



NUTS, NARS, and Speech

Dwane van der Sluis^(✉)

WiseWorks.AI, London, UK
ucabdv1@ucl.ac.uk

Abstract. To investigate whether “Intelligence is the capacity of an information-processing system to adapt to its environment while operating with insufficient knowledge and resources” [29], we look at utilising the non axiomatic reasoning system (NARS) for speech recognition. This article presents NUTS: raNdom dimensionality redUction non axiomaTic reasoning few Shot learner for perception. NUTS consists of naive dimensionality reduction, some pre-processing, and then non axiomatic reasoning (NARS). With only 2 training examples NUTS performs similarly to the Whisper Tiny model for discrete word identification.

Keywords: Non Axiomatic Reasoning · Perception · Few shot learning

1 Introduction

‘Artificial Intelligence’ now covers a wide range of tasks such as image recognition, speech recognition, game playing, and protein folding, each of which can be performed at, near, or beyond human level. Over time the term has drifted in meaning, away from a ‘thinking machine’, toward systems that often can only be applied to a single task, do not improve without further training, and take large amounts of resources to train and run. For example, GPT-3, a large language model, is estimated to have cost over 4.6 million dollars to train [19]. These models can be opaque, difficult to interpret, and unable to explain why a particular prediction was made, or unable to provide any guarantees in failure scenarios. Predicate Logic, on the other hand, is capable of robust and consistent decisions. One such predicate logic system, CYC [18], aims to encode all common human knowledge in a knowledge graph. This means CYC has the limitations of predicate logic, one being that all axioms (in the knowledge graph) be true and consistent, otherwise false statements can be derived. Another approach is Non Axiomatic Reasoning. The Non Axiomatic Reasoning System (NARS) performs reasoning that does not assign an objective value of truth to a statement, but instead assigns a subjective value. This subjective value is not fully trusted, and is revised over time as new information arrives. NARS has the advantages that it 1) can cope with holding conflicting information in its knowledge base 2) can

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-33469-6_31.

explain predictions, 3) requires less data for inference 4) explicitly implements logic choosing which concepts to remove rather than randomly forgetting.

While Open NARS for Applications (ONA) was designed with reasoning in mind, we choose to investigate its usefulness, and resource consumption, on perception, and in particular speech recognition. This is because as Peirce stated “abductive inference shades into perceptual judgement without any sharp line of demarcation between them.” [23], and advances in understanding one may shed light on the other. The integration of deep learning and logic reasoning is an open-research problem and it is considered to be the key for the development of real intelligent agents [21]. We narrow the focus of this paper to the dimensionality reduction, and logic needed to convert auditory sensory data into category labels, and the resources required in the Open NARS for Applications (ONA) software platform. First we give background, from the recent discussion around the definition of intelligence, and then how our human nervous system fused from 2 independent systems, perhaps leading to different characteristics of it. We then give a limited literature review, and then explain our method and experiments. In the last section we give and discuss our results.

2 Background

Until recently, and possibly still, industry (and maybe academia too) are interested in whether new tasks can be learnt by AI, and if so, can they be sold profitably to consumers. The developers are under no obligation to consider the environmental impact or safety concerns. That said, some do by choice, but there is little compulsion from a social or regulatory point of view. This is partly due to deep learning being “unreasonably good” [25] and partly due to no other known way of achieving the same level of performance. A focus on resource consumption was created by adding it explicitly into the definition of artificial intelligence [30]. Invitations to comment on this definition produced much discussion [31]. In this discussion it was pointed out that industry has existing finite resource limitations [17], which is true, however for most leader board tasks, resource consumption is not taken into account. It also appears that governments are hesitant to impose resource limitations on industry. A resource limitation is of interest as it prevents brute force approaches and opens up the possibility of investigating how fewer resources can be utilised over time. Brute force approaches can also encode an entire domain space, further limiting the conclusions that can be drawn. If a method requires fewer resources over time, as a task or operation is repeated, then this suggests a deeper (or perhaps more precise or over-fitted) understanding of that task or operation, which may be of interest in the investigation of intelligence. Wang’s definition of intelligence [30] separates skill (e.g. playing chess) from intelligence, and contains an assumption of insufficient knowledge and resources (AIKR). Under this definition learning a new skill to the same level of ability as some other method but with fewer resources or an insufficient understanding of the world (i.e. imperfect knowledge) is advantageous. Thus intelligence and skill are separate concepts.

