In Proceedings of the 44th Annual Meeting of the Cognitive Science Society, 2022.

Constructing Individualized Computational Models for Dementia Patients

Peggy Fidelman[1] (peggyf@cs.utexas.edu), Uli Grasemann[1] (uli@cs.utexas.edu), Claudia Peñaloza[2] (claudia_penaloza@ub.edu), Michael Scimeca[3] (mscimeca@bu.edu), Yakeel T. Quiroz[4] (yquiroz@mgh.harvard.edu), Swathi Kiran[3] (kirans@bu.edu) Risto Miikkulainen[1] (risto@cs.utexas.edu)

Abstract

Dementia is a common and debilitating condition that typically gives rise to increasing language impairment. There is a need to understand the nature of this impairment further so that therapies may be developed, particularly in the case of bilinguals. This paper extends BiLex, an existing computational model of bilingual lexical access, to simulate language decline in dementia. Six lesion types are evaluated for their ability to reproduce the pattern of decline in the semantic variant primary progressive aphasia (svPPA) subtype of dementia. Semantic memory lesions reproduce this pattern of decline best in monolinguals, and further suggest patterns that are likely to be found in longitudinal data from bilingual dementia patients in the future.

Keywords: personalized medicine; bilingualism; computational modeling; neural networks

Introduction

Dementia affects 55 million people worldwide, and approximately 10 million new cases are diagnosed each year (World Health Organization, 2021). Although there are therapies and medications that can slow its progression, there is not yet a known cure, nor any way to reverse the resulting decline. Most dementia patients experience increasing difficulties with language, which is often one of the most distressing symptoms because of the resulting loss of connection to others. Therapies that can slow the progression of these language difficulties are thus extremely beneficial for preserving quality of life in dementia patients.

An important but sometimes overlooked aspect of language behavior is that more than half the world's population speaks at least two languages (Grosjean, 2021). In the United States, census data indicates that bilingualism is increasing, and this may reflect a worldwide pattern (Grosjean, 2021). For individuals with language disorders who speak more than one language, determining the most effective treatment is more complex than in monolinguals because there are interactions between languages that affect the outcome of treatment. For example, treatment in one language may or may not improve ability in the individual's other language(s), and the effect may not be symmetrical nor consistent between different bilingual patients (Kiran, Sandberg, Gray, Ascenso, & Kester, 2013). There is thus an increasing need for therapies that can preserve language ability in bilinguals with dementia; it may also be possible to develop more effective therapies that take advantage of the patient's ability to speak multiple languages.

Obtaining sufficient data to guide this process is a major challenge, for three reasons. First, dementia patients, as well as their families and caregivers, face substantial additional burdens on their time due to the demands of living with dementia and caring for an affected person, which leaves little time for study participation. Second, bilinguals vary widely in terms of their proficiency in each language, and it is necessary to recruit study participants across the entire range. Third, data collection is extremely time-consuming. The process of testing to get a single data point can take hours, and in order to understand the nature of decline, data needs to be acquired at several points in time, with enough time in between so that meaningful decline may be observed. These requirements limit the amount of data that can be acquired, which in turn limits the rate at which potential treatments can be developed and evaluated.

A possible solution is to employ computational simulations. Computational models can be built based on current understanding of relevant neuroscience and psychology, and constrained with data on available human subjects. They can then be fit to new human subjects, making it possible to predict how various treatments may affect the decline.

As an instantiation of this approach, this paper presents a computational model of lexical access and semantic comprehension that can be applied to Spanish-English bilingual dementia patients. The model is an extension of BiLex, an existing computational model of bilingual lexical access (Peñaloza, Grasemann, Dekhtyar, Miikkulainen, & Kiran, 2019). While BiLex is a neural network model, it is different from deep learning models that require millions of parameters and training examples to construct. Instead it is a constrained model, incorporating principles from the neuroscience and language literature such as multiple maps and connections between them (Kroll & Stewart, 1994). These principles establish biases that make it possible to construct accurate models with training data from only a small number of subjects. BiLex was calibrated with data from 28 healthy Spanish-English bilinguals and five monolinguals, and shown to accurately simulate naming abilities across the full range of bilingual profiles (Peñaloza et al., 2019). In this paper, BiLex is extended with several possible pathologies, and shown able to account for the specific characteristics of two individual dementia patients.

