector—can imagine or conceive of whole, without approximation or simplification. Furthermore, if the autoencoder’s training was successful, the objects in the canon collectively exemplify an idealization or simplification of the objects of some worldly domain. Finally, and most strikingly, a trained autoencoder and its canon are effectively mathematically equivalent: not only are they roughly logically equivalent, it is also fast and easy to compute one from the other. In fact, merely autoencoding a small sample from the canon of trained autoencoder *S* is enough to accurately replicate or model *S*.

3. Imagine if you will the “hermeneutics of suspicion”<sup>2</sup>—the classical ‘90s kind of symptomatic or subversive academic reading—was a data-mining process that infers, from what is found and not found in the world constructed by a literary text, an organon (system of thought and feeling) that makes certain real-world phenomena unthinkable, invisible, foreclosed to the order of things. The critic would infer, from observation of the literary work’s selection of phenomena, a generative model of the work, finding what is repressed or marginalized in the text within ‘gaps’ in the generative model: states of the lifeworld that the generative model cannot generate. Pushing the process even further, an ambitious critic would go on to try to characterize *dimensions*—ways in which states of the world can be meaningfully different from each other—missing from the generative model. Contemporary cultural-materialist or ideology-sensitive readings are, as Rita Felksi argues in “After Suspicion,”<sup>3</sup> for the most part “post-suspicion”: recent

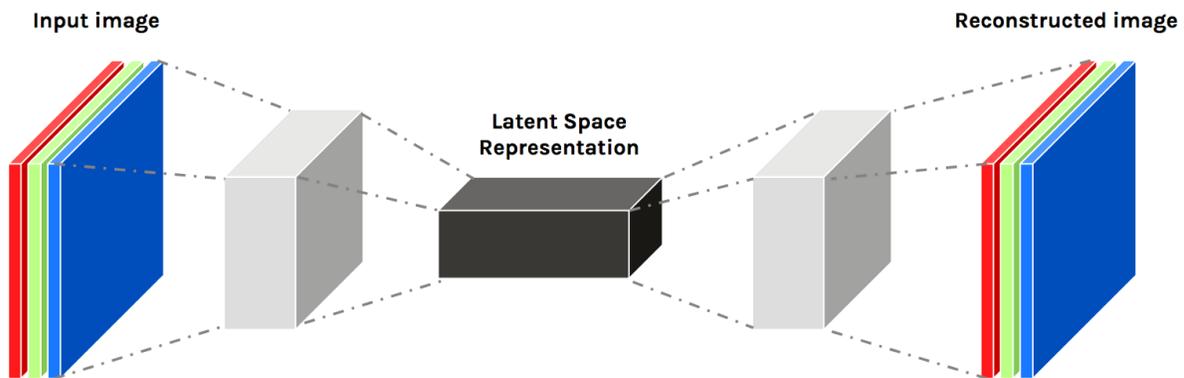
social-theoretic literary critics, especially those associated with the field of affect-studies, tend to differ from their predecessors in assigning reflexivity and agency to literary texts as the facilitators of the critical comparison between model and world. This modern turn places the framework of some recent social-theoretic readers—in particular, Jonathan Flatley and Sianne Ngai<sup>4</sup>—in a close alliance with our own. Specifically, Ngai’s landmark argument in *Ugly Feelings* that a work of literature can, through tone, represent a subject’s ideology—and so, both represent a structure of her subjectivity and touch upon the structure of the social-material conditions structuring her subjectivity—as strongly concordant with the proposition that systems of ‘respects of variation’ that we might define by the excess material reality that they marginalize (that is, defined as ‘ideology’) can be identically defined through the aesthetic unity of material realities they access best (that is, defined as ‘tone’). The *canon* of a trained autoencoder, we are proposing, recapitulates the ideology of a system of ‘respects of variation’ as a tone.

