A Neural Field Model of Word Recognition

Andrew P. Valenti (andrew.valenti@tufts.edu)

Bradley Oosterveld (bradley.oosterveld@tufts.edu)

Matthias Scheutz (matthias.scheutz@tufts.edu)
Tufts University Human-Robot Interaction Laboratory, 200 Boston Ave.
Medford, MA 02155

Abstract

We show how temporal and spatial information can be represented as stable patterns in a dynamical system. We describe a model in which category perception arises from the incremental recognition of temporal patterns from sequences of inputs and this is accomplished by decoding a pool of recurrently connected artificial neurons which is called a neural field. In an example application, we use these patterns to identify a set of words which share the word onset represented by the input sequence, consistent with the Marslen-Wilson COHORT model of word recognition. Similarly, we evaluate the extent to which information contained in the bottom-up sensory signal can be used to determine word boundaries. We suggest it is plausible that a neural field offers a naturalistic explanation of how perception arises in word processing.

Keywords: dynamic field theory, neural fields, connectionist model, word recognition, COHORT model, machine learning

Introduction

The brain encodes and processes sensory input acquired from the environment. Sensory input, regardless of modality, is encoded as spatiotemporal patterns, and a superior form of pattern processing has evolved in humans coinciding with the expansion of the neocortex. In this brain structure, several essential cognitive processes such as visual, auditory, and speech perception occur (Koch, 2004; Mattson, 2014). These processes include not only recognizing patterns, but also classifying them (Grossberg, 2005). During this processing, different sensory inputs which represent members of the same category are mapped to a singular representation for that category. In speech processing, for example, all pronunciations of the phoneme "ə", are mapped to the same pattern, allowing for invariance in speech perception across multiple speakers (Kleinschmidt & Jaeger, 2015). Consistent with these hypotheses, our model uses patterns of activation to represent sequences of states in the context of perceiving words; we modeled these states as equilibriums in a *neural field*.

The human neocortex consists of six layers of tissue containing approximately 10^{10} neurons. Columns of tissue can be represented mathematically as neural fields, which form patterns of activation through interaction with each other (Amari, 1977). These interactions between fields generate patterns of activation in a fashion that is believed to be similar to how sensory information is represented in the human neocortex (Amari, 1977; Brady, 2012). These patterns represent an encoding of spatial and temporal information from the brain's sensory input stream.

Each neuron in a neural field F (Figure 1) is connected to each of its neighbors with weights that create an on-center off-surround activation pattern, where the closest neighbors provide a positive influence on activation, further neighbors a negative influence, and the furthest no influence. If given no input and random initial conditions, the units of the field are guaranteed to quickly fall into a stable equilibrium state. Different equilibrium states of a field can be associated with different inputs, and thus the states of activation in a neural field can be used to store information by associating them with category labels (Valenti, Brady, Scheutz, Holcomb, & Pu, 2016).

In this work, we demonstrate a model of word perception using neural fields. Our research is not focused the initial interaction between perceptual signals and the sensory apparatus. We are instead interested in the processing of the output of such apparatuses, and how it can be used to constrain the patterns of activation in higher level cognitive processes, like lexical representation. Our model uses two neural fields, each representing a level of cognitive processing. Since sensory information unfolds over time as a continuous sequence, the input presented to the first neural field is a sequence of feature vectors which represent the letters of an artificial font. Sequences of output features from the first field representing letters are then presented as input to the second field which identifies likely word boundaries and classifies these letter sequences as words.

There are many theories about how patterns of activation in the lexicon are formed once the sensory information has been received (Dahan & Magnuson, 2006). This work focuses on the Marslen-Wilson (1987) COHORT model, which theorizes that information contained in the bottom-up perceptual signal can be exploited to determine which lexical items should be activated, and also used to identify perceptual characteristics such as word boundaries. To explore the extent to which this information is sufficient, we have developed a model where word onsets constrain the set of activated lexical entities such that word onsets activate lexical items with shared onsets. Our model thus makes predictions similarly to the COHORT model; the initial information contained in the sensory signal influences the activation of an initial word-cohort, allowing it to predict word boundaries in a higher level of processing.

