# Conceptual Alignment with Formal Concept Analysis

Uta Priss

Zentrum für erfolgreiches Lehren und Lernen, Ostfalia University
www.upriss.org.uk

**Abstract.** Endres et al. (2009) use concept lattices as a model for neural decoding. This paper continues with that idea and presents a thought experiment of how certain structures (associative concepts) and certain processes (conceptual alignment and filtering) might be used to model mental concepts. Several examples from the neuroscientific literature are discussed with respect to such a modelling with FCA.

## 1 Introduction

Formal Concept Analysis[1] (FCA) is a mathematical formalisation of concepts consisting of extensions and intensions which lead to a conceptual hierarchy (Ganter & Wille 1999). Because the natural language word "concept" refers to mental representations the question arises as to whether FCA can be used to model mental processes. There are a few papers that discuss neural applications of FCA (e.g. Endres et al. 2009) and also a few papers that discuss FCA and artificial neural networks, but so far not much research has ventured in that direction. Belohlavek (2000) shows that concept lattices can be learned by artificial neural networks of the type "bidirectional associative memory". In a similar manner, Kumar, Ishwarya & Loo (2015) use FCA for modelling cognitive functions relating to memory. Kuznetsov (2004) discusses how FCA can be used to model different types of artificial neural networks. There are a few more publications in similar directions but as far as we know, the neuroscience community does not yet appear to see FCA as a crucial component for modelling mental processes. Furthermore, the existing papers on neural applications focus on single lattices, not on networks of lattices.

In this paper we are providing some suggestions for further possible connections between neuroscience and FCA. This position paper is neither a mathematical nor a computational paper. But we believe that it would be feasible to implement a neural network simulation of what is described in this paper. Collecting neural data of the kind needed for this paper is difficult because such data tends to be either detailed but with a narrow focus or relating to the whole brain but not at a very detailed level. Furthermore if it was possible to collect detailed neural data on a larger scale, the sheer amount of data would be overwhelming and a challenge for FCA. It should be emphasised that FCA is just used as a model in this paper. We are not suggesting that concept lattices are actually computed in the brain.

[1] The mathematical details of FCA are not discussed in this paper. They can be found in the textbook by Ganter & Wille (1999).

This paper presents a thought experiment of how certain structures (associative concepts) and certain processes (conceptual alignment and filtering) could constitute a possible mechanism that might contribute to how mental concepts are formed. We expect many other mechanisms to occur in mental processes. The theory described in this paper is somewhat inspired by Barwise & Seligman's (1997) Information Flow Theory (IF). Barwise & Seligman describe a network of classifications (equivalent to concept lattices) that are connected through infomorphisms via a channel. In our opinion, a problem with IF is that once classifications and a channel exist, then infomorphisms can be observed, but the theory does not prescribe an algorithm for how to build a channel. In this paper associative concepts take on a similar role as an IF channel, but associative concepts are generated by the contributing concept lattices. Thus it is not necessary to build a channel in advance.

The next section recounts an example of an application of FCA to neural data and explores consequences from that example. Section 3 introduces the notion of "associative concepts" which is discussed with respect to neuroscientific evidence. Section 4 defines "conceptual alignment" as a lattice infomorphism and describes its relevance for associative concepts. The paper finishes with a conclusion.

## 2 Concept lattices in the visual cortex?

In this section we use the notion "natural language concept" for formal concepts which correspond to words of a natural language. For example, natural language concepts occur in the well known lattice of "bodies of water" (Wille 1984) which, for example, defines a puddle as "a temporary, stagnant, natural body of water". While lattices describing physical objects and their characteristics tend to contain at least some natural language concepts, we would assume that many FCA applications are too abstract to produce many (if any) natural language concepts. But we would argue that such concepts do occur in the example presented by Endres et al. (2009).

Endres et al. apply FCA to neuroscientific data consisting of stimuli that were presented to a macaque monkey (used as formal objects) and discretised spike trains from individual high-level neurons in its visual cortex (used as formal attributes). Traditionally such data is analysed with tree-based clustering methods. But the examples from Endres et al. show that using concept lattices is quite promising. Endres et al.'s lattices contain concepts corresponding to natural language concepts such as "face/frontal view of head", "sideview of head", "head from behind" and "round items" which indicates that such concepts might emerge within that area of the visual cortex of monkeys. The bottom concepts tend to relate to individuals, such as other monkeys and caregivers known to the monkey, and the top concepts correspond to abstract, general categories. In summary, we would argue that such lattices contain at least some natural language concepts.

