DANIEL DE MAGALHÃES CHADA

ARE YOU EXPERIENCED?

Contributions towards Experience Recognition, Cognition, and Decision Making

ARE YOU EXPERIENCED?

Contributions towards Experience Recognition, Cognition, and Decision Making

DANIEL DE MAGALHÃES CHADA

Submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Ph.D.) in Management
with emphasis on Behavioral and Decision Sciences

at the Escola Brasileira de Administração Pública e de Empresas — EBAPE
of the Fundação Getulio Vargas — FGV

Daniel de Magalhães Chada: *Are You Experienced?*, Contributions towards Experience Recognition, Cognition, and Decision Making, © October 2016

COMMITEE:
Fabio Caldieraro
Moacyr Alvim Horta Barbosa da Silva
Ross Gayler
Alexandre Mendes
Alexandre Linhares

LOCATION:
Rio de Janeiro

TIME FRAME:
October 2016

Ficha catalográfica elaborada pela Biblioteca Mario Henrique Simonsen/FGV

Chada, Daniel de Magalhães
Are you experienced?: contributions towards experience recognition, cognition, and decision making / Daniel de Magalhães Chada. – 2016.
97 f.

Tese (doutorado) - Escola Brasileira de Administração Pública e de Empresas, Centro de Formação Acadêmica e Pesquisa.
Orientador: Alexandre Linhares.
Inclui bibliografia.

1. Administração. 2. Ciência cognitiva. 3. Memória - Simulação por computador. 4. Processo decisório. 5. Sistema de reconhecimento de padrões. 6. Inteligência artificial distribuida. I. Linhares, Alexandre. II. Escola Brasileira de Administração Pública e de Empresas. Centro de Formação Acadêmica e Pesquisa. III. Título.

CDD – 658.403

DANIEL DE MAGALHÃES CHADA

ARE YOU EXPERIENCED? CONTRIBUTIONS TOWARDS EXPERIENCE RECOGNITION, COGNITION AND DECISION MAKING.

Tese apresentada ao Curso de Doutorado em Administração da Escola Brasileira de Administração Pública e de Empresas para obtenção do grau de Doutor em Administração.

Data da defesa: 08/12/2016.

ASSINATURA DOS MEMBROS DA BANCA EXAMINADORA

Alexandre Linhares
Orientador (a)

Fabio Caldieraro

Moacyr Alvim Horta Barbosa da Silva

Ross Gayler

Alexandre Mendes

Alexandre Mendes

Ohana means family.
Family means nobody gets left behind, or forgotten.

— Lilo & Stitch

Dedicated to the loving memory of Edson de Campos Chada.

1937 – 1994

ABSTRACT

This work is comprised of three independent contributions from the realm of cognitive modeling to management science. The first addresses Experience Recognition, a theory first introduced by Linhares and Freitas [91]. Here it is extended and better defined, and also its contribution to cognitive science and management science are discussed. The second contribution introduces a cognitive framework called Rotational Sparse Distributed Memory, and provides a sample application of its characteristics as a substrate for a highly relevant subject in management science: semantic networks. The final contribution applies Rotational Sparse Distributed Memory to modeling network motifs, dynamic flexibility and hierarchical organization, all highly impactful results in recent neuroscience literature. The relevance of a neuroscientific modeling approach towards a cognitive view of decision science are discussed.

Gratitude bestows reverence,
allowing us to encounter everyday epiphanies,
those transcendent moments of awe that change
forever how we experience life and the world.

— John Milton (1608 - 1674)

ACKNOWLEDGMENTS

Turns out this is the section upon which I put highest importance, but have the most trouble writing. Those who choose the academic path must always rely on the support of many, for it is a long and often arduous path, with delayed rewards. I know the following brief words cannot possibly do justice all who have shown me kindness, patience, and support throughout these years, but nonetheless I will make a humble effort to acknowledge them and the love they've shown me through their words and deeds.

