Implications of Developments in Machine Learning for People with Cognitive Disabilities

White Paper for Coleman Institute for Cognitive Disabilities

Clayton Lewis
Coleman-Turner Professor of Computer Science
University of Colorado Boulder
clayton.lewis@colorado.edu

November 2, 2018

Abstract: Machine learning is the process of training a computer program to reproduce and generalize the relationship between some kind of input and some kind of output. High profile successes for machine learning are attracting interest and investment. Can this technology be applied in ways that benefit people with cognitive disabilities and those who support them? What are the current limitations in the technology that affect these applications, and what are the prospects for overcoming them? What actions can maximize the benefits for people with cognitive disabilities?

1. Introduction

For many years the application of computers was based on programming by human beings. Humans needed to understand a problem, devise a solution to it, and express this solution as a sequence of instructions in a computer language. Recently, a very different approach has emerged, in which a computer system is trained to solve a problem, without any human programmer knowing or expressing a solution. This is machine learning.

In one common form of machine learning, training consists of the computer system processing a large collection of training examples, each of which specifies some input, and a desired output to be associated with that input. For example, the examples might consist of sentences in English, with associated sentences in Chinese. Having processed these examples, the system can then produce a sentence in Chinese when given a sentence in English, without having seen the sentence during training.

In another form, a computer system learns to take actions in a situation in which appropriate sequences of actions earn some kind of reward, such as a winning score in a game. This technique is called *reinforcement learning*. One can relate this to the form of training just mentioned, by considering that the system processes examples of states and associated actions, to learn what action to take in a given state. But rather than being told the correct actions, it determines that using the rewards it receives.

A striking example of success using machine learning is machine translation. While for many years machine translation was literally a joke, for some time now translations of practical value are freely available for many languages. The best translations are produced by machine learning systems, not by translation programs written by people. The machine learning systems are trained on large collections of translations created by human translators.

Another striking success story is that of the Alpha Zero game playing program, that learned to play three games, chess, shogi (Japanese chess), and Go (Silver et at., 2017). Alpha Zero learned to play chess and Go better than the strongest known human and machine players, in only several hours, with no human knowledge of how to play the games needed (beyond the rules). Its success with Go is especially notable, as for many years Go has been recognized as a much more demanding challenge than chess, for which machines attained world leading strength in 1997.

These successes rest on *deep learning*, the use of multiple layers of interconnected artificial neurons (see LeCun, Bengio & Hinton, 2015, who note that recent successes rest on much earlier ideas, rendered feasible by greatly increased computing capacity.) Differences between actual and desired outputs produce error signals that are propagated through the network of neurons, adjusting the connections among them so as to bring the output closer to what is needed.

What do these successes mean for people with cognitive disabilities, and those who support them? This white paper will discuss some key applications of machine learning that are potentially important for people with cognitive disabilities, consider the present state of the art, and discuss what challenges remain for successful implementation. It then presents a road map of possible actions for those interested in exploring and developing the potential of machine learning for our community.

2. Cognitive disabilities and technology

Because cognitive functions are extremely diverse and complex, differences among people in the operation of these functions are also diverse and complex. The Cognitive and Learning Disabilities Accessibility Task Force of the World Wide Web Consortium (https://w3c.github.io/coga/user-research/) lists ten categories of difference, including memory, executive function, reasoning, attention, language, and literacy. Each of these categories is further subdivided; for example, memory includes functions with different temporal characteristics (short term, long term, and working memory) as well as different content, for example memory for spatial patterns vs. sound patterns. Differences can be associated with diverse circumstances, including chromosomal variation, brain injury, effects of medications, aging, and many others.

As with other disabilities, the impact of these differences depends sensitively on the environment, and on the activities people wish to or are required to engage in. In general, people wish to be able to live interdependently as causal agents in their lives , m not having to depend on other people more than others do. Technology can modify the environment in such a way as to increase people's ability to live as they wish. This is considered the person-environment fit model that has come to define intellectual disability.https://www.apa.org/pi/disability/resources/publications/newsletter/2016/09/intellectual-disability-support.aspx   For example, a reminder system can make it easier for a person to remember to take medication, or to keep an appointment. Text can be read aloud for people who have difficulty decoding written language. In a series of studies, Bodine and collaborators at the Rehabilitation Engineering Research Center on Advancing Cognitive technologies, are seeking to relate the effectiveness of specific technologies to the functional capabilities of users with cognitive disabilities (http://www.ucdenver.edu/academics/colleges/Engineering/research/AssistiveTechnologyPartners/research/RERC/Projects/Pages/Projects.aspx).

Useful though it is, today's technology is limited in what it can do, compared with human assistants or caregivers. Thus some people have to rely on caregivers to keep them safe, by recognizing and avoiding risky situations, or to explain financial decisions.

As mentioned above, advances in machine learning are pushing back the boundaries of what technology can do. Can we expect these advances to enable people with cognitive disabilities to live more self-determined lives?

Progress on this has implications beyond the direct impact on people with disabilities. Current caregiving arrangements impose costs on family members and on social service programs, that may be reduced. Current caregivers might become more effective, and more available, able to serve more clients, if assisted by technology. These changes could help address the crisis in caregiving seen in most regions today: there are too few caregivers to meet the demand, and turnover is high (see for example https://acl.gov/news-and-events/news/aging-baby-boomer-generation-prompts-concerns-over-caregiver-shortage).

3. Application: Making text easier to understand.

People with cognitive disabilities may have more limited literacy skills. As a result, text that is easy to understand for many readers may be difficult for them. The Medicaid Reference Desk project (see Lewis & Ward, 2011) addressed this problem by using human editors to revise descriptions of Medicaid waivers to make them easier to understand for people with cognitive disabilities (who were often among the intended beneficiaries.)

Here is an example of a description from the Medicaid Reference Desk, first in its original form, and then as edited.

> *Original*: Occupational Therapy Services involve the treatment prescribed by a physician to develop, restore, or improve functional abilities related to self-help, adaptive behavior and sensory, motor, postural, and emotional development that have been limited by a physical injury, illness or other dysfunctional condition. It involves the use of purposeful activity interventions and adaptations to enhance functional performance.
>
> *As revised*: Occupational Therapy Services: Help with day-to-day activities and testing to see if you need special tools to help you eat, work and live in the community.

