



HAL
open science

Some thoughts about artificial stupidity and artificial dumbness

Jean Lieber, Jean-Guy Mailly, Pierre Marquis, Henri Prade, François Rollin

► **To cite this version:**

Jean Lieber, Jean-Guy Mailly, Pierre Marquis, Henri Prade, François Rollin. Some thoughts about artificial stupidity and artificial dumbness. *AI Communications*, 2023, 2023 (pre-press), pp.1–15. 10.3233/AIC-220322 . hal-04455911

HAL Id: hal-04455911

<https://inria.hal.science/hal-04455911v1>

Submitted on 13 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Some thoughts about artificial stupidity and artificial dumbness

Jean Lieber^{a,*}, Jean-Guy Mailly^{b,**}, Pierre Marquis^{c,***}, Henri Prade^{d,****} and François Rollin^{e,*****}

^a *Université de Lorraine, CNRS, Inria, LORIA, bâtiment Loria, B.P. 239, 54506 Vandœuvre-lès-Nancy, France*
E-mail: jean.lieber@loria.fr

^b *Université Paris Cité, LIPADE, 75006 Paris, France*
E-mail: jean-guy.mailly@u-paris.fr

^c *Université d'Artois, CNRS, CRIL, 62300 Lens, Institut Universitaire de France, France*
E-mail: marquis@cril.univ-artois.fr

^d *IRIT, CNRS, Université Paul Sabatier, 118 route de Narbonne, 31062 Toulouse cedex 09, France*
E-mail: prade@irit.fr

^e *Rollmops Unlimited Foolishness, France*
E-mail: frrollin@wanadoo.fr

Abstract. In a recently published book, the French writer and comedian François Rollin has discussed various aspects of the notion of stupidity, including artificial stupidity, the stupid counterpart of artificial intelligence. His claim is that a system of artificial stupidity is a system that provides wrong answers to any task it should solve, leading to absurd solutions in most cases. We believe that this claim is (at least partially) false and that designing artificial stupidity is not as trivial as it seems. In this article, we discuss *why* and *how* one could design a system of artificial stupidity. We believe that such a reflection on (artificial) stupidity can bring about some interesting insights about (artificial) intelligence.

Keywords: artificial stupidity, artificial intelligence

*Stupidity is often the ornament of beauty; it is it which gives to the eyes this dull limpidity of the blackish ponds and this oily calm of the tropical hours.*¹

Charles Baudelaire, *Journaux intimes* (1887)

1. Introduction

In a recently published book dedicated to stupidity, François Rollin² has written a chapter entitled “BA” (which stands for “Bêtise artificielle”), that can be translated into “AS”, for *artificial stupidity* [1]. Considering, as this chapter suggests, that AS is “hollow intelligence” implies a strong link between AS and AI, justifying a scientific interest by the AI community.

But how could AS be defined?

*Corresponding author. E-mail: jean.lieber@loria.fr.

**E-mail: jean-guy.mailly@u-paris.fr

***E-mail: marquis@cril.univ-artois.fr

****Corresponding author. E-mail: prade@irit.fr.

*****E-mail: frrollin@wanadoo.fr

¹The original Baudelaire’s sentence is “La bêtise est souvent l’ornement de la beauté ; c’est elle qui donne aux yeux cette limpidité morne des étangs noirâtres et ce calme huileux des heures tropicales.”

²According to Wikipedia, “François Rollin is a French actor, comedian, author and screenwriter.” He has also published about ten books for the French audience.

1 The thesis that François Rollin develops is that, while an AI system must perform sorting and adaptation on the
2 data, an AS system will skip one of these two types of operations (or both). From this, he concludes that an AS
3 system could be a system providing answers unrelated to the questions asked, and thus, that building such an AS
4 system should be easy. He illustrates this idea by some examples, such as the following (translated from French):

5 **Question:** How to solve a quadratic equation?

6 **Answer of the AS system:** By making all centenarians of the North of France eat six kilograms of artichoke
7 per day for eight weeks.

8 **Question:** How many stars are there in total in the Universe to the nearest 20%?

9 **Answer of the AS system:** 8.

10 This thesis has challenged us, and this article presents some thoughts on this subject. Stupidity is perhaps an even
11 more fascinating subject than intelligence, in view of the considerable number of works celebrating, questioning,
12 denouncing stupidity (such as, for example, [2–8]).

13 This article is organized as follows. Section 2 attempts to define the notion of stupidity. Then, some stupidities
14 produced by humans and by machines are discussed (Section 3). Afterwards, the two following issues related to AS
15 systems are considered:

- 16 • Why (knowingly) build an AS system?
- 17 • How to build such a system?

18 Sections 4 and 5 aim to address these two issues. After a tentative conclusion (Section 6), the first four authors leave
19 the floor to the fifth one: François Rollin himself.

20 This article is a revised and enriched translation of an article originally presented in the “Journées d’Intelligence
21 Artificielle Fondamentale” in 2022 [9].

22 **2. Trying to circumscribe the notions of stupidity and dumbness**

23 One possible way to approach “artificial stupidity” is to contrast it with “artificial intelligence” (AI). Before doing
24 that, in this section, we try to identify various forms of stupidity and foolishness in humans, to establish differences
25 between stupidity and dumbness.

26 *2.1. A terminology note*

27 This work is based on the reflections of French researchers (the authors of this article), with the objective of
28 studying the notion of *bêtise artificielle*, which we translate here as “artificial stupidity”. However, the French term
29 *bêtise* does not correspond exactly to the English term *stupidity*. The English and French vocabularies are both very
30 rich in terms to express various forms of impairments with respect to the behavior normally expected of a person
31 who reasons and makes decisions. Indeed, in French, the two main words *bêtise* and *stupidité* have slightly different
32 meanings, namely: *stupide*, the corresponding adjective, usually refers to behaviors where wrong conclusions are
33 drawn (while the right ones are easy to derive), while *bête*, the adjective counterpart of *bêtise*, rather applies to situ-
34 ations where people cannot conclude anything (although conclusions follow rather straightforwardly). Thus, *bêtise*
35 rather corresponds to English terms such as *dumbness* or *cluelessness*. It is also worth noting that *bête* in French
36 also means *beast* which suggests a comparison with animals (implicitly assumed to lack reasoning capabilities); in
37 English, if we except *donkey* (in French, *âne*) wrongly reputed to behave stupidly in both languages, there is no word
38 referring to animals for expressing some form of stupidity. *Foolishness* is another term that refers to the quality of
39 being unwise, stupid, or not showing good judgment, and has a rather broader meaning that encompasses craziness.
40
41
42
43
44
45
46
47
48
49
50
51

2.2. An anthropocentric viewpoint

Let us start by trying to understand the meaning of the term “artificial intelligence”. In such an attempt, one is faced with the tricky question of defining what “intelligence” means. To do this, we most often get around the difficulty by making a reference to the human being, which presents the form of intelligence that is probably the most accomplished, or at least the most polymorphous.

The Turing test [10] goes in this direction of such an imitation game: a process is seen as “intelligent” when the observer cannot predict whether he/she is interacting with a human or a machine.

Thus, the adjective “artificial” does not pose any specific concern in the elaboration of a definition (“artificial” only means that the intelligent process being analyzed is produced by a machine). Only the noun “intelligence” is questionable.