[1]The University of Texas at Austin
[2]University of Barcelona
[3]Boston University
[4]Massachusetts General Hospital, Harvard Medical School

These BiLex models provide value in two ways. First, they can be immediately useful in planning for care of individual patients. Knowing the level of communication possible in each language over time allows making better-informed decisions about living arrangements and degree of support required. Second, they can help accelerate scientific progress by allowing initial experimentation with new therapies in simulation. Scarce patient resources can then be spent on evaluating therapies that have already been found promising in computational studies. Thus, the work in this paper paves the way for better understanding and treatment of an important segment of dementia patients in the future.

Assessing Dementia

All types of dementia can give rise to language deficits. The nature and severity of these deficits vary between subtypes of dementia, as well as between individual patients. Some types of dementia, such as semantic variant primary progressive aphasia (svPPA), lead to deficits in semantic processing as well as severe word-finding difficulties. Other types, such as mild cognitive impairment (MCI), leave semantic processing intact and cause much milder difficulties with word-finding (Cummings, 2020).

One commonly used test of word-finding ability is the Boston Naming Test (BNT), which consists of 60 line drawings of objects which the subject is asked to name (Roth, 2011). This set of objects includes high-frequency and low-frequency items, leading to a range of difficulty. Shorter forms of the BNT exist, typically consisting of 15 or 30 items, where the full range of difficulty is still represented.

Semantic memory is commonly assessed with the Pyramids and Palm Trees test picture version (PAPT) (Howard & Patterson, 1992). In this test, subjects are shown one picture as a stimulus, and then asked to choose one of two other pictures based on relatedness. The example after which the test is named is a picture of a pyramid as the stimulus, and the subject is asked to choose between a picture of a palm tree and a picture of a fir tree. The full version of this test includes 52 items, but shorter versions (with e.g. 26 items) are sometimes used. Note that no words are involved in this assessment, enabling it to test semantic knowledge without requiring word-finding.

These two tests together indicate to what extent a patient's language difficulties stem from underlying semantic difficulties versus word-finding. Our model relies on these tests as a measure of a subject's naming and semantic abilities.

BiLex Model

BiLex is a computational model of the mental lexicon in bilinguals, inspired by the Revised Hierarchical Model of Kroll and Stewart (1994). It can predict naming performance accurately in healthy bilinguals given an individual's history of exposure to the two languages (Peñaloza et al., 2019). It can also be lesioned to match an individual's characteristic language deficits following a stroke, then used to accurately predict individualized treatment outcomes that would result from post-stroke language therapy in each of a patient's languages (Grasemann, Peñaloza, Dekhtyar, Miikkulainen, & Kiran, 2021). A clinical trial is currently underway to evaluate BiLex as a tool for clinicians to select the most effective treatment language in Spanish-English bilingual stroke patients; if successful, it would be the first computational model employed in this role. Given these prior studies, BiLex provides a promising foundation on which to build an individualized model of language decline in bilinguals with dementia.

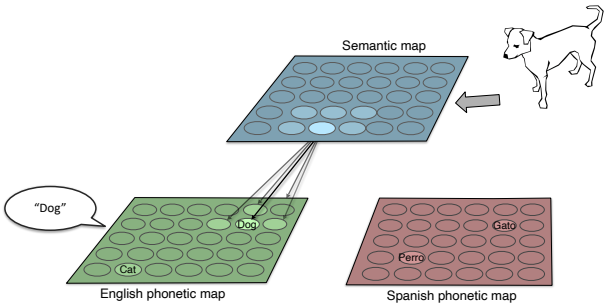

Figure 1: The BiLex model. The model consists of three SOMs: one for the semantic representations of words shared across languages, and two for the phonetic representations of words in each language. The model can be fit to the language history of individual patients, lesioned to model naming impairment in aphasia and dementia, and trained further to evaluate possible treatment options.