4. Autoencoders, we know, deal entirely in worlds rendered as sets of objects or phenomena. Whatever deeper worldly structures an autoencoder’s schema brings to the interpretation of an object, then, these structures are already at play, in some form, in the collective aesthetic of the objects they reign over.<sup>5</sup> I want to think about this aesthetically accessible, surface-accessible, world-making structure as the mathematical substrate of what writer/musician Ezra Koenig (via Elif Batuman) describes as “vibe”:

*It was during my research on the workings of charm and pop music that I stumbled on Internet Vibes (internetvibes.blogspot.com/), a blog that Ezra Koenig kept in 2005–6, with the goal of categorising as many “vibes” as possible. A “rain/grey/British vibe,” for example, incorporates the walk from a Barbour store (to look at wellington boots) to the Whitney Museum (to look at “some avant-garde shorts by Robert Beavers”), as well as the TV adaptation of Brideshead Revisited, the Scottish electronic duo Boards of Canada, “late 90s Radiohead/global anxiety/airports” and New Jersey. A “vibe” turns out to be something like “local colour,” with a historical dimension. What gives a vibe “authenticity” is its ability to evoke—using a small number of disparate elements—a certain time, place and milieu; a certain nexus of historic, geographic and cultural forces.<sup>6</sup>*

The meaning of a literary work like Dante’s “Inferno,” Beckett’s “Waiting for Godot,” or Stein’s *Tender Buttons*, we would like to say, lies at least partly in an aesthetic ‘vibe’ or a ‘style’ that we can sense when we consider all the myriad objects and phenomena that make up the imaginative landscape of the work as a kind of curated set. The meaning of Dante’s “Inferno,” let us say, lies in part in that certain *je ne sais quoi* that makes every

soul, demon, and machine in Dante’s vision of hell a good fit for Dante’s vision of hell. Similarly, the meaning of Beckett’s “Waiting for Godot” lies partly in what limits our space of thinkable things for Vladimir and Estragon to say and do to a small set of possibilities the play nearly exhausts. Part of the meaning of Stein’s *Tender Buttons* lies in the set of (possibly inherently linguistic) ‘tender buttons’—conforming objects and phenomena.<sup>7</sup>



Map of a trained autoencoder. All rights reserved.

5. The features or dimensions or ‘respects of variation’ of a trained autoencoder work very much like a fixed list of predicates with room to write-in for example ‘not’ or ‘somewhat’ or ‘solidly’ or ‘extremely’ next to each.<sup>8</sup> Within the context of the *feature function*, which produces ‘summaries’ of the input object, it is most natural to think of the ‘respects of variation’ as *descriptive* predicates. The features of a trained autoencoder take a rather different meaning if instead we center our thinking around the *decoder function*—the function that turns ‘summaries’ into reconstructions. From the viewpoint of the decoder function, a given list of feature-values is not a ‘summary’ that could apply to any number of closely related objects, but rather the (so to speak) DNA of a specific object. A given trained autoencoder’s features or ‘respects of variation’ are, from this

perspective, akin to a list of *imperative* predicates, structural techniques or principles to be applied by the constructor. For the decoder, the ‘generative formulae’ for objects in a trained autoencoder’s canon are lists of activation values that determine how intensely the construction process (the *decoder function*) applies each of the available structural techniques or principles.

6. It is a fundamental property of any trained autoencoder’s canon therefore that all the objects in the canon *align with a limited generative vocabulary*. The objects that make up the trained autoencoder’s actual worldly domain, by implication, roughly align or *approximately align* with that same limited generative vocabulary. These structural relations of alignment, I propose, are closely tied, and may have a strong relationship to certain concepts of aesthetic unity that commonly imply a unity of generative logic, as in both the intuitive and literary theoretic concepts of a ‘style’ or ‘vibe.’ To be a set that aligns with *some* logically possible generative vocabulary is hardly a ‘real’ structural or aesthetic property, given the infinity of logically possible generative vocabularies. To be a set that aligns with some (logically possible) *limited* generative vocabulary, on the other hand, is a robust intersubjective property.

7. By way of a powerful paraphrase, we might say that it means the objects that make up a trained autoencoder’s canon are *individually complex but collectively simple*. To better illustrate this concept (‘individually complex but collectively simple’), let us make a brief digression and describe a type of mathematical-visual art project, typically associated with late 20th century Hacker culture, known as a ‘64k Intro.’ In the artistic-mathematical subculture known as ‘demoscene,’ a ‘64k Intro’ is a lush, vast, and nuanced visual world that fits into 64 kilobytes of memory or fewer, less memory by a thousandfold than the standard memory requirements for a lush, robust, and nuanced visual world. In a 64k Intro, a hundred or so lines of code create a sensually complicated universe by, quite literally, using the esoteric affinities of surfaces with primordial Ideas. The code of a 64k Intro uses the smallest possible inventory of initial schemata to generate the most diverse concreta. The information-theoretical magic behind a 64k Intro is that, somewhat like a spatial fugue, these worlds are tapestries of interrelated self-similar patterns. From the topological level (architecture and camera movement) to the molecular level (the polygons and textures from which objects are built), everything in a 64k Intro is born of a ‘family resemblance’ of forms.