Representing State with a Neural Field

Our model is composed of two layers of neural fields. The structure of a single layer is shown in Figure 1. An Input vector (I) is fully connected to the neural field (F) by input

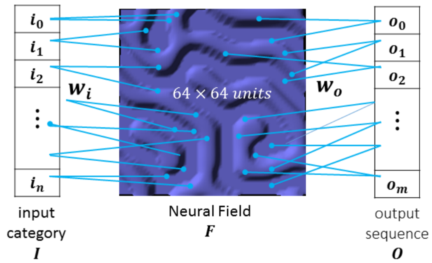

Figure 1: Single field design. Note: I and O are fully connected to F, but only a few connections are depicted here, for clarity.

weights (W_i). F is fully connected to an output vector (O) by output weights (W_o). In our model, such a layer (input, field, and output) can be interpreted as a cognitive processing layer, computing a specific function such as letter or word detection. These layers can be combined to represent a hierarchy of cognitive processes shown in Figure 3.

Our model is based on the following principles of dynamic field theory. Patterns can be stored as stable equilibrium states. A sequence can be "remembered" as a unique equilibrium, unrelated to any previously generated equilibrium, by calculating the sum of the pattern generated by the current input and the pattern representing the previous sequence. Fields converge to a stable equilibrium state after applying a finite series of "settling" operations after which the field ceases to change. Finally, fields can be forced out of a stable state into a target state by applying a finite series of operations which includes the target as its field input.

Field Dynamics

The input to a layer is received as a vector I, whose dimensionality is the number of discrete categories in the input domain. This input is used to first calculate the $n \times n$ matrix F_t (in our evaluation $n = 64$), which represents field activation at the given point in the sequence of input vectors. F_t is calculated using the following equations which are a variation of those widely used in dynamical systems (Amari, 1977).

$$D = W_i I \tag{1}$$

$$\sigma(x) = \begin{cases} \dfrac{x}{x+1} & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases} \tag{2}$$

$$S = W_{mh}\,\sigma(F_{t-1} + D) \tag{3}$$

$$F_t = \sigma(S + h + n) \tag{4}$$

As shown in Equation (1), the driver input, D, is generated by multiplying the input vector I by the input weights, W_i (dimensionality $n \times n \times \parallel I \parallel$). D is then added to the current field equilibrium, F_{t-1}, and the result is squashed to the range $[0,1]$ using Equation (2). This new equilibrium represents the sequence of input seen up to time t, plus the input at t. The result is multiplied by the within field weights, W_{mh}, which are defined using the Mexican hat function;[1] the result is the field influence term, S, Equation (3). Small bias h and noise n terms are added to the field influence, and the result is squashed again to produce the field activation F_t, Equation (4).

Settling to an Equilibrium State

Once a field has been updated from an input, a settling operation is applied resulting in convergence to an equilibrium state. This process is expressed in Equation (5), which is based on Equations (3) and (4) with the notable exception that the value of the field at the previous time step is not added to the input. This operation is repeated until a stable equilibrium is reached, determined by comparing the field activation at time t with its activation at $t-1$, repeating until the difference is below a small epsilon value.

$$S = W_{mh}\,\sigma(S + n) \tag{5}$$

Layer Output

The settled equilibrium is used to calculate the layer's output vector, O, whose dimensionality is the number of categories in the output domain.

$$O = tanh(F_t W_o) \tag{6}$$

The state of the field is multiplied by the output weights W_o ($n \times n \times \parallel O \parallel$), and their product is passed through the hyperbolic tangent activation function. The result is a vector whose values are normalized to the range $[-1,1]$.

