



Pamphlet **15**
→ September 2017

Literary **Lab**

Patterns and Interpretation

Franco Moretti

Pamphlets of the Stanford Literary Lab

ISSN 2164-1757 (online version)

Patterns and Interpretation

1. Abstraction

One thing for sure: digitization has completely changed the literary archive. People like me used to work on a few hundred nineteenth-century novels; today, we work on thousands of them; tomorrow, hundreds of thousands. This has had a major effect on literary history, obviously enough,¹ but also on critical methodology; because, when we work on 200,000 novels instead of 200, we are not doing the same thing, 1,000 times bigger; we are doing a different thing. The new scale changes our relationship to our object, and in fact *it changes the object itself*. “No one has ever seen the objects studied by contemporary historians”, Krzysztof Pomian once wrote, “and no one *could* ever have seen them [...] because they have no equivalent within lived experience”.² True. No one has a lived experience of demographic change, or of literacy rates, or of

– Figure 1.1




Figure 1.1 “Objects with no equivalent within lived experience”. From Sarah Allison, Marissa Gemma, Ryan Heuser, Franco Moretti, Amir Tevel, Irena Yamboliev, “Style at the Scale of the Sentence”, Literary Lab Pamphlet 5, 2013.

I will explain this chart in a minute; for now, let me just say that this is what literature has become in the new space of literary labs. These are still novels; but prepared for analysis in

1 See Mark Algee-Hewitt, Sarah Allison, Marissa Gemma, Ryan Heuser, Franco Moretti, Hannah Walser, “Canon/Archive. Large-scale Dynamics in the Literary Field”, Literary Lab Pamphlet 11, 2016.

2 Krzysztof Pomian, *L'Ordre du temps*, Paris 1984, p. 31.

a way that severs all connections with the lived experience of literature. Watching a play, listening to a poem, reading a novel: this is the lived, *concrete* experience of literature. **Figure 1.1**, is entirely abstract. It takes certain novelistic elements, it pulls them out of their context, and re-presents them in a completely different combination. In this case, the presentation is based on principal component analysis; it could be a trend, a map, a tree of distances – or many other things. But all of them, equally abstract. Alongside the explosion of the archive, then, and probably with even deeper consequences, this is the great novelty of computational criticism: *a redefinition of literature that foregrounds those traits that can be more easily abstracted, and hence programmed*. In this sense, algorithms have changed both *what* we study, and *how* we study it. Think of reading. For centuries, reading has been indispensable to the understanding of literature. In front of **Figure 1.1**, it is nothing. Nothing.

Just to be clear, it's not that we should stop reading books: reading is one of life's greatest pleasures, it would be insane to give it up. What is at stake is not reading, it's *the continuity between reading and (a certain kind of) knowledge*. I read books; but when I work in the Literary Lab they're not the basis of my work. Corpora are; ideally, those 200,000 novels. Here, size is again crucial, because a corpus is not "just like a text, only more of it".³ A text is a "communicative event": written by someone, in specific circumstances, to convey a specific meaning. But no one writes corpora; they are not "communicative events" – they are not events at all, they are artificial objects created by a researcher. A text is meant to address us; to "speak" to us. Corpora don't speak to us; which is to say, *they have no meaning in the usual sense of the word*.⁴

Now, meaning is not one of the things literary critics study; it is *the* thing. Here lies the great challenge of computational criticism: thinking about literature, removing meaning to the periphery of the picture. But of course this is also the great challenge *for* computational criticism: you discard meaning and replace it with – what? Dismembered in **Figure 1.1** is a sentence – "Yet when he arrived at Stone Court he could not see the change in Raffles without a shock" – which, when we encounter it in *Middlemarch*, triggers a chain of events which will turn a major character into an accomplice in someone's death. A very meaningful sentence. None of that remains visible in **Figure 1.1**. This is what we lose, with a computational approach. What do we gain?