It can be seen that the comparison to humans to characterize the term “intelligence” existed before the Turing test. In particular, for René Descartes [11, fifth part], “intelligence” (in the sense of reason) is the prerogative of human: It is what separates humans from animals. Therefore, stupidity would be a contrario characteristic of being an animal. The position of Descartes is debatable, of course: reason is just one form of intelligence and, from a biology viewpoint, humans are just animals among others. Today, in the field of AI, when bio-inspired algorithms implementing a form of collective intelligence are designed and studied, most of the time the inspiration comes from animal societies (such as colonies of ants [12]), not from human societies. Is it because the latter have too complex behaviors to be modeled in a satisfactory way or because the value added by the collective is too limited to be significant? Everyone can have a point of view on this issue.

2.3. To be dumb

Stupidity (in the sense of the ability to be stupid or rather of the inability to not be stupid) has often been defined as a deficiency of intellectual capabilities (in particular, as the etymology of “intelligence”³ and the design of the IQ tests shows, the lack of ability to make connections). Unsurprisingly, references to intellectual development in humans has been used as a benchmark to establish a grade of stupidity [13] in the 19th century, where the idiot is the one who cannot communicate by speech, the imbecile is the one who is unable to read and write, and the fool is the one who does not reason or behave in a normal way. As Eugène Marbeau wrote “Dumbness does not understand; stupidity understands in the wrong way.”⁴

For qualifying an agent (whether human or artificial) as stupid, it is obviously not enough that it produces a stupid answer from time to time (otherwise, every humans would be stupid), but that this agent persists in her/his/its errors, even when a continuous interaction with the environment provides her/him/it enough good arguments so that, in principle, these errors should not be repeated. From this viewpoint, artificial agents have a higher predisposition to stupidity than human agents, simply because the interactions with the environment is typically more reduced (some of them have neither sensors nor effectors) but also because their strength is to be able to repeat without getting tired (and very quickly) the same treatments and, consequently, the same errors. It can be noted that this idea of stupidity linked to the persistence in error is one of the points studied in detail in the work of François Rollin cited in the introduction, even if it is hardly put forward in the chapter on Artificial Stupidity of his book.

Stupidity takes very varied forms; it is probably vain to want to characterize precisely what it is to be stupid. However, this idea of persistence in error can be retained as one of its characteristic features. The immutability of beliefs takes us out of the human condition and of its capability to question them whenever necessary. It brings us closer once again to the animal, in particular, the ostrich (whose head in the sand policy could be expressed as “see nothing, know nothing”), or even to the super-human, since, as the famous Latin quotation suggests, *errare humanum est, perseverare diabolicum*. Leaving the last word to Eugène Marbeau: “Foolishness believes itself very clever and doubts nothing”⁵

³Intelligence is borrowed from Latin *intelligentia*, which is derived from *intellegere* that is composed of the prefix *inter-* (“between”) and the verb *legere* (“to pick, to choose, to read”).

⁴The original sentence is “La bêtise ne comprend pas ; la sottise comprend de travers.” (in Comtesse Diane’s *Livre d’or*, 1886).

⁵The original sentence is “La sottise se croit très habile, et ne doute de rien.” (*Pensées et maximes diverses*, 1906).

2.4. Doing stupid things

Defining what is a stupidity is probably simpler than defining what being stupid means. So, let us embark on an attempt to identify stupidity!

A stupid answer can be seen as the result of a process of reasoning or decision making that has gone awry; this result is simply wrong or inadequate, because it is based on information that is also wrong, because it is generated by using fallacious rules of reasoning or by taking shortcuts in the reasoning, or because it is based on an irrational decision model. Assuming from the start the existence of an “intelligent” process that might have been appropriate allows us to separate stupidity from nonsense: the former would be a deviation from the result of the intelligent process whereas the latter would be produced by a totally irrelevant process. Thus, to the question “How much is 12345 times 67890?”, an answer like “Stéphanie de Monaco” would not only be stupid, but absurd.

To illustrate the production of stupid conclusions as a result of an inadequate intelligent process, consider the following two examples (one related to reasoning, the other to decision making):

*Cats are mortal, Socrates is mortal, therefore Socrates is a cat.*⁶

Eugène Ionesco, *Rhinoceros*.

Here, the information used is correct, it is the inference that is not: in this case, it is abductive, not deductive, and therefore can lead to erroneous conclusions (as is the case here).

Here is another example (a dialogue between two friends):

- *I read in a magazine that my alcohol consumption (one bottle of Bordeaux per day) is excessive and could seriously damage my health.*
- *I guess so. What did you do then?*
- *I have stopped reading.*

In this case, the decision taken is unexpected and probably inappropriate.

Deciding whether the result is incorrect or inadequate requires *comparison to a standard*. Depending on a certain “distance” to a reference point, the result in question may or may not be considered appropriate and the process that produced it will be considered (more or less) “smart” or “dumb”. The distance in question may then arouse various feelings in the person measuring it, from anger to laughter, or even wonder.

Note that the standard used to measure the adequacy of the result produced by the reasoning or decision-making process under consideration may be *universal* (mathematical logic allows, for example, the separation of valid reasoning from invalid reasoning) or it may not be (there is no single, consensual model to characterize what reasonable decision making is). Indeed, this norm is often *context-dependent* (thus non universal), in particular for the person who receives the message and decides to classify it as stupid; and this is a function of his or her own abilities, so that what is stupid for one individual is not necessarily so for another. To take an example for which there is no universal norm, consider the problem of preparing a glass of Scotch whiskey. Serving it “on the rocks” (i.e., with ice) is viewed as stupid by some and as normal by others.

Let us take an example for which there is a universal norm. It is again the question of multiplying 12345 by 67890. There is only one true result: the product is 838102050. The result ‘-1’ produced by individual A could be considered stupid by anyone with little mathematical background, the result 838102051 produced by individual B, also erroneous, will be considered as a calculation error by some and as stupidity by others (we can observe that as 67890 ends with 0, the product sought must also end with 0), and the result 838012050 produced by individual C, just as false, will probably be perceived as less stupid. It is this “distance” from the expected result (which may also be incorrect) that causes a simple error to acquire the status of stupidity.

It can be noticed that for this kind of tasks, humans, because of their cognitive limitations, are often dumber than simple machines such as a simple pocket calculator. It has been known for a long time that the capability of humans to correctly reason with numbers (and, in particular, with probabilities) is reduced. In his book [14], the American mathematician John Allen Paulos has introduced the neologism “innumeracy” to designate this inability (innumeracy is to numbers as illiteracy is to letters). An illustration of this idea in the context of the pandemic

⁶The original sentence is “Tous les chats sont mortels, Socrate est mortel, donc Socrate est un chat.”

1 period that started in 2020 can be given. During this period, some people were refusing to be vaccinated because
2 of the (proven) fact that there are more severe cases of COVID among the vaccinated ones than among the non-
3 vaccinated ones. This fact had been seen as a (fallacious) argument that would support the idea that the probability
4 of contracting a severe form of COVID when being vaccinated would be higher than the probability of contracting
5 a severe form of COVID when not being vaccinated (obviously, the error comes from the fact that the numbers of
6 vaccinated and non-vaccinated persons are not the same!). Such errors of reasoning can unfortunately lead to tragic
7 outcomes, as it happens in other settings such as the legal framework (see the fascinating book [15]).

8 More generally, it has been shown that human decision making frequently appears as irrational. Indeed, although
9 *rationality postulate systems* [16] are sets of postulate that are widely accepted as reasonable in decision making, it
10 has been shown that most human decision makers have made choices that deviate from these postulates in experi-
11 mental settings [17].