8. Remarkably—and also, perhaps, trivially—the relationship between succinct expressibility and depth of pattern that we see in 64k Intros provably holds for any informational, cognitive, or semiotic system. A deeply conceptually useful, though often

technically unwieldy, measure of ‘depth of pattern’ used in information theory is ‘Kolmogorov complexity’: the Kolmogorov complexity of an object is the length of the shortest possible description (in a given semiotic system) that can fully specify it.<sup>9</sup> Lower Kolmogorov complexity generically means stronger pattern. A low Kolmogorov complexity—i.e. short minimum description length—for an object relative to a given semiotic system implies the existence of deep patterns in the object, or a close relationship between the object and the basic concepts of the semiotic system.

9. When all the objects in a given set  $C$  have low Kolmogorov+ complexity relative to a given semiotic system  $S$ , we will say the semiotic system  $S$  is a *schema* for  $C$ . If  $S$  is a given trained autoencoder’s generative language (formally, *decoder function*), and  $C$  the canon of this trained autoencoder  $C$ , for example, then  $S$  is a schema for  $C$ . Importantly, any schema  $S$  is in itself a semiotic object, and itself has a Kolmogorov complexity relative to our own present semiotic system, and so the ‘real’—that is, relative to our own semiotic system—efficacy of  $S$  as a schema for an object  $c$  in  $C$  is measured by the sum of the Kolmogorov+ complexity of  $c$  relative to  $S$  and the Kolmogorov complexity of  $S$ . Because one only needs to learn a language once to use it to create however many sets of sentences one wishes, though, when we consider the efficacy of  $S$  as a schema for *multiple* objects  $c_1, c_2, c_3$  in  $C$  we do not repeatedly add the Kolmogorov complexity of  $S$  to the respective Kolmogorov+ complexities of  $c_1, c_2, c_3$  relative to  $S$  and sum up, but instead add the Kolmogorov complexity of  $S$  just once to the sum of the respective Kolmogorov+ complexities of  $c_1, c_2, c_3$  relative to  $S$ . The canon of a trained autoencoder, we suggested, comprises objects that are individually complex but collectively simple. Another way to say this is that as we consider larger and larger collections of objects from a trained autoencoder’s canon  $C$ , specifying the relevant objects using our own semiotic system, we quickly reach a point whereupon the shortest path to specifying the collected objects is to first establish the trained autoencoder’s generative language  $S$ , then succinctly specify the objects using  $S$ .

10. Suppose that when a person grasps a style or vibe in a set of worldly phenomena, part of what she grasps can be compared to the formulae of autoencoder trained on this collection. The canon of this abstract trained autoencoder, then, would be an idealization of the worldly set, intensifying the worldly set’s own internal logic. Going the other way around, we might consider the idea that when the imaginative landscape of a literary work possesses a strong unity of style, the aesthetic unity of the artifactual collection is potentially an idealization of a looser, weaker aesthetic unity between the objects or phenomena associated with a real-world domain that the work of art encodes.

In the autoencoder case, we know to treat the artifactual collection of objects or phenomena—the trained autoencoder’s canon, mathematically equivalent to the trained autoencoder itself—as a systemic, structural gestalt representation of a worldly set whose vibe it idealizes. Applying the same thinking to the literary case, we might speculate that a *dense vibe* in the imaginative landscape associated with a work of art potentially acts as a structural representation of a *loose vibe* of the collective objects and phenomena of a real-world domain. I would offer, similarly, that the ‘dense aesthetic structure’ in question thus potentially provides a schema for interpreting the objects and phenomena of a real-world domain in accordance with a ‘systemic gestalt’ given through the imaginative landscape of the literary work.