The focus of this paper is on the cognitive processes that underlie speech recognition. It is assumed that these are similar to the cognitive processes that underlie other forms of perception. However, as stated earlier, the line between perceptual judgement and abductive inference has no clear demarcation [23]. This may be because the physiology underpinning these functions has different origins. Genetic patterning studies suggest that the ‘blastoporal nervous system’, which coordinates feeding movements and locomotion, and the ‘apical nervous system’, which controls general body physiology, evolved separately in our ancestors more than six hundred million years ago [2, 28], and subsequently fused. This may help explain why we (humans) are still aware of differences in different parts of our nervous system, being aware of our cognition around feeding and locomotion, but have little to no awareness of our ‘apical nervous system’. There are arguments that perception and cognition are unified [3, 6, 10, 13, 16, 26, 27], and arguments against, that is, for modularity [9, 22]. The debate between modularity and unified is beyond the scope of this paper. However, the important point for the purposes of this paper is that the mechanisms of perception are not fully available to us. We do not know, for example, how we recognise objects or how we understand speech. Speech is temporal in nature and involves nuanced differentiation between acoustically similar sounds, (for example b in bright, and f in fright).

Model performance over the last few decades has steadily improved, however it is computationally expensive. Current state-of-the-art models rely on a low level acoustic model, followed by a language model. The acoustic model converts a sound wave into an encoded representation of a sound, and the language model gives the probabilities of the next word, given the last few words, along with the acoustic encoding. As said, language models, like GPT-3 require a large amount of data to train, which conflicts with how children learn a new word with very few examples [5]. Wang’s definition of intelligence [30] is based on the idea that intelligence is about making the most of the resources that are available and that it is not always possible to know everything that is going on in the world. So, someone who is able to learn new skills quickly and efficiently, even if they don’t have a lot of knowledge about those skills initially, would be considered intelligent. One approach to understanding the mechanisms of intelligence is to consider the different ways in which it can be measured. One way to measure intelligence may be by looking at someone or some system’s ability to learn new skills, and then measure the quality of that skill, as well as the energy consumed to learn it and perform it once learnt. Speech recognition is one such potential skill.

Generally large dimensionality reduction is needed to convert perceived inputs into symbols that logic can be applied on. One approach [4] is to cluster the inputs in the feature space before similarity and difference are calculated and used as input into a deep network that is trained. Another approach is to pre-process with a DL model specifically trained for that modality, i.e. YOLO for computer vision, and use the generated labels. The generated labels form a lower bound on the resolution of the logic system, i.e. if the labels are ‘dogs’ and ‘cats’, it would be difficult for a logic system to learn of a new breed of dog. If

the logic system uses scalar output features of a DL model, e.g. the bounding box of the cat, x by y pixels, and if only far away cats are seen, and then later a closer (and therefore much bigger) cat is seen, scaling issues can be created, as the system may not see the full range of sizes immediately, and needs to recalibrate previous observations when the scale is readjusted on seeing a much larger, or smaller example. In speech these challenges are exhibited in the form of the dimensionality reduction from 16k samples per second to ~ 2 words per second, with uneven speaking speed. Speech recognition has traditionally used labelled data sets which cost USD50/hr to hand label, limiting the training data set size into the 10,00 h or less range. The resultant systems have low generalisability with many recent state of the art systems reporting $< 5\%$ word error rates, which collapses into the 30–40% range when used on other, but similar datasets [24]. The exact costs of training models such as Whisper [24] with 1.6 billion parameters, are unknown, but the 175 billion parameter GPT-3 ($109\times$ larger), also by OpenAI, is estimated at 4.6 million USD. If the training costs were a constant multiplier of the number of parameters (they are not), the training of Whisper could be in the order of magnitude of 40k USD. The training data set of Whisper was 168k hours of 16 KHz speech. This equates to a data set size of 19 TB, approximately half the estimated size of GPT3 about 45 TB of training data. The recent success of attention in other domains has also been applied to speech. Andrade et al. [1] developed a 202K parameter neural attention model, we will refer to as ANAM, which also targeted at the Speech Command dataset.