Figure 1.1 was part of a study of the stylistics of nineteenth-century novelistic sentences, in which, at a certain point, we compared sentences in which the independent clause came first, and the dependent clause followed ("I opened the door, as soon as the bell rang": abbreviated to ICDC), to those where the order was reversed ("As soon as the bell rang, I opened the door": abbreviated to DCIC).⁵ We were particularly interested in the latter type,

3 Elena Tognini Bonelli, "Theoretical overview of the evolution of corpus linguistics", in Anne O'Keeffe and Michael McCarthy, eds., *The Routledge Handbook of Corpus Linguistics*, Routledge, New York 2010, p. 19.

4 This is not a literary critic's idea of meaning; if anything, a strong link between meaning, context, and use is more typical of anthropologists like Geertz, or historians of political thought like Skinner and Pocock.

5 Sarah Allison, Marissa Gemma, Ryan Heuser, Franco Moretti, Amir Tevel, Irena Yamboliev, "Style at the Scale of the

because a sentence that begins with a dependent clause can't really stop there (As soon as the bell rang ... what?), you have to go on, it's a very *narrative* arrangement, and we wanted to know whether it differed from the other type, not just in syntactical order, but in semantic content as well. So, we broke the sentences down into four classes of clauses,⁶ calculated their most distinctive words, and used principal component analysis to visualize the results. **Figure 1.1** is the conclusion of the process: the four clauses are the purple vectors, and are surrounded by the nouns that most characterize them; the color of the words indicates their raw frequency, their size their statistical distinctiveness, and the good separation among the clauses suggests that syntax and semantics are indeed correlated. Good. And now?

2. Patterns

What we did – what most literary researchers would do, in these circumstances – was to look for signs of an unusual configuration. And luckily, in the upper right quadrant, around the dependent clause of the dependent-independent sequence, a cluster of related terms emerged: “drawing room”, “home”, “house”, “door”, “hall”, “church”, “building”, “gate”, “town”, “road”, “street”, “palace”, “yards”, “slope”, and “park” (**Figure 2.1**). In the upper left quadrant, around the independent clause of the same type of sentence, there was a different, but equally consistent cluster: “feelings”, “jealousy”, “indignation”, “despair”, “admiration”, “fancy”, “interest”, “memory”, and “tears” (**Figure 2.2**). From what looked like chaos, a pattern had




Figure 2.1 Dependent clauses and their main semantic cluster

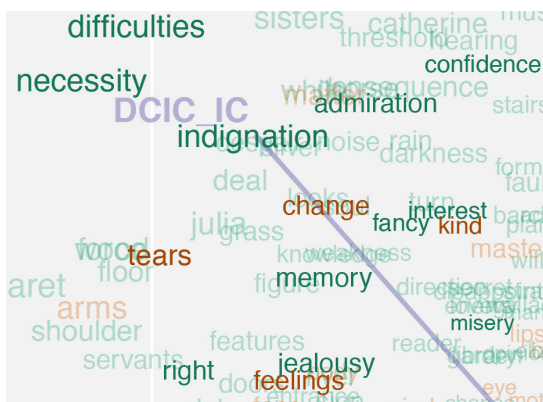


Figure 2.2 Independent clauses and their main semantic cluster

Sentence”, Literary Lab Pamphlet 5, 2013, especially pp. 5-10.

⁶ To continue using the previous examples, the four classes and their abbreviations are as follows: “I opened the door, as soon as the bell rang”, independent clause of the ICDC sequence, or ICDC_IC; “I opened the door, as soon as the bell rang”, dependent clause of the ICDC sequence, or ICDC_DC; “As soon as the bell rang, I opened the door”, independent clause of the DCIC sequence, or DCIC_IC; “As soon as the bell rang, I opened the door”, dependent clause of the DCIC sequence, or DCIC_DC.

emerged (Figure 2.3). We remove meaning to the periphery and replace it with what, I asked a minute ago. With this: with patterns. Instead of reading, pattern recognition.