I would like to thank my friend . . . then girlfriend . . . then fiancée . . . then wife Michelle. She shared every step of this journey side by side with me, through thick and thin. Her loyalty and commitment make me the luckiest man alive. More than that, the patience, kindness and love she shows me everyday light the way to my being a better man. I love you. I will also be forever grateful for the ray of light you brought to my life: my step-son Dan. To share in his discovery of the world is the greatest joy I have found, and I can only hope to be there at every new step.

I must also thank my mother, Lúcia, whose tirelessness, hard work and dedication to her family are forever ingrained in my mind as the meaning of the word 'nobility'. This Ph.D. would not have been possible without her unwavering support.

I am unaware of how many advisees can honestly say they found great friends in their advisors, but I did in Alexandre Linhares. His capacity to shape an endless flow of raw, chaotic creativity into wildly innovative work will forever astound me. His optimisim and trust in my abilities were undying, even when I occasionally 'disappeared'. He trusted my vision that the results would be there and, If I can say I have lived up to his expectations, then I can be proud of this work. Thanks captain, and thanks for making it all look easy and fun.

I thank Kaillen Givigi, the lab's protector and defender. Without her management and aid, navigating the different hurdles of graduate school would have certainly ended in ruin, for myself and many other members of the aptly named cognitive "psych-os" lab.

I thank my colleagues-in-arms: Diego Faveri and Bernardo Guelber Fajardo: the hard times were made easier, and the good times were made better by being able to share them with you.

I thank prof. Tom Griffiths and all the cocosci lab members at Berkeley, especially Jess Hamrick, Mike Pacer, Josh Abbott and Falk Lieder. I never knew I could learn so much about so much in so little time. The year I spent with them was truly a trial by fire for my academic endeavors, and I've never met a group so kind or welcoming. I thank them for making me feel like one of their own, though I was (and am) undeserving.

I must thank the staff at EBAPE/FGV for their tireless professionalism and ready technical support. The Getulio Vargas Foundation is a genuinely unique institution in Brazil, and it was a privilege to have access to such a collection of the country's (and world's) finest academic minds.

Last, yet surely far from least, I thank my friend, colleague and coauthor Marcelo Brogliato for the countless enlightening conversations in our shared passion of science, computation and cognition. His feedback was and remains instrumental in the development of my ideas and research. Research is born from ideas, and ideas grow from being shared, discussed and challenged. If the quality of an idea is passing the crucible of a strong argument, mine must be indeed strong for passing the test of our discussions. I pray my ideas may always have the opportunity of meeting the challenge of your own, my friend.

CONTENTS

LIST OF FIGURES

ACRONYMS

AI artificial intelligence

BSC Binary Spatter Code

BStr bit string

FFI feedforward inhibition

HL hard location

HRR Holographic Reduced Representation

PSR Predictive State Representation

RSDM Rotational Sparse Distributed Memory

SDM Sparse Distributed Memory

TPN Tensor Product Network

VSA Vector Symbolic Architecture

XOR exclusive-or

XR Experience Recognition

Part I

DEFINING EXPERIENCE RECOGNITION

" Two roads diverged in a yellow wood,
And sorry I could not travel both
And be one traveler, long I stood
And looked down one as far as I could
To where it bent in the undergrowth;

Then took the other, as just as fair,
And having perhaps the better claim,
Because it was grassy and wanted wear;
Though as for that the passing there
Had worn them really about the same,

And both that morning equally lay
In leaves no step had trodden black.
Oh, I kept the first for another day!
Yet knowing how way leads on to way,
I doubted if I should ever come back.

I shall be telling this with a sigh
Somewhere ages and ages hence:
Two roads diverged in a wood, and I—
I took the one less traveled by,
And that has made all the difference."

— Robert Frost (1874–1963)

1 INTRODUCTION

A CENTRAL PRECEPT OF MANAGEMENT SCIENCE IS TO DEVELOP a comprehensive theory of human decision-making. The choices and mechanisms which underlie the rationale of decision-makers are an essencial aspect of management. While the 'rational-decision actor' has been successful in modeling, and in bringing valuable insights into a number of decision scenarios, several studies have made clear that humans depart from rationality (for recent summaries, see [7, 8, 130]).