We can see a number of changes in this example. Rare words like “dysfunctional” and “postural” are gone. The first sentence in the original has four complex conjunctions: “develop, restore, or improve”, “self-help, adaptive behavior, and ... development”, “sensory, motor, postural, and emotional”, “physical injury, illness, or other dysfunctional condition”. The replacement sentence has only one conjunction, “eat, work, and live”. The obscure phrases “purposeful activity interventions” and “enhance functional performance” are replaced by “help” and “help you eat, work, and live in the community”.

Underlying these changes, one can see that the editors have understood what the original passage *means*, not just what it says. For example, there is no mention of “testing” in the original passage, but the editors know that for a physician to prescribe something, they have to do some form of testing. There is no mention of “community” in the original passage, either, but the editors know that a very important aspect of the abstract concept “functional performance” is being able to “live in the community”, for the people for whom this service is intended.

Unfortunately this human editorial process was apparently not sustainable, as the resource is not available today. Can natural language processing technology, based on machine learning, help with this problem?

One can envision two kinds of tools. One possible tool could be used by a content provider, like the Medicaid Reference Desk project, to produce more comprehensible versions of text automatically. A second tool, perhaps even more useful, could be used by a consumer to view any existing text, and would present an edited version. The second tool could be personalized, so that (for example) vocabulary would be tailored to the needs of an individual reader.

Some text simplification tools are available today, for example from TextHelp (toolbar from Chrome browser, and IBM (Content Clarifier). These can be used either by content creators or by readers. How do these tools stack up against the need?

We lack data on the effectiveness of these tools for readers with cognitive disabilities, so our assessment will be uncertain (though see Djamasbi, 2017). However, there appear to be a number of areas in which improvement should be sought.

Weakness in handling vocabulary. A system may lack a meaningful replacement for unfamiliar words. Work by Callison-Burch and collaborators is enabling words to be replaced by phrases, rather than by single words, and also phrases, rather than single words to be replaced (Pavlick & Callison-Burch, 2016a). For example, “hypertension” can be replaced by “high blood pressure”. or “medical practitioner” by “doctor”.

Coherence. Existing simplification techniques produce text that does not read smoothly. For this reason the Newsela service, that provides news feeds for schools, edited for different grade levels, uses human editing (Dan Cogan-Drew, personal communication, May, 2018).

Personalization. Ideally, a simplification tool should respect the reader’s vocabulary, rewording only vocabulary that the reader does not understand. While most research on simplification does not respond to this situation, Lee & Yeung (2018) explore a process in which one of four graded vocabularies is assigned to a reader, based on the reader’s responses on a 40-item vocabulary test.

Content augmentation. Often a text isn’t easy to understand because it assumes background knowledge that a reader lacks. Content Clarifier partially addresses this issue by enabling a reader to request supplementary information about things mentioned in a passage. But it seems likely that integrating this information into a rewritten passage would be better.

Use of situation model. Psychologist Walter Kintsch argued that human readers construct a *situation model* when reading: a mental description of the situation that the text deals with. In doing this readers routinely augment the content of the text itself with information drawn from their knowledge of the world, including inferences (for review see Zwaan & Radvansky, 1998). This can be seen in the Medicaid Reference Desk example, when the editors refer to “testing”, not mentioned in the original passage, but supplied by a knowledgeable reader.

3.1 Related progress in NLP Research

NLP researchers are attacking all of these problems. For example, Yejin Choi of the University of Washington, with collaborators, is working on increasing coherence in generated text by maintaining a kind of agenda of content to be expressed (Kiddon, Zettlemoyer, & Choi, 2016).

Choi and collaborators are also working on aspects of the situation model, by training systems to construct explanations of the events described in a text. This includes making inferences about the motives of the people who figure in a text, even though the motives are not mentioned (Rashkin et al., 2018). Choi is also working on processing and generating text using an abstract meaning representation (AMR) (Konstas et al., 2017), an approach that might be used to represent a situation model for a text , and perhaps, readers' background knowledge.

While progress is being made, there remain substantial obstacles. One is the lack of good measures of understandability, whether based on human judgement or on automated techniques. Without these, it is difficult to evaluate and compare new approaches. (Popular readability formulae, while useful for comparing the likely understandability of naturally occurring texts, are not suitable for comparing simplification methods, because they are based on factors that are correlated with understandability, but not important in themselves. For example, naturally occurring texts with shorter sentences are likely to be more comprehensible, but simply breaking up long sentences in a text often results in a text that is *less* comprehensible; see Redish, 2000). Xu et al. (2016) report progress on this matter, though they are still using some correlational measures.

A second issue might be called the Monkey's Paw, problem, after the W.W. Jacobs story in which a sinister talisman grants wishes in a way that is consistent with the wishes but is horrible in the means used. NLP systems sometimes perform a task as specified, but not in the way researchers would wish, or had expected. For example, a challenge competition was devised that was intended to promote progress in understanding the meaning of stories, by requiring systems to discriminate continuations that human readers found plausible or implausible. One of the best performing entries didn't even read the stories, finding that the continuations themselves contained sufficient cues to do well on the task (Schwartz et al. 2017). This system accomplished the explicit goal of the task, but avoided the processing researchers had hoped to observe. Pavlick & Callison-Burch (2016b) show that near-human accuracy in a test of inferences humans make when reading can be obtained by simply counting in the training materials whether inferences that feature a given adjective are more often accepted or rejected, regardless of context. They also note that this performance is considerably better than that given by any available, more sophisticated method, showing that artificial systems have some way to go to represent human-like language ability. Pavlick discusses these issues in a 2018 talk, available at http://www.ipam.ucla.edu/abstract/?tid=14546&pcode=DLT2018 .

4. Application: Artificial cognitive assistants.

Human beings assist one another under many conditions. What are the prospects for artificial agents, trained using machine learning, playing this role? We can distinguish two forms this assistance could take. In one form, a person wanting assistance initiates an interaction with an artificial agent. The agent then has to interpret the request the person is making, together with relevant context, and then to carry out appropriate actions. We'll call such an agent *passive*, since it does nothing until activated by a person. In a second form, an artificial agent identifies a situation in which a person needs assistance, without the person needing to be aware of the need. We'll call such an agent *active*, since it does not have to be activated by a request from a person.

A particular application of an active agent could be as part of a remote caregiving service. Today, human caregivers monitor a client's living space, watching for signs of danger or distress that the client might not recognize, or might not be able to deal with. Could an active artificial agent learn to play this role, helping a human caregiver notice situations that need attention?

Another application for an active agent could be in coaching "soft" job skills, for example teamwork, decision-making, and communication. Having or lacking these skills is often crucial for employment (US Department of Labor, 2012). Could an active artificial agent monitor a person's interactions in the workplace, and provide suggestions for improvement?

