Natural and Artificial Intelligence; Natural and Artificial Language

Diana Santos
Linguateca
University of Oslo, Norway

Abstract

This text starts by discussing what it means to be intelligent for humans and machines, what is the purpose of language, and how is human language fundamentally different from artificial languages. It presents the issue of values as one inescapable property of human language, and of human categorization in general, after reviewing five distinctive caracteristics of natural language. Then it proceeds to discuss static word embeddings, raising two questions: is the wisdom of the crowd an appropriate justification for using the underlying large text collections? And have the differences between languages been taken into account when intrinsically evaluating Portuguese word embeddings?

2012 ACM Subject Classification Computing methodologies → Machine learning; Computing methodologies → Natural language processing

Keywords and phrases Artificial Intelligence, Natural Language Processing

Digital Object Identifier 10.4230/OASIcs.SLATE.2021.1

Category Invited Talk

Acknowledgements I want to thank Alberto Simões for inviting me to SLATE.

1 Introduction

This invited talk gave me the opportunity to reflect on artificial intelligence (AI) and natural language processing (NLP) after ca. 35 years of having been introduced to both in my student years in the 80'es.

In fact, I have worked for more than 35 years in natural language processsing, and was a student of João Pavão Martins, who belonged to the first generation of artificial intelligence scholars in Portugal. After having launched Linguateca in the end of the 90's, with the goal of fostering R&D in the computational processing of Portuguese, I have taught at the Faculty of Humanities of the University of Oslo for the last ten years.

Therefore I thought it appropriate to share with you some thoughts about natural language and artificial intelligence, widening the scope of influence to philosopy, psychology, statistics, medical history and sociology, while somehow following up a previous discussion of the specificity of natural language as opposed to artificial languages, which I prepared for PROPOR and SBLP in 2006 [17].

2 A note on terminology

SLATE has a very interesting approach to language seen from a computer science view: it divides the field among human and computer languages, or better, in a tripartite way: computer-computer languages, human-computer languages, and human-human languages.

10th Symposium on Languages, Applications and Technologies (SLATE 2021).
Editors: Ricardo Queirós, Mário Pinto, Alberto Simões, Filipe Portela, and Maria João Pereira; Article No. 1; pp. 1:1–1:11
OpenAccess Series in Informatics
OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

However, there's no denying that *natural language* is by far the most common term when dealing with human language. And, considering the four terms in my title, the most common of all is obviously *artificial intelligence*. Interestingly, the only call in SLATE which mentioned intelligence was CCL, which used instead *computer intelligence*.

Now, the purpose of my title was exactly to point out that, no matter the fact that *natural* and *artificial* are antonyms, natural language does not work as expected, that is, clustering natural intelligence with natural language, and artificial intelligence with artificial languages. Quite the contrary, the pseudo-term *natural intelligence* does not exist (except perhaps as a pun): one always contrasts *artificial intelligence* with *human intelligence*. And in computer science one generally uses just *languages* to refer to programming languages. I would say that *artificial languages* are rather those natural languages invented by people, like Esperanto.

So, and despite the apparent symmetry and plausible analogy among the two pairs, in reality natural language is much more complicated and less predictable. It is the terms *natural language* and *artificial intelligence* which are linked: often natural language processing is considered a part of AI, or even the hallmark of AI – remember Turing's test.

Anyway, terms have a history, and evolve, as language in general does. Therefore they are not defined once and for all and, in fact, they mean very different things forty years ago and now.

Let me then start trying to discuss intelligence.

3 A linguistic inquiry of intelligence

While the concept of intelligence is something of the realm of philosophy or psychology, the linguistic approach to meaning is to investigate how the word is used, and to what it is applied, as Firth aptly expressed in his dictum "a word is defined by the company it keeps".

I will be using in this talk English and Portuguese, so I should hasten to say that the meaning depends on language:

- in English, there is *intelligent* and *clever* and *smart*, *cunny*, etc.
- in Portuguese, there is *inteligente* and *esperto* and *sagaz*, *perspicaz*, *astuto*, etc.

It is rather the clusters represented by these quasi-synonyms that have some sort of translation relation. The particular shades of meaning in each language are language-specific, and virtually impossible to precisely compare.