So, what are the conditions that neurons must meet so that the resulting lattices contain natural language concepts? In Endres et al.'s examples the higher level concepts all tend to have large extensions which cannot easily be characterised by an obvious single feature. This means that individual neurons respond to a wide range of stimuli. The criteria that neurons use for their activations are most likely fuzzy, but appear not to be totally random. The fact that some concepts can be labelled (such as "round" or "face") shows that there must be some neurons which react to all stimuli of a certain kind. Thus, there must be a precise core group for the activation of many of the neurons. Nevertheless each neuron also reacts to some other stimuli which are not in the core group. The core groups of different neurons appear to overlap to a large degree leading to closely related subconcept clusters in the lattices. But no two neurons have exactly the same activation pattern. The lower concepts indicate that all neurons react to important, most frequently encountered stimuli. Nevertheless, because the bottom concept has an empty extension, for each stimulus there appears to be at least one neuron that does not react to it.

In summary, the conditions for neurons are that each neuron appears to have a special activation focus yet also some other stimuli it reacts to. Important stimuli lead to the activation of most, but not all neurons. There are large overlaps between the sets of stimuli of different neurons but none are exactly the same. Thus natural language concepts appear to be formed not only in higher association or language processing areas, but for example already in the visual cortex. As mentioned before, FCA is just used as a (very suitable) model - we are not suggesting that lattices are actually computed in the brain. FCA suggests a kind of naturally occurring subconcept-superconcept hierarchy corresponding to the relationships between sets of stimuli and sets of neurons.

Kriegeskorte et al. (2008) show that categorical object representations in the inferior temporal cortex of humans and macaque monkeys are similar. Thus, it can be assumed that the functionality of brains of humans and other monkeys is similar in this respect and similar concept lattices as the ones described by Endres at al. (2009) could in theory be constructed for human data. Unfortunately such data is not normally available for the following reasons: Retrieving data from individual neurons involves invasive methods which cannot normally be performed on humans. Non-invasive neural imaging methods are ethically preferable because they can be performed on humans, but they tend to collect data from brain regions instead from individual cells. Furthermore any type of current neuroscientific data collection method tends to be tedious and expensive. Therefore most studies only focus on very specific aspects of cognition. Last but not least, a fair amount of processing is required in order to discretise any data. Therefore detailed, large scale neuron/stimulus data sets are not available. Nevertheless apart from neuroscience, robotics may be a possible application area for this kind of analysis. It may be possible to collect or generate data from artificial neural networks of humanoid robots or to reverse engineer large scale datasets that are available (knowledge graphs, ontologies, thesauri, web-based data collections, and so on) in order construct hypothetical concept lattices that contain natural language concepts in a similar manner as in the examples of Endres et al. (2009).

## 3 Associative concepts

### 3.1 Results from neuroscience

The previous section describes how concept lattices may be relevant for modelling individual brain areas. This raises the question as to how concepts can be modelled that combine information from different brain areas. Bauer & Just (2019) provide a review of the neuroscientific literature on the topic of how conceptual knowledge is represented in the brain. This section summarises that article but with one modification: we use the term “associative concept” for concepts in the sense of Bauer & Just in order to distinguish them from “formal concepts” in the sense of FCA. According to Bauer & Just, associative concepts are spatially distributed neural representations. This means that multiple brain areas are involved in any associative concept. For example an associative concept of “knife” activates visual areas (what a knife looks like), sensory areas (what a knife feels like), motor areas (how a knife is held and used) and association areas (further information about knifes). From a modelling view, the different areas can be considered dimensions. The formal concepts mentioned in the previous section would then be a component of an associative concept in one dimension. Bauer & Just state that in addition to the dimensions themselves, the relations between the dimensions are also represented in “specific high-order brain areas or convergence zones”. Thus, associative concepts form relations which are represented or computed somewhere.

According to Bauer & Just, associative concepts can be elicited by pictures or words. The picture or word of an associative concept is stored in sensory brain regions separately from the concept itself. Apparently it does not matter whether an associative concept is elicited by a picture or by a word, although pictures evoke potentially more detailed descriptions whereas words evoke the most generic properties of an associative concept. Neural representations of associative concepts appear to be largely independent of the language which is used to describe them and, furthermore, fine-grained activation patterns for associative concepts appear to be common across different people. More concrete associative concepts correspond to more activation in sensorimotor areas whereas more abstract concepts evoke other verbal concepts. Abstract scientific concepts repurpose neural structures originally evolved for more general purposes. Clearly, this is a very reduced and simplified summary of Bauer & Just’s already quite condensed review article. But if structures within single brain areas are modelled as concept lattices as suggested in the previous section, associative concepts then require some sort of combination or further processing of concept lattices as discussed in the next section.