The COHORT Model Using Neural Fields

The Marslen-Wilson (1987) COHORT model of spoken word recognition suggests that the real-time constraints of a speech signal influence how bottom-up information is used to determine which items in the mental lexicon become activated. According to the model, on each new input onset only the cohort of possible values remains activated; a cohort is the set of all lexical items that share an onset. A decision is reached only when one possible value remains in the cohort. Consistent with the original version of COHORT (which assumed a highly categorized, abstract string of phonemes rather than feature vectors as input), our model uses only the features present in the bottom-up sensory input to develop the cohort.

The neural fields in our model simulate what happens when sensory information makes contact with the mental lexicon. It is assumed that, once past the sensory apparatus, information flow through differentiated neural architectures, e.g., visual or auditory, is represented in the same fashion. In our evaluation we chose to focus on visual information, and the following sections describe how it is handled by the model.

[1] A Mexican Hat function where D is the Euclidean distance between two units in the field: $W_{mh} = e^{\frac{-D^2}{r^2 d}} \cdot (cos(\frac{\pi D}{2r}) - z) \cdot (\frac{1}{1-z}), r = 4.5, z = 0.15$

Model Input

The input to our model is a sequence of feature vectors which represent the output of the visual perceptual system. As in the Interactive Activation model of word recognition (McClelland & Rumelhart, 1981), visual features are extracted from the raw input, sequences of letters, by separating each letter into a set of component features. These features can be thought of as the pen strokes used to write the letter. For simplicity, we have chosen the font used by Rumelhart and Siple (1974) which is shown in Figure 2. Sequences of feature vectors are generated from a letter by arbitrarily circumnavigating the font clockwise from the outermost feature, spiraling inward. For example, the letter "R" is represented by the sequence [0, 1, 4, 5, 8, 9, 12].

In our implementation, the model receives visual features as input. As mentioned earlier, the input features could also represent information from other sensory modalities. For example, the input features could also come from a speech recognizer and instead represent sequences of phonemes. At the cognitive processing level, the model's basic results do not depend on what the input represents. With different input features there are obviously low-level feature processing issues (e.g., different types of variance in the input) which are outside the scope of this work.

(a) Labeled Segments (b) Letters

Figure 2: 14-segment display using the letter font.

Model Architecture

The architecture of our model is shown in Figure 3. It uses two neural fields, each representing a stage of cognitive processing: *F1* which is a letter detector, and *F2* which is a word detector. $F1$ and $F2$ are connected via a Send Gate which controls when information from $F1$ is sent to $F2$.

Before the model can be used, it must first be initialized. This initialization generates the associations between input and neural field equilibrium states which are used to detect sequence boundaries. During this initialization the model is presented with sequences of input features and the boundaries between them that represent meaningful units. Perceptrons are trained to detect these boundaries based on the equilibrium states which represent them.

Once initialized, information flows through the model as shown in Figure 3. Letter segment sequences previously generated by the visual recognizer flow as input to $F1$; letters detected by $F1$ flow as letter sequences to the word detector $F2$. Perception arises as the generated pattern predicts a set, or cohort, of letter or word candidates. As new features are fed incrementally into the model, a new pattern is generated and each field's perceptron updates its prediction, removing candidates from its cohort. Letter or word recognition occurs when there is a single candidate left in the respective cohort. When recognition occurs, the send gate is opened sending the output perceptron's value to the next layer as input. In the case of the letter layer, the output is sent as input to the word layer; in the case of the word layer, the output value of the perceptron is used by the decoder perceptron to generate results interpretable by a human.

Model Initialization

The characteristic theory of the COHORT model is instantiated in the neural field model during initialization. This initialization forms the associations between input features and neural field equilibrium states used during perception. This initialization is composed of three steps: (1) Initial equilibrium generation (2) W_i training (3) W_o training.

The initial values of the weight matrices used in the model ($F1$: W_i, W_o and $F2$: W_i, W_o) are chosen randomly from a truncated normal distribution with a standard deviation of: $\frac{2}{\sqrt{n_{inputs}}}$ Using this particular standard deviation helps the training to converge more quickly (Géron, 2017).