Negotiation methods, strategic decisions, leadership styles and policy-making are all subject to the minds which generate them. Are managers capable of rational decision-making under complex scenarios? [7, 50]. What forms of contextual influence may be present? [10]. What cognitive processes occur in decision-makers' sense-making and understanding? [88, 141] What role do analogies play?[49, 69]. How does memory's intrinsic characteristics sway one's perception of a given situation? [92].

Cognitive science, the study of human information-processing, is slowly filling the void between the need for formal models of human behavior and the numerous shortcomings of the rational model. The computational tools generated through the field's explorations into human behavior can provide new paradigms of pattern recognition, exploratory data analysis, information retrieval, data visualization and decision support, all crucial to the aim of decision science.

For example, in the field of business strategy, Gavetti et al. [49, 51] postulate that choice, in novel environments at least, is guided by analogy-making. More generally, case studies are a highly popular tool in business education and their value often hinges on the understanding of one situation in terms of another. The literature provides a number of examples of analogical and contextual reasoning:

1. The chain "Toys 'R' Us", launched in the 1950s, was tied to the vision and success of supermarkets (the chain was effectively called "Baby Furniture and Toy Supermarket" at a particular point). Afterwards, the launch of the office store "Staples" was based on similar reasoning: "Could we create a Toys R Us for office supplies?" [49]

2. In the 1980s the largest European car-maker decided to invest heavily in the U.S. market by introducing a large range of cars that were bestsellers in Europe. Before this decision, the car-maker had 63% of the im-

ported car market in the U.S. But the car-maker was called Volkswagen. The American consumer's experience with Volkswagen consisted of the Beetle, an inexpensive and odd-looking car first sold in 1938. American consumers rejected the idea of a large, well-built, modern-looking, powerful and expensive Volkswagen. To make matters worse, the company decided to withdraw the Beetle from the market, and its share of the imported car market in the U.S. dramatically fell from 63% to less than 4%. The exact same cars were being sold in Europe and in the U.S; the only difference was in the consumers' experiences of what a "Volkswagen" meant. A twenty-thousand-dollar Volkswagen seemed, to Americans, like a practical joke. Similarly, "the new Honda" to an American consumer meant a new car model; to the Japanese, it meant a new motorcycle [121].

3. Uber, a driver services startup based in Silicon Valley has experienced stratospheric growth between 2010 and 2016 [3, 136, 148]. Their smartphone 'app' allows clients to place a trip request, which is sent to a nearby driver, who is registered on the Uber network of drivers. The company's phenomenal success has given rise to a host of other startups using essentially the same business model, but for services other than transportation. This became known among the entrepreneur community and media as the "we're the Uber for X" [40, 144] model, where a client makes an ad-hoc order for X (service or product), which is then quickly delivered to their location.

Decision-makers often obtain strategic insights by understanding one situation via analogies with another; however, analogy-making is just one of the subjects from cognitive science that have crossed the bridge to management science. Neural networks, which can be placed alongside models of nonlinear regression in general, and more recently Deep Learning architectures and other models of large-scale parallel processing, have found use in a number of more traditional management science domains. Notable among them are credit-risk evaluation [112], the understanding of new product development [100], and consumer targeting [79].

Furthering the relevance of this cross-disciplinary exploration, it is widely known that algorithms developed for cognitive models or explored in cognitive science (such as neural networks, deep learning, reinforcement learning and Bayesian learning) are in use today powering the most advanced tools used in countless aspects of management and administration, including targeted marketing and recom-

mender systems [14], predictive modeling in business strategy [101], data-driven policy-making [122, 137, 150], financial systems [31, 104], and many other areas which are vital to business and management.

At this moment the reader may contend that fields such as applied analytics adopt a “whatever works” approach, and any system would be adopted if it yielded positive results. Therefore, any overlap between its models and those of cognitive science would be, deep-down, incidental. However, this same bottom-line view must also be seen as signaling, in the same sense, that applications of cognitive architectures are either solving problems that other algorithms cannot, or are solving them better. In both cases their relevance in applied problems, be it intentional or not, remains.