### 3.2 Modelling with FCA

In Section 2 we have argued that individual brain areas might be modelled as concept lattices. Section 3.1 poses the question as to how associative concepts can be formed by combining the individual lattices and possibly other information. Different functional brain areas have different biological configurations (size, shape, type of neurons, connectivity). Thus, it cannot be assumed that they all function in the same manner and should be modelled in exactly the same manner. Other models than FCA might be more suitable for planing, logical thinking, motor functions, emotions and so on. Therefore the definition below does not state that dimensions must be lattices, but they can be. Whereas the formal concepts discussed in Section 2 are most likely part of the unconscious realm of thought, associative concepts as discussed in Section 3.1 are evoked by pictures or words and therefore belong to the realm of conscious thought. Therefore associative concepts are a type of mental construct which is formed in some high level functional areas by relying on information provided by lower areas. Based on the information from the previous sections, we suggest the following definition of associative concepts:

**Definition 1.** *For a many-valued formal context with a set $C$ of objects and sets of attributes $D_1, ..., D_n$: the elements of $C$ are called* associative concepts *if the $d \in D_i$ themselves are concepts (modelled using FCA or some other theory) and a common null element $d_\perp$ exists (i.e., $\forall_{1 \leq i \leq n} : d_\perp \in D_i$). For each $c_i \in C$ the non-null entries in its row form a vector $(d_{i_1}, d_{i_2}, ..., d_{i_k})$ of length $k \leq n$ with* components $d_{i_j}$ *and* signature $(D_{i_1}, D_{i_2}, ..., D_{i_k})$.

The definition indicates that some dimensions can be concept lattices or substructures of concept lattices. Using standard FCA techniques for many-valued contexts (such as conceptual scaling or pattern structures) concept lattices can be computed in any case. But we are arguing below that also other approaches might be of interest for associative concepts. Dimensions are not meant to correspond to functional brain areas but to subareas which contribute one specific aspect to an associative concept. Not every associative concept uses all of the dimensions. For example, colour concepts most likely do not activate areas related to movement or touch.

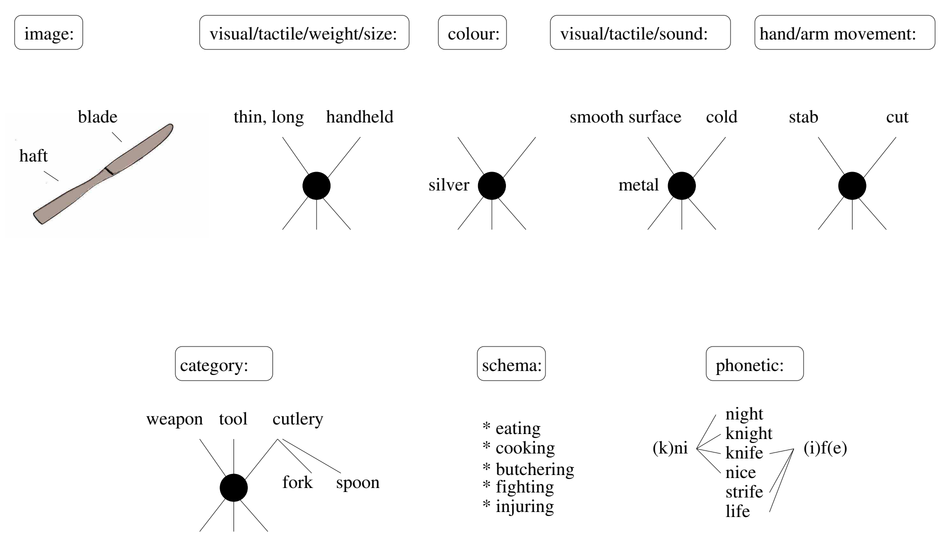


**Fig. 1.** An associative concept of "knife"

Fig. 1 shows a possible associative concept for the word "knife". It consists of a visual dimension that contains a prototypical picture of a knife. This dimension also represents part-whole relationships. In accordance with Section 2, this dimension could be modelled as a concept lattice where the pictures are formal objects corresponding to visual stimuli. The reason why some of the other dimensions are labelled with combinations ("visual/tactile/sound") is explained further below. The different dimensions can also be viewed as "facets" or "aspects". They refer to different physical characteristics of knifes: their colour, material, shape, size and typical usages. The characteristics are described from an embodied perspective. For example "handheld" is meant to represent a physical characteristic of something that can be held in a hand because of its size, weight and shape. While the dimensions at the top half of Fig. 1 are perceptual and formal concepts as described in Section 2, the dimensions at the bottom half are at a higher functional level, such as categories, schemata (typical usage situations) or referring to a representational form. If someone is asked to provide a definition of "knife" they will most likely invoke a category ("A knife is a weapon that ...", " ... a tool for ...", "a type of cutlery"). Understanding a word requires phonetic processing of the consonants and vowels. The schema and phonetic dimensions are not modelled as lattices in this example. Schemata could be modelled as conceptual graphs; phonetic processing could be modelled as a neural network that processes sequential inputs in a probabilistic manner. Fig. 1 is just meant as a possible example for illustrating what associative concepts might be like.