POETRY

thing which no one living  
prizes but the dead see  
and ask among themselves,  
What do we remember that was shaped  
as this thing  
is shaped? while their eyes  
fill  
with tears. By which  
and by the weak wash of crimson  
colors it, the rose  
is predicated  
WILLIAM CARLOS WILLIAMS

90

Excerpt from William Carlos Williams, Paterson, 1927

II. It is logically possible to share a trained autoencoder’s formula directly, by listing the substrate of a neural network bit by bit, but it is a pretty bad idea to try: the computations involved in autoencoding, let alone in any abstractly autoencoding-like

bio-cognitive processes, are mathematically intractable and conceptually oblique. If what a person grasps in grasping the ‘aesthetic unity’ or vibe of some collection of phenomena is, even in part, that this collection of phenomena can be approximated using a limited generative language, then we cannot hope to express or share what we grasped in its abstract form. One mathematical fact about neural nets that neural-netty creatures like us *can* easily use, however, is the practical identity between a trained autoencoder and its canon: if grasping a loose worldly vibe has the form of a trained autoencoder, we should expect to share our vibe-insight with each other by intersubjectively constructing an appropriate set of idealized phenomena. At the same time, we should expect that the ‘idea’ that our constructed set of idealized phenomena expresses is essentially impossible to paraphrase or separate from its expressive form, despite its worldly subject matter.

12. A vibe is therefore, in this sense, an *abstractum that cannot be separated from its concreta*. The above phrasing tellingly, if unintentionally, echoes and inverts a certain formula of the “romantic theory of the symbol”—as given, for example, in Goethe’s definition of a symbol as “a living and momentary revelation of the inscrutable” in a particular, wherein “the idea remains eternally and infinitely active and inaccessible [wirksam und unerreichbar] in the image, and even if expressed in all languages would still remain inexpressible [selbst in allen Sprachen ausgesprochen, doch unausprechlich bliebe].”<sup>10</sup> The relationship of our literary-philosophical trope of a ‘vibe’ to the romantic literary-philosophical trope of ‘the Symbol’ is even clearer when considering Yeats’s more pithy paraphrase a century later, at the end of the romantic symbol’s long trans-European journey from very early German romanticism to very late English Symbolism: “A symbol is indeed the only possible expression of some invisible essence, a transparent lamp about a spiritual flame.”<sup>11</sup>

13. A question therefore brings itself to mind: does the idea of an abstractum that cannot be separated from its concreta simply reaffirm the Goethe/Yeats theory of the symbol from the opposite direction, positing a type of abstractum (a ‘structure of feelings’) that can only be expressed in a particular, rather than a type of particular (a ‘symbol’) that singularly expresses an abstraction? Not really, I would argue; indeed, I would say the difference between the two is key to the elective affinity between vibe and specifically *Modernist ars poetica*.

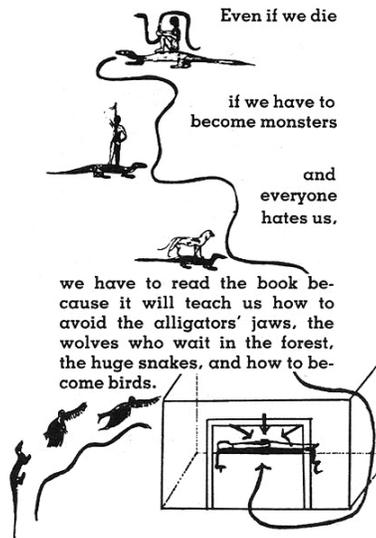
14. Despite its oh so many continuities with Symbolism and romanticism, the era of Pound, Eliot, Joyce, and Stein is marked by the ascendancy of a certain *materialist* reorientation of the Symbolist/romantic tradition. One relevant sense of ‘materialist’ is the sense that Daniel Albright explores in his study of Modernist poetic theory’s

borrowings from chemistry and physics, but a broader relevant sense of ‘materialist’ is closer to ‘not-Platonist,’<sup>12</sup> or to ‘immanent’ in the Deleuzian sense. Recalling Joyce’s and Zukofsky’s Aristotle fandom, and perhaps observing that William Carlos Williams’s “no ideas but in things”<sup>13</sup> is about as close as one can get to ‘universalia in re’ in English, we might even risk calling it an *Aristotelian* reorientation of the Symbolist tradition, both in aesthetic theory and in aesthetic practice.