Consistent with Kroll and Stewart's model, BiLex consists of three interconnected maps: a phonetic map for each language and a shared map that represents semantic concepts. Each map is implemented as a self-organizing map (SOM; Kohonen, 1982), and all three maps are fully connected via bidirectional associative connections, as depicted in Figure 1.

BiLex is trained with a corpus of words in English and their direct translations to Spanish (638 concrete nouns in the simulations in this paper). Words are represented by two types of vectors in BiLex. One is an encoding of the semantic features of a word, where each position in the vector represents a semantic feature (e.g. "can fly"). The other type is an encoding of the phonetic features of the word. Since BiLex is a bilingual model, each word has three representations: one semantic vector and two phonetic vectors (Spanish and English).

BiLex training follows the normal SOM training process (Kohonen, 1982). All three maps are trained simultaneously, and the associative connections between the maps are strengthened via Hebbian learning, so that connections between simultaneously activated SOM nodes are strengthened proportionally to the activation of the two nodes they connect. Training occurs throughout the entire simulated lifetime, and the training hyperparameters vary to capture the way language learning capacity declines with age in humans: the learning rate and neighborhood size are decreased over time, and random noise is added to the associative connections throughout the entire simulated lifetime to capture lan-

guage attrition effects. The exact values and rate of change for these hyperparameters were fit previously to healthy human data using evolutionary optimization (Peñaloza et al., 2019; Grasemann, Miikkulainen, Peñaloza, Dekhtyar, & Kiran, 2019).

After training, BiLex's performance is evaluated using simulated BNT tests in both English and Spanish, as well as a simulated PAPT test, similarly to testing of patients.

Extending BiLex to Model Dementia

Dementia and stroke both give rise to language deficits resulting from neurological damage, but the nature of the damage and the resulting deficits differ significantly. Whereas stroke typically consists of a single event in which brain tissue is damaged, dementia is progressive in nature. Dementia also arises from several different brain abnormalities which give rise to different patterns of language decline. Therefore, it is useful to evaluate several different ways of lesioning BiLex progressively, so that it is possible to model as many of the different subtypes of dementia as possible.

Lesion Types

This study evaluates six possible lesion types targeting the functional areas in BiLex analogous to the brain areas and functions that are affected by dementia pathology. These include the semantic map (analogous to semantic memory; Lesions 1, 2, and 3), propagation between the semantic and phonetic maps (analogous to connections between semantic memory and phonetic memory; Lesions 4 and 6), and the input to the semantic map (analogous to the combined sensory and cognitive inputs; Lesion 5).

Candidate Lesion 1: Random Deletion of Neurons in the Semantic Map Nodes in the semantic map are randomly chosen to be marked as lesioned, meaning that that they will no longer activate in response to the input nor propagate any activation to phonetic maps. Once a node is marked as lesioned, it will remain that way for the rest of the simulated lifetime. Beginning at the simulated age of onset, progressively more nodes are added to the lesioned set over time.

Candidate Lesion 2: Focused Deletion of Neurons in the Semantic Map This lesion is identical to Lesion 1, except that the next positions to be lesioned are chosen so that they are adjacent to other lesioned nodes. The position of the first node to be lesioned is chosen at random.

Candidate Lesion 3: Blurring of Semantic Features in the Semantic Map At each time step during the lesioning phase, the weight of each feature of each node in the semantic map is changed to be slightly closer to the average of its neighboring nodes' weights for that feature. This lesion is implemented using a pooling operation over neighbors, as is commonly used in some layers of a convolutional neural network. It is intended to simulate the "conceptual averaging" observed in semantic dementia patients.

Candidate Lesion 4: Deletion of Associative Connections Connections from the semantic map to the two phonetic maps are removed at random. Once a connection is removed, it no longer propagates activation from the semantic map to the phonetic map, and it can no longer be strengthened via the Hebbian learning process.