Figure 2.3 A pattern emerges

Great keyword of our times, pattern; a recent article uses it 74 times. But what exactly does it mean? Certainly, not what it used to when it entered the English language: derived from the French *patron* – master, owner, boss – pattern initially indicated, the OED tells us, “something shaped or designed to serve as a model from which a thing is to be made” (1324); “an example or model to be imitated” (c1450); “an example, an instance, esp. one taken as typical, representative, or eminent” (1555). Originally, “pattern” was thus a *normative* concept: people created patterns, and then imposed them onto the world. Then, about a century ago, a 180° shift occurred: we no longer brought patterns into the world, we found them there. Patterns acquired an empirical, independent existence: “a regular and intelligible form or sequence discernible in certain actions or situations; esp. one on which the prediction of successive or future events may be based (1883); “an arrangement or relationship of elements, esp. one which indicates or implies an underlying causative process other than chance” (1900).⁷ And it was at the moment these “objective” features emerged, incidentally, that the frequency of “pattern” suddenly increased (Figure 2.4).




Figure 2.4 “Pattern” in English, 1800-2000. Source: Google Books Ngram Viewer, consulted on April 21, 2016

⁷In the new meaning of the term, the earlier subjective element hasn’t entirely disappeared, but it has shrunk to the merely subordinate function of “recognition”, in which subjectivity is still needed to extricate the pattern from its surroundings (to “discern” it, as in the 1883 entry), but it no longer imposes a model “onto” the world, as had been the case earlier.

Pattern as a “relationship of elements”: like the one we had found between semantics and syntax in our study of the sentence. A correlation: another great keyword of our times. A few years ago, the editor of *Wired* magazine, Chris Anderson, used this notion to declare “the scientific method obsolete”:

Scientists are trained to recognize that correlation is not causation, that [...] you must understand the underlying mechanisms – the model – that connect the two. But faced with massive data, this approach to science [...] is becoming obsolete. [...] Petabytes allow us to say: “Correlation is enough” [...] Correlation supersedes causation, and science can advance even without coherent models, unified theories, or really any mechanistic explanation at all.⁸

This is a wonderful statement, because it expresses with absolute candor the ambivalence towards knowledge that is typical of our times. Correlations, yes; but no theories, no models, not even explanations. Knowledge, without explanations? It’s the same anti-intellectualism as that generated by the industrial revolution two centuries ago, and for the same reason: knowledge is good – to make machines work. Period. Useful knowledge, the Victorians called it; engineers, not scientists. The same in Silicon Valley today, this El Dorado of engineers where Anderson has in the meantime reinvented himself as a drone manufacturer.

Now, It’s not that pattern recognition doesn’t matter. It does – as a beginning. What *really* matters, though, are all those things Anderson had declared to be obsolete: finding the underlying mechanism; establishing causation; *explaining* why patterns are there in the first place, and what do they do. And this requires, as a first step, a further leap into abstraction. If the image in **Figure 1.1** seemed abstract, well, *it wasn’t abstract enough*: we must take all the “home” and “door” and “road”, all the “tears” and “jealousy” and “fancy” of **Figures 2.2-2.3**, and reduce them to a purely *formal* relationship: space in dependent clauses – emotions in independent ones (**Figure 2.5**).

From the chaos of data of **Figure 1.1**, through the emergence of patterns in

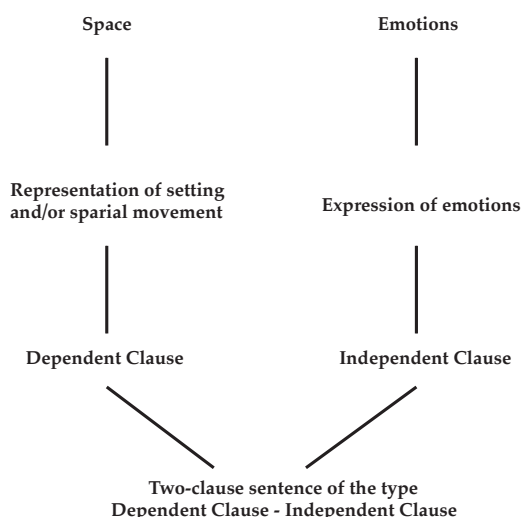


Figure 2.5. The form behind the pattern

Here are a few instances of this abstract structure, all drawn from the corpus of “Style at the Scale of the Sentence”: “When Peter perceived the village, he burst into a shout of joy” (Radcliffe); “When the day came round for my return to the scene of the deed of violence, my terrors reached their height” (Dickens); “When Deronda met Gwendolen and Grandcourt on the staircase, his mind was seriously preoccupied” (Eliot); “Yet when he arrived at Stone Court he could not see the change in Raffles without a shock” (Eliot).