Machine learning researchers are generally pessimistic about creating artificial cognitive assistants today. They identify two key challenges, the challenge of data, and the challenge of understanding.

The *data* challenge arises from the fact that machine learning systems generally require a large number of examples of whatever the system needs to observe, together with an appropriate response for each example. Because the space of situations in which an artificial cognitive assistant may have to function is very large, complex, and diverse, collecting a useful corpus of examples is challenging.

There are some ways to temper this challenge, however. One is to use *unsupervised* learning, in which examples need not be associated with appropriate responses, which greatly simplifies the problem. The system itself determines the goodness of possible responses, often by testing the quality of a model of the situation it constructs. A good model will be able to predict what will happen next, or may allow the details of a situation to be reconstructed from a simplified representation. No human input is needed to do this.

Izo & Grimson (2007) describe a surveillance camera that learns to detect unusual behavior in a parking lot by building a model of the behaviors that it sees. In the parking lot example, the camera learns that objects of a certain size move along certain paths, while objects of a different size move along different paths. We know, but the camera does not, that these objects are cars and pedestrians. After operating for some period of time, things that do not fit its model can be flagged as unusual. For example, if a car sized object moves into an area where those objects don't normally move, this can be noticed. This could be explored as an approach to creating active cognitive assistants, for remote caregiving, or for monitoring workplace activity, alerting a job coach of something unusual occurring.

Another approach to the challenge of data is to pretrain a system to handle a wide range of different situations, and then train it specifically to work in a given situation of interest. The hoped for result of this *metalearning* procedure (Schmidhuber, 1987; Finn, Abbeel, & Levine, 2017; many others) is that the pretraining enables the system to learn what to do in the situation of interest with very few examples. The pretraining may or may not employ special processing of the data to prepare the ground for subsequent learning.

Extensive pretraining would be expensive, but it would be shared among many situations of interest. For example, a remote caregiving monitor might be trained (expensively) in many different client homes; it might then quickly and cheaply learn how to monitor events for a new client in a new home.

An issue in training systems to cope with a very wide range of tasks is forgetting: changes to a system needed to handle a new task may interfere with what has been learned about previous tasks. Achille et al. (2018) report progress in devising representations that allow new tasks to be learned without disrupting old ones.

The data challenge affects a great many applications of machine learning, so we can expect continued research focus on it. For example, DARPA includes this challenge as a focus of its $2B AI Next campaign (https://www.darpa.mil/work-with-us/ai-next-campaign).

The challenge of *understanding* can be illustrated using the surveillance camera example. While the camera can identify unusual behaviors, it has no understanding at all of what it is seeing. As mentioned, while its model distinguishes objects of different sizes, it has no information about what the objects are. Nor does it have any interpretation of the various trajectories that are described in its model, for example that some trajectories of small objects (pedestrians) lead to a building entrance.

These lacks imply differences between what the camera can do and what a human observer would do. In observing what they know is a pedestrian, a human observer would attend to movements not just because they are unusual, but because they suggest something about the intent of the pedestrian. Moving towards a window of a building would be interpreted differently from moving toward an equally unusual location on the outside of the parking lot. In remote caregiving the camera would similarly lack the ability to distinguish unusual but unimportant events from important ones. Note that such a camera could still be useful as an aid to a human observer, even though it could not substitute for a human observer.

Another impact of lack of understanding arises for an active agent determining when a client needs assistance. Given appropriate training data, an agent could be able to detect visible signs of distress, such as facial expressions or vocalizations. But it could not make inferences about distress based on failed attempts to accomplish a goal, as a human observer could, without being trained on particular goals, and particular failure behaviors.

Some progress towards some form of understanding may be suggested by work on *time contrastive networks* (Sermanet et al., 2017). This technique uses video recordings of the same action, taken from different viewpoints, and trains a system to identify different views of the same action, from views of different actions taken from the same viewpoint. Because views from different viewpoints will be different in background, and other irrelevant visual features, while views from the same viewpoint will tend to share these features, the approach pushes the system to identify what the essential features of an action are, while ignoring irrelevant details. While this approach is a long way from supporting the forms of understanding we're considering here, it perhaps at least suggests how trained systems can see beneath superficial details to grasp something closer to meaning.

Approaching more closely our aims, Fern et al. (2014) describe an artificial assistant that can learn a person's goals by observing the person's actions. They have the explicit aim of creating "a domain-independent framework" that can be used to create artificial assistants for different situations, without extensive human programming.

The challenge of understanding, like the data challenge, is widely recognized. DARPA has recently announced a program aimed directly at it, Machine Common Sense (see https://www.darpa.mil/news-events/2018-10-11), and other research programs will continue to address it.

4.1 Managing context and state.

Besides the challenges of data and understanding, there are other challenges to be met in creating useful artificial agents, whether active or passive. One is that assistance is commonly conditioned on the context in which assistance is needed. That is, what an assistant must do is determined not just be the content of a client's request, or (for an active agent) by the signs that assistance is needed, but also by the specifics of the situation. For example, if a client needs help navigating, the assistance depends on where the client is, as well as on where they want to go. If a client needs help remembering what gifts they have given, to avoid giving a friend a duplicate gift, the assistance depends on the agent having stored information about past gifts, and being able to retrieve that information. The Appendix includes a number of examples of services of this kind, drawn from the experience of a brain injury survivor.

In any particular case, like providing navigation advice, one can imagine creating a specialized service that can consult location information, for example on a smartphone, and can respond to a client's request using that information. Machine learning could be used to train a system to provide useful advice, based on various forms of client requests, and various situations seen in maps. Providing simple prompts or recommendations as opposed to automation allows for greater development of self-determination skills, a key outcome identified in many programs and services. Further, a collection of such services could be developed, over time, that could provide useful support in a range of common situations. Saadi Lahlou's *Installation Theory* (2018) suggests that our knowledge of how to behave may be largely a mosaic of this kind of special knowledge.

The data challenge means that such developments would require gathering a good deal of example data, and the understanding challenge means that, even for a relatively well defined task like navigation, the advice provided would sometimes fall short of what a human assistant might give. For example, if I ask how to get to a distant grocery store, a human assistant might tell me that there's a similar store on the next block. Today, an artificial agent would need to be especially trained to handle that and indefinitely many other special cases.