If all dimensions were concept lattices, then a direct product of the lattices could be calculated and an associative concept would just be a concept in the product lattice. But considering that there are billions of neurons in the human brain and single lattices already tend to have many more concepts than attributes (which correspond to neurons in Section 2), the number of concepts in such a direct product would be huge. Furthermore, although humans generate subconcept-superconcept relationships in their minds (of the sort: "a poodle is a dog"), there is no evidence that a long chain of such relationships is ever generated. Most people who are not trained biologists would generate a sequence such as "poodle" $\rightarrow$ "dog" $\rightarrow$ "mammal" $\rightarrow$ "animal" $\rightarrow$ "thing". At the level of "thing" the relationships tend to become circular, for example, by defining a thing as an object and then an object as a thing. Such circular top-level definitions can be found in any dictionary. Thus, it is questionable whether associative concepts form a lattice at all. In our opinion it is more likely that there are some local subconcept-superconcept-like relationships amongst associative concepts but also many other semantic relationships. The hierarchical relationships are generated ad hoc when an associative concept is thought about and never extend very far. Using Miller's (1956) famous number, maybe at most 7 associative concepts and their relationships are ever considered at once. Such relationships will be locally meaningful and free of contradictions, but globally potentially circular and contradictory. Thus, we would argue that there is no reason to calculate a global structure that encompasses all possible associative concepts.

Instead of computing a global structure, it is of interest to observe the changes occurring within and amongst the dimensions if one navigates from one associative concept to another. A train of thought might correspond to navigating between associative concepts. For example, navigating from "knife" to "cutlery" only changes components in some dimensions (image, movement, category, phonetic) while the other components all stay the same. Apart from the phonetic dimension, the changes in the other dimensions are all small. It is to be expected that there is a mutual interaction between associative concepts and their components both while an associative concept is thought about but also while the brain is learning new concepts. Somehow selecting one component can influence how other components are selected and modifying a component in one dimension might modify components in other dimensions. In our opinion, the interesting questions about associative concepts are how they interact, how the different dimensions shape and influence each other and how they contribute to and are modified by other associative concepts. Section 4 discusses this further.

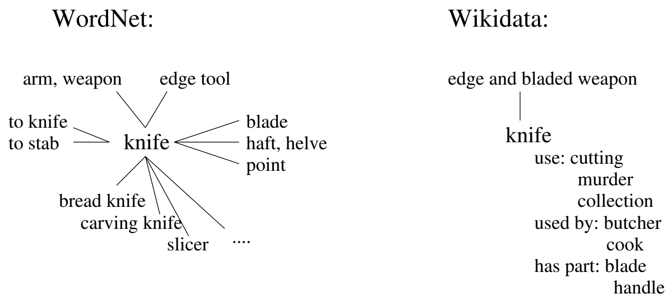


**Fig. 2.** Associations of "knife" in lexical or semantic databases

Because, as mentioned before, it is difficult to obtain neural data, it is of interest to use databases that are available, such as lexicons, dictionaries, ontologies or knowledge bases. Unfortunately, their structures are somewhat different from Fig. 1. Fig. 2 shows entries for "knife" in a lexical database (WordNet[2]) and a knowledge base (Wikidata[3]). The focus of WordNet and Wikidata is more on verbal relationships and ignores some of the more basic perceptual dimensions. It might be possible to extract perceptual dimensions but this would require some additional processing, for example, supported by formal ontologies. Lexical databases might be more suitable for more abstract associative concepts (such as "democracy") which are according to Bauer & Just (2019) represented in a more verbal form in the brain as well.