In any case, and restricting now our look at the words *intelligent* and *inteligente*, it is not hard to see that people use these words in a much broader way than dictionary definitions would foresee. I give just some random examples, to illustrate that its use is very dependent on context, and that intelligence can be assigned to a device or to the person who devised it, and it can even simply indicate a better solution compared to a previous one:

- Another classic example of such parallel communications is the device that we will use as the example in this section, the ***intelligent liquid crystal display (LCD) module***. (*Introduction to Mechatronic Design*)
- newspaper parlance: *sistema de semáforos inteligentes...* (intelligent traffic lights)
- gadgets *smart phones* or techniques *smart queries*, or even environments *casas inteligentes* (intelligent homes)

Anyway, the most important point I would like to make about the concept of intelligence when it appears in *human inteligence* and *artificial intelligence*, is that, again, the meaning of the compounds does not necessarily imply that we are talking about the **same** concept of *intelligence*. One could of course point to the paralellism of human vs. artificial/machine and conclude that the difference is who possesses/displays intelligence. But it is not hard to see that compounds can also completely change the meaning of the noun, like in *human nature* vs. *tropical nature*.

And this reminds me of one of the old lessons of AI: it is not by copying birds that mankind managed to fly. It was by understanding the physical principles birds use.

For most of the activities we use the word *intelligent*, either machines or humans are considered intelligent when they do it, but most often than not, not both. Let me give some examples:

- RoboCup – soccer playing
  While this is a display of intelligent robots, who has ever used the word intelligent when describing how one's child is good at soccer?
- Chess or Go playing
  Human champions of chess or Go are usually considered highly intelligent, but computers have surpassed them by (mainly) the capacity to foresee thousands of outcomes.
- Encyclopaedic answer
  The possibility of knowing a lot of facts in the past was a hallmark of a learned, intelligent person, but the access of a computer to more facts than a human can store has given a computer the lead (cf. the Jeopardy contest)
- Poetry writing
  Poets like Camões or Shakespeare are considered geniuses, so their poetic intelligence is high, although sensibility and way with words are ingredients equally necessary. So far machine poetry is considered funny and sometimes interesting, but not really intelligent.
- Lying to protect other's feelings
  Emotional intelligence like the one displayed by people who choose different versions of a story, or even lie, something that most humans do, is so far outside the realm of computer communication.
- Visually identifying a cancerous tissue
  This is an activity that required specific training for radiologists and other medical staff, and that has been partially taken over by image recognition systems built based on machine learning over large amounts of data. However, an apparently better-than-human system got it totally wrong with given images produced by a different vendor, something that would not fool a human. That forced us to caution "blind" machine learning. See [11] for more information on this.
- Recognizing people in the street
  These capabilities are different from human to human, but may be called social intelligence. The recent surveillance attempts by authoritarian governments using AI systems to identify populations also call into attention this ability.
- Identifying a dialect
  A related ability, that of being able to detect a specific dialect, which is being appropriated by intelligent computers, can also be considered a sign of human intelligence.

The examples could be multiplied at will. Their point is basically that not all these activities are equally praised – and are obviously differently implemented – in humans and in computers.

I would like to here support a new paradigm called hybrid intelligence (HI), beautifully presented by Frank van Harmelen in his 2020 keynote at IC3K, entitled "Hybrid Intelligence: AI systems that collaborate with people, instead of replacing them"[24]. He suggests augmenting human intellect and capabilities instead of replacing them, with the aim of achieving goals unreachable by either humans or machines alone. Van Harmelen argues that AI is unaware of norms and values; reasons; and contexts, and that it is absolutely necessary to have explanations in human society, so that decisions can be disputed. And explanations, he argues, need to be grounded on values, norms, motives, committments and goals.