12 Obviously, if machines largely surpass humans when it comes to making calculations, their intelligence is much
13 less polymorphic: they do not know how to do anything else! In fact, the intelligence that is manifest at the execution
14 of these programs comes more from the humans that have designed these programs than from the machine. In
15 particular, a pocket calculator has no more conscience of the operations that it computes and of what is a number
16 than a clock has of the time, or than a facial recognition system has of what is a face and what are relevant parts of
17 a face.

18 One should also keep in mind that humans can perfectly consider as stupid some reasoning processes that are not,
19 since they correspond perfectly to a universal norm, when the latter exists (e.g. for deductive reasoning). Thus, the
20 validity of an inference depends only on whether its conclusion is true in all cases where its premises are true and
21 absolutely not on the truth of the premises and conclusion involved in the inference. In particular, when the premises
22 are inconsistent, any proposition can be deduced from them. The great logician Bertrand Russell gave this famous
23 example: “If $2 + 2 = 5$ then I am the Pope.” Indeed, it is known that the same value can be subtracted from the
24 two members of an equality, while preserving its validity, and this can be applied to the value 3 and the equality
25 $2 + 2 = 5$. Therefore, from premise $2 + 2 = 5$, it can be deduced that $1 = 2$. Now, the Pope and I are two persons.
26 So, if $1 = 2$, then we are the same person: I am the Pope! The formal validity of this reasoning in classical logic can
27 be established as follows: first, every logical consequence of any statement α must hold when α holds; second, $\alpha \vee \beta$
28 is a logical consequence of α whatever β ; third, when α is a contradiction, $\alpha \vee \beta$ is equivalent to β . Overall, every
29 β holds under any contradiction α . Accordingly, the reasoning followed by Bertrand Russell is valid, it respects the
30 canons of classical logic and yet it appears as meaningless, in particular because its premises α ($2 + 2 = 5$) and its
31 conclusion β (“I am the Pope”) have no connection.
32
33
34

35 **3. Stupidity is the most shared thing in the world⁷**

36
37 This section presents some interesting examples of stupidities: a human example of stupidity that seems chal-
38 lenging to capture by an AI system, an AI system built on very basic principles but that did have some success, and
39 sophisticated AIs producing (sometimes) stupidities (because of misuses or not).
40
41

42 *3.1. A nice example of human stupidity*

43
44 There is a story going around about a simple math problem that a student would have solved in a surprising,
45 creative, but wrong way (a nice piece of stupidity). This problem (after some reformulation) was the computation of
46
47

$$48 \lim_{x \rightarrow 0} \frac{1}{5x^2}$$

49
50 ⁷This title is inspired by the sentence of Descartes “Common sense is the most shared thing in the world” (original sentence: “Le bon sens est
51 la chose du monde la mieux partagée” [11]).

1 It seems that the student has remembered the solution of a similar problem: 1

$$\lim_{x \rightarrow 0} \frac{1}{8x^2} = +\infty$$

2
3
4
5
6 In order to give the reader the opportunity to imagine the student’s solution, the latter is given in appendix A: the remainder of this section is based on the assumption that the reader has read this Appendix. 6

7
8 Now, the following question can be raised: what competences an AI system requires to generate an answer such that this one in an “unintentional” way (i.e., an AI system that is not designed to generate stupid answers but generates some despite of its designer). We imagine that a system capable of generating the student’s stupidity must at the same time work on a representation of mathematical expressions (by a tree structure, for example) and on a visual representation, with links between these representations. From our perspective, building such an AI system would not be simple! 7
8
9
10
11
12
13
14

15 3.2. *Stupid AIs* 15

16
17 An AI program, or simply an AI, does only what it is programmed to do. One easily understands that a sophisticated program can do things that will be considered impressive, with regard to problem solving or image recognition, showing some kind of (artificial) “intelligence”. However, even very simple programs can fool us. This is the case, in particular, with the program ELIZA [18] which was able, in the mid-1960s, to simulate an interview with a therapist by simply rephrasing the “patient”’s sentences into questions, and then restarting the dialogue with preconceived sentences appearing to be related with the patient’s words. It is clear that when ELIZA was stating “I understand” because it was unable to provide a better answer in the context of the dialogue, it was bluffing its interlocutor, even if it was actually a quite stupid system. 17
18
19
20
21
22
23
24
25

26 This question of AIs that may seem “realistic” to their users, while relying on very simple (or even simplistic) techniques, has been relevant for a long time in video games. As stated in [19, Section 1.2.2], in the first era of video games, AI games were mainly based on predefined scripts. This limits the possible behaviors of non-playable characters (NPCs, which are the characters of the game controlled by the machine). The result of this is that, despite the huge progress made in the domain of AI for video games, there are still many players who complain about the stupidity of NPCs,⁸ including professional players who can identify the scripted phases of the game.⁹ As a consequence, an experienced player can anticipate the behavior of the AI. 26
27
28
29
30
31
32
33

34 3.3. *AI systems that produce stupid outcomes (or are stupidly used)* 34

35
36 *The Dong Mingzhu case, or Stupid and Mean AI.* Dong Mingzhu is a Chinese businesswoman, leading an important company in the field of air conditioning. To promote her company, she has organized a publicity campaign. Her portrait has been printed on city buses. However, in the People’s Republic of China, video surveillance is generalized. The large number of cameras and AI algorithms dedicated to face recognition allow the authorities to detect citizens who commit infractions or incivilities, such as crossing a road when it is forbidden to pedestrians because vehicles are passing. A social credit system is used, and when some incivility is noticed, the person who has committed it loses points; after losing too many points, this person can be denied some rights such as using public transportation. Dong Mingzhu quickly lost her points because each time a bus with her face on it was going in front of a security camera, the AI system concluded wrongly that Dong Mingzhu had crossed the road while it was forbidden.¹⁰ This scenario shows at the same time the great quality of the facial recognition algorithms (they are not intrinsically stupid) and also the huge stupidity produced by the AI system, coming from the lack of context-sensitivity in the decision-making. 36
37
38
39
40
41
42
43
44
45
46
47
48

49 ⁸<https://www.vice.com/fr/article/xya8kw/le-concept-de-vraie-ia-dans-les-jeux-vidos-nest-quun-outil-marketing> 49

50 ⁹<https://www.sofoot.com/rafsou-le-script-existe-dans-fifa-ce-n-est-pas-une-legende-484027.html> 50

51 ¹⁰<https://www.bbc.com/news/technology-46357004> 51

1 *When natural language processing gets wrong.* To be able to perform automatically (i.e., algorithmically) and in
2 a satisfactory way language processing tasks, such as dialogues based on animated conversational agents (or other
3 *bots*), summarizing or translating from one language to another, requires an ability to understand, which is not
4 acquired by AI systems that remain devoid of common-sense knowledge about the world in which we live.

5 Solving anaphora (in particular, satisfying pronominal references) is a difficult problem that requires knowledge
6 of the world (and the ability to reason about it) in order to be solved correctly. This problem is on the basis of
7 Winograd schemas [20], which can be used as variants of the Turing test.

8 Ask your personal assistant (a robot): “Could you please go shopping for me at the book store and buy the last
9 book by François Rollin, and if they have pencils, get 10.” If a short time later your robot comes back with 10 copies
10 of the last book by François Rollin, you should not be so surprised. Indeed, it is probable that they had pencils.
11 The mistake here comes from the fact that, for you, “10” refers to the pencils and not to the last book by François
12 Rollin, while your robot solves the anaphora differently: for it, “10” refers to the number of copies of the last book
13 by François Rollin to be bought, and not to the pencils.

