

# What Does Perspectivism Mean? An Ethical and Methodological Countercriticism

**Mathieu Valette**

ERTIM, Institut National des Langues et Civilisations Orientales (Inalco, Paris)  
2 rue de Lille, 75007 Paris, France  
mvalette@inalco.fr

## Abstract

In this paper, we address the epistemological and ethical break of perspectivism in NLP. First, we propose to consider data annotation from the point of view of the scientific management of annotation work - which is part of the automation process inherent in NLP, in order to ideologically situate the perspectivist paradigm. We then analyze some of the concepts of perspectivism (in particular, truth). Finally, based on this analysis, we formulate a set of proposals aimed at overcoming the observed limitations of corpus annotation in general and perspectivism in particular.

**Keywords:** perspectivism, postmodernism, ethics, annotated corpus, corpus linguistics, statistical analysis of textual data

## 1. Introduction

The dynamism of perspectivist work (Basile et al., 2021; Cabitza et al., 2023) attests to the fact that the construction of datasets in NLP, as in AI, is a research challenge that is more topical than ever. The inclusion of variation and *situated* interpretation in a field of research still strongly marked by a referentialist approach is good news. Of course, the days are long gone when all semantic phenomena were represented by means of ontologies, often general, stable and unchanging, mostly constructed in a top-down or onomasiological manner (Fellbaum, 1998). In this old paradigm, abundance and variety were a *problem* (polysemy problem for linguists, disambiguation problem for NLP scientists), whereas they are inherent in semiotic activity, i.e. the production of signs. Probabilistic approaches have won out over the dominant ontological paradigm.

The NLP community is increasingly asking itself about the political and sociological biases present in the datasets it processes (Feng et al., 2023), especially as these datasets are, as they grow, increasingly opaque. By connecting the issue with dataset human annotation, the perspectivist paradigm links a general scientific question (the process of qualifying an object) to an ethical question (the under-representativeness of minority sensibilities, in particular), paving the way for a more general discussion on the ethics of NLP and AI.

We aim to address the epistemological and ethical break of perspectivism in NLP. First, we propose to consider data annotation from the theoretical point of view of *scientific management* (Taylor, 1911) of annotation work which participates in the automation process inherent in NLP, in order to ideologically situate the perspectivist paradigm. We then analyze some of the concepts of perspectivism (in particular, truth). Finally, based on this analysis, we make a set of proposals aimed at overcoming the observed limits of perspectivism.

## 2. Annotation as Scientific Management of Work

The aim of NLP is automation, i.e. the elimination of the human component in the processing of textual data. NLP consists in setting up processes to obtain an output result from a set of input data. Among the proposed system improvements, reducing human intervention appears to be almost as important as improving raw performance, computed on the basis of well-known metrics (f-score, accuracy, etc.). Tasks performed by humans are traditionally described as "manual". However, "manual" has two antonyms: automatic, of course, but also intellectual, in which case it's not out of the question for "manual" to have a depreciatory connotation to describe a non-intellectual or low-intellectual task.

Manual work is indeed the stumbling block of NLP. A manual task par excellence is, paradoxically, reading and interpreting datasets. The NLP scientist must neither read nor interpret their dataset, firstly for practical reasons (the dataset is too large) but also for methodological reasons: they must keep their distance from it, and this distancing constitutes a strategy of objectification. In most cases, ethical standards for corpus annotation require that annotators do not participate in algorithm design. Knowledge of the corpus would induce a bias in algorithmic choices and bias the results. Interpretation tasks are therefore outsourced for methodological reasons.

Manual annotation (or partially automated annotation for the most robust tasks, such as POS tagging and named entity recognition) is the only form of text interpretation available in NLP (as we'll see in paragraph 4, it's not the only one possible). NLP scientists entrust annotation tasks to various third parties. Depending on the resources deployed, the stakes involved and the type of task, these annotation (and therefore reading) tasks may be delegated to experts (linguists, doctors, lawyers), or to less qualified individuals (interns, students) or to subcontractors (such as the most famous of them all:

Amazon Mechanical Turk) who provide little or no guarantee of the annotators' expertise - in short, they deskill the annotation work (Cohen et al., 2016).