The example of keeping track of gifts (and many other examples shown in the Appendix) shows that assistance depends on handling *state*, information that changes over time, and must be kept up to date. In the example the state would be something like a list of people and associated gifts. The assistant would need to be able to update this information, when a gift is given, and to be able to hold a conversation that refers to the list. As discussed earlier, current Natural Language Processing technology is best at analyzing relationships within a text, and weak at connecting a text to outside information, like the gift list.

Researchers are making some progress on this. Eric and Manning (2017) report a system that can be trained to converse about a collection of key-value pairs (a common data representation) in three assistance domains, scheduling, weather, and simple navigation. It outperforms earlier systems, while falling well short of human performance. He et al. (2017) report progress on a different task that requires conversation about a body of stored information; notably, this system attains some success managing information that was not available at training time.

4.2. Instructability.

The gift tracking example suggests a further challenge for the development of artificial assistants. While getting navigation advice is a common need, and might be supported by a specially trained agent, gift tracking seems like something only a few people would request. A human assistant can easily help: the client describes that is needed, and the assistant can readily figure out that to do, and do it. That is, a human assistant can be *instructed* to perform a new task, for which they have not been trained. The instruction is brief, and does not involve providing thousands of examples on which to train.

This problem is called *interactive task learning*, referring to the need for a system to learn how to perform a task from an interaction with a client (Laird et al., 2017). Interestingly, workers in this field argue that the deep learning approaches that dominate the other research we have discussed are not appropriate to this domain (though others would point to deep learning work that addresses related tasks, like imitation). Instead, they employ methods with an earlier history in artificial intelligence and cognitive science, relying on structures of symbols rather than networks of artificial neurons. Progress is being made, but researchers note that limitations in natural language processing, and managing background knowledge, as we've discussed earlier, are serious. Kirk, Mininger, & Laird (2016) report success in combining verbal descriptions of aspects of tasks with visual demonstrations. Their system is capable of learning simple games, puzzles, and procedures from instruction, rather than from large collections of examples. Considering examples like the gift tracking assistant, however, a leading researcher in this field says, "We're only scratching the surface."

As mentioned earlier, this work draws on a different, older line of thought in artificial intelligence research, than current deep learning methods. But proponents of this work point out that many practical applications of artificial intelligence today, in self-driving cars, for example, rely heavily on these methods, and not on deep learning. Further, deep learning systems often have poor *explainability*, compared to systems using symbolic methods. That is, it is usually difficult to say much about how a deep learning system chose a particular response, and thence to understand what might be needed to repair mistakes. Lack of explainability can also make it hard to predict how a system will respond in a novel situation, since how its responses are derived is obscure. This is a problem for systems that have to be trusted in new situations.

Smolensky and collaborators are pursuing a program of research that may unify these older and newer lines of work, showing how symbolic structures can be represented and managed within a neural network framework (Cho, Goldrick, & Smolensky, 2017; Huang et al., 2017, 2018). In the longer term this work may succeed in overcoming some of the limitations both of current machine learning based approaches to NLP, and of traditional linguistic theorizing, providing another perspective on the challenges noted by Pavlick in her 2018 talk.

4.3 Interacting with a cognitive assistant.

As artificial cognitive assistants become more capable, how might people with cognitive disabilities interact with them? The gift tracking application discussed above is a task that was actually supported by a remote human assistant, hired by a person with a brain injury. All interaction with the assistant, who was in a distant country, was carried on by phone and email. This suggests that a passive artificial agent could be useful if available in the same way to a client. Interestingly, this client suggests that they would prefer to interact with an artificial rather than a human assistant, because of the awkwardness of interacting with another person about cognitive challenges.

Another model is suggested by Aira (https://aira.io/), an assistance service for blind people. The service allows a blind client to establish voice communication with a trained human agent, who can also see through a camera that the client has, mounted in a headset or on their phone. Analysis of Aira transactions suggests that significant parts of many of them could be automated (for example, many transactions are requests to read text.) A version of the Aira service might be devised for sighted clients, and could provide a good interaction model for passive cognitive assistants. The value of passive cognitive assistants is that they allow the user to direct the perceived need of support and interaction promoting self-determination. There is concern that over automation of systems may inhibit already vulnerable populations from exerting control in their own environments.

4.4 The prospect of direct brain connection?

In the longer term, many technologists expect that it may be possible to connect artificial agents directly to a client's brain. The envisioned connection goes beyond today's brain computer interfaces in that it would be bidirectional, That is, today's brain computer interfaces can communicate signals from the brain to an external system, but future interfaces may communicate signals from an external system, an active or passive cognitive assistant, to a client's brain. The OpenWater startup says on its Web site: "With read/write ability - we may be able to upload/download and augment our memories, thoughts, and emotions with a ski-hat form factor, non-invasively. (https://www.openwater.cc/technology)." Some neuroscientists think this unlikely. Some feel that our current understanding of brain mechanisms may leave out essential processes, being based so far only on easy to observe phenomena like action potentials. We may have farther to go to understand how information is represented and processed in the brain than we now think.

Another possible limitation is individual difference: if a brain connection can be established for one person, the large engineering investment needed to do this may need to be repeated for each subsequent person, because of differences in how information is represented in different brains. But some results suggest that the representations created by different people are similar; see Mason and Just (2016).

Because of the potential implications of this form of interaction, it seems appropriate to give some attention to it before it becomes a reality. To be concrete, we can consider two use cases. In the first, a hypothetical brain connection supplies a person with navigation knowledge. While today, I know how to cycle from my home to my office, for other destinations I must consult my phone for directions. With a brain-connected navigation system I would know how to reach any destination, just as I reach my office, with no need to consult any outside source. In a sense, my phone, a passive assistant, would be replaced by an active assistant, connected to my brain, that would detect my need for navigational knowledge, and supply it.

In a second use case, a system would improve a person's ability to recognize risky behavior, something difficult for some people. Rather than relying on an external assistant to monitor their behavior and provide cautionary advice, a person could have a brain-connected system that would lead them to make appropriate choices themselves.

Such scenarios raise a number of potential concerns, especially about autonomy. Would a person whose judgement of risk is artificially modified feel less free? While the matter is of course very unclear, there is some reason to think not. Psychologist Jonathan Haidt argues that our conscious selves have little or no access to the processes by which we make decisions and choose actions (Haidt, 2001). Rather, our conscious selves observe what we do, and seek to provide an explanation of it, generally an explanation that makes our actions seem rational. If Haidt is right, our conscious selves would not be aware of the existence of a brain-connected system influencing our choices. We would continue to explain our behavior, maybe having an easier time of it, with fewer risky actions to rationalize. Indeed, the person with the brain-connected system might feel *more* free than a person accompanied by a human or artificial assistant that kept advising them not to do something they were starting to do. In the latter case the signs of reduced autonomy are visible.