3 Literature Review

In the 1970’s speech recognition relied on hand crafted features. This changed as end to end differential systems were developed and new SOTA were reached [24]. These systems lack interpretability, while not important for speech recognition, are of interest if the features triggering decisions can be exposed and validated. With concept whitening [7] it is possible to concentrate (grounded) meaning in single neurons to aid interpretability, but requires category labels, which may not be available at training time, and adds complication as category labels “need to address topics like the representation of concepts, the strength of membership in a category, mechanisms for forming new concepts and the relation between a concept and the outside world” [32]. Another approach is Deep Logic Models which integrate deep learning and logic reasoning in an end-to-end differentiable architecture [21]. This work leads onto Relational Reasoning Networks R2Ns [20] which perform relational reasoning in the latent space of a deep learner architecture. However they suffer an explosion in memory needed as the number of possible ground atoms grows polynomially on the arity of the considered relations. This underpins the useful implication of the AIKR. Shanaha [26] explores perception of objects via computer vision, with abduction, but does not describe how resource intensive the work was or if resource limits were reached. Johansson [14] investigates learning match to sample and the relationships more, less and opposite. Noting the advantages that AGI has, allowing us to experiment

with an agents internals, and giving the example where NARS layer 6 (variables for generalisation) is shown to be needed for the work Hammer [11] does on a system designed to only process data from perception, i.e. has no predefined knowledge. Generalised identity matching, where a new example is matched to a reference sample has been shown to be possible in NARS [15], and further that the derived identity concept could generalise to novel situations. In unpublished work Durisek [8] postulated speech recognition leveraging phonemes may be possible with NARS.

4 Method

For simplicity we attempt to identify single whole spoken words, which has been the focus of much research. We use a standard data set, the Speech Command v2 [33], which contains 35 single word commands, 0–9, back, forward and other confusing words (bird, bed). Each word had over 3000 recordings, each of 1 s in duration or less. As in Whisper [24], we take 16 bit, 16 KHz audio, on which 80 bin MEL (logarithmic) spectrum was calculated every 10 ms. This produced 8000 (80×100) energy intensity values per second, which were normalised. Utterances shorter than 1 s were padded. This reduced the input dimension from 16000 to 8000, and is a standard pre-processing step in speech recognition. These 8000 values needed to be restructured to be passed into ONA. Data was presented to ONA in the form of Narsese¹ statements. As a simple example, we encode three examples as Narsese instances, A, B and C with n properties each. The strength of the property relationship to the instance was encoded in the truth value, i.e. a property with a strength of 0.9 would be encoded as:

$$\langle \{A\} \rightarrow [p1] \rangle .\%0.9\% \quad (1)$$

meaning “‘{A}’ has the property ‘p1’ with strength of 0.9”. It was then asserted that {A} is a LABEL and {B} is a LABEL. e.g.

$$\langle \{A\} \rightarrow LABEL \rangle . \quad (2)$$

And then the system was queried to see if C was labelled correctly:

$$\langle \{C\} \rightarrow LABEL \rangle ? \quad (3)$$

After a grid search we set the number of labelled examples per class to 2 (+ unlabelled example = 3). We note this is the smallest number that allows similarity to be exploited. With this setup and synthetic and real data we attempted to answer the following questions:

- RQ1: Can non axiomatic reasoning, which can cope with conflicting information, be leveraged to perform speech recognition?
- RQ2: Is there a computational or performance?

For a baselines we used 1) general speech recognition pre-trained Whisper models, and 2) the earlier mentioned ANAM model. We expected these to produce state-of-the-art results, at the cost of larger computation.

¹ For Narsese see <https://cis.temple.edu/~pwang/NARS-Intro.html>.

Experiment 1 NARS, Computational Complexity. We expressed data in the same manner as above, encoding each real world utterance as 8000 properties, using the energy in each bin as the strength of a property as Narsese statements. Energy values below 0.5 posed a problem, as they expressed absence of a feature in the data, e.g. the word ‘Moo’, should not have high frequency ‘s’, or ‘t’ sound in it. To enable ONA to track the absence of something, we negate the property name, i.e. [mel16x9] becomes [NOTmel16x9], and subtract the truth value from 0.5, so that a low truth value 0.1, becomes 0.9. E.g.

$$\langle \{U_1\} \rightarrow [\text{mel16x9}] \rangle .\%0.1\% \quad (4)$$

is replaced with

$$\langle \{U_1\} \rightarrow [\text{NOTmel16x9}] \rangle .\%0.9\% \quad (5)$$

We took 3 random utterances of ‘one’ (from the 3893 possible), generated the 8000 values for each, then encoded these as Narsese statements. We then asserted $\langle \{U_1\} \rightarrow \text{one} \rangle$. For utterances 1 and 2, and then queried ONA to see if utterance 3 is similar to the label $\langle \{U_n\} \rightarrow \text{one} \rangle$?. All performance tests used a 64 GB AMD Ryzen 5 3600 6-Core Processor running Ubuntu (no GPU).