Annotation campaigns are based on NLP's industrial concepts of automation, optimization and rationalization. Thus, the annotation task must be understood in terms of the scientific management of work also known as Taylorism, and thus the division of labor. Without going into the details of this division, we can distinguish three roles: (i) the project manager selects the corpus, defines the sequence of operations, controls the process, is in charge of the annotation guide, and the methodological and logistical choices, (ii) the curator supervises the annotation tasks and is the interface between the project manager and the annotators. He trains them in the annotation guidelines, checks the quality of annotations and acts as an expert; (iii) the annotators, who have no expert status, comply with the annotation guidelines. Finally, annotators work on data samples and do not have an overall view of the dataset or even of the processing line, which they may even be completely unaware of (Bontcheva, 2010).

In another industrial context, the annotator would correspond to the unskilled laborer working on a production-line. Until perspectivist proposals overturned the paradigm, the annotator was suspected of being, structurally, the weak link in the processing chain (Bontcheva, 2010): he or she may make careless mistakes or fail to understand guidelines. The whole evaluation system (inter-annotator agreement, Cohen's Kappa...) of the annotation task is based on this structural weakness, and determines annotator recruitment processes. Admitting that the annotator has an intrinsically low confidence rating deskills their work and forces the project manager to compensate by multiplying the number of annotators, who are recruited at low cost (which further lowers the level of requirement) or without pay, following practices such as incidental crowdsourcing (Park et al., 2019) or gamification.

A Marxian reading of these annotation campaigns is necessary: it shows a process entirely controlled by computer scientists who own the technological production apparatus (the ability to build up large-scale datasets, machine learning algorithms) and who buy the labor power of annotators-proletarians. This vision in terms of class rule may seem inappropriate when we think of IT, but on the one hand, we can only observe the incredible rise in power of the industrial players in the digital sector, proportional to their technological domination and the human costs it generates. On the other hand, the production of resources, i.e. the creation of value through annotation, is necessarily linked to the production means.

### 3. The Ethics of Perspectivism in Debate

We will not discuss here the scientific value-adding virtues of perspectivism, which consists in considering noise as information and therefore error as a positive value (Basile 2021; Cabitza et al., 2019; Cabitza et al., 2023; Sachdeva et al., 2022; Kralj Novak et al., 2022). We wish to raise two issues, one ethical, the other methodological, which can be linked in terms of solutions.

#### 3.1 Perspectivism Has no Impact on Work Management and Labor Conditions

We could be content to see only the technical and scientific contribution of perspectivism, i.e. the enrichment of data, but the Perspectivist Manifesto explicitly adopts an ethical stance<sup>1</sup>, both in its general argument and in the datasets available (Measuring Hate Speech, Pejorative Language in Social Media, Work and Job-Related Well-Being), so it's legitimate to discuss the ethics of perspectivism itself, since technical means are never neutral and, as we've seen, organize work.

The Manifesto forcefully denounces aggregation and harmonization as forms of obliteration of annotators' personal opinions and disregard for cultural background. The authors rightly observe that harmonization, in particular, can take place at the end of deliberative phases, where relationships of domination can be established to the prejudice of minority opinions (Noble, 2012). What we note, however, is that in no case does the Manifesto denounce the alienation of the task. *Crowdsourcing is still production-line work*<sup>2</sup>. Not only do these working conditions potentially expose annotators to abnormally repeated harmful content (Steiger et al., 2021), but also, we are not sure that many annotators declare that corpus annotation is fulfilling work and leads to well-being at work (contrary to what academic computer scientists might say about their high-qualified work). If we exclude the incidental crowdsourcing or gamification techniques already mentioned, annotation is a tedious, repetitive task, socially unrewarded, solitary by method and, finally, poorly paid (Fort et al., 2011; Gray and Suri, 2019). This last point is soberly mentioned - but not discussed - in the Manifesto in a footnote.