It seems clear that the ethical status of brain-connected systems is affected by whether the connection is voluntary. A person might choose to have a system installed that would help them avoid behaviors they have learned to regret. But can we be confident that connection would remain voluntary? If a brain-connected system could reduce behaviors that cause serious accidents, for example, might not its use be mandated, by the justice system or by the insurance system? Or would it deny people with disabilities the “dignity of risk” that allows for learning and opportunity (http://mn.gov/mnddc/ada-legacy/pdf/The_Dignity_of_Risk.pdf)?

5. Application: Use of machine learning to optimize administrative and other services.

A general task well suited to machine learning is *classification*, determining whether an example of some kind falls in a desired category or not. Many administrative functions can also be thought of as classification: is a person eligible for a service or not; is a job applicant qualified for a job or not. One can train a machine learning system by giving it a large collection of examples: descriptions of people who are eligible or ineligible for a service, or descriptions of job applicants and hiring outcomes, and so on. Such a system could perhaps perform these tasks as well as or better than human administrators, at lower cost. Methods related to machine learning can also potentially be used to evaluate social services, given descriptions of services and associated outcome data. A recent French report (https://www.aiforhumanity.fr/pdfs/MissionVillani_Report_ENG-VF.pdf) calls for investment in these kinds of applications, suggesting that they could reduce costs and improve quality for financially stretched social service agencies.

Two challenges emerge in attaining these hoped for benefits. First, experience has shown that systems trained on data can exhibit behavior that is strongly biased, in socially undesirable ways (see Josh Lovejoy, “Fair is not the default”, https://design.google/library/fair-not-default/). A notorious example is a system that predicts recidivism, as part of a process of advising judges on sentencing in criminal cases. Because black offenders in the sample used to train the system were more likely to reoffend, the system recommended longer sentences for blacks. Most people would agree that the simple fact of someone being black should not result in longer sentences, yet that is what a seemingly objective system trained on actual data produces.

The problem is deeper than might at first appear. If one simply removes race as a descriptor in the example, the system will likely still discriminate in undesirable ways, using variables that are strongly correlated with race, such as (in some localities) zip code.

Another form of bias can arise if some class of people is not represented in the sample used to train a system. If a system for screening job applications is trained on a sample of applicants who do not have disabilities, for example, the system will likely do a bad job of evaluating applicants with disabilities. Even if some people with disabilities are included in the training data, the trained system may ignore factors that are important in evaluating their applications, because, for most applicants, these factors aren’t very important. Jutta Treviranus has suggested that systems design should focus not on the typical cases in a sample, as statistical methods do, but the exceptional cases (https://medium.com/ontariodigital/if-you-want-the-best-design-ask-strangers-to-help-e37bdb73567). Treviranus argues that supporting the exceptional cases in a design is the only way to attain truly robust and inclusive coverage, as opposed to relying upon the statistical “average”. Designs focussed on typical cases inevitably fail in situations that are too distant from the typical. She argues further that efforts to expand coverage by creating different categories of users, for example different disability types, will still fail when exceptional cases are considered.

A second challenge is that systems trained on data lack the means to recognize exceptions that human administrators would recognize. If missing required appointments is a disqualification for a given service, an automated system will likely enforce the disqualification regardless of what a human administrator would accept as extenuating circumstances. Virginia Eubanks (2018) describes a situation in which a service client who was hospitalized had service eligibility revoked by an automated system, because of missing appointments. Perhaps a system trained on a sufficiently broad range of cases, including exceptions, could do better, but it’s the nature of exceptions that they are rare, and can differ from previous exceptions in hard to foresee ways.

This challenge is another outcropping of the understanding challenge discussed earlier, a broad issue identified by many machine learning researchers: lack of artificial “common sense”, or “artificial general intelligence”. That is, while machine learning has produced very successful systems for performing narrowly defined tasks, they are so far unable to adapt their behavior to deal appropriately with tasks that are slightly different, or to recognize the need to deviate from learned behaviors, as a human would do.

5.1 Privacy concerns

Any use of data about people creates privacy risks. People with disabilities may be especially concerned for two reasons. First, they may feel that they can more easily be identified from data collected about them, because they are not typical. Second, they may feel that loss of privacy is more damaging, because of the possibility of discrimination, for example in employment.

These risks create a dilemma. On the one hand, collecting data about a person's speech patterns, for example, makes it possible for a speech recognition system to better recognize what they are saying. But these data might allow a third party to determine that the person has a disability.

The fact that data about people have commercial value has led to the emergence of *data brokers*, entities that buy and sell data. This means that if a person allows data about them to be collected, with a good purpose in mind, the data may easily become available for other purposes.

Thus there are problems on both ends of a possible spectrum of policies on data collection. Some fear that restrictions on data collection could slow or prevent the development of useful services for people with disabilities, while unrestricted data collection could expose people to loss of privacy and potential discrimination. Can we craft privacy policy and law that is tailored to accommodate genuine advances in accessibility without opening the door to the kind of abuses described above?

6. Other applications

The applications of machine learning discussed here are only some of those that may emerge, with implications for people with cognitive disabilities. Machine learning is in everyday use to improve speech recognition, in personal assistant technologies like Amazon Alexa, Google Assistant, or Microsoft's Cortana. Extensions of this work may assist people with speech and language disabilities. Machine learning may also have potential for customizing educational interactions, for example choosing lesson materials adapted to an individual learner's capabilities. Improvements in assessing people's functional strengths and weaknesses may also be possible.

People with cognitive disabilities may benefit from systems that help others participate more meaningfully in their chosen environments. . The same technology that may help people with disabilities develop soft skills for the workplace may help their coworkers be more supportive, and accepting of diversity .

7. Reflections

AI and robotics pioneer Rodney Brooks cautions against overenthusiasm about the prospects for artificial intelligence, in Technology Review:
https://www.technologyreview.com/s/609048/the-seven-deadly-sins-of-ai-predictions/
One of his cautions focuses on the limitations to understanding that we've discussed:

> [S]uppose a person tells us that a particular photo shows people playing Frisbee in the park. We naturally assume that this person can answer questions like What is the shape of a Frisbee? Roughly how far can a person throw a Frisbee? Can a person eat a Frisbee? Roughly how many people play Frisbee at once? Can a three-month-old person play Frisbee? Is today's weather suitable for playing Frisbee?
>
> Computers that can label images like "people playing Frisbee in a park" have no chance of answering those questions. Besides the fact that they can only label more images and cannot answer questions at all, they have no idea what a person is, that parks are usually outside, that people have ages, that weather is anything more than how it makes a photo look, etc.