15. For the Modernist aesthetic theorist, the philosophical burden on poetics partly shifts from the broadly Platonist burden of explaining how concreta could rise up to reach an otherwise inexpressible abstract idea, to the broadly Aristotelian burden of explaining how a set of concreta is (or can be) an abstract idea. Where Coleridge looked to the Imagination<sup>14</sup> as the faculty that vertically connects the world of things to the world of ideas for example, William Carlos Williams looked to the Imagination as the faculty that horizontally connects *things* to create a *world*. From a broadly Aristotelian point of view, the Poundian/Eliotian —or, less canonically but more accurately, Steinian—operation wherein poetry explicitly arranges or aggregates objects in accordance with new, unfamiliar partitions<sup>15</sup> is precisely what it means to fully and directly represent abstracta: an abstractum *just is* the collective affinity of the objects in a class. In fact, in “New Work for the Theory of Universals,” the premier contemporary scholastic materialist David Lewis formally proposes that universals are simply ‘natural classes,’ metaphysically identical to sets of objects that possess internal structural affinity.

16. By way of an example of a literary work’s production of a ‘horizontal’ *symbol* as described above, we might consider the imaginative landscape of Franz Kafka’s corpus. It is not very outrageous, I believe, to offer that it operates as just this kind of aesthetic schema for the unity or the affinity of a collection of real world phenomena. A reader of Kafka learns to see a kind of Kafkaesque aesthetic at play in the experience of going to the bank, in the experience of being broken-up with, in the experience of waking up in a daze, in the experience of being lost in a foreign city, or in the experience of a police interrogation—in part by learning that surprisingly many of the real life nuances of these experiences can be well-approximated in a literary world whose constructs are all fully bound to the aesthetic rules of Kafkaen construction. We learn to grasp a Kafkaesque aesthetic logic in certain worldly phenomena, in other words, partly by learning that the pure Kafkaesque aesthetic logic of Kafka’s literary world can generate a surprisingly good likeness of these worldly phenomena.

17. This minor brush with Kafka, and with the inevitable ‘Kafkaesque,’ also provides us with a good occasion to remark an interesting relationship between ambient meaning, literary polyvalence, and processes of concept-learning. Let us take the late French Symbolist and early Parisian avant-garde concept of ‘polyvalence’ to include both phenomena of collage, hybridity, and polyphony, where the heterogeneous multiplicity is on the page, and phenomena of indeterminacy, undecidability, and ambiguity where the heterogeneous multiplicity emerges in the readerly process. On the view suggested here, a vibe-coherent polyvalent literary object functions as a nearly-minimal concrete model of the abstract structure shared by the disparate experiences, objects, or phenomena spanned by the polyvalent object, allowing us to unify these various worldly phenomena under a predicate, e.g., the ‘Kafkaesque.’ The paradigmatic cases of this cognitive work are, inevitably, those that have rendered themselves invisible by their own thoroughness of impact, where the lexicalization of the aesthetically generated concept obscures the aesthetic process that constitutively underlies it: we effortlessly predicate a certain personal or institutional predicament as ‘Kafkaesque,’ a certain worldly conversation as ‘Pinteresque,’ a certain worldly puzzle as ‘Borgesian.’ (I’m still waiting for ‘Ackeresque’<sup>16</sup> to make it into circulation and finally name contemporary life, but Athena’s owl flies only at dusk and so on.)



Excerpt from Kathy Acker, *Blood and Guts in High School*, New York, Grove Press, 1984

18. Perhaps the best conceptual bridge from the raw ‘aesthetic unity’ that we associated with an autoencoder’s canon to a kind of systemic gestalt modeling of reality that we associate with the computational form of a trained autoencoder is what we might call the relation of *comparability* between all objects in a trained autoencoder’s canon. The global aesthetic unity of the objects in a set fit for autoencoding, I propose, is not just technically but conceptually and phenomenologically inseparable from the global intercomparability of the manifold’s objects, and the global intercomparability of the manifold’s objects is not just technically but conceptually and phenomenologically inseparable from the representation of a *system*.