Candidate Lesion 5: Deletion of Features from the Semantic Input Vector Positions in the incoming semantic vector, each of which corresponds to a semantic feature, are chosen at random to be added to the set of lesioned features. These positions in the input vector will be ignored subsequently when choosing the semantic map node that most closely matches the input vector.

Candidate Lesion 6: Abnormal Spread of Activation in the Semantic Map During normal operation of BiLex, the neighborhood size in all three maps decreases gradually as the simulated age increases. This process is commonly used with SOMs to allow the maps to organize at a global scale early in the training process and then to refine their organization at a local scale later according to more fine-grained similarities in the data. In BiLex, the neighborhood size also effectively determines how many neurons adjacent to the winning neuron in a given map can propagate their activation to the other maps via the associative connections. In this lesion, semantic map activation and subsequent propagation to the phonetic maps is perturbed by progressively *increasing* the neighborhood size in the semantic map during the decline. Thus, progressively more nodes in the semantic map are allowed to propagate activation to the phonetic maps than in a healthy model, causing the activation in the phonetic maps to become less focused.

Lesion Progression

When applying any of these candidate lesions to BiLex, two choices must be made: when to start the lesioning process, and the rate at which it should progress. Because the age of onset and the rate of progression of dementia can both vary among patients, even within the same subtype of dementia, these are choices that must be tailored to each individual patient.

The choice of when to start the lesioning process is informed by the timing of the patient's diagnosis and when the symptoms were first noticed. However, this information does not provide an exact answer, as neurological damage occurs for some time before it is possible to formally diagnose dementia. To deduce the appropriate rate of progression, some indication of the individual patient's rate of decline is needed. The most direct way is to obtain BNT and PAPT scores at two timepoints with enough time in between so that the scores can decline meaningfully.

Given such information, it is possible to find the best lesioning start time S and rate of progression P using the following search process: for each combination of S and P within a range of plausible values, a BiLex model is trained according to the language exposure history and with lesion-

ing starting at S and progressing at rate P. Each combination of S and P is evaluated for fit to the patient's actual scores (according to the sum of the squared distances) and the S and P with the best average fit (over the available time points and languages) is found.

Predicting Bilingual Dementia

The above process can be repeated with all six lesion types, evaluating how well each one of them explains the patterns of decline in the two languages. It will then be possible to identify the lesions that provide the best way to model dementia computationally. However, longitudinal BNT and PAPT data do not yet exist on bilingual dementia patients. Therefore, a two-step process is employed in this paper:

First, the six lesion types are implemented in models of healthy bilinguals, and their effect on the decline of the two languages characterized. Second, these lesions are evaluated based on how well they predict the decline in monolingual dementia patients, for which data does exist. Thus, the second experiment allows identifying the best lesions; the first experiment shows what the predicted patterns for that lesion are. These experiments are described in the next two sections.

Experiment 1: Qualitative Comparison of Candidate Lesions

Language exposure history data from two healthy bilingual controls whose data were collected as part of the initial tuning of the BiLex model (Peñaloza et al., 2019) were used to generate predictions of the shape of decline. These two subjects were chosen because they represent different degrees of bilingualism: one is significantly more proficient in Spanish than English at the age when lesioning begins, and the other has approximately equal abilities in both languages.

For each combination of candidate lesion and human subject, 20 BiLex models were trained with the subject's language exposure history data and lesioned beginning 11 simulated years before the last data point. The rate of progression P was chosen so that full decline happens in 10 simulated years: At that point, no unlesioned neurons remain in the semantic map (Lesions 1 and 2), the weights are fully averaged (Lesion 3), no unlesioned associative connections remain (Lesion 4), no unlesioned positions in the semantic input vector remain (Lesion 5), and the propagation neighborhood contains the entire semantic map (Lesion 6).