⁸ Chris Anderson, “The End of Theory: The Data Deluge Makes the Scientific Method Obsolete”: http://archive.wired.com/science/discoveries/magazine/16-07/pb_theory/

Figures 2.1-3, to the clarity of form of **Figure 2.5**. Distillation. And it's not just a process of epistemological clarification: it's also the discovery of *an actual causal mechanism*. If there is a pattern in the data, *it's because behind it there is a form which repeats itself over and over again*. Form is the repeatable element of literature, as I wrote elsewhere, and the pattern is *nothing but* this repetition.⁹ Spatial terms keep occurring in dependent clauses, emotional ones in independent ones, and eventually we notice what is going on; we see – we “recognize” the pattern. But what we see is just the shadow of the form of **Figure 2.5**. *Patterns are the shadows of forms over data*. If you don't grasp the form, your hands remain empty.

3. Interpretation

The correlation between “home” and “road” and dependent clauses is – to use a word contemporary criticism is allergic to – a *fact*: meaning, that any investigation of these clauses would find the same quantities of homes and roads as we did.¹⁰ That those words form a cluster, however, and that the cluster means “space” – these are not facts, but *interpretations*. Someone else may set “home” and “road” in opposition, as “inside” and “outside”; I cluster them together, as two versions of “space”. Here, subjective decisions clearly return to be dominant. Algorithms generate new facts, whose interpretation continues however to rely on a different hermeneutic tradition.

I began this pamphlet by saying that computational criticism was replacing meaning with patterns; now, I am interpreting patterns on the basis of the meaning of words. But I don't think this is a contradiction, it's that the term “meaning” can be used in (at least) two very different ways. In interpreting the pattern emerging from **Figures 2.1-3**, I was referring to the “dictionary” meaning of words, which the founder of modern hermeneutics, Schleiermacher, would usually indicate with the term *Bedeutung*; earlier on, in discussing texts as communicative events, I was referring to meaning in a specific context, which Schleiermacher would usually indicate with the term *Sinn*.¹¹ “Significance” is a possible English translation of *Sinn*, leaving “meaning” for the more abstract dictionary use; “senso” and “significato” are used in Italian in a similar way, and comparable pairs exist in other languages. But the point is not what words we use, it's being aware that our reflection on meaning may take two very different directions: dictionary meaning, or meaning-in-context; large aggregates, or the “given

⁹ See “The Slaughterhouse of Literature” [2000], now in *Distant Reading*, Verso, London 2013.

¹⁰ A fact because inter-subjective, therefore, but also in the sense of *factum*: because the correlation we had found was not a “given”, but the result of a series of operations: the identification of specific sentences, their division into clauses, the establishment of their most distinctive words, the decision to focus on nouns, the choice of a quantitative threshold, then of principal component analysis ... Anyone repeating this series of steps would find the same values; the steps themselves, though, are dictated by a specific, and subjective, research interest.

¹¹ “Some term what a word is thought to mean “in and of itself” its meaning [*Bedeutung*] and what the word is thought to mean in a given context its “sense” [*Sinn*].” Friedrich Schleiermacher, *Hermeneutics: the Handwritten Manuscripts*, Heinz Kimmerle, ed., The Scholars Press, Missoula, Montana 1977 , p. 117.

statement”, as Schleiermacher put it in the page where he defined “the rules of the art of interpretation” [111].

Schleiermacher’s claim for this art is well-known, but it’s so wonderfully shameless that I want to repeat it: ‘The task [of interpretation] is to be formulated as follows: “To understand the text at first as well as and then even better than its author.” ’ [112] This is, incidentally, the only justification for the existence of literary critics; without us darkness would prevail, fortunately we are here, understanding literature better than its authors. But why do we understand better? Here lies Schleiermacher’s stroke of genius: the reason is not our superior knowledge, but *our ignorance*: “we have no direct knowledge of what was in the author’s mind”, and so we’re forced “to become aware of many things of which [the author] himself may have been unconscious”. The author knows many things that, so to speak, he doesn’t know he knows; an interpreter can only know them by making them explicit: that’s what understanding a text “better than its author” means. And it’s extremely elegant, but with a big obstacle for quantitative and computational criticism: for Schleiermacher, the aim of this type of interpretation consists in understanding *individual texts in their “individuality”*,¹² and in relation to the “intention” (however defined) of the author. But corpora have neither individuality nor intentions. So?