4 What is intelligence; what is language

I would briefly suggest four criteria for intelligence, namely

- Learning
  One of the most mentioned properties of intelligence is learning. An intelligent person learns from failure. But is learning enough to become intelligent? If we do not know what a machine has learned, we cannot predict, or explain, unexpected failures. If one does not critically choose from whom to learn, one can learn things that are wrong, incorrect, even dangerous if one has no critical sense...
- Knowledgeability
  At least for a person, the more she knows, clearly the more intelligently she can behave. Computers can "know" many more facts than humans, but there are several problems with facts: they depend on theory, they are not consensual, and it requires intelligence to generalize over them. Humans are extremely good at creating generalizations and decide, with incomplete information. Plus, knowing is knowing where to find. Who are the authorities. Who to trust.
- Alternative worlds
  Humans very easily devise alternative worlds: if / as if. We often decide based on imagining different realities, computing consequences that are just thought, even pursuing "impossible" paths. Humans can define intensional concepts, computers so far only extensional ones.
- Context awareness
  An intelligent being/device/system changes behaviour depending on the context, reacting to the environment/situation in a proper way. This is perhaps the property that has been endowed most artificial systems, but to a certain extent only. Human reactions embody assumptions, and beliefs. Humans are able to revise and change their beliefs, and to reason with incomplete knowledge. One of the properties that allow us to do it are emotions [20]. See for example the paradigm of affective computing proposed by Picard [12] to endow machines with some emotions.

But let me point out that, anyway and pace Turing, mimicking a human is different from being a human.

If we look at language – and let me emphasise again that we only talk of *natural language* when we process it with computers, therefore doing AI – the first remarkable thing is that when humans devise languages (programming languages) for computers (which is called language engineering), they endow them with properties very different from those they use in human languages.

In PROPOR 2006 I had a keynote [17] on what distinguished between natural and artificial languages, arguing for the following distinctive properties of atural language:

1. Metaphorical nature
2. Context dependency
3. Reference to implicit knowledge
4. Vagueness
5. Dynamic character (evolution and learnability)

Here I would like to add a sixth characteristic that I believe is extremely important, namely that natural language embodies **values**.

To argue for this, I start by discussing what natural language is for, borrowing heavily from Ellis [2] and Steiner [22]: language is a pre-requisite, or the way, for humans to understand the world (through characterization), language allows one to do things (with others), and language is instrumental to create a shared community, as well as to put others (the onew that do not speak our language) outside.

Human language is a source of power, as sociolinguists and sociologists have argued for a long time.

But another characteristic that is not so much discussed, and which may at first look like a truism, is that it is human-centered, in that the values that constitute it are all relative to Man.

Let me give two examples of natural language concepts that illustrate this. Take first disease. As scholars of medical studies have pointed out, disease is a concept related to humans. As Sedgwick [21] puts it:

> There are no illnesses or diseases in nature. (...) The medical enterprise is from its inception value-loaded; it is not simply an applied biology, but a biology applied in accordance with the dictates of social interest.
> (...) All illness, whether conceived in localized bodily terms or within a larger view of human functioning, expresses both a social value-judgment (contrasting a person's condition with certain understood and accepted norms) and an attempt at explanation (with a view to controlling the disvalued condition).

Only the states which have an undesired effect for the goals Man pursues receive this description. In other words, a disease is something natural that is considered bad for humans (or pets or crops).

Take now the concept of weed: it is a plant that is considered bad for human gardens (or for the concept that humans have of gardens, or of plants in general). We know that this is not a biological property of a plant, it is a value that humans attach to it.

Generalizing, the concept of good or bad is something that pervades our language. Values are essential to communicate among humans. And they are absent from computer languages.

Although we can generalize to language in general (not only human language) that the purpose of language is knowledge representation, and communication, in order to do things with others, and to inform or disinform others, it is only among humans that values are shared and communicated.

As a side remark, and to highlight the importante of values for humans, consider another keynote at SLATE, on "What Programming Language Design Taught Me About Life". In the abstract, Pitman [13] states:

> I came to see languages as much more complex entities than mere functional behavior or stylized syntax. Languages are about community and shared values – and not just the kinds of values that get returned from a function call. The choices a language designer makes will attract certain users and alienate others

So, even in the (in principle, value-free) design of programming languages (to communicate with computers) the issue of values is paramount. (And it is (also?) with humans that language engineering is preoccupied with, not (only?) computers.)

Following the path of looking into other keynotes at SLATE, in last year's on "How Humans Succeed While Failing to Communicate", Graat [6] concludes:

> The task of making a computer understand human communication therefore seems to be the hardest thing to do.