As can be seen in Figure 2, the six candidate lesions result in different patterns of decline. Lesions 4 and 6 do not affect PAPT score at all; Lesion 3 affects PAPT score before any decline in BNT scores is observable; and Lesions 1, 2, and 5 affect PAPT scores and BNT scores simultaneously, but with varying differences in impairment between them. The shape of the decline in BNT scores also varies between candidate lesions, with some showing steeper decline early, some showing steeper decline later, and others showing approximately linear decline throughout the entire lesioning period. The second experiment then identifies which ones of these patterns are the best match with those observed in patients.

Experiment 2: Matching Candidate Lesions with Patient Data

Models with different lesions were fit to the data of two English-speaking monolingual svPPA patients included in the study of Flurie et al. (2020). For each patient, four timepoints of naming accuracy scores were reported, spanning a period of 20-23 months, with six to eight months between timepoints. PAPT scores exist for each patient at the first and last timepoint. The age of each patient and the time since onset were also included.

These data were collected as part of a study investigating the effectiveness of a maintenance-based treatment intended to preserve naming ability on the specific words included in the treatment. Some of the words included in the treatment were also words that appear in the BNT. To avoid this confound, the model's naming accuracy was measured only on words that were not included in the treatment.

Both patients have significantly impaired naming and PAPT scores, with naming more severely affected, and both naming and PAPT scores continue to decline over the course of the disease in both cases. This pattern immediately allows elimination of Lesions 4 and 6, neither of which causes PAPT scores to decline (Figure 2). Similarly, Lesion 3 can be eliminated because its PAPT score first declines steeply and then plateaus (and its naming scores do not fall to the patients' level until its PAPT scores have plateaued at a significantly lower level than those of both patients).

For each of the remaining three candidate lesions, an exhaustive search was run over a range of values for P and S, subject to the constraint that S cannot be later than the patient's reported age of onset. For each combination of lesion type, P, and S, 40 BiLex models were trained with language exposure history set to 100% English and 0% Spanish, with lesioning starting at simulated age S and progressing at rate P.

Figure 3 shows the best fits of S and P for each of the three candidate lesions. Lesions 1 and 2 provide the best match with the patient data. They both involve progressive deletion of semantic map neurons, thus suggesting a likely source of pathology in svPPA. Patterns for these lesions in Figure 2 then predict how this type of dementia is likely to progress in bilingual patients. More specifically, more balanced bilinguals are predicted to decline similarly to the top two plots on the right column, and bilinguals with one dominant language similarly to the top two plots on the left.

Discussion and Future Work

Two of the six candidate lesions fit the data of monolingual svPPA patients well. The patterns of decline in the two languages shown in Figure 2 then constitute predictions on svPPA decline on bilingual patients, and similar predictions can be generated for future patients with different language histories. As longitudinal data become available for bilingual svPPA patients, these predictions can be tested. Accurate predictions then build confidence that BiLex can be used to evaluate treatment options for individual patients in the future.