Psychologist Gary Marcus provides a more detailed critique, including many of the themes we've considered here (Marcus, 2018):

> Deep learning systems work less well when there are limited amounts of training data available, or when the test set differs importantly from the training set, or when the space of examples is broad and filled with novelty. And some problems cannot, given real-world limitations, be thought of as classification problems at all. Open-ended natural language understanding, for example, should not be thought of as a classifier mapping between a large finite set of sentences and large, finite set of sentences, but rather a mapping between a potentially infinite range of input sentences and an equally vast array of meanings, many never previously encountered. In a problem like that, deep learning becomes a square peg slammed into a round hole, a crude approximation when there must be a solution elsewhere.

It is of course uncertain what the results will be, as enormous research investments are deployed to attack these limitations; indeed many researchers believe that Marcus's critique already no longer accurately reflects the state of the field. Anyway, Marcus himself ends his critique on an optimistic note, though anticipating not the ultimate success of current machine learning methods, but the emergence of "entirely new methods".

On the other hand, there may be more life in the existing methods than Marcus (and Brooks) suppose. Trinh and Le (2018) report success in some of the commonsense reasoning tasks for which deep learning methods have been weak. Their method relies entirely on regularities detected in a large corpus of text, with no human supervision, to determine (for example) whether “it” in the sentence “The trophy doesn’t fit in the suitcase because it is too big" refers to "the suitcase" or "the trophy". This kind of choice is easy for people, but (because it depends on knowledge about suitcases and fitting) has been difficult for artificial systems. Evidently the required knowledge is in fact available in a sufficiently large volume of text, and can be extracted and used automatically. Sap et al. (2019) also report progress in supporting commonsense reasoning, using a large corpus of if-them relationships collected using crowdsourcing.

So we may expect continued progress in deep learning, as well, perhaps, as significant new ideas. Besides awaiting (and encouraging) these developments, our community should consider how more limited capabilities may be useful in the applications important to us.

8. Roadmap

Here are some actions that may help realize benefits from machine learning to people with cognitive disabilities and those who support them. If you are interested in participating, or if you know of existing efforts that address these opportunities, please email ColemanInstitute@cu.edu .

8.1 Pilot development projects

Remote caregiving. The shortage of caregivers means that many people do not have access to the care they want and need. Remote caregiving services, like RestAssured, may help meet the need. Can machine learning be used to make remote caregiving more effective, and more available? What data collection would be needed to enable this? Can remote companion care also be explored to address the loneliness epidemic impacting aging populations and people with disabilities?

Remote personal assistants. The Aira system provides remote support to blind people, by connecting a trained human agent to a camera carried by the client. Can a similar service benefit sighted people with cognitive disabilities? Using machine learning to automate some common interactions for blind clients seems feasible; are there similar opportunities for clients with cognitive disabilities? Navigation could be a promising area to explore.

Conversational assistants. Systems like Amazon Alexa, Google Assistant, and Microsoft Cortana are supporting a widening range of applications, including for people with sensory and motor disabilities, allowing people to carry out useful interactions by voice. Are there opportunities to extend this technology, using machine learning, to provide support for people with cognitive disabilities? Could such supports also provide value to consumers generally, potentially increasing the investment available to develop them?

8.2 Policy projects

Safe data sharing. Our community has a role to play in defining an appropriate balance between the value of data collection in improving supports for people with disabilities, and the associated privacy risks. Can workable opt-in policies be developed, accompanied by confidence in how data would be used and shared? Can cross-industry agreements be brokered, allowing data to be shared among companies for the benefit of consumers, while not exposing the data to unwanted uses? Can educational programs be developed for people with cognitive disabilities to enhance awareness and understanding of data ownership and data sharing?

Ethics in Artificial Intelligence. Concerns about possible negative effects of artificial intelligence, including machine learning, have led to the formation of many interest groups, including the Center for Human Compatible AI (Berkeley), the Partnership on AI (https://www.partnershiponai.org/ ; includes Amazon, Apple, Facebook, Google, IBM, Microsoft, ...), the Future of Life Institute (https://futureoflife.org/; Oxford) and many more. Our community needs to participate in these discussions.

Cross disability cooperation. Many of the issues and opportunities considered here, including those involving personal assistance and privacy, concern people with disabilities of all kinds. We should establish communication with other advocacy groups to consider these matters.

8.3 Research Initiatives

NLP Benchmarking. Progress in automatically transforming text to make it more intelligible would be aided by examples that show the kind of transformations that are needed, such as the materials prepared for the Medicaid Reference Desk project. Also, a good way to evaluate such transformations, for example, a protocol for testing comprehensibility, would also be useful.

Direct brain interfaces. The practicality of direct brain interfaces is a matter of controversy. So is their desirability, and their ethical status. It seems prudent to promote discussion of their possible impact, and how their possible development might best be shaped.

ACKNOWLEDGEMENTS

Preparation of this paper was supported by the Coleman Institute for Cognitive Disabilities, University of Colorado. A draft was provided as background for a workshop on this topic, held on October 2, 2018, in Broomfield, Colorado, in conjunction with the Institute's annual conference. This version reflects some of the discussion at the workshop, and many suggestions provided by the participants. The participants were: Andrew Begel, Cathy Bodine, Kelly Buckland, Isabel Bush, Henry Claypool, Bill Coleman, Ed Cutrell, Pam Gregory, Shaun Kane, Anirudh Koul, John Laird, Hobson Lane, Chris Langston, Clayton Lewis, Sue-Ann Ma, Jim Martin, Kyle Nesbit, Derek Nord, Blake Reid, Rob Rupert, Maarten Sap, Will Scott, Sue Swenson, Shea Tanis, Jutta Treviranus, Nancy Ward, Dustin Wright, and Nida Zada. The discussion was held under Chatham House rules, meaning that remarks were not to be attributed to individuals. Amie Lulinski and Aruna Gunda provided valuable assistance during the workshop.