Experiment 2 Nalifier, NARS, Synthetic Data. We take the same method as experiment 1, but this time pass the statements into a python pre-processor, Nalifier.py [12], that suppresses certain Narsese statements, and synthesises other Narsese statements, which are in turn passed into ONA. The Nalifier has several functions, if the statement received consists of an instance property statement e.g. $\langle \text{utterance}_n \rightarrow [\text{property}_p] \rangle$. It collects all the properties for this new instance, all the properties for all other instances in its memory and starts comparing them to find the closest. If an instance is found that the current instance is similar to, it synthesises and emits new narsese. The new narsese is passed into ONA (or more specifically NAL, the executable of ONA), and interpreted. The success criteria is the same as experiment 1, we check to see if the unknown instance is labelled correctly.

Experiment 3 - NUTS. We now introduce ‘NUTS’ : raNdom dimensional-ity redUction non axiomaTic reasoning few Shot learner for perception. NUTS consists of four modules, dimensionality reduction, conversion into narsese, a narsese preprocessor (the Nalifier), and open NARS for applications (ONA). We used a random projection without sparsity, to reduce dimensionality, specifically we pass the input 16k samples through MEL encoding, producing 8000 values. These 8000 values were multiplied by a randomly generated $8000 * D$ matrix, reducing the dimensions to D . These D values were then used to generate narsese as before, which is passed into the Nalifier which filters and generates narsese, which is passed into ONA. Each class was tested in turn with the negative classes consisted of the remaining 34 words in the speech command dataset. The number of learning examples of each word could be varied, along with ONA’s setting for the size of the AIKR, the size of the reduced dimensionality space², and the

² A grid search showed 4 dimensions was reasonable.

number of repeats. The matrix used for reduction was re-generated before each run. Success was measured as the proportions of runs where the correct “is a” relationship is identified.

5 Results

Experiment 1 NARS, Computational Complexity. As mentioned, baselines were OpenAI’s whisper model, and Andrade’s et al’s ANAM. Whisper was tested on 100 random utterances from each of the 35 words in the Standard Commands data set, comparisons were case insensitive excluding punctuation. As seen in Table 1, Whisper tiny model took an average of 0.8 sec per inference (including encoding) with a performance of 58%. ONA was unable to accurately identify the unknown utterance as being similar to anything in memory. This may have been because the full structure of the speech was not exploited, but we wished to avoid manual feature engineering. Analysis of the derived statements showed that the instances were considered similar to the properties rather than the instances, this was unexpected.

Table 1. Performance & Baselines: Whisper Tiny, Large, and ANAM

	Exp1	Exp2	NUTS	Large	Tiny	ANAM
Vocabulary Size	1	2	35	50257	50257	35
Training Samples	2	2	105	1e9 (est)	1e9 (est)	84843
Input Dimensions (1 sec audio)	16000	16000	16000	16000	16000	16000
Intermediate Dimensions	8000	200	4	8000	8000	9600
Inference Time (sec)	0.05	2615.00	0.02	43	0.80	0.095
Training Time (sec)	19	5700	16			7200
Performance Accuracy	0%		64%	68%	58%	93%

Experiment 2 Nalifier, NARS, Synthetic Data. The Nalifier took considerable time to execute, to load and ‘train’ 2 instances with 2000 properties each, took 95 min. To load, encode and perform inference on a new example took an additional 43 min. This version of the Nalifier contains a $O(n^2)$ algorithm which executes each time a new property was observed for an instance. After 3 instances, each with 2000 properties were added into NARS, (first pre-processed by the Nalifier), NARS successfully determined that instance C, the unlabelled instance, was similar to instance A. This showed that speech recognition is possible with NARS, and that the Nalifier is needed. We did not attempt 8000 properties, or measure accuracy due to execution time.

Experiment 3 NUTS. We were surprised randomly reducing dimensions worked, even for small numbers of training examples. For these experiments the best performance was obtained with 4 dimensions, when the unknown class was labelled 64% of the time, compared to 2.8% for random performance, see Table 1. This compared favourably with Whisper Tiny’s 58%³, but far below the ANAM’s state of the art 94%. Training was label and compute efficient, only needing 2 training samples per class, and inference time was 0.02 sec (including encoding and dimensionality reduction), far below that experiment 2’s 43 min, showing that dimensionality reduction is the source of the computational efficiency.