Initial Equilibrium Generation: A "seed" equilibrium is generated to represent each unique input feature a field will receive. For the letter detector, $F1$, 15 equilibriums are generated to represent each of the 14 possible letter visual features, plus an equilibrium to represent the beginning of a sequence when no input has been presented yet. For the word detector, $F2$, 27 equilibriums are generated, for the 26 letters in the English alphabet and one more for the initial sate. These initial equilibriums are generated by a variation of Equation 1 where I is the product of a one-hot vector whose 1-bit corresponds to the ordinal value of the letter segment in the range $[0,15]$ or letter in the range $[0,27]$ and the randomly drawn W_i for the given field.

W_i **Training:** For a feature vector (i.e., letter segments for the first field, letters for the second), a set of weights (W_i) is trained which will reliably reproduce the initial equilibrium associated with that feature vector. The model's operation assumes that a *settled* field generated from new input can be added to the current settled field to produce a new equilibrium representing the input sequence seen thus far. Without trained driver weights, an *unsettled* equilibrium would be added, violating a core model assumption. The training of W_i, uses a version of the perceptron learning rule, Equation 7, to train the single layer perceptron whose activation is found by multiplying the driver input by W_i,

$$\Delta W_i = \eta I(Target - IW_i) \tag{7}$$

where *Target* is the seed equilibrium for the category and η is a learning rate. Training proceeds until $Target - IW_i < 0.0001$. This approach is a variation of Hebbian learning, a biologically plausible mechanism for learning associations

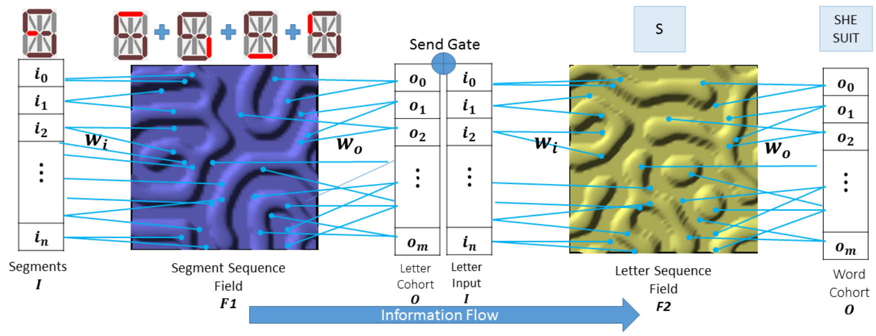

Figure 3: Neural Field based Cohort Model architecture.

between neurons (Laszlo & Plaut, 2012).

W_o **Training:** The output weights W_o map a field's current equilibrium to the output domain relevant to its cognitive layer (i.e., letters cohort or words cohort). For a given cognitive layer the output vector O represents the members of the cohort that are currently active. O is the size of the lexicon of known labels at the given cognitive level, each of its elements representing a member of the lexicon. A member of the lexicon is considered activated if the value of its corresponding element in O is above an activation threshold. For example, the letter segment sequence [0, 1] is the prefix of the letters {A, B, D, O, P, Q, R} (see Figure 2a).

The weights W_o are trained so that the same value of O can be calculated every time a corresponding equilibrium is present in F. The weights are updated using Equation 8,

$$\Delta W_o = \eta F(Target - FW_o) \tag{8}$$

where $Target$ is the vector in the output domain (e.g., letters or words) indicating cohort membership of the lexical entries whose onset is represented by the equilibrium of the neural field F. Training proceeds until $Target - FW_o < 0.001$.

Detecting a New Sequence

An equilibrium for a sequence is generated by adding the seed equilibrium of the new element of the sequence to the current sequence equilibrium. For the first segment in a letter, its seed equilibrium is added to a default value (i.e., the 15^{th} seed equilibrium). This process is repeated for each segment of the letter and after each addition, the field is settled. A challenge in implementing the COHORT model is detecting when a new sequence begins. In our model, an input sequence is "remembered" as an equilibrium whose value is the sum of the equilibriums seen as input thus far. Thus, when the end of the sequence is detected there must be some way of resetting the field so that the next sequence is not affected by the previous sequence. To do this, each field has a reset gate whose purpose is to detect the conditions under which the field should be reset to its default equilibrium.