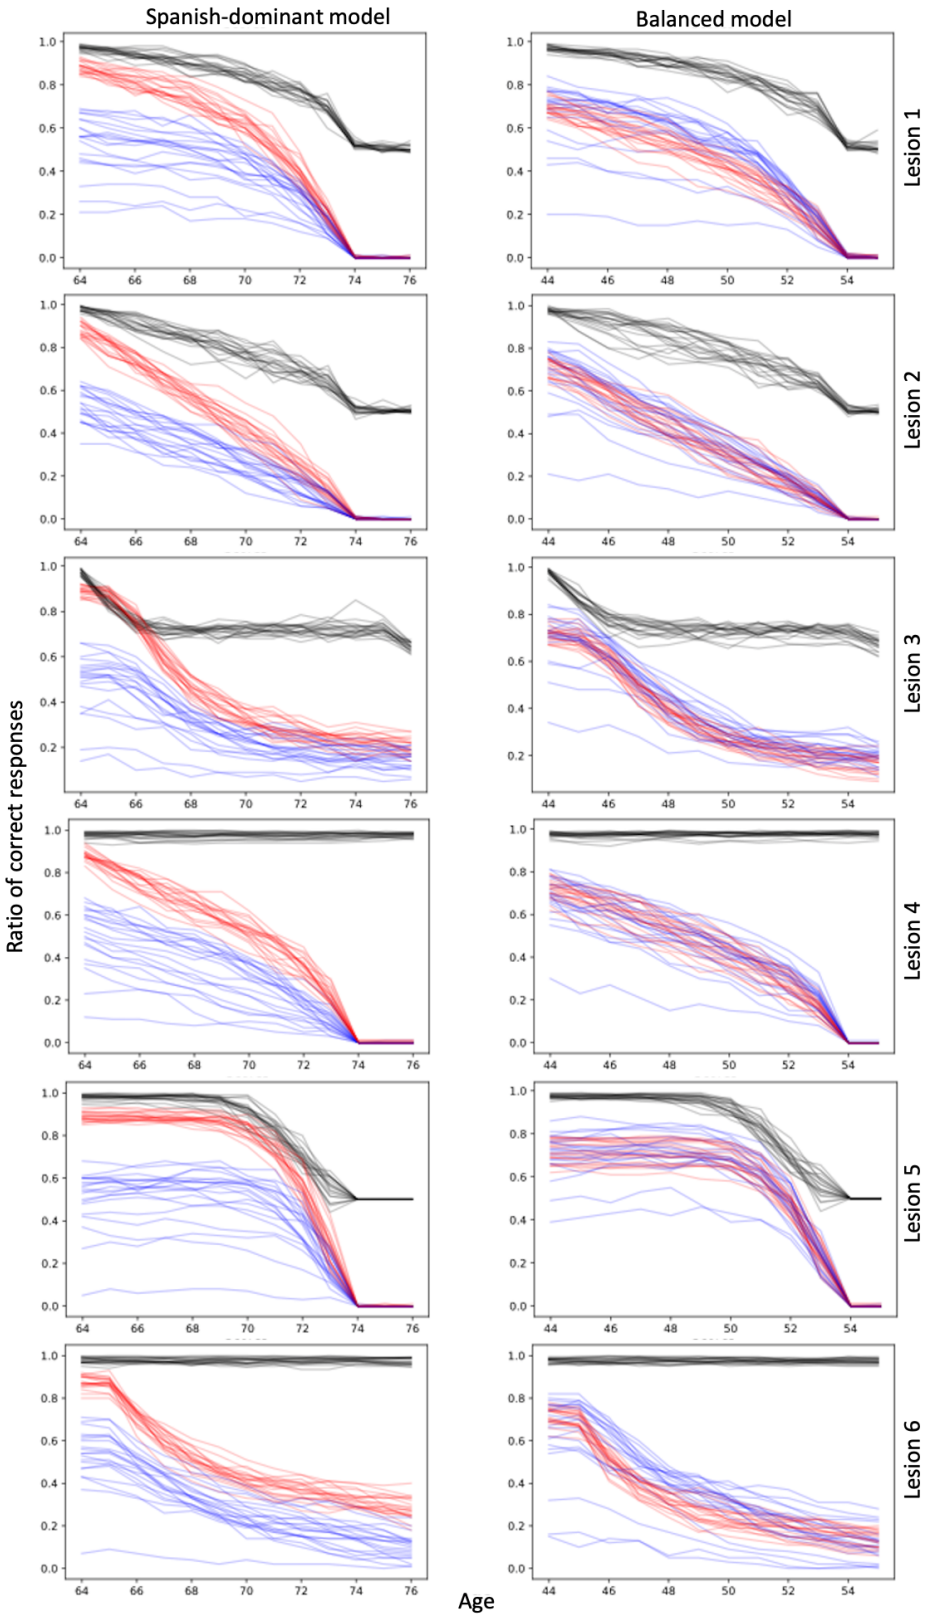

Figure 2: Decline observed with the six candidate lesions in a Spanish-dominant model (left) and a balanced bilingual model (right). Black lines represent PAPT scores, red lines represent Spanish BNT scores, and blue lines represent English BNT scores. Candidate lesions are arranged according to their indices, with Lesion 1 at the top of the figure and Lesion 6 at the bottom. The lesions result in different patterns of decline; Lesions 1 and 2 provide the best match with existing monolingual data (Figure 3), predicting that the pattern in the top two rows will be observed in bilingual longitudinal data in the future.