My answer to this is that maybe it is not necessary that computers understand our communication, maybe the right choice is to communicate differently) with them (as we so far have done). And communicate other things.

Because, I would argue that to assign value is something absolutely human: good and bad do not exist in nature or reality.

In order to evaluate, you have to compare with something else. Usually, human judgement.

But – and this is a highly relevant detail – not all judgements are consensual. All of us are aware of ethical paradoxes, different legal opinions, etc.

In fact, and even in a more general sense, cultures have been defined (by Delfim Santos [16]) as different rankings of values.

The bottom line is that human language always includes values, and these values are inherently human.

5 Word embeddings and the wisdom of the crowd

I turn now to one specific technology used in NLP, which one may say has come to dominate NLP in the last years: word embeddings, a form of representing context based on co-occurrence. Based on machine learning over big text collections (the crowd) – see [18] for looking critically at size.

I am obviously not the first one who looks critically at this technology. In fact, an excellent presentation this year by Rada Mihalcea [10] has voiced at least the following concerns: unpredictable, unstable, crowd-dependent, human-ununderstandable, climate-unfriendly, corporation-owned.

My impression is that often one uses word embeddings in a way aptly described by the Portuguese expression *atirar o barro à parede* (let's see if it sticks, if it works), without even providing a rationale for using them.

But another criticism I want to raise here is that people "playing" with word embeddings do not take different languages seriously.

Let us first investigate what underlies the use of word embeddings: the assumption that the larger the set of texts one uses, the more one (system or person) learns – in other words, that quantity leads to quality. This is actually backed by a scientific observation done by statisticians (Galton [3]) more than one hundred years ago, with the name *the wisdom of the crowd*, which states that the median estimate of a group can be more accurate than the estimates of individual experts.

But the application of this "law" has some flaws, as I will proceed to argue.

First, it is not the size itself of the crowd that is the relevant factor: one has to ask the right crowd. What if had asked SLATE's audience two simple questions, one dealing with the meaning of a Norwegian expression, and the other about the family of an Angolan politician? Assuming that no one in the audience was aquainted with the politics of Angola, neither knew Norwegian, no matter how many answers I would get, I would not trust them, compared with those provided by one single Norwegian speaker, or one single Angolan historian.

Second, if one (person) reads a lot of texts, s/he will be able to understand that some texts oppose the others, or are based on others, and s/he will use her or his intelligence to make sense of what s/he read. Not taking everything at face value, or considered equally true or reliable.

Last but not least, recent studies in psychology have shown that social pressure undermines the effect of the wisdom of the crowd. The crowds studied by Galton had been independently asked, and had no idea of what the other respondents had answered.

Lorenz et al. [9] contend, after performing some interesting experiments, that

> [a]lthough groups are initially "wise," knowledge about estimates of others narrows the diversity of opinions to such an extent that it undermines the wisdom of crowd effect in three different ways.
> - The "social influence effect" diminishes the diversity of the crowd without improvements of its collective error.
> - The "range reduction effect" moves the position of the truth to peripheral regions of the range of estimates so that the crowd becomes less reliable in providing expertise for external observers.
> - The "confidence effect" boosts individuals' confidence after convergence of their estimates despite lack of improved accuracy.

So, if we use (static) word embeddings in Portuguese, what is the crowd? The next tables show the size of some of the most common WEs (Table 1)[1], and how many words/entries they share with each other (Table 2).

Table 1 How many words have embeddings; removing words with numbers.

	Size in words	without numbers
nilc	929,606	910,215
nlx	873,910	752,001
pt-lkb	202,001	201,877
cc	2,000,000	1,665,247
base	1,052,405	984,226
lemas	1,613,937	1,374,196

Table 2 Removing words with numbers, how many words are shared.

	nilc	nlx	pt-lkb	cc	todosbase	todoslemas
nilc	–	296,157	75,285	380,252	**536,720**	158,813
nlx	296,157	–	58,716	**596,091**	231,931	304,249
pt-lkb	75,286	58,726	–	70,217	**75,311**	65,900
cc	380,252	**596,091**	70,217	–	365,048	456,284
base	**536,720**	233,795	75,301	281,097	–	314,314
lemas	158,813	304,249	65,900	**390,415**	281,097	–

In Tables 3 and 4 one can appreciate the 15 closest words to the Portuguese word *inteligência* given by some of these models, as an illustration of what they can do, and of what they may not do.