#### 3.2 Perspectivism Mistakes Sincerity for Truth

It's tempting to draw a parallel between the perspectivist paradigm in AI and its philosophical homonym (Leibniz, Nietzsche, but especially Deleuze) (Astor, 2020). The subjectivist relativism of the perspectivist paradigm is akin to the postmodern assumption that there is not just one truth, but many truths, and that all truths are equal. Many authors see this as a worrying drift, particularly in science and politics, as it leads to *post-truth politics* (Holzem, ed. 2019) and to pseudoscientific or negationist positions,

<sup>1</sup> "Aggregation and harmonization destroy any personal opinion, nuance, and rich linguistic knowledge that come as a result of the different cultural and demographic background of the annotators" (<https://pdai.info/>).

<sup>2</sup> In fact, the term "annotator" is sometimes replaced by "worker" (Aroyo and Welty, 2015).

which in turn provide the breeding ground for totalitarianism (Rastier, 2019).

Of course, the perspectivist paradigm does not claim to offer alternative scientific truths, but several interpretations that reflect the opinions of the annotators, and it is not uncommon for the concept of “truth”, sometimes renamed “ground truth” to be relativized, whether in the Manifesto or in *position papers* (Basile et al. 2021; Planck, 2022; Cabitza et al, 2023). For instance, Aroyo and Welthy (2015), pervert the traditional concept of truth with a leitmotiv that systematically contradicts the rational thinking that underpins modern science (“truth is a lie”, the “antiquated ideal of truth”...). They even propose a “new theory of truth”: “Crowd truth is the embodiment of a new theory of truth that rejects the fallacy of a single truth for semantic interpretation, based on the intuition that human interpretation is subjective and that measuring annotations on the same objects of interpretation (in our examples, sentences) across a crowd will provide a useful representation of their subjectivity and the range of reasonable interpretations” (Aroyo and Welthy 2015: 21). The label “theory of truth” is a misuse of language, firstly because their purpose is limited to corpus annotation and not to a scientific definition of truth (only the 1st of the 7 myths they identify calls into doubt the concept of truth, the other 6 myths are purely methodological criticisms), and secondly, because the authors confuse “truth” with opinion, or some concept we could design as “sincerity”.

The confusion maintained between sincerity and truth in almost all the scientific production of the perspectivist paradigm is not just a problem of terminological inaccuracy. It is an epistemological confusion with methodological consequences. What is valued in the perspectivist paradigm is the sincerity of annotations. Truthfulness can be evoked, for example, in the annotation of linguistic norms, or in the perspective of establishing linguistic norms (e.g. POS tagging of poorly endowed languages leads to the identification of variations that are not variations of personal sensitivities but of norm perception). Sincerity is a psychological concept and is not a matter for the annotated texts, but for the annotators themselves, their subjectivity and the resulting interpretation.

Interpretations are not just a matter of the annotator and the sample to be annotated going head to head. Numerous interpretative biases could be cited and, in order to inventory and model them, we could draw on millennia of bibliography, from Aristotle's Rhetoric (what is the ethos of the speaker?) to R. Jakobson's functions of language. We could also ask about the material conditions of the device (how was the corpus constituted? for what purpose? how was the task described? What is the intertext of the sample to be annotated?) as well as the *Dasein* of the annotator (What is his or her psychological state at the moment of annotation? is he or she happy? unhappy? worried about the future? How long has he been working? is he tired? hungry? etc. etc.).

All these questions seem trivial, but from the moment we address the sincerity of the annotator – a fortiori if we aim to make it a truth – it seems necessary to ask the question of their psychological condition. From the point of view of a psychologist, this condition could be as relevant as the sampling recommended by (Cabitza et al. 2023: 6885) “both in regard to their origin and culture as well as to their expertise and skills” – a necessary condition that could involve sociologists capable of correctly sampling the team of annotators. The more extra-linguistic criteria we include in the constitution of the dataset, the more we have to mobilize the corresponding sciences.