## 4 Conceptual Alignment

### 4.1 Synesthesia as a form of conceptual alignment

The previous section suggests that the components of an associative concept may influence each other. Ludwig, Adachi & Matsuzawa (2011) observe that not just humans but also chimpanzees associate high pitched sounds with high luminant colours. They argue that such a relationship of synesthesia is not learned but "constitutes a basic feature of the primate sensory system". In our opinion this is an example of *alignment* amongst visual, tactile, auditory or motoric conceptual dimensions which we define as follows:

[2] https://wordnet.princeton.edu/
[3] https://www.wikidata.org

**Definition 2.** Conceptual alignment *is an order-preserving partial isomorphism which maps formal concepts of one concept lattice onto formal concepts of another concept lattice. Two concept lattices are in* total conceptual alignment *if a total order-preserving isomorphism exists between them.*

A partial isomorphism is only a mapping between substructures of the lattices. If two concept lattices that are not aligned at all are combined by concatenating their formal contexts, then the combination lattice is in the worst case the direct product of the initial lattices. Contrary, if two totally aligned concept lattices are combined into one lattice (by concatenating their reduced and isomorphically ordered contexts), then the combination is also isomorphic to the initial lattices. Thus, a larger degree of conceptual alignment reduces the complexity of combining lattices. It follows that if two dimensions which belong to an associative concept are strongly aligned, then it is possible to combine them into one dimension (without changing the original lattices much) in order to reduce the number of dimensions of the associative concept.

When a baby first explores the world, it simultaneously uses multiple senses while it is building its conceptual hierarchies. For example, a round shape can be simultaneously seen and felt. That is the reason why in Fig. 1 we labelled some of the dimensions as "visual/tactile/weight/size" and so on. Fig. 3 shows how that dimension is potentially itself constructed from further lower-level dimensions. For small physical objects that can be held in one hand, a mental, conceptual representation of shape should be the same no matter whether the object is seen or felt. Such an object may have additional visual features (colour) and tactile features (temperature) but the joint visual-tactile features should have a similar conceptual representation. Therefore we would argue that there are either lattices for parts of the visual and tactile areas that are strongly aligned (as in the left hand side of Fig. 3) or only one joint lattice is built for visual and tactile inputs. A similar situation occurs for weight and size (right hand side of Fig. 3). For most everyday objects there is a correlation between weight and size. Thus, the two lattices should be aligned. There are exceptions: a lead ball is heavy, a polystyrene ball is lightweight. But picking up such an object produces a surprise. Therefore exceptions must be dealt with in a different manner, but in non-exceptional cases, a correspondence between weight and size can be expected. There may also be some connections between the dimensions of shape and of size/weight. But while the concepts for weight/size are graded and interval-based (heavy/large, medium, light/small), concepts for shape are feature-based and not graded. Thus, lattices for weight/size and shape cannot be conceptually aligned apart from a few connections.

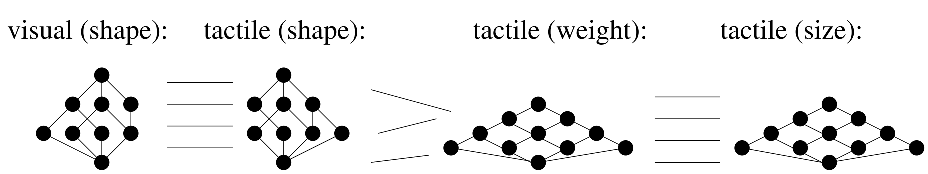


**Fig. 3.** Conceptual alignment

The synesthetic connections observed by Ludwig, Adachi & Matsuzawa (2011) for chimpanzees are not between vision and touch as in Fig. 3 but between vision and sound which might seem surprising. But we would argue that they are actually also based in physics[4]. From a physical viewpoint dark colours and soft surfaces collect heat whereas bright colours and smooth surfaces reflect light and heat. This even extends to sounds: soft surfaces produce a muffled sound when tapped whereas solid/smooth surfaces produce a sharp sound. Therefore synesthetic connections can develop between these perceptual dimensions reflecting what babies (human or chimpanzee) observe when they explore physical objects. It should be mentioned that this account of the relationship between perception and physics is related to the well-known idea of "embodied cognition" (Varela, Thompson, & Rosch, 1992). But the neural basis of embodied cognition is still being investigated (Caramazza et al. 2014). Our suggestion is that low level synesthesia is related to conceptual alignment. It can be speculated that conceptual alignment is such a strong driving factor of early developmental learning processes that it even overgeneralises and creates more connections than what is actually observed by a baby. This could be an explanation of unusual synesthesia in some people.