I thank the DIAGRAM Center of Benetech for providing a stimulating and welcoming environment in Palo Alto, where much of the background research was done. I am grateful to the following people, who kindly contributed their time to discuss the topic with me during the preparation of the original draft, and provided essential ideas and suggestions: Sina Bahram, Yejin Choi, Henry Claypool, Andreas Girgensohn, Pam Gregory, Jonathan Grudin, He He, Katherine Hermann, Jess Holbrook, Tomas Izo, Dan Jurafsky, Kirill Kireyev, John Laird,Charles LaPierre, Hobson Lane, Saadi Lahlou, Pat Langley, Chris Langston, Owen Lewis, Rainer Lienhart, Josh Lovejoy, Jim Martin, Mark Nitzberg, Aude Oliva, Caroline Pantofaru, Roy Pea, Marc Pickett, T V Raman, Alexis Roos, Rob Rupert, Franco Salvetti, Paul Schroeder, Jeff Shrager, Paul Smolensky, Jutta Treviranus, Vincent Vanhoucke, Nancy Ward, Jason White, Dustin Wright, Yinfei Yang, and Mohammed Yousouf. Of course none of them, nor any of the workshop participants, is responsible for the opinions and distortions in the paper.

REFERENCES

Achille, A., Eccles, T., Matthey, L., Burgess, C. P., Watters, N., Lerchner, A., & Higgins, I. (2018). Life-Long Disentangled Representation Learning with Cross-Domain Latent Homologies. *arXiv preprint arXiv:1808.06508*.

Bosselut, A., Levy, O., Holtzman, A., Ennis, C., Fox, D., & Choi, Y. (2017). Simulating action dynamics with neural process networks. *arXiv preprint arXiv:1711.05313*.

Cho, P. W., Goldrick, M., & Smolensky, P. (2017). Incremental parsing in a continuous dynamical system: Sentence processing in Gradient Symbolic Computation. *Linguistics Vanguard*, *3*(1).

Djamasbi, S. (2017). *Improving Manual and Automated Text Simplification* (Doctoral dissertation, UMass Medical School).

Eric, M., & Manning, C. D. (2017). Key-value retrieval networks for task-oriented dialogue. *arXiv preprint arXiv:1705.05414*.

Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.

Fern, A., Natarajan, S., Judah, K., & Tadepalli, P. (2014). A decision-theoretic model of assistance. *Journal of Artificial Intelligence Research*, 50, 71-104.

Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *arXiv preprint arXiv:1703.03400*.

Haidt, J. (2001). The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological review*, *108*(4), 814.

He, H., Balakrishnan, A., Eric, M., & Liang, P. (2017). Learning symmetric collaborative dialogue agents with dynamic knowledge graph embeddings. *arXiv preprint arXiv:1704.07130*.

Huang, Q., Smolensky, P., He, X., Deng, L., & Wu, D. (2017). A Neural-Symbolic Approach to Natural Language Tasks. *arXiv preprint arXiv:1710.11475*.

Izo, T., & Grimson, W. E. L. (2007, September). Unsupervised modeling of object tracks for fast anomaly detection. In *Image Processing, 2007. ICIP 2007. IEEE International Conference on* (Vol. 4, pp. IV-529). IEEE.

Kiddon, C., Zettlemoyer, L., & Choi, Y. (2016). Globally coherent text generation with neural checklist models. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing* (pp. 329-339).

Kirk, J., Mininger, A., & Laird, J. (2016). Learning task goals interactively with visual demonstrations. *Biologically Inspired Cognitive Architectures*, *18*, 1-8.

Konstas, I., Iyer, S., Yatskar, M., Choi, Y., & Zettlemoyer, L. (2017). Neural amr: Sequence-to-sequence models for parsing and generation. *arXiv preprint arXiv:1704.08381*.
Lahlou, S. (2018). *Installation theory: the societal construction and regulation of behaviour*. Cambridge University Press.

Laird, J. E., Gluck, K., Anderson, J., Forbus, K. D., Jenkins, O. C., Lebiere, C., Salvucci, D., Scheutz, M., Thomaz, A., Trafton, G., Wray, R. E. , Mohan, S. & Kirk, J. (2017). Interactive task learning. *IEEE Intelligent Systems*, *32*(4), 6-21.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *nature*, *521*(7553), 436.

Lee, J., & Yeung, C. Y. (2018). Personalizing Lexical Simplification. In *Proceedings of the 27th International Conference on Computational Linguistics* (pp. 224-232).

Lewis, C., & Ward, N. (2011, July). Opportunities in cloud computing for people with cognitive disabilities: designer and user perspective. In *International Conference on Universal Access in Human-Computer Interaction* (pp. 326-331). Springer, Berlin, Heidelberg.

Marcus, G. (2018). Deep learning: A critical appraisal. *arXiv preprint arXiv:1801.00631*.

Mason, R. A., & Just, M. A. (2016). Neural representations of physics concepts. *Psychological science*, *27*(6), 904-913.

Pavlick, E., & Callison-Burch, C. (2016a). Simple PPDB: A paraphrase database for simplification. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)* (Vol. 2, pp. 143-148).

Pavlick, E., & Callison-Burch, C. (2016b). Most" babies" are" little" and most" problems" are" huge": Compositional Entailment in Adjective-Nouns. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (Vol. 1, pp. 2164-2173).

Rashkin, H., Sap, M., Allaway, E., Smith, N. A., & Choi, Y. (2018). Event2Mind: Commonsense Inference on Events, Intents, and Reactions. *arXiv preprint arXiv:1805.06939*.

Redish, J. (2000). Readability formulas have even more limitations than Klare discusses. *ACM Journal of Computer Documentation (JCD)*, *24*(3), 132-137.

Sap, M., LeBras, R., Allaway, E., Bhagavatula, C., Lourie, N., Rashkin, H., Roof, B., Smith, N. & Choi, Y. (2019). ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning. AAAI. Online at https://arxiv.org/abs/1811.00146 .

Schmidhuber, J. (1987). *Evolutionary principles in self-referential learning, or on learning how to learn: the meta-meta-... hook* (Doctoral dissertation, Technische Universität München).

Schwartz, R., Sap, M., Konstas, I., Zilles, L., Choi, Y., & Smith, N. A. (2017). The effect of different writing tasks on linguistic style: A case study of the ROC story cloze task. *arXiv preprint arXiv:1702.01841*.

Sermanet, P., Lynch, C., Chebotar, Y., Hsu, J., Jang, E., Schaal, S., & Levine, S. (2017). Time-Contrastive Networks: Self-Supervised Learning from Video. *arXiv preprint arXiv:1704.06888*.

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., & Hassabis, D. (2017). Mastering chess and shogi by self-play with a general reinforcement learning algorithm. *arXiv preprint arXiv:1712.01815*.

Trinh, T. H., & Le, Q. V. (2018). A Simple Method for Commonsense Reasoning. *arXiv preprint arXiv:1806.02847*.