To sum up: noting the methodological and ethical limits to immanent corpus annotation, the perspectivist paradigm proposes to abandon corpus annotation in favor of annotation of annotators (by sampling) and their sincere perceptions of the dataset. But the perspectivist literature only partially solves the ethical problems it raises, it only notes some of these problems and, rather than fighting discrimination, it makes it visible, measurable and computable by tagging variations. The real ethical problem with NLP is the management of work.

#### 4. Constructing Rather Than Annotating

Let's change perspective.

What if the problem wasn't the inclusion of annotators, but the very principle of corpus annotation? Let's try to justify our change of perspective by 3 main proposals:

##### **A. Give a philological value to the dataset, i.e. turn it into a corpus**

With the widespread use of machine learning, the concept of the corpus as a built set of texts designed for a specific task has been greatly devalued in NLP, in favor of the dataset, i.e. a collection of data often “scraped” from the Internet with little preliminary characterization (mainly some sources and keywords), and whose main characterization comes from the annotation itself. Moreover, with deep learning, datasets are now quantified in terms of gigabytes rather than linguistic units (words, sentences, texts). However, the construction of a corpus is the first scientific act involved in the definition and establishment of the object of science. Consequently, building a corpus is a high-level activity (so much so that it can take years for a linguistics PhD student!). Our first proposal for a change of perspective is this: we need to give value to the construction of the corpus (Dusserre and Padró, 2017). Constructing a corpus requires setting up a task, identifying and examining sources, selecting texts, verifying them, characterizing them: author, type of discourse, textual genre, etc. (Biber and Conrad, 2009). Generally speaking, the reuse of datasets for new tasks, commonly accepted as good scientific practice, is not always acceptable when we're talking about corpora (a problem that also arises for gold standard corpora). Indeed, if we think in terms of corpus rather than data, sharing is only acceptable

if the new task aims to pursue or verify the objectives of the previous task. If the aim is to recycle a dataset that is more or less suitable for a new task, the corpus is downgraded to a dataset, because the intention is inherent in the corpus.

Just as the concept of corpus is underused in NLP, that of text is also poorly understood. A text is not just a collection of sentences or a bag of words, it's a semiotic object produced in a particular social and cultural context, corresponding to an enunciative project and containing interpretative rules, which are conditioned by its intertextuality (Mayaffe, 2002) (i.e. the set of texts linked by a text, for instance, in the case of an interlocution, usual on the social web). Not taking intertextuality into account when collecting texts is to deprive future annotators of their interpretative clues, because the corpus is not a resource but a multi-scale contextualization of observable phenomena (Mayaffe and Viprey, 2008). Text is the first interpretable semantic unit if we have to select a first level of annotation.

### **B. Focus on coarse-grained annotations or on intrinsically annotated corpus**

The more fine-grained the annotation, the greater the number of annotations, the greater the risk of variation. If variation is a quality in the perspectivist paradigm, this is without considering the ethical biases we discussed earlier: variation can indeed be an effect of arduous working conditions. A coarse-grained annotation is one that requires a longer, more reflective - less reflex - less manual, more objective intellectual work of interpretation. For example, when it concerns hate speech, it can be interesting to collect all the tweets of an author identified as a habitual hater, rather than a sample of his or her explicitly hateful tweets, which allows annotators to safely distance themselves from hateful content. Moreover, hateful sentiments are not necessarily expressed in hateful words (Eensoo et al., 2015). Fine-grained annotation is unfortunately confused in NLP with word-level annotation, but the word is only a minimal semiotic unit carrying lexical meaning, not text sense.

An ultimate coarse-grained annotation would be to use corpora that are intrinsically annotated, i.e. corpora that, by their very constitution, already contain metadata that can be used as annotations (e.g. gender declared, age, opinion, city, etc.), as is the case, for example, with corpora of polarized comments with ratings. One of the most widely cited NLP articles in sentiment analysis, (Pang and Lee, 2002), uses this type of corpus. The use of intrinsically annotated corpora makes it possible to concentrate efforts on higher-level (and therefore more skilled, better-paid) annotation tasks.