14 Beyond assuming common-sense knowledge about the world, speech acts, as they are achieved by human beings,
15 are also impacted by a number of conventions ruling the way communication between people usually takes place in
16 talk exchanges. A chatbot that neglects those conventions may easily lead to provide stupid statements. As a matter
17 of illustration, let us consider the following dialog:

- 18 • *The chatbot:* “I have just been informed that your friend Elsa had a baby”.
- 19 • *You:* “Great! A boy or a girl?”
- 20 • *The chatbot:* “Yes”.

21 Although the response provided by the chatbot is perfectly correct from a logical point of view, it is clearly off the
22 mark. The problem comes from the fact that focusing on the locutionary content of the question is not enough. The
23 question has an illocutionary function: what you really want to know is the birth sex of the baby. Answering “Yes”
24 also does not comply with one of the Gricean maxims of conversation [21] that describe specific rational principles
25 observed by people who follow the cooperative principle in pursuit of effective communication. Indeed, it can be
26 assumed that everyone knows that the birth sex of any person is male or female. Thus, answering “Yes” is at odds
27 with Grice’s maxim of quantity, which requires that the response should give as much information as is needed, but
28 no more.

29 The automatic translation tools that are available online are wonderful tools for *shallow* translations. They are not
30 stupid at all, but it is possible to have fun using them as stupidity generators because of anaphoras that are wrongly
31 solved, but because of other problems as well.

32 For instance, let us try to translate in French the following sentence (which gives an argument whose justification
33 is quite questionable) with an automatic translation tool:

34
35 The beginning always precedes the end. It is normal, “beginning” starts with
36 a “b”, and “end” starts with an “e”.

37
38 When we translate it into French, we obtain:

39
40 Le début précède toujours la fin. C’est normal, le “début” commence par
41 un “b” et la “fin” commence par un “e”.

42
43 This sentence means that the French word “début” starts with a “b”, and the word “fin” starts with a “e”, which
44 is obviously not the case. After translating it, the justification of the argument is thus even more questionable. This
45 comes from the fact that the AI algorithm does not distinguish between *de re* and *de dicto* understanding of the
46 sentence, despite the quotation marks.

47 Of course, translating is a very difficult task, even for human beings (see, e.g., [22]); as a general rule, for a
48 somewhat complex text, no ideal, perfect translation exists, and thus the nonsense that is produced does not reflect a
49 great distance from a universal standard, but rather a difference from what is generally expected.

50 Among the sentences that are hard to translate are ambiguous sentences, especially those containing syntactic am-
51 biguities. Syntactic ambiguity, also called structural ambiguity (the scholarly term is “amphibology”), is a situation

1 in which a sentence may be interpreted in more than one way due to the ambiguous structure of the sentence. Thus, 1
2 the sentence “British left waffles on Falklands” (a classic of its kind) is ambiguous, since its meaning is not at all 2
3 the same depending on whether one considers that the verb is “waffles” or that the verb is “left”. Did the British 3
4 political left rambles indecisively about Falkland Island policy, or did the British physically leave waffles on the 4
5 Falkland Islands? Taken outside of a context that makes it possible to eliminate one of the two possible hypotheses, 5
6 this sentence is not really translatable since its very meaning is unknown. It is funny to watch how machine trans- 6
7 lation tools do with these kinds of sentence. For this, once translated into French, the result is surprising: “Gaufres 7
8 britanniques à gauche sur les Malouines”. This result does not correspond to any of the two meanings presented 8
9 above, but to a third one (at least as unlikely), which (back to English) would be given by “British waffles on the left 9
10 on the Falkland Islands”.

11 Many more examples of translations failing to preserve what is expected when achieved using automatic transla- 11
12 tion tools could be found. Since AI algorithms do not understand the texts they process, the possible hidden meaning 12
13 of these texts remains inaccessible, even when taking them into account is essential, such as in puns or other jokes 13
14 based on words. For example, the pun “How do you cut the ocean in half? With a see-saw!” once translated into 14
15 French becomes “Comment couper l’océan en deux ? Avec une balançoire !” The homophony between “see” and 15
16 “sea”, which is at the core of the pun, is fully lost in translation. In the process, the sentence loses its status as a 16
17 pun, as one can hardly identify what could be funny in the answer. Once translated into French, it becomes mainly 17
18 non-sensical, just as if the answer in English had been “With a banana!”

19 Here, it is important to highlight the fact that the errors produced by natural language processing systems, in 19
20 particular for translation, can be considered as stupid only according to a non-universal norm (as discussed in 20
21 section 2.4), in this case, from the context of human understanding the natural language used (or natural languages 21
22 for the case of translation) as well as puns in this (or these) language(s). 22
23
24

25 **4. Why would one build an AS system?**

26
27 The actual construction of an AS system can be seen as a rather futile exercise. Nevertheless, this section brings 27
28 together some reasons why one might want to embark on such an adventure. 28
29

30 *4.1. In order to pass the Turing test*

31
32 As we have seen, the Turing test aims at verifying whether an automatic system could fool humans by making 32
33 them believe that it is human itself, under certain conditions of interaction [10]. If such a system is too clever when 33
34 answering a question, it may fail the test. To take an example from section 2.4, if a machine interlocutor answers 34
35 the question “ $12345 \times 67890 = ?$ ” very quickly and correctly, he would fail in its deception. An answer that is far 35
36 enough from the result without being completely nonsense could fool the human examiner.¹¹ 36

37 Therefore, with the aim of deception induced by the Turing test, in order to develop a computer system that would 37
38 pass this test, it would have to include both AI and AS. Therefore, accepting the Turing test as a characterization of 38
39 intelligence involves accepting stupidity as a part of intelligence. 39
40

41 *4.2. Deliberate foolishness*

42
43 One could imagine an AI system that would deliberately produce nonsense, perhaps in the spirit of [23], which 43
44 notices that video games players usually find less enjoyable a game where the opponent (i.e. the AI algorithm 44
45 controlling the NPCs) is “too intelligent” and plays optimally, which makes it very hard (or even impossible) to 45
46 defeat. A slightly stupid opponent may then be more realistic, and more enjoyable (of course, a very stupid opponent 46
47 is also less enjoyable since then winning the game would become too easy). Finding the good level of stupidity in 47
48 this context is challenging. 48

49 One can imagine at least two other reasons for conceiving AS systems: 49
50

51 ¹¹Another answer from the machine might be: “You wouldn’t be, by any chance, a computer trying to pass its Turing test?” 51

- to lighten the atmosphere, to make the interlocutor laugh, if it is a pet robot.
- to deceive, or to disconcert a caller who, for example, would get angry.

These are obviously hard challenges for AI! Understanding and automatically generating funny stories is a very complicated problem since humor is a matter of belief revision (the punchline forces listeners of a joke to revise the context set up at the beginning of the joke) [24, 25]. It is also a matter of incongruity and violation of norms (incongruity often makes people laugh) [26, 27]. Saying stupid things deliberately to make people laugh requires to be smart enough. Obviously a program able to do such things would be a very special AS system, able to say stupid things on purpose.

4.3. Creative stupidity

If separating the true from the false is the Holy Grail of the scientific process and is indispensable to walk with confidence in this life (to paraphrase René Descartes), one should not be afraid to make mistakes. Foolishness is fruitful. We learn by doing errors.