19. In the phenomenology of reading, we experience this (so to speak) ‘sameness of difference’ as primary, and the ‘aesthetic unity’ of a literary work’s imaginative landscape as derived. A literary work’s ‘style’ or ‘vibe,’ is, at first, an invariant structure of the very transformations and transitions that make up the work’s narrative and rhetorical *movement*. As we read Georg Büchner’s ‘Lenz,’ for instance, plot moves, and the lyrical

processes of Lenz's psyche revolve their gears, and Lenz shifts material and social sites, and every change consolidates and clarifies the higher-order constancy of mood. A given literary work's invariant style or vibe, we argued, is the aesthetic correlate of a literary work's internal *space of possibilities*. This space of possibilities is, from the reader's point of view, an extrapolation from the *space of transformations* that encodes the logic of the work's narrative, lyrical, and rhetorical 'difference engine.' Or, more prosaically: no less than it means a capacity to judge whether a set of objects or phenomena does or does not collectively possess a given style, to grasp a 'style' or 'vibe' should mean a capacity to judge the difference between two (style-conforming) objects in relation to its framework. 20. Learning to sense a system, and learning to sense in relation to a system—learning to see a style, and learning to see in relation to a style—are, autoencoders or no autoencoders, more or less one and the same thing. If the above is right, and an 'aesthetic unity' of the kind associated with a 'style' or 'vibe' is immediately a sensible representation of a logic of difference or change, functional access to the data-analysis capacities of a trained autoencoder's feature function and abstract lower-dimensional representation-space follows, in the *very* long run, even from appropriate 'style perception' or 'vibe perception' alone, since the totality of representation-space distances between input-space points logically fixes the feature function. More practically, access to representation-space difference and even to representation-space distance alone is—if the representation-space is based upon a strong *lossy compression schema* for the domain—practicably sufficient for powerful 'transductive'<sup>17</sup> learning of concrete classification and prediction skills in the domain. When we grasp the loose 'vibe' of a real-life, worldly domain via its idealization as the 'style' or 'vibe' of an ambient literary work, then, we are plausibly doing at least as much 'cognitive mapping' as there is to be found in the distance metric of a strong lossy compression schema.

21. One reason the mathematical-cognitive trope of autoencoding matters, I would argue, is that it describes the bare, first act of treating a collection of objects or phenomena as a set of *states of a system* rather than a bare collection of objects or phenomena—the minimal, ambient systematization that raises *stuff* to the level of *things*, raises *things* to the level of *world*, raises *one-thing-after-another* to the level of *experience*. (And, equally, the minimal, ambient systematization that erases nonconforming *stuff* on the authority of *things*, marginalizes nonconforming *things* to make a world, degenerates experience into false consciousness.)<sup>18</sup>

22. In relating the input-space points of a set's manifold to points in the lower dimensional internal space of the manifold, an autoencoder's model makes the fundamental distinction between phenomena and noumena that turns the input-space points of the manifold into a system's range of visible states rather than a mere arbitrary set of phenomena. The parallel 'aesthetic unity' in a world or in a work of art—what we have called its 'vibe'—is arguably, in this sense, something like a maximally 'virtual' variant of Heideggerian mood ('Stimmung'). If a mood is a 'presumed view of the total picture' (Flatley) that conditions any specific attitude toward any particular thing, the aesthetic unity that associates the collected objects or phenomena of a world or work with a space of possibilities that gives its individual objects or phenomena meaning by relating them to a totality is sensible cognition of (something like) the Stimmung of a system—and much like Stimmung, it is the "precondition for, and medium of"<sup>19</sup> all more specific operations of subjectivity. What an autoencoding gives is something like the system's basic system-hood, its primordial having-a-way-about-it. How it vibes.

---

## Footnotes

1. 'Vanilla' autoencoders, as described here, are antiques in deep learning (DL) research terms. Contemporary variants like autoencoder generative adversarial networks (GANs), however, have performed exceptionally in 2017.
2. The term originally comes from Paul Ricoeur, in reference to Marx and Freud. Colloquially, it has come to name the academic reading practices of mainstream Anglo-American critical theory at the turn of the 21st century. See Paul Ricoeur. *Freud and Philosophy. An Essay on Interpretation*. New Haven, CT: Yale University Press, 1977. Print.
3. Rita Felski. "After Suspicion." *Profession*. 2009. 28-25. Print.
4. See Sianne Ngai. *Ugly Feelings*. Cambridge, MA: Harvard University Press, 2007. Print. And Jonathan Flatley. *Affective Mapping*. Cambridge, MA: Harvard University Press, 2008. Print.
5. Compare with Trisha Low: "The idea is that all this ethereal, feminine language is really concrete, or a sort of sublime mass of flesh that can really press down on certain kinds of structures that produced it in the first place. Like tar. Well, I guess I'm not secretly a structuralist anymore because I've said I'm secretly a structuralist so many times that people just know. But I'm interested in the way that somatic disturbances can press up against templates or structures, which make them more visible. Or not even necessarily more visible, but which produce a tension between what you feel is the fleshy part and what you feel is the structure underneath. The two are still indivisible though." (Sarah Gerard. Interview with Trisha Low. "Trisha Low by Sarah Gerard." *BOMB Magazine* 3 June 2014. Web.)
6. Elif Batuman. "What Am I Doing Here." *The Guardian* April 2008. Web.
7. The same goes, I would say, for meaning in the works of Modernists like Alfred Jarry, Virginia