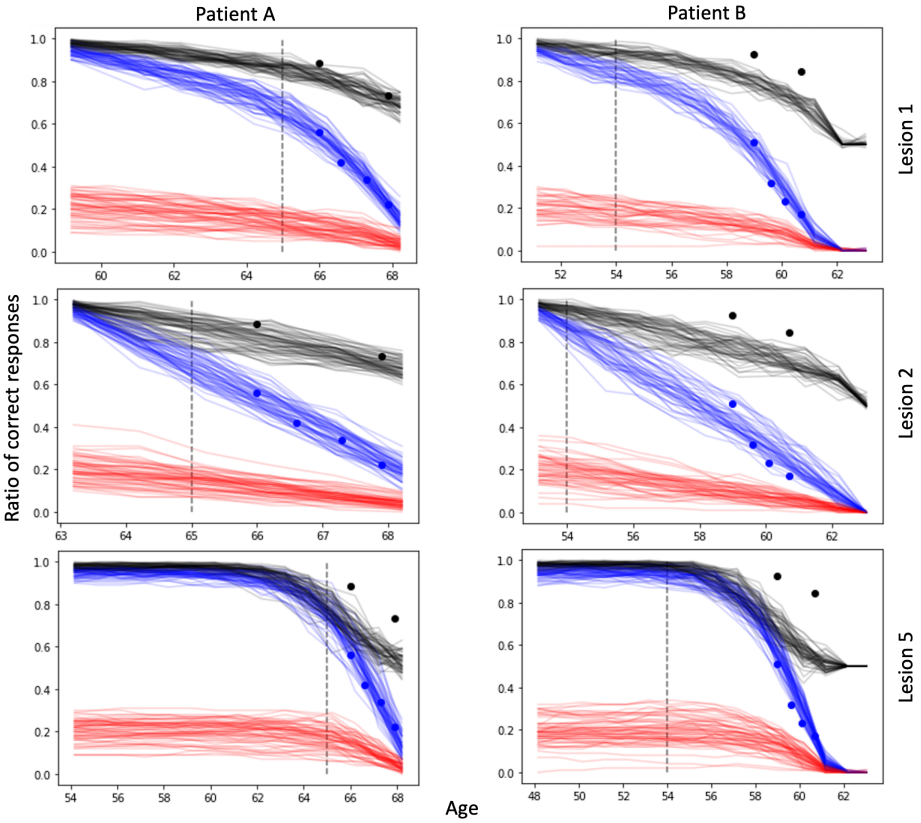

Figure 3: Best fits of candidate lesions to data from two svPPA patients in the study of Flurie et al. (2020). Fits for Patient A are on the left and fits for Patient B on the right. Lesions 1, 2, and 5 are shown since they are the only ones that result in a qualitative match to patient data. Black dots represent actual patient PAPT scores, and blue dots represent actual patient naming scores. Similarly, black lines show the model's PAPT scores, and blue lines show the model's English naming scores. The dashed vertical line is the patient's reported age of onset. Lesions 1 and 2 provide the best match with data, suggesting that these lesions can be used to predict the patterns of bilingual decline as well, as shown in Figure 2.

It is likely that other subtypes of dementia will be best modeled with different types of lesions. For instance, MCI patients characteristically are not impaired on PAPT (Taler & Phillips, 2008). Therefore, Lesions 4 and 6 may be appropriate for modeling MCI, despite not being good choices for modeling svPPA. In future work, we plan to fit the current candidate lesions to the pattern of decline seen in MCI. This process can be repeated for other subtypes of dementia as longitudinal data become available for them.