Mistakes are, in fact, a source of creation. History is full of scenarios where blunders have led to successes, such as the “bêtises of Cambrai”, the tarte Tatin, or leaving the culinary field while remaining in the realm of serendipity, the discovery of the vulcanization process by Charles Goodyear in the XIXth century. Of course, in order for an error to be fruitful, intelligence is then required to recognize the existence of a potential to be exploited in the situation created by the error. Thus, an error in a mathematical proof will be fruitful if its discovery can lead to a new approach to the problem.

In humans, since early childhood, learning goes through phases of trial and error, through abusive generalizations that education, teaching, and reasoning allow (sometimes) to correct. Reinforcement learning implements this idea in AI: it goes through the exploration of situations that are deemed unsatisfactory or not. Observe that to be creative, reinforcement learning should probably allow a certain permissiveness in the estimation of the unsatisfactory character of a situation.

The same is true for the machine: inductive generalization, like abduction, analogy, or other forms of common-sense reasoning, does not preserve the truth. Learning, when it is not limited to rote learning (which does not offer the ability to take into account new situations), presupposes the risk of being wrong. The algorithms implemented in machines to learn aim at minimizing this risk in an empirical way, by relying on the available situations. They cannot completely eliminate it.

In the machine, it should be noted that the error committed can be more significant when the data is too small or of poor quality. This is the question of bias in the data used to learn, which can lead AIs to reproduce reprehensible behaviors,¹² or to generate predictions that are nothing but self-fulfilling prophecies, as in the case of the predictor *PredPol*. This system has been used by the Los Angeles Police Department to determine the areas to be patrolled; obviously, the LAPD could observe more crimes, misdemeanors, and infractions where it was sent to patrol than where it was not (and therefore was not in a position to observe them)! Therefore, the original bias in the data has led *PredPol* to predict more crimes and infractions in the areas under which LAPD was more present and to predict almost no such behaviors in the areas for which no data was available. Moreover, these predictions have tended to reproduce, or even strengthen, the bias in the data.

For stupidity to be creative, in humans and machines, the risk of being wrong must be accepted, but it must also be modulated and controlled by the implementation of mechanisms that allow us to question our beliefs and the decisions they lead to. This is essential in order to avoid the aforementioned *perseverare diabolicum* and thus strive not to remain stupid. Thus, a lot of work in AI has turned over the last 40 years to this multifaceted problem, including nonmonotonic reasoning (where conclusions change according to the hypotheses made, whether implicit or not), belief revision (roughly speaking, how to question one’s beliefs in order to change them sufficiently for restoring consistency without changing them too much?), or even learning techniques based on the decreasing of an error rate (such as the backpropagation algorithm for neural nets).

¹²See for instance <https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>

5. Designing an AS system

Suppose that one wants to build an AS system. For this purpose, two large fields of AI are considered first: section 5.1 considers the artificial intelligence/stupidity systems based on artificial reasoning, and section 5.2, the AI and AS systems based on supervised learning. Then two more specific domains of AI are considered. Case-based reasoning (CBR) can be considered at the intersection of the study of reasoning systems and supervised learning systems and is studied from an AS perspective in Section 5.3. Abstract argumentation is another domain of AI that can be viewed from this perspective (Section 5.4). The section ends with section 5.5 that presents a non-exhaustive list of AI domains and how they could be treated as AS domains.

5.1. Reasoning foolishly

To commit errors in reasoning is undoubtedly to reason foolishly (even if it is within everyone's reach). Deductive reasoning is largely based on the following scheme $\frac{p \rightarrow q \quad p}{q}$ (if p is true and if p implies q , then q is true). Allowing other schemes, such as, for example, $\frac{p \rightarrow q \quad q}{p}$ (which has an abductive flavor), or $\frac{p \rightarrow q \quad p}{\neg q}$, is more risky! According to the first, which does not offer guarantees of deduction, the observation of a consequence of p suggests that p is true, which is not stupid! Whereas the second scheme leads to incoherence, for sure. One can consult [28] for a study of fallacious arguments. A bad analogy, based on an overly hazardous parallel, likewise leads to absurdity: "The hedgehog is like a toothbrush, it has a stiff coat, and the toothbrush curls when scared".

Improperly using deduction also leads to absurdities:

- *Cheap apartments are rare.*
- *Everything that is rare is expensive.*
- *So cheap apartments are expensive.*

The problem here is that the second statement has exceptions (as the first one indicates), whereas it is treated as if it did not. The real stupidity may then be not accepting to abandon a false or absurd conclusion.

In the above situations, stupidity is due to the production of erroneous conclusions. On the contrary, dumbness may lead to the lack of conclusions from premises that would lead to obvious conclusions for most people. Drawing wrong conclusions and failing to conclude (when it is simple) are two different types of deficiencies that reveal unintelligent behavior.

If an AS system wants to look silly, producing false conclusions is not enough! It has to make it perceptible how it arrives at them. The top of silly reasoning is perhaps arriving at a correct conclusion by false inferences... unless this is a sign of intellectual virtuosity!

In a similar way, a perfectly correct reasoning process could be perceived as stupid simply because it is not clever, because it is too complex, because it requires complicated calculations, whereas there are much simpler approaches (as an example, we can think of the famous problem of the cyclists and the fly, see, e.g., <https://www.othot.com/blog/the-complexity-of-simplicity> and [29]).

Finally, if stupidity is also, as has been said, an inability to see, to create links between facts or statements of those facts, then, one can use ideas of systems with limited reasoning capacity [30] in the design of an AS system. For example, the system may be aware that p and $p \rightarrow q$ are true, without being able to form their conjunction and conclude q , e.g. [31].

5.2. Learning stupidities

A supervised learning system is based on examples where an example is an ordered pair $(x, y) \in \mathcal{P} \times \mathcal{S}$ with \mathcal{P} , the space of problems (or inputs), and \mathcal{S} , the space of solutions (or outputs).

In order to simplify the forthcoming discussion, it is assumed that the relation "has for solution" is functional and the underlying function is denoted by F : the non-noisy examples (x, y) of a supervised learning system verify $y = F(x)$. Thus, the objective of such a system is to approximate the function F : if \mathcal{TS} denotes the training set and

F_{TS} the function learned by a given supervised learning system, the error is computed by estimating the difference between the functions F and F_{TS} .

In order to build a supervised learning system for AS, one idea is to apply an AI technique of supervised learning, using stupid examples. So, following this idea, the issue that must be addressed is how to constitute TSS, the training set of stupid examples. Now, it is assumed that a *necessary* condition for (x, y) to be a stupid example is that (x, y) is an error, that is, $y \neq F(x)$. This raises two questions:

- Is the condition $y \neq F(x)$ *sufficient* to learn stupidity?
- If not, what other conditions should be stated?

These questions are not answered in this article: we only consider the necessary condition and examine what taking it as sufficient involves in several frameworks.

First, consider the binary classification (i.e. $|\mathcal{S}| = 2$) and let \bar{y} denote the element of \mathcal{S} that is different from $y \in \mathcal{S}$. Then TSS can be built from a training set TS (that is meant to be used by an AI system): TSS would be the set of (x, \bar{y}) where $(x, y) \in TS$. So, the binary classification system will learn a function F_{TSS} that, provided that the learning is successful, would be wrong most of the time, meaning that the function $\overline{F_{TSS}} : x \in \mathcal{P} \mapsto \overline{F_{TSS}(x)} \in \mathcal{S}$ would give a correct learning function:¹³ asking a stupid person what to do in case of a binary choice could be useful to know what not to do! For instance, if an image classification system is trained to distinguish photos of naked mole rats and photos of bicycles, with a stupid training set, it would take most of the time the photo of a naked mole rat for a photo of a bike, and vice versa.