The default equilibrium is the starting state to which subsequent equilibriums are added. The model hypothesizes that a reset signal represents constraints arriving top-down from higher cognitive processing levels (e.g., syntactic, semantic, pragmatic) as well as bottom-up from the features contained in the input data.

The Send Gate

Each layer of the model has an associated Send Gate which controls the information that it sends to the next highest level of cognitive processing. The first layer's send gate connects the letter detector field to the word detector field and the second layer's determines the overall output of the model. In different configuration of the model, the second layer's send gate could connect to a third field and so on. Send Gate processing is the same for every layer (refer to Figure 4). First, the input features A are presented, and the field is updated B. The cohort is then calculated C and evaluated by the Reset Gate D. The field may or may not then be reset to its default state. The cohort is calculated and if it has shrunken to one member, the Send Gate E opens.

Notice that the Send Gate only opens when the cohort has shrunk to one member. Thus we must ensure that the state of the cohort is reset so that a new sequence can be subsequently recognized,otherwise a feature that is repeated across category boundaries will not be recognized. The Send Gate behavior models the recognition point prediction of the COHORT model which states that word recognition occurs as soon as sufficient information is received such that all other candidates are eliminated (Marslen-Wilson, 1987).

Model Evaluation

The primary goal of our research was to determine whether neural fields are a plausible way to model word perception. Prior research theorizes that humans represent word forms as categories, abstracted away from variability (Dahan & Magnuson, 2006) and it is this view that our model seeks to explore. There are several well-known cognitive models (e.g., COHORT, TRACE, Neighborhood Word Activation) whose theories make different predictions once the input signal make contact with the lexicon. We chose the COHORT model as a starting point and explore whether two of its predictions can emerge from our neural field model: word-initial cohort and the identification of word boundaries. The operation of the model is summarized as follows:

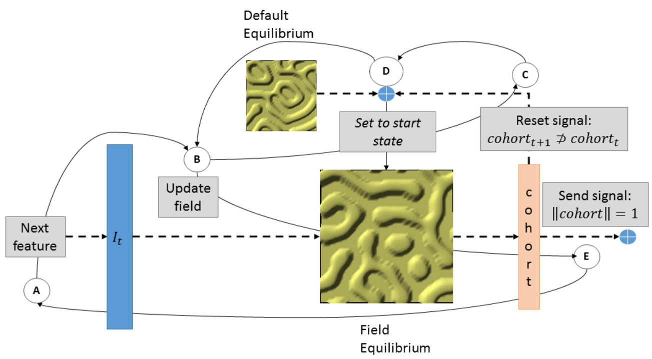

Figure 4: Reset/Send Signal Processing. If the reset gate (D) is open, it will set the field to a default equilibrium; otherwise it will update the field to the sum of the equilibrium of the current field and the equilibrium of the new feature. Send gate processing (E) takes place *after* the reset gate is processed.

1. Each feature (i.e., letter segment) of sensory input is converted to a pattern.
2. Sequences are generated by adding the current input's pattern to the previous input's pattern.
3. Perception arises as the generated pattern predicts a set, or "cohort", of letter or word candidates. A perceptron is trained to decode the pattern and interpret the prediction.
4. As new features are fed into the model, a new pattern is generated and the perceptron updates its prediction, removing candidates from the cohort.
5. Letter or word recognition occurs when there is one candidate left in the cohort.
6. New categories are recognized at the point when either the cohort is empty or when new candidates are added.