We thus argue that a process of alignment occurs between some of the concept lattices that constitute dimensions of an associative concept as shown in Fig. 3. We further hypothesise that the associative concepts themselves may be what drives the process of conceptual alignment. Therefore, if a baby repeatedly conjures sets of associative concepts which appear to have some regular combination of components then the lattices may have a tendency to arrange themselves so that their corresponding concepts align. For example, if two domains can be modelled with intervals, then these might attempt to align themselves. Furthermore if there is a subconcept-superconcept relationship for two associative concepts in two of the lattices whereas in a third lattice only the subconcept exists but not the superconcept, then maybe that concept is added to that lattice so that the subconcept-superconcept relationship holds in all three lattices. We propose the following principle:

**Principle 1** Reduction through Alignment: *if two dimensions of an associative concept are partially aligned, they may have a tendency to increase their alignment so that they can be combined into one dimension thus reducing the number of dimensions of the associative concept.*

This process does not need to lead to total alignment. Lattices of perceptual dimensions are most likely never totally aligned because different types of perception always have subtle differences. In the example in Fig. 1, there is a relationship between "silver" and "metal" (because metals tend to have a silver or copper colour) but such relationships cannot be found between all kinds of colours and materials. There is also a relationship between the weight/size of an object and the types of hand/arm movements that can be performed with an object.

[4] Some forms of synesthesia may not be based in physics. For example, some people associate colours with graphemes (i.e. letters or numbers). Since graphemes are learned at a later stage than basic shapes, it is likely that higher level brain areas are involved in that type of synesthesia. Thus, it is a different type of synesthesia than the one discussed in this paper.

### 4.2 Lack of conceptual alignment

Further evidence for alignment between low level perceptual concepts arises from conditions where such alignment fails to occur. Autism is a potential example for such a condition. Autism involves impairment in social interaction but also unusual perceptual abilities and deficits most likely due to cerebellar abnormalities (Fatemi et al. 2012) which appear to obstruct fast neural communication between certain brain areas, for example, motor and vision. In our terminology, conceptual alignment of low level perceptual brain areas may be lacking. Asperger's syndrome and high functioning autism (HFA) are milder forms of the condition. In this paper we are using the "speech-in-noise" problem as an example which refers to difficulties in understanding speech in noisy situations, for example, in a restaurant setting. Adults with HFA tend to have more difficulties with speech-in-noise understanding than people without autism (i.e. neurotypical or NT people) because their brains have difficulties filtering the relevant sound from the background noise (Smith & Bennetto 2007). Lip reading improves speech-in-noise understanding for NT people but not for adults with HFA (Smith & Bennetto 2007).

Fig. 4 shows a modelling of the conceptual alignment required for speech-in-noise understanding. The upper half of the figure shows potential lattices involved in speech-in-noise perception in an NT person. By using "theory of mind" which involves understanding and mapping someone else's actions, a connection is established between the other person's lip movements and the person's own lip movements. Concepts relating to seeing consonants being formed are aligned with concepts of oral muscle movement. For HFA people lip reading does not help because they tend to have problems with "theory of mind" cognition (Baron-Cohen et al. 1997). Thus, the connection labelled "A" in Fig. 4 is difficult for people with HFA. Furthermore, because they generally have a problem with the alignment of perceptual inputs from different channels, the connection labelled "B" is also difficult for them. Therefore, in the case of autism, the low level perceptual lattices are not aligned. In fact it is known that people with autism perceive and produce too much pitch in spoken language (Nadig & Shaw 2012) which in NT people is only relevant for emotional connotations and not for phoneme perception. DePape et al. (2012) show that people with HFA have "less specialisation for native speech sound categories than controls". Or in other words they have too many categories for speech sounds. But people with HFA tend to produce less diverse mouth movements than NT adults (Parish-Morris et al. 2018). Translated into our terminology, in autism, the lattices of pronunciation related motor neurons have too few concepts for consonants whereas the phoneme perception lattices have too many concepts. Therefore they are not aligned. Both lattices contain irrelevant concepts pertaining to pitch because of the insufficient connection A. Because speech understanding without noise is similar to NT people, HFA people must be using higher level processing for language understanding that overcomes low level difficulties but without ever leading to an alignment of the lower level lattices. The lower level deficiency only becomes apparent in speech-in-noise situations.

NT people also find it more difficult to understand speech in a noisy situation if the speech is not in their native language but in a second language that they learned at a later age. Because this depends on the type of noise it cannot be solely due to better semantic