When the size of \mathcal{S} grows, things may get different. Let us consider the case of regression, where $\mathcal{P} = \mathbb{R}^m$ and $\mathcal{S} = \mathbb{R}^n$, a setting used in many numerical machine learning techniques, such as multilayer perceptrons. Given $(x, y) \in TS$, how will $(x, y') \in TSS$ be chosen with the sole condition $y' \neq y$? If this is through a one-to-one correspondence $\varphi : y \in \mathbb{R}^n \mapsto y' \in \mathbb{R}^n$, this would amount to a situation similar to the one of binary classification as mentioned above ($\varphi^{-1} \circ F_{TSS}$ could be an approximation of F as good as F_{TS}). If the computation of y' is done randomly and regardless of the value of y then F_{TSS} would only be representative of the random generator on \mathcal{S} : this falls back to the hypothesis of François Rollin of an AS system that answers any random thing to a given question (as exemplified in the introduction). If the computation of y' is done randomly but stays close to y (e.g., for $n = 1$, y' is chosen randomly in $[y - C, y + C] \setminus \{y\}$, where $C > 0$ is a constant that is relatively “small”), this amounts to the framework of learning with noisy examples, with additive noise.

These few abstract examples illustrate the fact that the necessary condition “to be stupid is to be wrong” is not always sufficient to obtain a learning system in AS that is neither a “reverse” AI learning system nor a system that just gives nonsense answers.

5.3. SBR (Stupidity-Based Reasoning)

SBR would be case-based reasoning (CBR) producing stupid solutions. CBR consists in solving (or trying to solve) a target problem with the help of a case base, where a case is a representation of a problem-solving episode [32, 33]. In many applications, a case consists of an ordered pair (x, y) and thus coincides with the notion of example in supervised learning and is sometimes associated with additional pieces of information about the relation between x and y . For many applications, a session of CBR consists in selecting in the case base a case that is considered to be similar to the target problem (case retrieval step) and to modify it in order to solve the target problem (case adaptation step).

Two complementary ways to build an SBR system can be considered.

First, stupid supervised learning, as briefly studied in the previous section, can be a source of inspiration for building a base of stupid cases. This meets the recent issue of CBR using positive and negative cases. These negative cases can be compared to (even be assimilated to?) stupid cases. Negative cases can be acquired during the use of a CBR system: these are the cases corresponding to failures of this system. Using positive and negative cases for the learning of adaptation knowledge has improved the results of carrying out this learning with only positive cases [34].

¹³Under some symmetry assumptions, it should be possible to show that $\overline{F_{TSS}}$ and F_{TS} should have the same error expected value.

1 This tends to suggest that if both positive and negative experiences are memorized, with a distinction between the
2 two kinds, this can be useful for improving adaptation in CBR.

3 A second way to turn CBR into SBR would consist in changing the reasoning steps (as this is the case in the
4 general framework of Section 5.1), for example by selecting from the case base the case that is *least* similar to the
5 target problem: if the adaptation manages to modify the case-recipe of the *gratin dauphinois*¹⁴ in order to solve the
6 problem “I would like a recipe for blackcurrant sorbet”, the result could be surprising (and useful when you have
7 guests who linger too long). Another way to have an SBR system would be to use an irrelevant distance function (or
8 similarity measure) for case retrieval. For example, for the CBR system presented in [35], the use of a straightforward
9 distance function between problems would result in a stupid retrieval process that involves the generation of stupid
10 solutions. It is also possible to imagine a relevant case retrieval followed by a stupid adaptation: the example in
11 Section 3.1 illustrates this idea (for an SBR system solving limit computation problems).

12 It should be noted that these ideas are in line with the one presented by François Rollin of an AS system that
13 would abstain from sorting (here: making a stupid case retrieval) or adapting (here: making a stupid adaptation).

15 5.4. Argument-Based Stupidity

17 Although we have already mentioned the issue of reasoning foolishly in Section 5.1, we now turn to the topic of
18 formal argumentation, which is an instantiation of the general question of reasoning. In AI, formal argumentation is a
19 field dealing with the representation of conflicting information and the definition of non-trivial reasoning approaches
20 in presence of this conflicting information. Classically, reasoning is based on notions of logic and rhetoric: an
21 argument is made of a set of reliable pieces of information (the support) and one piece of information that can be
22 deduced from this support (the claim). If an argument a contradicts an argument b , and nothing contradicts a , then
23 we consider that a is acceptable and b is not. Argument-based reasoning approaches are based on various principles,
24 so it seems quite easy to define a system of argument-based stupidity as an argumentation system that would fail to
25 satisfy some of these basic principles.

26 Let us start with abstract argumentation [36]. In this context, we only use the relations between arguments to
27 determine which ones are acceptable, without any need to focus on the exact nature of the arguments (in particular,
28 their internal logical structure). An abstract argumentation framework is then a directed graph $F = \langle A, R \rangle$, where the
29 nodes A represent arguments, and the edges $R \subseteq A \times A$ represent the attacks, i.e. the notion of “counter-argument”.
30 Reasoning with such a graph is done thanks to the notion of extension, which are sets of jointly acceptable arguments,
31 corresponding to a possible solution of the problem represented by the argumentation graph. There are various
32 ways to define an extension, depending on the properties expected from this set of arguments. Among the classical
33 approaches, we focus here only (as a matter of exemplification) on the stable semantics: a set of arguments $E \subseteq A$ is
34 a stable extension if and only if

- 35 (1) it is conflict-free (formally, $\forall a, b \in E, (a, b) \notin R$);
- 36 (2) it is a set that attacks its complement ($\forall a \in A \setminus E, \exists b \in E$ such that $(b, a) \in R$).

38 It is thus possible to define a system of argument-based stupidity as a system returning sets of arguments that
39 attack their complement, but without satisfying the conflict-freeness principle. Let us illustrate this notion with the
40 following simple scenario, where the dialogue:¹⁵

- 41 – *Charles Lytton stole the jewel!* (a)
- 42 – *It’s impossible, he has a broken leg.* (b)
- 43 – *He’s probably faking his injury.* (c)

44 is modeled by an argumentation framework $F = \langle A, R \rangle$, where $A = \{a, b, c\}$ and $R = \{(c, b), (b, a)\}$. While stable
45 semantics requires to accept only $\{a, c\}$, a stupid version can lead to accept $\{b, c\}$: this set attacks its complement,
46 but it is not conflict-free. Stupid versions of other semantics for abstract argumentation frameworks can be defined
47 similarly by choosing sets of arguments that only satisfy a part of the required properties. We note that, in some
48 contexts, such as the existence of information about the priorities to apply between arguments, some of the basic
49

50 ¹⁴“Gratin dauphinois is a French dish of sliced potatoes baked in milk or cream” according to Wikipedia.

51 ¹⁵Any similarity with a detective comedy from the 1960’s is absolutely non-fortuitous.