### **C. Use computer-aided corpus analysis**

"One could determine the different ages of a science by the technique of its measuring tools " said the philosopher Bachelard (1938: 216, our translation). It seems astonishing that corpus annotation is still today an irreducibly manual task when numerous "distant

reading" tools exist and have been used for over 40 years now, first in the Statistical Analysis of Textual Data and then Digital Humanities community (e.g. Compagno, eds., 2018; lezzi, eds., 2018; Lebart et al. 2019). This is indicative of a stubbornly marked divide between the NLP scientists and the humanities. Yet, humanities provide several bottom-up analytic methods, and tools and heuristics for corpus description. Textometrics tools (Heiden, 2010), combined with ad hoc text semantic theories (Pincemin, 2010; Rastier, 2018), can be used to generate annotations that can then feed philologically built datasets, for instance for opinion mining (Eensoo et al., 2015; Valette, 2018; Baiocchi, 2019). Finally, work combining semiotic theory and word embeddings aims to formally model the 'reader' (Sanna and Compagno, 2020) – a relevant idea in the context of perspectivism.

By combining statistical measurement, corpus analysis and semantic theorization, computer-aided corpus analysis creates an (objectifying) distance between annotator and *empirie*. What's more, this distance protects annotators from prejudicial content to which they might be brutally subjected, without the necessary reflexivity to withstand it. What's more, the proposed corpus analysis reclassifies the creation of value as intellectual work. This twofold distancing, by theory and by tool, leads to an objectification of the phenomena studied. Objectivity could wrongly be seen as contrary to the perspectivist paradigm. On the contrary, it is a construct that emanates from experiences, not subjectivities: As the philosopher Bitbol (2014) says: "*Objectification means focusing attention on what, in experience, can be shared. This presupposes an education, which begins with the transmission of a common language and an ethic of truth*" (our translation, our emphasis).

## **5. Conclusion**

The perspectivist paradigm renews work on the production of learning data for machine learning. It is based on both the methodological obstacles observed in the community and an ethical position that appears to be inspired by the DEI (Diversity, Equity, and Inclusion) ethical management framework. In brief, the aim of this short position paper was to sketch a discussion about the ethical and methodological proposals of perspectivism which is mainly based on a managerial background, by using both a Marxian-based ethical and a methodological criticism inspired by Humanities proposals.

Our conclusion is that an ethical approach to manual annotation of corpora, today organized according to the principles of production-line work, would be to reclassify it as scientific corpus analysis, in order to revalue this necessary interpretative work.

## **6. References**

Aroyo, L. and Welty, C. (2015). Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1): 15–24.