Data

The TIMIT corpus (Garofolo, Lamel, Fisher, Fiscus, & Pallet, 1993) provides a set of 10 phonetically rich sentences spoken by 630 speakers of eight major dialects of American English which are annotated at the word and phoneme level. The annotations of the corpus were used as a set of naturally occurring sequences to train the model's letter and word detectors. The text of the corpus was used to create feature vectors, as described in *Model Input*, which was presented to the $F1$ as a sequence of letter segments.

Results

The entirety of the TIMIT training set was pre-processed by the visual feature recognizer and its output, an unbroken sequence of letter segments was presented as input to the model. In the first experiment, the model was artificially reset to a default state at the end of every word so that errors in the perception of one word did not affect the perception of other words. This was done to verify correct operation of the model. For 100% of the words in the lexicon, the activation matched the ground truth for every letter segment in that word. Furthermore the model generated the correct cohort (when one existed) of letters for every letter segment sequence and of words for every letter sequence. In a separate experiment (Valenti, Oosterveld, & Scheutz, 2017), the model was not reset and in 82.5% of cases, the model detected the word level transition, suggesting that bottom-up information alone is insufficient to detect word boundaries.

Discussion

The model uses the structure of the data to represent top-down cues which are simulated through a "forced reset" when the start of a new letter or word is detected from the structure of the input data set. This is not ideal but allows the model to continue processing when the bottom-up cues alone are insufficient. One alternative to the forced reset would be to train a detector to recognize likely word boundaries in a training corpus, using it to augment the existing cohort-based reset mechanism. Consider the following sequence of letters without any explicit separation (we could have equally used a letter segment sequence, but that would have been harder to visualize):

shewashedyourdarksuitingreasywashwaterallyear

Humans can usually distinguish each letter sequence of a word and consequently recognize each word of the target sentence; however it is not as straightforward for a computer model to do so. Without further information constraints, a naïve model might correctly reject all sequences of letters that form non-words (e.g., *shew*) but erroneously recognize legal words such as *suiting*, resulting in a syntactically implausible reading of the sentence. Our model attempts to discern sequence boundaries by exploiting the cohort dynamics when processing letters. As a sequence of letters is read into the model, a cohort of possible words is initially formed which shrinks in size until only a single word candidate is left; this is the word's recognition point. If a shrinking cohort begins to grow again when a new letter is added to the sequence,decided this might indicate the start of a new word sequence and that the model should reset the field (*D* in Figure 4).

Since the present design does not model higher level cognitive processes, we abstract over all those that might be relevant to detecting a category boundary and combine them into one signal per field called *forcedReset*. Specifically, the data is preprocessed by the visual feature recognizer so that the letter segments have been grouped into sequences by letter. This roughly corresponds to how the higher areas of the visual cortex constrain lower area feature sequences during perception (Friston, 2005). The model uses this information to force the letter detector field to reset at the start of every new sequence. Likewise, the model uses the word size as a top-down cue to force a reset in the $F2$ word detector. During evaluation, the percentage of times the model accurately detects a word boundary using only the bottom-up signal is calculated; the forced reset ensures the model can continue processing when there is insufficient bottom-up information.

Future Directions

The first version of COHORT assumed input to be an abstract phoneme string. Thus, we arbitrarily chose to present visual input to the neural field as an unambiguous, noiseless sequence of letter segments which made it easier to visualize the model's operation in its graphical user interface. Real-world data is noisy yet perception still arises from these cognitive "noisy channels". Developing a design that incorporates noisy channels is key to understanding situated cognitive processes. Similarly, the input is invariant. In the speech perception domain, humans can usually recognize what is being said regardless of the speaker's accent, gender, etc. The model design needs to incorporate the ability to map varying input to invariant representations in order to simulate human performance in most perception domains. The model uses bottom-up information contained in the input signal to determine word boundaries, which is insufficient for 100% accuracy. Training an additional perceptron on a large speech corpus such as TIMIT should allow the model to statistically learn when a word boundary is likely to occur and this can be as a top-down cue to be added to the reset signal and improve its accuracy. Lastly, human cognitive language processing in the auditory and visual domains is often studied using electro physiological measures such as Event-related Potentials (ERPs). We have previously demonstrated a mapping of a single neural field model's dynamics to an ERP component (Valenti et al., 2016). Extending this to multiple fields could lead to models of human cognitive performance under varying cognitive workloads.