We conclude that perception, specifically speech recognition is possible with NARS. However we note performance collapsed certain words such as Bird, and Bed, yet ANAM’s confusion matrix shows it is possible to distinguish them. This may be due to a limitation of NARS, or information loss in the dimensionality reduction. Figure 1 shows the overall performance of a random generated matrix, a random word, and reduced dimensions 2–10, repeated 3500 times (100 times per class). Figure 2 shows performance increases with the number of examples, raising from 64% at 2 examples to 90% at 20.




Fig. 1. Accuracy as a function of the reduced dimension embedding. Number of examples per class = 2.




Fig. 2. Accuracy as a function of the number of examples. Reduced dimensions = 4.

6 Discussion

RQ1: Can non axiomatic reasoning, which can cope with conflicting information, be leveraged to perform speech recognition? Yes, we demonstrated that NARS (along with the Nalifier and dimensionality reduction), can interpret speech data in a meaningful way, obtaining 64% accuracy with only 2 training examples on a 35 class problem.

³ Whisper leverages language models greatly improving multiword performance.

RQ2: Is there a computational advantage? A crude comparison of inference time on CPU suggests they are in a similar order of magnitude. But as Whisper runs efficiently on GPU, and NUTS is a mix of C and Python, a strict comparison of counts of each operation type, could not be completed in the time, and it is left as further work.

We started by discussing the term ‘Artificial Intelligence’ and resource consumption. While we did not produce a system with same or better performance with fewer resources, if we had done so would it be more ‘Intelligent’? We would argue it would not, as the catalyst is not part of the described algorithm. We suspect building ‘intelligent’ systems (as in ‘thinking machines’) will remain elusive until the terms skill and intelligence are dis-entangled, and the catalyst for improvement is isolated and automated.

Acknowledgement. We thank reviewers and Parker Lamb for their comments.

References

1. de Andrade, D.C., Leo, S., Viana, M., Bernkopf, C.: A neural attention model for speech command recognition. ArXiv abs/1808.08929 (2018)
2. Arendt, D., Tosches, M., Marlow, H.: From nerve net to nerve ring, nerve cord and brain evolution of the nervous system. *Nat. Rev. Neurosci.* **17**, 61–72 (2016). <https://doi.org/10.1038/nrn.2015.15>
3. Barsalou, L.W.: Perceptual symbol systems. *Behav. Brain Sci.* (1999). <https://doi.org/10.1017/S0140525X99002149>
4. Blazek, P.J., Lin, M.M.: Explainable neural networks that simulate reasoning. *Nat. Comput. Sci.* **1**(9), 607–618 (2021)
5. Bloom, P.: *How Children Learn the Meanings of Words*. MIT press, Cambridge (2002)
6. Chalmers, D.J., French, R.M., Hofstadter, D.R.: High-level perception, representation, and analogy: a critique of artificial intelligence methodology. *J. Exp. Theor. Artif. Intell.* **4**(3), 185–211 (1992). <https://doi.org/10.1080/09528139208953747>
7. Chen, Z., Bei, Y., Rudin, C.: Concept whitening for interpretable image recognition. *Nat. Mach. Intell.* **2**(12), 772–782 (2020). <https://doi.org/10.48550/arXiv.2002.01650>
8. Durisek, P.: Speech recognition using nars (2014). <https://www.applied-nars.com/articles/speech-recognition-using-nars>
9. Fodor, J.A.: *The Modularity of Mind*. MIT press, Cambridge (1983). <https://doi.org/10.7551/mitpress/4737.001.0001>
10. Goldstone, R.L., Barsalou, L.W.: Reuniting perception and conception. *Cognition* **65**(2–3), 231–262 (1998). [https://doi.org/10.1016/s0010-0277\(97\)00047-4](https://doi.org/10.1016/s0010-0277(97)00047-4)
11. Hammer, P.: Adaptive Neuro-Symbolic Network Agent, pp. 80–90. Springer, Heidelberg (2019). https://doi.org/10.1007/978-3-030-27005-6_8
12. Hammer, P., Lofthouse, T., Wang, P.: The OpenNARS implementation of the non-axiomatic reasoning system. In: Steunebrink, B., Wang, P., Goertzel, B. (eds.) *AGI-2016. LNCS (LNAI)*, vol. 9782, pp. 160–170. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41649-6_16
13. Jarvilehto, T.: Efferent influences on receptors in knowledge formation. *Psychology* **9**(1) (1998)