Woolf, Franz Kafka, Maeterlinck, Raymond Roussel, Ezra Pound, TS Eliot, Robert Musil, Andrei Bely, Viktor Shklovsky, Walter Benjamin, Vladimir Khlebnikov, Daniil Kharms, Yukio Mishima, Harold Pinter, John Ashbery, Nathalie Saurraute, Haroldo de Campos, Samuel R. Delany, Kathy Acker, or Robbe-Grillet, and of staple 'proto-Modernist' anchors like Georg Büchner, Herman Melville, Comte de Lautreamont or Emily Dickinson, as well as parts of later Johann Wolfgang von Goethe, Charlotte Brontë, later Anton Chekhov, and later Gustav Flaubert.

8. More formally, we proposed to understand the features of a trained autoencoder as analogous to a fixed list of predicates with room to write-in a real-valued numerical grade from 0 to 9 next to each, where 0 means 'not at all' and 9 means 'extremely.'

9. In the normal definition of Kolmogorov complexity, the 'semiotic system' in question must be Turing-complete: that is, the semiotic system in question must be capable of describing a universal Turing-machine. (Our own 'default' semiotic system—that is, the semiotic system of the human subject currently communicating with you, the reader—is of course Turing-complete, since we can think about, describe, and build universal Turing-machines.) In the coming discussion, we will lift this restriction, in order to allow us to also talk about the Kolmogorov complexity of certain sets relative to more limited semiotic systems—semiotic systems like a given trained autoencoder's 'generative vocabulary' (decoder function). The purpose of this deviation is to save space we would otherwise have to devote to the fidgety technical concepts of conditional Kolmogorov complexity and of upper bounds on Kolmogorov complexity. We take this liberty because unlike most other mathematical concepts, the concept of Kolmogorov complexity does not have a preexisting one-size-fits-all fully formal definition, and always calls for a measure of customization to the purposes of a given discussion. For the sake of propriety, we will mark each instance of this 'off brand' application of the concept of Kolmogorov complexity as 'Kolmogorov+ complexity.'

10. Johann Wolfgang von Goethe. Quoted in Miguel Beistegui. *Aesthetics After Metaphysics: From Mimesis to Metaphor*. New York: Routledge, 2012. Print.

11. W. B. Yeats. "William Blake and His Illustrations to The Divine Comedy." *Ideas of Good and Evil*. 1903.

12. We will leave aside the question whether Plato was, himself, a Platonist in this sense.

13. William Carlos Williams. *Paterson*. 1927.

14. Coleridge and William Carlos Williams both take their concept of Imagination from Kant.

15. A partition is the division of a set into non-overlapping subsets.

16. See Kathy Acker. *Blood and Guts in High School*. New York: Grove Press, 1984. Print.

17. A. Gammerman, V. Vovk, and V. Vapnik. "Learning by Transduction." *UAI'98 Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*. San Francisco: Morgan Kaufmann Publishers, 1998. 148-155. Print.

18. See Adrian Piper. "Xenophobia and Kantian Rationalism." *Philosophical Forum* XXIV 1-3 (Fall-Spring 1992-93). 188-232. Print. Piper discusses xenophobia as "a special case of a more general cognitive phenomenon, namely the disposition to resist the intrusion of anomalous data of any kind into a conceptual scheme whose internal rational coherence is necessary for preserving a unified and rationally integrated self."

19. See Martin Heidegger. *The Fundamental Concepts of Metaphysics. World, Finitude, Solitude*.

Bloomington: Indiana University Press, 1983. Print.

---

Peli Grietzer recently finished his PhD in mathematically informed literary theory at Harvard Comparative Literature and the HUJI Einstein Institute of Mathematics. His dissertation borrows mathematical forms from deep learning theory to model the ontology of 'ambient' phenomena.