- Astor, D. (2020). *Perspectivisme. Leibniz, Nietzsche, Whitehead, Deleuze*. Phd Thesis, Institut Polytechnique de Paris: 471 p.
- Bachelard, G. (1938). *La formation de l'esprit scientifique*, Paris: Vrin.
- Baiocchi, M. (2019). *Diversité et recommandation : une investigation sur l'apport de la fouille d'opinions pour la distinction d'articles d'opinion dans une controverse médiatique*, PhD thesis, Université de Montréal: 509 p.
- Biber, D. Conrad, S. (2009). *Register, genre, and style*. Cambridge: Cambridge University Press.
- Bitbol, M. (2014). L'expérience d'objectiver Ou comment vivre en première personne la possibilité de la troisième. In N. Depraz (eds), *Première, deuxième, et troisième personne*, Presses Universitaires de Rouen.
- Cohen, K.B., Fort, K., Adda, G., Zhou, S., Farri, D. (2016) Ethical issues in corpus linguistics and annotation: pay per hit does not affect effective hourly rate for linguistic resource development on Amazon Mechanical Turk. In *LREC Int Conf Lang Resour Eval*. W40: 8-12.
- Compagno, D. eds. (2018). *Quantitative Semiotic Analysis*, Springer International Publishing: Berlin. 189 p.
- Dusserre, E., Padró, M. (2017). Bigger does not mean better! We prefer specificity. In *Proceedings of the 12th International Conference on Computational Semantics (IWCS) — Short papers*.
- Fellbaum, Chr. eds. (1998). *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press.
- Eensoo, E., Valette, M. (2015). Une méthodologie de sémantique de corpus appliquée à des tâches de fouille d'opinion et d'analyse des sentiments : étude sur l'impact de marqueurs dialogiques et dialectiques dans l'expression de la subjectivité. *TALN'2015*. Caen (France).
- Fort, K. Adda, G., Bretonnel Cohen, K. (2011). Amazon Mechanical Turk: Gold Mine or Coal Mine? *Computational Linguistics*, 37 (2): 413–420.
- Gray, M. L., Suri, S., (2019). *Ghost Work. How to Stop Silicon Valley from Building a New Global Underclass*. Boston: Houghton Mifflin Harcourt, 288 p.
- Heiden Serge. (2010). The TXM Platform: Building Open-Source Textual Analysis Software Compatible with the TEI Encoding Scheme. In *24th Pacific Asia Conference on Language, Information and Computation*. Sendai, Japan.
- Holzem, M. (ed) (2019). *Les sciences contre la post-vérité : Vérités citoyennes*. Vulaines-sur-Seine: Éditions du Croquant. 174 p
- Iezzi, D. F. Celardo, L., Misuraca, M., eds. (2018). *Proceedings of the 14th International Conference on Statistical Analyses of textual Data (JADT'18)*, UniversItalia, 2 vols.
- Lebart, L., Pincemin, B., Poudat, C. (2019). *Analyse des données textuelles*. Presses de l'Université du Québec. Mesure et évaluation. 510 p.
- Mayaffre, D. (2002). Les corpus réflexifs : entre architextualité et hypertextualité. In *Corpus*, 1: 51-69.
- Mayaffre, D., Viprey, J.-M., eds. (2012). *La cooccurrence, du fait statistique au fait textuel, Corpus*, 11.
- Noble, J. A. (2012). Minority voices of crowdsourcing: Why we should pay attention to every member of the crowd. In *Proceedings of the ACM 2012 conference on computer supported cooperative work companion*: 179-182.
- Planck, B. (2022) The “problem” fo human label variation: on ground truth in data, modeling and evaluation. Planck, 2022, In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*: 10671-10682.
- Park, J., Krishna, R., Khadpe, P., Fei-Fei, L., & Bernstein, M. (2019). AI-based request augmentation to increase crowdsourcing participation. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7(1): 115-124).
- Pincemin B. (2010). Semántica interpretativa y textometría. In Duteil-Mouguel C. and Cárdenas V. eds., *Semántica e interpretación, Tópicos del Seminario*, 23, Enero-junio 2010: 15-55.
- Rastier, F. (2018). Computer-Assisted Interpretation of Semiotic Corpora. In Compagno, D. eds., *Quantitative Semiotic Analysis*. Springer International Publishing: Berlin: 123-139.
- Rastier, F. (2019). Autour de la « post-vérité », de menaçantes convergences, in Holzem, M., eds. (2019). *Les sciences contre la post-vérité : Vérités citoyennes*. Vulaines-sur-Seine: Éditions du Croquant: 23-57.
- Sanna, L., Compagno, D. (2020). Implementing Eco's Model Reader with Word Embeddings. An Experiment on Facebook Ideological Bots. *Proceedings of the 15th International Conference on Statistical Analyses of textual Data (JADT'20)*, Toulouse, France.
- Steiger, M., Bharucha, T., Venkatagiri, S., Riedl, M.J., Lease, M. (2021). The Psychological Well-Being of Content Moderators: The Emotional Labor of Commercial Moderation and Avenues for Improving Support. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing System (CHI '21)*: 1-14.
- Taylor, F. W. (1919). *The Principles of Scientific Management*. Harper & Brothers – via Internet Archive (Prelinger Library)
- Valette, M. (2018). Elements of a Corpus Semantics for Humanities. Application to the Classification of Subjective Texts, In D. Compagno, eds., *Quantitative Semiotic Analysis*, Springer International Publishing: Berlin: 141-150.