Conclusion

We explore the cognitive process of word recognition by creating a dynamacist model of the COHORT theory of Marslen-Wilson (Marslen-Wilson, 1987). This theory describes how sensory input is mapped to a specific word from a person's mental lexicon. Whereas Marslen-Wilson predicted the identification of a word cohort from which a unique word is selected and recognized he did not address how it might arise functionally from the input signal nor did he specify an implementation of the model. Moreover, we know of only one implementation of COHORT (Johnson & Pugh, 1994); it too conceives of encoding the input as patterns from which a cohort emerges and resolves. However it does not discuss the underlying algorithm for this process nor how it was trained, so it is difficult to assess its plausibility. In contrast, the presented model provides a general way to encode sequences in patterns and to find positions within those sequences which is applicable to any type of sensory information unfolding over time.

References

Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological cybernetics*, *27*(2), 77–87.

Brady, M. C. (2012). *A field-based artificial neural network with cerebellar model for complex motor sequence learning*. Unpublished doctoral dissertation, Indiana University.

Dahan, D., & Magnuson, J. S. (2006). Spoken Word Recognition. In M. J. Traxler & M. A. Gemsbacher (Eds.), *Handbook of Psycholinguistics.*

Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society London, Series B, Biological Sciences*, *360*(1456), 815–836.

Garofolo, J. S., Lamel, L. F., Fisher, W. M., Fiscus, J. G., & Pallet, D. S. (1993). DARPA TIMIT acoustic-phonetic continuous speech corpus CD-ROM, NIST speech disc 1-1.1. *NASA STI/Recon technical report n*, *93*.

Géron, A. (2017). *Hands-on Machine Learning with Scikit-Learn & TensorFlow* (1st ed.). Sebastopol, CA: O'Reilly.

Grossberg, S. (2005). Adaptive resonance theory. In L. Nadel (Ed.), *The encyclopedia of cognitive science* (1st ed.). Wiley.

Johnson, N. F., & Pugh, K. R. (1994). A cohort model of visual word recognition. *Cognitive Psychology*, *26*, 240–346.

Kleinschmidt, D. F., & Jaeger, T. F. (2015). Robust speech perception: Recognize the familiar, generalize to the similar, adapt to the novel. *Psychological Review*, *122*(2), 148–203.

Koch, C. (2004). *The Quest for Consciousness: a Neurobiological Approach* (1st ed.). Englewood, CO: Roberts and Company Publishers.

Laszlo, S., & Plaut, D. (2012). A neurally plausible parallel distributed processing model of event-related potential reading data. *Brain and Language*, *120*, 271–281.

Marslen-Wilson, W. D. (1987). Functional parallelism in spoken word-recognition. *Cognition*, *25*, 71–102.

Mattson, M. (2014). Superior pattern processing is the essence of the evolved human brain. *Frontiers in Neuroscience*, *8*(265). doi: 10.3389/fnins.2014.00265

McClelland, J., & Rumelhart, D. (1981). An Interactive Activation Model of Contextual Effects in letter perception: Part 1. *Psychology Review*, *88*(5), 375–407.

Rumelhart, D., & Siple, P. (1974). Process of recognizing tachistoscopically presented words. *Psychology Review*, *81*, 99–118.

Valenti, A. P., Brady, M. C., Scheutz, M. J., Holcomb, P. J., & Pu, H. (2016). A neural field model of word repetition effects in early time-course ERPs in spoken word perception. In A.Papafragou, D. Grodner, D. Mirman, & J. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 2765–2770).

Valenti, A. P., Oosterveld, B., & Scheutz, M. J. (2017, Sep. 6-8). A neural field model of sequence perception. In *Cognitive Computational Neuroscience Conference 2017.* New York, NY.