U.S. Department of Labor, Office of Disability Employment Policy (2012) Soft Skills to Pay the Bills – Mastering Soft Skills for Workplace Success. Online at http://www.dol.gov/odep/topics/youth/softskills/

Xu, W., Napoles, C., Pavlick, E., Chen, Q., & Callison-Burch, C. (2016). Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, *4*, 401-415.

Zwaan, R. A., & Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological bulletin*, *123*(2), 162.

APPENDIX

Example tasks for a personal assistant

Note: These examples are provided by a brain injury survivor, and include details of that person's situation, so as to make the examples concrete and realistic. Some, but not all, of these tasks were delegated to a remote human assistant in the past.

About me:

I have worked hard to create a life that maximize my strengths and minimize my weakness;
I am not able to hold a job;
I take art classes, show and sell my art (mostly sculptures and paintings);
I thrive on social relationships and they are very important to me. I hate doing "case work" on myself;
I have a terrible sense of time-I am lousy at estimating how long a task will take and even worse at scheduling time to accomplish a task;
Even though I hate being late, I find myself chronically late. I function best on routine;
Little kinks in my schedule throw me off disproportionately;
My long-term memory (pre-accident) is excellent. My long-term memory the accident is impaired;
My short-term memory is severely impaired.

Manage Prescriptions:

Prescriptions at CVS (brick and mortar);
Prescriptions at Walgreens (brick and mortar);
Prescriptions through CVS' Mail Order Pharmacy (sent to me at my home address);
Prescriptions through Walgreens' Mail Order Pharmacy (sent to me at my home address);
Prescriptions through CVS' Specialty Mail Order Pharmacy (some are sent to me at my

home, some are sent to my neurologist’s office);
Prescriptions through Walgreens’ Specialty Mail Order Pharmacy (some are sent to me at my home, some are sent to my neurologist’s office);
Manage auto-refill requests
Ascertain if the brick and mortar pharmacy has appropriately filed with my insurance;
If there is a delay in filling my prescription, determine if it is an insurance issue an expiration issue, or availability issue;

Note: My insurance does not pay for all of my prescriptions. Some of my prescriptions are only valid for 30 days and my neurologist post-dates prescriptions. If I attempt to fill a prescription after the 30 day period, the pharmacy is unable to fill the prescription. This triggers having to contact my neurologist and explain what happened, and have him write a new prescription. I take mega-doses of some medications and most pharmacies do not keep that level of a medication in stock. Often times, the medications cannot be put on auto-refill. When I request a refill, pharmacy will have to order the medication. The medication has to be filled in full. If not, any remaining pills cannot be filled.

Manage My Calendar:

Medical appointments;
- Neurologist every three months;
- Gastroenterologist every three months;
- Sleep doctor every three months;
- Primary Care Physician every six months;
- Opthamologist every 12 months;
- Endocrinologist every 12 months;

Dental appointments every 6 months;
Hair appointments every 3 months;
Classes (at five different institutions)
Social events
- Double Dates (weekly)
- Date Night with my husband (weekly)
- Birthdays

Neighborhood club meetings;

Notes: When I see someone on my treatment team, I need to bring a copy of any new medical records from other services or tests. There are certain people that I take out to lunch for their birthday. This includes looking at my availability, look at their availability, selecting an appropriate restaurant. The neighborhood club has officer meetings monthly, general meetings monthly, and specific planning meetings.

Manage/Schedule My Tasks/Responsibilities:

½ hour per week to organize medications in my pill dispenser.

My pill dispenser system is based on 30 days; Be sure that I have all of the medications to fill my pill dispenser;

3 hours per week for mail:

Throw out trash mail;

Shred data-sensitive mail;

Prepare birthday and anniversary cards to my important people;

Renew subscriptions/membership;

1 hour every Sunday to review upcoming week:

Review my appointments;

Review my husband's week;

Review my children's schedule/upcoming week;

1 hour per week (to be performed between Monday evening and Wednesday evening) to organize supplies to bring to Thursday's class;

2 hours per week to research new restaurants, art shows, events of interest;

2 hours per week for laundry;

3 hours per week for groceries;

2 hours per week to organize and put art supplies away.

Track Important College Application Deadlines:

Create a spreadsheet for each child (twins);

Populate fields for best ACT and SAT math/language scores;

Populate fields for their class rank, AP classes and GPAs;

Populate fields for colleges/universities of interests;

Indicate if they visited or want to visit a campus;

Populate fields for each college/university of average ACT/SAT score;

Populate fields indicating if a college/university accepts the common application and if it requires an essay in addition to the common application essay;

Populate fields that indicate a college/university's early decision, early action, regular application dates, when each school notifies students of acceptance, waitlist or denial status;

As we go through this year, other important dates will arise, such as dormitory preference, and so add fields to the spreadsheet as we learn the process.

Track My Health:

When I visit professionals on my treatment team, they will ask me how I am doing. I usually don't remember and need assistance with creating and maintaining a list:
List migraine frequency, duration, severity, with or without auras, which medication taken and its efficacy;
List frequency of vertigo, and falls due to vertigo;
List incidents of high photosensitivity (for me);
List incidents of high hyperacusis (for me);
List incidents of double vision;
List incidents of slurred speech, stuttering and/or aphasia;

Track My Art Activity:

Maintain a list of art projects started and any plans for that project;
Maintain a list of all completed pieces of art;
List art shows that I am attempting to enter:
Any shows that I have submitted a piece for consideration;

Show registration fee;
Notification date of acceptance/denial in show;
If a piece is accepted, where my piece should be delivered/mailed;
Date and time to deliver the art;
Date and time to pick up the art if not sold;
Price of the piece;
Size, matting and framing for the piece;

Gallery:

List of all pieces of mine they have and my bottom-line price (meaning how much I must take home from any sale);
Size, matting and framing of piece;
Media used in the piece;

Sales:

If a piece is purchased, who bought it and for how much;

Outreach/Marketing

Email list of people who are interested in knowing of my shows;

Do I have enough business cards?
Did I post information on Facebook, Instagram or my webpage?
My webpage (I need my personal assistant to build my website. I built and managed 3 personal websites before, but I no longer have the skills to do so)
My Instagram (I can do this myself, but I do not remember to do it)

Classes:

Alert me when the next term class catalogue is published;
List classes that I have taken and classes or instructors I want to take/hav
Maintain a supply list for each class as well as studio and/or modeling fees;
Add classes to my calendar;
If a class or workshop is in a new-to-me place, find out what the parking situation is.