1 properties such as conflict-freeness can be violated [37–39]. Does it mean that these extended frameworks are stupid? 1
2 We leave it to the reader to decide. 2

3 If we do not only focus on the relations between arguments, but also on the method used to build arguments 3
4 from structured knowledge (generally using some logical formalism [40]), then there are also several ways to argue 4
5 stupidly. In summary, there are two steps in the building of structured argumentation frameworks from logical 5
6 knowledge: the identification of arguments and the identification of attacks towards them. Assume that we are 6
7 facing someone stating the following sentence: “I like chocolate, so I am the King of England.” This is equivalent 7
8 to considering that $(\{a\}, b)$ is an argument (where a represents the fact that I love chocolate and b the fact that I am 8
9 the King of England), even if there is no reason to assume that b is a consequence of a . It is already quite stupid to 9
10 claim to be a monarch just because one loves chocolate, and this kind of reasoning can obviously lead to more severe 10
11 problems if argumentation is applied to more sensitive subjects. Even in the case where arguments are “correctly” 11
12 defined, a way to build a system of argument-based stupidity is to fail the second step, that is, the identification of 12
13 attacks. Ignoring some attacks or, on the contrary, adding some attacks that should not exist, will obviously change 13
14 the outcome of the reasoning. 14

15 One can reasonably¹⁶ wonder whether an unlimited system of argument-based stupidity would not lead to chaotic 15
16 reasoning, without particular interest. As mentioned in the introduction, the stupidity of reasoning can be defined 16
17 thanks to a distance function to a correct expected result. So a “good” system of argument-based stupidity might 17
18 be a system that incorrectly identifies some predefined ratio of arguments and attacks and that selects acceptable 18
19 arguments that are somehow related to a correct solution. 19
20

21 5.5. AS in other fields of AI 21

22
23 All fields of AI could be investigated under the lens of AS. Below is presented a short and non-exhaustive inven- 23
24 tory of “from AI to AS” exercises. 24

25 For decision-making, the following question could be raised: “How to make a stupid decision?” For example, in 25
26 a multi-criteria decision, one can decide from an irrelevant criterion: buying a car because it is red when it has been 26
27 decided to buy it mainly on ecological criteria, on price, and on performance, can appear to be stupid. 27

28 For unsupervised learning, a question could be “How to stupidly categorize a domain?” For example, humans 28
29 can be separated into two categories according to the number of vowels of the first name of the mothers of their 29
30 fathers. If such categories appear to be completely useless, this would confirm that this unsupervised learning would 30
31 be stupid. 31

32 In the framework of multi-agent systems, a question would be: “How to make stupidity emerge from a group of 32
33 agents that are individually non-stupid?” In this context, the five fundamental laws of human stupidity proposed by 33
34 Carlo Cipolla can be pointed out [8]. This theory states, in its third law, that a stupid (human) agent is an agent 34
35 who harms another agent (or a group of other agents) without benefiting from it or, even worse, by being harmful 35
36 to herself/himself. Figuratively, this could be seen as someone who shoots her/his own foot while shooting someone 36
37 else! These laws have given birth to experiments aiming at estimating the validity of Cipolla’s theory in the context 37
38 of simulations using artificial agents [41], but also to work that aims to give a biophysical interpretation to this 38
39 theory [42]. 39
40

41 6. Tentative conclusion 41

42
43
44 This article was born out of a reflection on the notion of AS introduced by François Rollin in a recent book. The 44
45 objective here is to gather some thoughts on the question of stupidity seen from the point of view of AI. 45

46 The article does not have the pretension to cover all the issues related to artificial stupidity but raises the fol- 46
47 lowing questions between these still ill-defined notions of intelligence and stupidity (whether or not they are arti- 47
48 ficial): can the latter be seen as complementary of the former? or, conversely, do we consider that there can be no 48
49 [artificial] stupidity without (some) [artificial] intelligence? 49
50

51 ¹⁶Is it really so reasonable? 51

1 Another question or, perhaps, another way of asking this question, is what would be an AS system: an AI system
2 when it goes wrong? Or an AI system designed to produce nonsense? The second answer may be a problem if we
3 consider stupidity as unintentional, but one can avoid this difficulty by stating that the wish (that the system is stupid)
4 comes from the system’s developer, while the system itself would not be conscious of its stupid reasoning, exactly
5 like an AI system is not conscious (until proven otherwise) that it is (sometimes) intelligent. Finally, we do not feel
6 exempt from being stupid by the mere fact of having written this article (which contains at least 7 stupid statements:
7 Can you find them?), but we take comfort in appreciating the charm, if not the beauty of stupidity.

8 This work offers several perspectives of study, partly mentioned in the article. One of them would be the attempt
9 to formally model stupidity. Stupidity is not only a production of errors; it has to be considered as a deviation from
10 a standard that needs to be characterized. This standard can be that of an observer judging whether there is stupidity
11 or not. It can also be an estimation of the consequences of stupidity (for example, by a utility function).

12 This article raises more issues than it actually solves (does it solve any?). In fact, it can be seen as an invitation
13 to the AI community to consider this field of research from an unusual point of view. And let us not forget Gustave
14 Flaubert, who wrote “Yes, stupidity consists of wanting to conclude.”¹⁷ Let us be wary of definitive conclusions!

15
16
17 All sections of this article, with the exception of the last one, were written by its first four authors (who neverthe-
18 less cited the fifth author in the introduction). The last section is written by the fifth author.

21 7. A few words from François Rollin

22
23
24 From the reflections and questionings exposed above, I am afraid that I have to deduce that it is even more difficult
25 to produce Artificial Stupidity than to produce Artificial Intelligence, as long as the objective is to generate authentic
26 stupidity, good real home-grown stupidity, and not simply absurdity or nonsense.

27 This conclusion would be counterintuitive if stupidity was considered a simple failure. In fact, it is easier to fail
28 than to succeed: in cooking, in music, in literature, in medicine, in chess, in love, and in other fields, it takes real
29 competence to do well, while simple incompetence is enough to do badly. When it comes to intelligence, it seems
30 to be the opposite, at least if the desired outcome is good real stupidity.

31 This surprising difficulty can only support the necessity and urgency of scholarly and studious work on AS. The
32 adventure has started, it is a good thing, now it has to be completed, and it is not up to me: I am far too stupid for
33 that.

37 Appendix A. Solution to the exercise of Section 3.1

38
39
40 The solution proposed by the student was $\pm\infty$.

43 Acknowledgements

44
45
46 The authors are thankful to Anaëlle Wilczynski and Zied Bouraoui, organizers of the *Journées d’Intelligence*
47 *Artificielle Fondamentale*, for having given them the opportunity to submit this article. We also thank the anonymous
48 reviewers for their work. Finally, they wish to thank Mathieu d’Aquin, for a proofreading of this article.

50
51 ¹⁷Flaubert’s original sentence is “Oui, la bêtise consiste à vouloir conclure.” (in a letter to Louis Bouilhet, September the 4th, 1850).