Why is interpretation necessary in the first place? Because of the “many things” that an author “knows”, but leaves implicit. And why does an author do that? Schleiermacher doesn’t really say, but we can conjecture: because these things were learned *by doing*, not by explicit declarative knowledge. Know-how, more than know-that; practices and conventions which were widely shared in the author’s world. *This* is what the interpreter must “become aware of”: not hidden individual secrets, but habits that were socially ubiquitous, and that escaped attention, purloined-letter-like, precisely *because* of their visibility. What the interpreter recovers, in other words, is *the social within an individual oeuvre*.

Here, the aim of Schleiermacher’s project and the means necessary to achieve it emerge as somewhat at odds with each other. The aim is the understanding of individual texts; the means, the making explicit of social practices. The latter are necessary to the hermeneutic enterprise, but in a purely instrumental way; in themselves, they are not interesting. On this, computational criticism clearly reverses the hierarchy: conventions matter for us much more than any individual text: they matter *as such*, unlike what happens in the hermeneutic project. Here, the disagreement is clear. In another sense, though, our work can be seen as the *completion* of hermeneutics, because the “abstract objects” computation produces – these objects no one experiences directly, but which we all somehow know how to take into account – are exactly what hermeneutics wants to raise to the level of consciousness. We

12 “Technical interpretation: The language and its determining power disappear and seem to be merely an organ of the person, in the service of his individuality” [161]

could make Schleiermacher's motto our own: quantitative hermeneutics is the art of understanding conventions – forms, genres, styles, practices – better than their society ever did.¹³

4. Form, History, Explanation

More on this at the end. First, a quick step back to [Figure 2.5](#). There is no logical relationship between space and emotions; so, why did they “click”, and appear so frequently together? As always with form, the explanation is to be found in history. Around 1800, as nation-building and early industrialization were violently reshaping European geography, novelists ceased to treat space as a mere “container” of the story – as a sort of “box” where events unfolded on their own – and turned it into a force that actively *shaped* events: the “demonic” power of the *milieu*, as Auerbach put it in *Mimesis*. Social space had changed too much for literary space not to change as well. And likewise for emotions: around the same time, their expression on stage and on the written page was drawn towards that “rhetoric of excess” which, as Peter Brooks has shown in *The Melodramatic Imagination*, was a major symbolic consequence of the French revolution.

Why space and emotions? Because they had both been over-heated by contemporary events, and combining them in the micro-unit of a two-clause sentence was a way to bring some unity into historical experience. This is what form *always* does: it selects some elements of the given world, combines them, and creates a *model* of that world. Space has become so powerful, it generates events – and emotional intensity is itself one of the most significant events of human existence: this is what (the structure of) that sentence says. And it seems to have all originated with Gothic novels, where space and emotions were both sensational and inextricable: castles and terror, dungeons and madness, caves and despair ... Totally implausible. But, once the connection was in place, it could be endlessly re-functionalized – form, the repeatable element in literature – and adapted to the changing needs of the age. First came historical novels, where the space of internal peripheries provided an asylum for the grand passions banished from the modern world; later, it was the turn of the metropolises, and of their peculiar pathologies (Paris and ambition, London and eccentricity, Petersburg and radicalism, Madrid and folly); then the provinces, and the deadly dangers of modern boredom; while the great naturalistic cycles of the last quarter century endowed specific regions, like Verga's Sicily or Hardy's Wessex, with tragic significance. With each generation, the space-emotion complex was reinvented, joining a specific articulation of the

13 Of course, it's not like the computational approach has invented the study of conventions: genre theory and stylistics have existed for a very long time. But our abstract objects are finally the right objects of study for this type of investigation; or at least, better than their predecessors. They allow to shift from “type thinking” to “population thinking”, to use Ernst Mayr's concepts. What we've done, in other words, is creating the right laboratory environment for what had so far largely remained speculative theories.