But the question remains: who should answer? Should one use the crowd, that is, a lot of different people who wrote different texts in different contexts and take the average/median? Or should texts about a particular subject be used when one is interested in that particular subject?

[1] See, respectively, [8, 14, 7, 4] for nilc, nlx, cc, and pt-lkb and [19] for base and lemas.

It has been argued by Tshitoyan et al. in an interesting letter to *Nature* [23], that

> models trained on the set of all Wikipedia articles (about ten times more text than our corpus) perform substantially worse on materials science analogies. Contrary to what might seem like the conventional machine learning mantra, throwing more data at the problem is not always the solution. Instead, the quality and domain-specificity of the corpus determine the utility of the embeddings for domain-specific tasks.

Here we are back to one of the most fundamental issues of (at least old) artificial intelligence: aren't we confusing language knowledge with world knowledge? Domain vs. general knowledge?

What do word embeddings encode? World knowledge, or language knowledge? Should they be used for expert tasks, or for general language? AI has always excelled in specific domains, domain expertise, and not in general-purpose reasoning or language.

If we look at the closest words to the Portuguese word *inteligência* given by different ways of employing the several WEs available, we see that they are widely different depending on the method, and also depending on the textual base. (The last line gives the range of similarities.)

Table 3 The closest words in models with 300 dimensions: *lemas* indicates that the word embeddings were created after lemmatizing the corpus, and *lemasmwe* after lemmatizing and connecting multiword expressions, both done by PALAVRAS [1].

base, w2v	lemas, glove	lemasmwe, fasttext	cc, fasttext
intuição	senso	ininteligência	inteligencia
imaginação	habilidade	contra-inteligência	Inteligência
sabedoria	mente	desinteligência	ainteligência
criatividade	criatividade	Inteligência	perspicácia
sagacidade	experiência	inteligência-do-cinema	inteligência.A
habilidade	talento	inteligibilidade	deinteligência
perspicácia	certo	inteligXncia	intelecto
astúcia	capaz	deligência	contra-inteligência
intelecto	imaginação	inteligencia	intelegência
destreza	excelência	inteligência	super-inteligência
0.68-0.57	0.85-0.76	0.95-0.85	0.77-0.62

While it becomes clear to the naked eye that fasttext finds the morphologically closest words, while glove and word2vec the semantically closest, and that the common crawl (cc) embeddings embed a lot of noise (misspellings and tokenization errors), there are all sorts of choices that provide different results. Further examples can be seen in Table 4.

In order to illustrate the same question now done to a specifialized set of texts, namely literary texts in Portuguese, I also present the results in Table 5.

But, of course, we are not going to evaluate a set of embeddings based on one word. There are two standard ways of evaluating (static) embeddings: intrinsic and extrinsic. The intrinsic approaches for Portuguese have been mainly the use of a set of analogies translated from English (the google set), and this is something I would like to fiercely criticize – assuming that word embeddings for Portuguese are conceived as a reflection of the language, and not primarily of the world.

In my opinion, it is wrong to use translated analogies in at least three counts:

Table 4 The closest words in further models with 300 dimensions.

nilc, w2v	nilc, glove	nilc, fasttext	nlx, w2v
inteligencia	habilidade	inteligênciaX	criatividade
inteligências	intuição	geointeligência	intuição
astúcia	criatividade	super-inteligência	imaginação
imint	força	contra-inteligência	sagacidade
criptológica	sabedoria	desinteligência	perspicácia
engenhosidade	senso	ciberinteligência	argúcia
laboriosidade	sensibilidade	foto-inteligência	sensibilidade
sagacidade	imaginação	contrainteligência	testreza
intuição	capacidade	superinteligência	lucidez
imaginação	talento	deligência	inventividade
0.68-0.54	0.54-0.44	0.96-0.88	0.71-0.66

Table 5 The closest words in further models with 300 dimensions, trained on 50 millions of literary text, mainly 19th century.