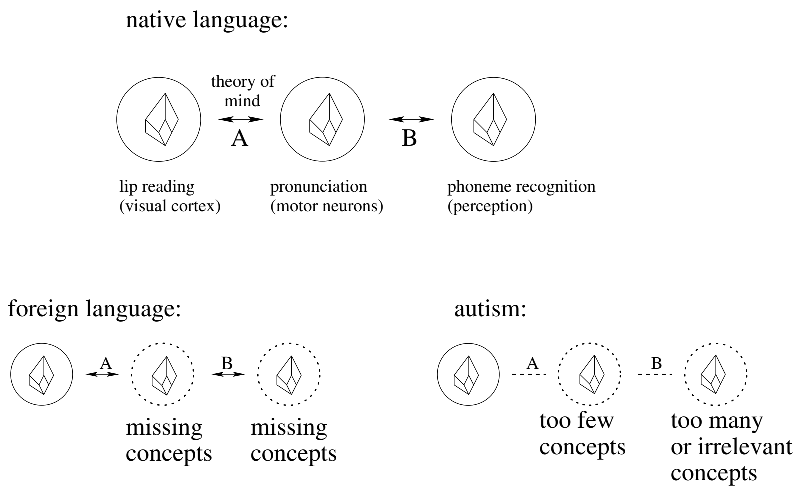


**Fig. 4.** Speech-in-noise processing for people with and without autism

networks in the native language (Broersma & Scharenborg 2010). For NT people connection A is not affected and there is also no reason to assume that they have difficulties in mapping muscle movements for producing consonants onto the sounds produced by consonants (connection B). Their problem is most likely that they are missing concepts both for their motor neurons as well for their phoneme perception. Apparently such concepts can only be learned at an early age. For example, Japanese people have a problem hearing the difference between "l" and "r" in English and tend to pronounce "l" instead of "r". Thus, they are missing the concept of "r" both as a phoneme as well as a pronunciation. In some cases, higher level processing can support speech understanding and production. For example, with adequate training, it is possible for movie actors to learn to pronounce a foreign language with a high degree of accuracy. But this involves higher levels of concentration. Another example for high level processing is the fact that most people can hear the tonal differences of Chinese words if their attention is drawn to them. But associating (lower level) meaning differences with tonal differences may be impossible to learn at a later age. In summary, we are arguing that the notion of "conceptual alignment" can serve as a model for explaining low level perceptual processing. In particular, a lack of alignment may cause deficits which cannot be fully compensated for by higher level processes.

### 4.3 Filtering

We suspect that associative concepts and their dimensions mutually influence each other: if concepts from different dimensions co-occur then they give rise to an associative concept that combines them more solidly. One the other hand once the associative concept has been formed, it encourages its dimensions to conceptually align with each other if at all possible. From the viewpoint of the associative concept it is advantageous to have fewer dimensions because that reduces its cognitive load. We believe that apart from alignment there is a second possibility of reducing dimensions as described in the following definition and principle:

**Definition 3.** *The* strength *of a dimension is the degree to which it is a result of combinations of other conceptually aligned dimensions.*

**Principle 2** Filtering: *associative concepts prefer stronger dimensions. Very weak dimensions may be omitted from an associative concept thus reducing the number of dimensions.*

Thus, both the Principle of Reduction through Alignment as well as the Principle of Filtering lead to a reduction of dimensions and therefore reduce the computational load related to an associative concept. Low-level conceptual alignment as in Fig. 3 is advantageous because it strengthens the individual dimensions and reduces dimensions. Dimensions that do not fit at all are considered noise and are filtered out. There are many known examples of mental filtering. For example, there tend to be many continuity errors in films[5], such as objects that are suddenly in a different place or changes in clothing of the actors in the middle of a scene, which most people do not notice. With respect to linguistic processing, people tend to overlook missing letters or words in written language and do not hear certain consonant variations in spoken language. We would argue that this is a consequence of alignment (which overwrites erroneous values) and filtering which ensures that an unusual perceptual input is not added as an extra dimension to an associative concept but instead omitted. Not everything can be filtered out, though. Inputs that interfere with normal processing (for example spoken language with foreign accents or unusual rhythms such as stuttering) or contradict expectations (for example an incorrect note in a piece of music) are not filtered out but instead consciously noticed. The same holds for observations that attract emotions (such as curse words). Thus, as mentioned before, exceptions must be processed differently from non-exceptional perceptive inputs.

With respect to the speech-in-noise perception mentioned in the previous section, it is known that people with HFA have problems filtering. In our terminology this means that people with HFA perceive many simultaneous, conflicting associative concepts of few dimensions instead of one single coherent associative concept with many dimensions. In summary, this section describes how conceptual alignment may be related to the formation of associative concepts and how associative concepts may attempt to reduce the number of their dimensions via alignment and filtering.

## 5 Conclusion

This paper starts with an example of an application of FCA to neural data by Endres et al. (2009) which describes potential concept lattices for single perceptional brain areas. The paper then argues that individual lower level lattices are combined as dimensions in higher level associative concepts. Relationships of mutual influence exist between

[5] https://en.wikipedia.org/wiki/Continuity_(fiction)

associative concepts and their component dimensions, for example, in the form of conceptual alignment.