nation-state to a different reaction to modern life. Chronotopes, saturated with feelings. No wonder writers kept returning to the structure of [Figure 2.5](#).¹⁴

This was a crazily accelerated summary of course; but I wanted to give an idea of how recognizing the form behind the pattern allows us to unify separate historical givens in a single explanation. If this pamphlet began by moving away from the lived experience of literature towards abstraction, form has emerged as simultaneously *the apex* of the process, and the turning point that allows to *reverse the direction*, and return from abstraction to literary history. Here, form functions very much like Max Weber's concept of the ideal-type: "a mental construct" that "cannot be found empirically anywhere in reality", but that, once constructed, can be used "for *comparison with and measurement of reality*".¹⁵ This is exactly how we should think of literary form: a mental construct which we will never find *as such* in individual works, but which we can use to "measure" their relationships. Form will *never* explain a single text, and is the *only* thing that can explain a series of them.

5. Noise

And yet, the return to "reality", to repeat Weber's term, may happen in more than one way. Let's ask ourselves: why do we look for patterns to begin with, and are so happy when we find them? Because they reveal some kind of order, and we *want* to find order, especially when we are confronted with large masses of data. "It is not by chance", wrote a great precursor of computational criticism, Leo Spitzer, "that the 'philological circle' was discovered by a theologian,




Figure 5.1. Patterns are real, but never perfect

who was wont to harmonize the discordant, to retrace the beauty of God in this world".¹⁶ The beauty of God in the world: order in the strongest possible sense. But the problem is, there is no beauty of God in patterns. There can't be, because patterns lie halfway between two domains – the chaos of empirical data, and the clarity of conceptual form – that are too "discordant" to be truly "harmonized". Patterns somehow bridge the gulf between them, *making form visible within data*, and it's a small miracle. But, small: right in the middle of the spatial cluster of "town" and "road" and "door" you find ... "god" (and a little lower there's a "bishop"; plus "sounds", "questions", "match", "hero", "letter", etcetera: [Figure 5.1](#)). The pattern is real,

¹⁴ If this scenario is correct, the space-emotions complex should persist throughout the nineteenth century, while its specific semantic clusters should change every generation or so. Unfortunately, the corpus of the original study is not large enough to test this prediction, which remains therefore, for the time being, just a hypothesis.

¹⁵ Max Weber, "Objectivity" in social science and social policy, 1905, in *The Methodology of Social Sciences*, Glencoe, IL, The Free Press, pp. 90, 97.

¹⁶ Leo Spitzer, "Linguistics and Literary History", 1948, in *Representative Essays*, Stanford UP, 1988, p. 32 (emphasis mine).

but not perfect: its borders are porous; its space, full of un-harmonized discordance. And the question is: what should we do with this discordance?

For the logic of research, discordance – disorder, noise, chaos, call it what you will – is fundamentally an obstacle that we must learn to ignore: in order to “see” the space-emotion pattern, we must somehow *not* see god, bishop, and company. They are there, we know they are there, but we don’t look at them. We have put on blinders, to use Weber’s great metaphor in *Science as a Vocation*. But once we have seen the pattern, we can take off the blinders, and then, it’s impossible not to notice how *small* the region of order actually is. Form is everywhere encircled by noise: a maelstrom of semantic options that have kept whirling around, without ever crystallizing into stable structures. Failed conventions. Failed styles. Noise. Now: what if we took this noise *as itself an object of knowledge*? Not an obstacle to interpretation, but its aim. What would interpretation become – what would a *hermeneutics of noise* be like? This is the ultimate challenge of digitization to literary theory. Traditional interpretation, once it had removed noise from the horizon of research, could legitimately keep it out of the picture forever. The alliance of algorithms and archive returns noise in front of our eyes again, and again, and again. “The times have been / That, when the brains were out, the man would die, / And there an end; but now they rise again.” Computation takes all the Banquos who have been slaughtered in the process of literary selection, and brings them back to life. Speaking to these ghosts, will be our task.