words, w2v	words, glove	words, fasttext	lemas, w2v	lemas, glove
ciência	ciência	ininteligência	espírito	espírito
compreensão	sensibilidade	desinteligência	intelectual	talento
capacidade	compreensão	Inteligência	talento	'pírito
penetração	espírito	intelligência	inteligente	entendimento
perspicácia	energia	deligência	'pírito	capacidade
espírito	humana	inteligências	ciência	conhecimento
instrução	imaginação	vigência	capacidade	ciência
intelectual	capacidade	desinteligências	compreensão	compreensão
sensibilidade	superior	consCiência	aptidão	prático
concepção	experiência	conciência	faculdade	bastante
0.57-0.47	0.59-0.48	0.97-0.83	0.62-0.50	0.60-0.46

First, because it from the start only caters for what is common in both languages. If there were an interesting feature of Portuguese one wanted to assess that was not present in English, it would never pop up from translation.

Second, because it gives unwarranted importance to things, words and concepts that are key in English but marginal in Portuguese. The most discussed example, that of queen and king, is striking in that respect: no doubt that this is a much more relevant concept in English than in Portuguese. And capitals of US states are certainly part and parcel of the average American's knowledge, but very far away from what one would expect a Portuguese-speaking man in the street to know.

Finally, even the structure of concepts so basic as family are different in the two languages, given that *pais* means both parents and fathers, and *irmãos* can be translated by siblings or brothers. The translation of analogies involving these lexical items will not work, which means that also translation-related phenomena, like errors or translationese, will impact the "Portuguese" analogies.

There is, fortunately, one other set of evaluation data, created from Portuguese lexical resources, TALES [5]. But I hope that the NLP community for Portuguese will join efforts and develop many more sets of Portuguese-inspired evaluation data to evaluate NLP of

Portuguese (not necessarily only or even specifically for word embeddings), because different languages (and cultures) embody different data, categories, and assumptions, and reducing everything to English would amount to an epistemicide of terrible proportions (for the concept of epistemicide, cf. the sociologist Boaventura Sousa Santos [15]).

6 Concluding remarks

To conclude, let me restate my most important messages here:

Values are an extremely important feature of any human language. Human language is fundamentally human centered, on values towards mankind.

Artificial intelligence as a way to replace human intelligence is doomed to failure, because of the absence of values. Hybrid intelligence, namely cooperation between humans and machines, is what we should aim for.

Different (human) languages and different cultures, embodying different values and ways of seeing the world, should not be dismissed as a nuisance. Rather, they are (also) a manifestation of what is human, and what is intelligence.

References

1. Eckhard Bick. PALAVRAS, a Constraint Grammar-based Parsing System for Portuguese. In Tony Berber Sardinha and Thelma de Lurdes São Bento Ferreira, editors, *Working with Portuguese Corpora*, pages 279–302. Bloomsbury Academic, 2014.
2. John M. Ellis. *Language, Thought and Logic*. Northwestern University Press, 1993.
3. Francis Galton. Vox populi. *Nature*, 75(7):450–451, 1907.
4. Hugo Gonçalo Oliveira. Learning Word Embeddings from Portuguese Lexical-Semantic Knowledge Base. In Aline Villavicencio, Viviane P. Moreira, Alberto Abad, Helena de Medeiros Caseli, Pablo Gamallo, Carlos Ramisch, Hugo Gonçalo Oliveira, and Gustavo Henrique Paetzold, editors, *Computational Processing of the Portuguese Language - 13th International Conference, PROPOR 2018, Canela, Brazil, September 24-26, 2018, Proceedings*, pages 265–271. Springer, Cham, 2018.
5. Hugo Gonçalo Oliveira, Tiago Sousa, and Ana Alves. TALES: Test set of Portuguese Lexical-Semantic Relations for Assessing Word Embeddings. In *Proceedings of the Workshop on Hybrid Intelligence for Natural Language Processing Tasks HI4NLP (co-located at ECAI-2020)*, pages 41–47, 2020.
6. Jang Graat. How Humans Succeed While Failing to Communicate, 2020. Abstract of keynote at SLATE'20.
7. Edouard Grave, Piotr Bojanowski, Prakhar Gupt, Armand Joulin, and Tomas Mikolov. Learning Word Vectors for 157 Languages. In *Proceedings of LREC 2018*, pages 3483–3487, 2018.
8. Nathan S. Hartmann, Erick R. Fonseca, Christopher D. Shulby, Marcos V. Treviso, Jéssica S. Rodrigues, and Sandra M. Aluísio. Portuguese Word Embeddings: Evaluating on Word Analogies and Natural Language Tasks. In *Proceedings of Symposium in Information and Human Language Technology, Uberlândia, MG, Brazil, October 2–5, 2017*, 2017.
9. Jan Lorenz, Heiko Rauhutb, Frank Schweitzera, and Dirk Helbing. How social influence can undermine the wisdom of crowd effect. *PNAS - Proceedings of the National Academy of Sciences of the United States of America*, 108(22), 2011. `doi:10.1073/pnas.1008636108`.
10. Rada Mihalcea. The Ups and Downs of Word Embeddings, 2021. Seminar 19 May 2021. URL: `https://www.youtube.com/watch?v=33XtLnPDOC0`.
11. Luke Oakden-Rayner, Jared Dunnmo, Gustavo Carneiro, and Christopher Ré. Hidden stratification causes clinically meaningful failures in machine learning for medical imaging. In *Proceedings ACM Conference on Health Inference Learning*, page 151–159, 2020. `doi:10.1145/3368555.3384468`.