14. Johansson, R.: Scientific progress in agi from the perspective of contemporary behavioral psychology. In: NARS AGI 2020 Workshop, 2020 (2020)
15. Johansson, R., Lofthouse, T., Hammer, P.: Generalized identity matching in nars. In: Artificial General Intelligence. Springer, Heidelberg (2022). https://doi.org/10.1007/978-3-031-19907-3_23
16. Lakoff, G., Johnson, M.: Philosophy in the flesh, the embodied mind and its challenge to western thought (1999). <https://doi.org/10.1590/S0102-44502001000100008>
17. Legg, S.: A review of “on defining artificial intelligence.” J. Artif. General Intell. **11**(2), 45–46 (2020). <https://doi.org/10.2478/jagi-2019-0003>
18. Lenat, D., Guha, R.V.: Cyc: a midterm report. AI Maga. **11**(3), 32–32 (1990). <https://doi.org/10.1609/aimag.v11i3.842>
19. Li, C.: Openai’s gpt-3 language model: A technical overview. blog (2020). <https://lambdalabs.com/blog/demystifying-gpt-3>
20. Marra, G., Diligenti, M., Giannini, F., Maggini, M.: Learning representations for sub-symbolic reasoning. ArXiv abs/2106.00393 (2021).arXiv:2106.00393v2
21. Marra, G., Giannini, F., Diligenti, M., Gori, M.: Integrating learning and reasoning with deep logic models. In: Brefeld, U., Fromont, E., Hotho, A., Knobbe, A., Maathuis, M., Robardet, C. (eds.) ECML PKDD 2019. LNCS (LNAI), vol. 11907, pp. 517–532. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-46147-8_31
22. Prinz, J.: Is the mind really modular. Contemp. Debates Cogn. Sci. **14**, 22–36 (2006)
23. Project, P.: The Essential Peirce, Volume 2: Selected Philosophical Writings (1893–1913). Essential Peirce. Indiana University Press, Bloomington (1998)
24. Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I.: Robust speech recognition via large-scale weak supervision. Technical report, OpenAI (2022). <https://doi.org/10.48550/arXiv.2212.04356>
25. Sejnowski, T.J.: The unreasonable effectiveness of deep learning in artificial intelligence. Proc. Natl. Acad. Sci. **117**(48) (2020). <https://doi.org/10.1073/pnas.1907373117>
26. Shanahan, M.: Perception as abduction: turning sensor data into meaningful representation. Cogn. Sci. **29**(1), 103–134 (2005). https://doi.org/10.1207/s15516709cog2901_5
27. Shimojo, S., Shams, L.: Sensory modalities are not separate modalities: plasticity and interactions. Curr. Opin. Neurobiol. **11**(4), 505–509 (2001). [https://doi.org/10.1016/s0959-4388\(00\)00241-5](https://doi.org/10.1016/s0959-4388(00)00241-5)
28. Tosches, M.A., Arendt, D.: The bilaterian forebrain: an evolutionary chimaera. Curr. Opin. Neurobiol. **23**(6), 1080–1089 (2013). <https://doi.org/10.1016/j.comb.2013.09.005>
29. Wang, P.: Non-axiomatic reasoning system: exploring the essence of intelligence. Indiana University (1995)
30. Wang, P.: On defining artificial intelligence. J. Artif. Gener. Intell. **10**(2), 1–37 (2019). <https://doi.org/10.2478/jagi-2019-0002>
31. Wang, P., Hahm, C., Hammer, P.: A model of unified perception and cognition. Front. Artif. Intell. **5** (2022). <https://doi.org/10.3389/frai.2022.806403>
32. Wang, P., Hofstadter, D.: A logic of categorization. J. Exp. Theor. Artif. Intell. **18**(2), 193–213 (2006). <https://doi.org/10.1080/09528130600557549>
33. Warden, P.: Speech commands: a dataset for limited-vocabulary speech recognition. arXiv preprint [arXiv:1804.03209](https://arxiv.org/abs/1804.03209) (2018). <https://doi.org/10.48550/arXiv.1804.03209>