Another possibility for the future is to combine some of the lesion types to achieve a more precise fit to patient data. For example, Lesion 4 (which does not affect PAPT scores) can be combined with either Lesion 1 or Lesion 2 (both of which lead to PAPT scores slightly below the patients' actual scores in Experiment 2) to achieve a closer fit to the data from the two monolingual patients reported in Experiment 2. However, allowing lesion types to be combined weakens the constrained nature of the model and increases the danger of overfitting the small amount of patient data available to us at this time. As more data become available, it will be possible to investigate how lesion types can be combined to generalize across all the available data.

In the current model, lesions were assumed to progress at a steady rate from the time they begin until the end of the simulated lifetime. Variable rates may be more accurate, especially rates that either increase or decrease steadily (unconstrained variation would raise the risk of overfitting). If longitudinal data with more than a few time points becomes available, this possibility can be evaluated as well.

Conclusion

While it is important to be able to take into account bilingualism in understanding and treating dementia, it is difficult to obtain comprehensive enough data to cover the variety of patients. Computational modeling can be a crucial tool, but the models have to be principled in order to be sufficiently constrained with the available data. BiLex, a self-organizing map model with associative connections, is based on neuroscience and psychological principles that make it possible to fit it to individual patients. With longitudinal data on monolingual patients and healthy data on bilinguals, it is possible to identify the candidate lesions that predict the likely patterns of decline in bilingual dementia. Once verified with bilingual dementia data, these models can serve as a foundation for identifying the most effective treatments in the future.

Acknowledgments

We would like to thank Jamie Reilly, Maurice Flurie, and Molly Ungrady for sharing the patient data which made Experiment 2 possible, for taking the time to explain the details of their testing process, and for valuable discussions that inspired some of the lesion designs. This research was supported in part by NIH under grant 3U01DC014922-05S1.

References

Cummings, L. (2020). *Language in dementia*. Cambridge: Cambridge University Press.

Flurie, M., Ungrady, M., & Reilly, J. (2020). Evaluating a maintenance-based treatment approach to preventing lexical dropout in progressive anomia. *Journal of Speech, Language, and Hearing Research*, *63:12*, 4082–4095.

Grasemann, U., Miikkulainen, R., Peñaloza, C., Dekhtyar, M., & Kiran, S. (2019). Evolutionary optimization of neural-network models of human behavior. *Proceedings of the International Conference on Cognitive Modeling*.

Grasemann, U., Peñaloza, C., Dekhtyar, M., Miikkulainen, R., & Kiran, S. (2021). Predicting language treatment response in bilingual aphasia using neural network-based patient models. *Scientific Reports*, *11*, 10497.

Grosjean, F. (2021). *Life as a bilingual: Knowing and using two or more languages*. Cambridge: Cambridge University Press.

Howard, D., & Patterson, K. E. (1992). *The pyramids and palm trees test.* Pearson Assessments, https://www.pearsonassessments.com.

Kiran, S., Sandberg, C., Gray, T., Ascenso, E., & Kester, E. (2013). Rehabilitation in bilingual aphasia: Evidence for within- and between-language generalization. *American Journal of Speech-Language Pathology*, *22*, S298–S309.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, *43*, 59–69.

Kroll, J. F., & Stewart, E. (1994). Category interference in translation and picture naming: Evidence for asymmetric connections between bilingual memory representations. *Journal of Memory and Language*, *33:2*, 149–174.

Peñaloza, C., Grasemann, U., Dekhtyar, M., Miikkulainen, R., & Kiran, S. (2019). Bilex: A computational approach to the effects of age of acquisition and language exposure on bilingual lexical access. *Brain and Language*, *195*, 104643.

Roth, C. (2011). Boston naming test. In J. S. Kreutzer, J. DeLuca, & B. Caplan (Eds.), *Encyclopedia of clinical neuropsychology* (pp. 430–433). New York, NY: Springer New York.

Taler, V., & Phillips, N. A. (2008). Language performance in alzheimer's disease and mild cognitive impairment: A comparative review. *Journal of Clinical and Experimental Neuropsychology*, *30:5*, 501–556.

World Health Organization. (2021). https://www.who.int/news-room/fact-sheets/detail/dementia.