12 Rosalind Picard. *Affective Computing*. MIT Press, 1997.

13 Kent Pitman. What Programming Language Design Taught Me About Life, 2018. Abstract of keynote at SLATE'18.

14 João Rodrigues, António Branco, Steven Neale, and João Silva. LX-DSemVectors: Distributional Semantics Models for the Portuguese Language. In João Silva, Ricardo Ribeiro, Paulo Quaresma, André Adami, and António Branco, editors, *Computational Processing of the Portuguese Language: 12th International Conference, PROPOR 2016, Tomar, Portugal, July 13-15, 2016, Proceedings*, pages 259–270. Springer, 2016.

15 Boaventura de Sousa Santos. *Epistemologies of the South. Justice against Epistemicide*. Paradigm Publishers, 2014.

16 Delfim Santos. Sobre o conceito de cultura. In *Obras Completas, vol. III (1953-1966)*, pages 543–550. Gulbenkian, 2011. 3.a edição, revista e ampliada. Organização e notas: Cristiana de Soveral.

17 Diana Santos. What is natural language? Differences compared to artificial languages, and consequences for natural language processing, 2006. Invited lecture, SBLP2006 and PROPOR'2006, Itatiaia, RJ, Brazil, 15 May 2006. URL: `https://www.linguateca.pt/Diana/download/SantosPalestraSBLPPropor2006.pdf`.

18 Diana Santos. Grandes quantidades de informação: um olhar crítico, 2021. Invited lecture, HD-Rio 2020/2021, 14 April 2021. URL: `https://www.youtube.com/watch?v=Qi-3QzP0NxM0`.

19 Diana Santos. Palavras pulverizadas: algumas experiências, 2021. URL: `https://www.linguateca.pt/Diana/download/valoresPalPulv.pdf`.

20 Diana Santos and Belinda Maia. Language, emotion, and the emotions: A computational introduction. *Language and Linguistics compass*, 12(5), 2018. `doi:10.1111/lnc3.12279`.

21 Peter Sedgwick. Illness – mental and otherwise. *Hastings Center Studies*, 1(3):19–40, 1973.

22 George Steiner. *After Babel: aspects of language and translation*. Oxford University Press, 1992. 1st edition 1975.

23 Vahe Tshitoyan, John Dagdelen, Leigh Weston, Alexander Dunn, Ziqin Ron, Olga Kononova, Kristin A. Persson, Gerbrand Ceder, and Anubhav Jain. Unsupervised word embeddings capture latent knowledge from materials science literature. *Nature*, 571:95—98, 2019.

24 Frank van Harmelen. Hybrid Intelligence: AI systems that collaborate with people, instead of replacing them, 2020. Keynote at IC3K 2020. URL: `http://www.ic3k.org/Documents/Previous_Invited_Speakers/2020/IC3K_2020_KS_3_Presentation.pdf`.