As a thought experiment, the paper poses many questions for further research. We believe that in particular questions about different types of learning are of interest for concept lattices. In early developmental phases the concept lattices must be built. This could be modelled with bidirectional associative memory as suggested by Belohlavek (2000). But learning does not stop at that point. Later in life conceptual structures are still modified, but slowly and incrementally. In our opinion the mutual influences that exist between associative concepts and their dimensions provide a model for such kind of learning. We hope that this paper might encourage further research in these directions.

## References


1. Baron-Cohen, S., Jolliffe, T., Mortimore, C., & Robertson, M. (1997). *Another advanced test of theory of mind: Evidence from very high functioning adults with autism or Asperger syndrome.* Journal of Child psychology and Psychiatry, 38, 7, p. 813-822.
2. Barwise, Jon; Seligman, Jerry (1997). *Information Flow. The Logic of Distributed Systems.* Cambridge University Press.
3. Bauer, Andrew J. & Just, Marcel A. (2019). *Neural representations of concept knowledge.* In: Zubicaray, & Schiller (Eds.), The Oxford Handbook of Neurolinguistics.
4. Belohlavek, R. (2000). *Representation of concept lattices by bidirectional associative memories.* Neural Computation, 12, 10, p. 2279-2290
5. Broersma, M., & Scharenborg, O. (2010). *Native and non-native listeners' perception of English consonants in different types of noise.* Speech Communication, 52, 11-12, p. 980-995.
6. Caramazza, A., Anzellotti, S., Strnad, L., & Lingnau, A. (2014). *Embodied cognition and mirror neurons: a critical assessment.* Annual review of neuroscience, 37, p. 1-15.
7. DePape, A., Hall, G., Tillmann, B., & Trainor, L. (2012). *Auditory Processing in High-Functioning Adolescents with Autism Spectrum Disorder.* PLoS ONE 7, 9, e44084. https://doi.org/10.1371/journal.pone.0044084
8. Endres, Dominik M.; Foldiak, Peter; Priss, Uta (2009). *An Application of Formal Concept Analysis to Semantic Neural Decoding.* Annals of Mathematics and Artificial Intelligence, 57, 3, Springer-Verlag, p. 233-248.
9. Fatemi, S. H., Aldinger, K. A., Ashwood, P., Bauman, M. L., Blaha, C. D., Blatt, G. J., ... & Estes, A. M. (2012). *Consensus paper: pathological role of the cerebellum in autism.* The Cerebellum, 11, 3, p. 777-807.
10. Ganter, Bernhard; Wille, Rudolf (1999). *Formal Concept Analysis. Mathematical Foundations.* Berlin-Heidelberg-New York: Springer.
11. Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., & Bandettini, P. A. (2008). *Matching categorical object representations in inferior temporal cortex of man and monkey.* Neuron, 60, 6, p. 11261141.
12. Kumar, C. A., Ishwarya, M. S., & Loo, C. K. (2015). Formal concept analysis approach to cognitive functionalities of bidirectional associative memory. Biologically Inspired Cognitive Architectures, 12, p. 20-33.
13. Kuznetsov, S. O. (2004). *Machine learning and formal concept analysis.* In International Conference on Formal Concept Analysis, Springer, p. 287-312.
14. Ludwig, V,; Adachi, I; Matsuzawa, T. (2011). *Visuoauditory mappings between high luminance and high pitch are shared by chimpanzees (Pan troglodytes) and humans*. Proceedings of the National Academy of Sciences, 108, 51, p. 20661-20665.

15. Nadig, A., & Shaw, H. (2012). *Acoustic and perceptual measurement of expressive prosody in high-functioning autism: Increased pitch range and what it means to listeners.* Journal of Autism and Developmental Disorders, 42, 4, p. 499-511.
16. Miller, G. A. (1956). *The magical number seven, plus or minus two: Some limits on our capacity for processing information.* Psychological review, 63, 2, p. 81.
17. Parish-Morris, J., Sariyanidi, E., Zampella, C., Bartley, G. K., Ferguson, E., Pallathra, A. A., ... & Brodkin, E. S. (2018). *Oral-Motor and Lexical Diversity During Naturalistic Conversations in Adults with Autism Spectrum Disorder.* In Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic, p. 147-157.
18. Smith, E. G., & Bennetto, L. (2007). *Audiovisual speech integration and lipreading in autism.* Journal of Child Psychology and Psychiatry, 48, 8, p. 813-821.
19. Varela, F.; Thompson, E.; & Rosch, E. (1992). *The Embodied Mind: Cognitive Science and Human Experience*, MIT Press.
20. Wille, Rudolf (1984). *Liniendiagramme hierarchischer Begriffssysteme.* Studien zur Klassifikation. Indeks Verlag.