References

- [1] F. Rollin, *Suis-je bête ! — L'héroïque Professeur Rollin foudroie la bêtise avec ruse et modestie*, Presses Universitaires de France, 2020, Illustrations by D. Goossens, foreword by A. Astier.
- [2] Jean Paul, *Das Lob der Dummheit*, 1782, *Eulogy of stupidity*, no English translation known; French and Spanish translations: *Éloge de la Bêtise*, Jose Corti, 1993, and *Elogio de la estupidez*, Sequitur, 2012.
- [3] R. Musil, On stupidity, *PN Review* 5(3) (1978), 5, *Über die Dummheit*, Reclam Verlag, 1937.
- [4] L. Latzarus, *Éloge de la Bêtise*, Hachette, 1925.
- [5] G. Bechtel and J. Carrière, *Dictionnaire de la Bêtise et des erreurs de jugement, suivi du Livre des Bizarres*, Robert Laffont, Bouquins, 1965, 1991, 2014.
- [6] H. Frankfurt, *On Bullshit*, Princeton University Press, 2005.
- [7] P. Denis, *Éloge de la Bêtise*, Presses Universitaires de France - PUF, 2001.
- [8] C.M. Cipolla, *The basic laws of human stupidity*, Bologna: il Mulino, 2011.
- [9] J. Lieber, J.- Mailly, P. Marquis, H. Prade and F. Rollin, Quelques réflexions autour de la notion de bêtise artificielle, in: *16èmes Journées d'Intelligence Artificielle Fondamentale, JIAF 2022, Saint-Étienne, France, June 30 - July 1, 2022*, Z. Bouraoui and A. Wilczynski, eds, 2022, pp. 50–60. <https://hal.archives-ouvertes.fr/hal-03781382/document>.
- [10] A. Turing, Computing Machinery and Intelligence, *Mind* LIX(236) (1950), 433–460.
- [11] R. Descartes, *Discours de la méthode*, 1637, Discourse on the Method, Optics, Geometry and Meteorology, translation by P. J. Olscamp, Revised edition: Hackett, Indianapolis, 2001.
- [12] E. Bonabeau, M. Dorigo and G. Théraulaz, *Swarm Intelligence: From Natural to Artificial Systems*, Oxford University Press, 1999.
- [13] C. Godin, Does stupidity exist?, *Le philosophe* 42(2) (2014), 27–38, The author develops his point in his *Encyclopédie Conceptuelle et Thématique de la Philosophie* (Editions Champ Vallon, 2018) in chapter 96 on intelligence which includes a section III on stupidity and a section IV on artificial intelligence.
- [14] J.A. Paulos, *Innumeracy: Mathematical Illiteracy and Its Consequences*, Hill and Wang, 1988.
- [15] L. Schneps and C. Colmez, *Math on trial: how numbers get used and abused in the courtroom*, Basic Books New York, 2013.
- [16] L.J. Savage, *The Foundations of Statistics*, John Wiley and sons Inc., New-York, 1954.
- [17] D. Bouyssou, D. Dubois, M. Pirlot and H. Prade (eds), *Decision Making Process: Concepts and Methods*, Wiley, 2009.
- [18] J. Weizenbaum, ELIZA - A computer program for the study of natural language communication between man and machine, *Communications of the ACM* 9(1) (1966), 36–45.
- [19] G.N. Yannakakis and J. Togelius, *Artificial Intelligence and Games*, Springer, 2018.
- [20] T. Winograd, Understanding natural language, *Cognitive psychology* 3(1) (1972), 1–191.
- [21] P.H. Grice, *Studies in the Way of Words*, Harvard University Press, 1989.
- [22] U. Eco, *Experiences in translation*, University of Toronto Press, 2008.
- [23] L. Lidén, Artificial stupidity: The art of intentional mistakes, in: *AI Game Programming Wisdom*, Vol. 2, S. Rabin, ed., Charles River Media Rockland, MA, 2003, pp. 41–48.
- [24] G. Ritchie, *The Comprehension of Jokes: A Cognitive Science Framework*, CRC Press, 2018.
- [25] F. Dupin de Saint-Cyr and H. Prade, La compréhension des histoires drôles : une affaire de révision de croyances, *Revue Ouverte d'Intelligence Artificielle* 3(3–4) (2022), 253–285.
- [26] A. Nijholt, "All the world's a stage": incongruity humour revisited, *Ann. Math. Artif. Intell.* 88(5–6) (2020), 405–438.
- [27] F. Dupin de Saint-Cyr and H. Prade, Belief revision and incongruity: is it a joke?, *J. of Applied Non-Classical Logics* (2023), to appear.
- [28] P. Bisquert, F. Dupin de Saint-Cyr and P. Besnard, Assessing arguments with schemes and fallacies, in: *Proc. 15th Int. Conf. on Logic Programming and Nonmonotonic Reasoning (LPNMR'19), Philadelphia, June 3-7*, M. Balduccini, Y. Lierler and S. Woltran, eds, LNCS, Vol. 11481, Springer, 2019, pp. 61–74.
- [29] P.R. Halmos, The legend of John von Neumann, *Amer. Math. Monthly* 80 (1973), 382–394.
- [30] R. Fagin and J.Y. Halpern, Belief, Awareness, and Limited Reasoning., *Artif. Intell.* 34(1) (1987), 39–76.
- [31] H. Prade, Handling (un)awareness and related issues in possibilistic logic: A preliminary discussion, in: *Proc. of NMR 2006*, 2006, pp. 219–225.
- [32] C.K. Riesbeck and R.C. Schank, *Inside Case-Based Reasoning*, Lawrence Erlbaum Associates, Inc., Hillsdale, New Jersey, 1989.
- [33] M.M. Richter and R.O. Weber, *Case-based reasoning, a textbook*, Springer, 2013.
- [34] J. Lieber and E. Nauer, Adaptation knowledge discovery using positive and negative cases, in: *Case-Based Reasoning Research and Development, proceedings of the 29th ICCBR*, A.A. Sánchez-Ruiz and M.W. Floyd, eds, Springer, 2021, pp. 140–155.
- [35] É. Astier, H. Iopeti, J. Lieber, H. Mathieu Steinbach and L. Yvoz, Case-based cleaning of text images, in: *Case-Based Reasoning Research and Development, 31st International ICCBR*, S. Massie and S. Chakraborti, eds, Aberdeen, UK, 2023, pp. 344–358.
- [36] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artif. Intell.* 77(2) (1995), 321–358.
- [37] L. Amgoud and C. Cayrol, A Reasoning Model Based on the Production of Acceptable Arguments, *Ann. Math. Artif. Intell.* 34(1–3) (2002), 197–215.
- [38] K. Atkinson and T.J.M. Bench-Capon, Value-based Argumentation, *FLAP* 8(6) (2021), 1543–1588.
- [39] J. Rossit, J.-G. Mailly, Y. Dimopoulos and P. Moraitis, United we stand: Accruals in strength-based argumentation, *Argument Comput.* 12(1) (2021), 87–113.

1	[40] P. Besnard and A. Hunter, <i>Elements of Argumentation</i> , MIT Press, 2008. ISBN 978-0-262-02643-7.	1
2	[41] A.G.B. Tettamanzi and C. da Costa Pereira, Testing Carlo Cipolla's Laws of Human Stupidity with Agent-Based Modeling, in: <i>2014 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), Warsaw, Poland, August 11-14, 2014 - Volume III</i> , IEEE Computer Society, 2014, pp. 246–253. doi:10.1109/WI-IAT.2014.174.	2
3		3
4	[42] I. Perissi and U. Bardi, The Sixth Law of Stupidity: A Biophysical Interpretation of Carlo Cipolla's Stupidity Laws, <i>Systems</i> 9 (3) (2021),	4
5	57.	5
6		6
7		7
8		8
9		9
10		10
11		11
12		12
13		13
14		14
15		15
16		16
17		17
18		18
19		19
20		20
21		21
22		22
23		23
24		24
25		25
26		26
27		27
28		28
29		29
30		30
31		31
32		32
33		33
34		34
35		35
36		36
37		37
38		38
39		39
40		40
41		41
42		42
43		43
44		44
45		45
46		46
47		47
48		48
49		49
50		50
51		51