Current market events [24, 64, 103, 119, 129] show how the technology market has clearly signaled that cognitive architectures are a worthy investment. Better models of cognitive function may provide essential knowledge in what guides decision-making, memory, analogy-making and attention, among others. Any system which can provide better personal recommendations; provide strategies for better communication with customers or employees; find ways of holding clients’ attention; or induce better learning, is a relevant contribution towards management science. Here we shall explore the underlying bases which may allow for such a system to exist.

At first, the skeptical reader might argue that parts of this thesis, especially Chapter 4, are not usually within the scope of management. This is a matter of perspective. In Silicon Valley, in different corporations such as Google, Facebook, Dropbox, etc., I have participated in the same discussions with high-level researchers and managers and, from their perspective, these are amongst the pivotal problems in management today. The study of cognitive-based architectures such as Deep Learning and techniques such as Hyperdimensional Computing and Random Indexing are now highly funded within all major companies in the world (not only those involved in software technology). As entrepreneur and venture capitalist Marc Andreessen says, "Software is eating the world" [5]. The cutting edge problems of computational cognitive science *have become*, in these environments, management problems. So it is worth repeating that this is a matter of perspective: there are organizations in which these are not management problems, and there are organizations in which these are.

This thesis consists of three independent contributions which seek to explore the underpinnings of cognitive modeling theory and architecture. These, as will be argued, may provide new insights not only into better computational systems as useful tools for management science, but also a clearer understanding of the human cognitive apparatus.

The following chapter explores and refines a number of propositions collectively called Experience Recognition (XR). Indeed, the ap-

plications of a framework guided by the tenets of XR, a framework first outlined by Linhares and Freitas [91], expanded in Linhares and Chada [89], and further explored in this text, can potentially have an impact in online marketing and consumer behavior models, along with cognitive modeling and computer science; areas of research which are now more than ever becoming intertwined. As it is further explored herein, XR may prove a bridge towards better decision-support, marketing, and management tools. The chapter closes in an exposition of inadequacies of current systems, which, if corroborated, an XR-based system possibly could alleviate.

Chapter 3 introduces RSDM, a cognitive architecture, based on a highly influential memory model first presented by Kanerva [72], called SDM. An application of RSDM, which reproduces a construct from a well-founded psychological theory called *spreading activation*, is presented. It is argued that RSDM has the means to assuage a number of diverging views in cognitive science. The chapter also asserts that, through spreading activation, RSDM takes a first step as an applied cognitive model in management science, with potential application in recommender systems and other relevant areas.

Chapter 4 builds on the previous chapter, presenting results on how RSDM can provide a model for neuronal micro-circuitry and dynamic neuronal reconfiguration, argued to be central mechanisms in executive function [48] (which includes working memory, attentional control, inhibitory control and cognitive flexibiliy) and learning — thus providing a psycho-neurologically plausible architecture that can reproduce both low-level and high-level cognitive functions while also able to model the prolific network-theory interpretations of brain function.

2 DEFINING EXPERIENCE RECOGNITION

NO THEORY IS BORN FULLY GROWN AND PERFECLTY SHAPED. It must face the trials of experience and time; it must be shaped and grow from debate, divergence, difficulty and discussion. A good theory is not only robust, it is antifragile [138]. This chapter oulines the tenets of Experience Recognition (XR), as initiated by previous works [89, 91], and expands on its definition, further clarifying the theory and exposing potential applications thereof. Its relevance towards management and decision-science are discussed and defended.

2.1 INTRODUCTION: THE MIND AS A PATTERN RECOGNITION DEVICE

If we look at the human mind as a pattern-recognition device, what is the nature of its pattern-recognizing? And how does it differ from the majority of pattern-recognition methods we have collectively devised over the decades?

The study of (static) pattern recognition generally requires a database of known patterns, be it a training set or pre-established hypothesis space. Thus, a system (or theoretical model), given a new pattern, is faced with the task of classifying it against the database store. While the field of pattern recognition has brought deeply impactful contributions towards solving some central problems of modern science, this work argues that, towards the goal of modeling human cognition, one should favor a similar, but in many ways distinct approach. This work argues the case that pattern recognition as it is currently defined, is not a sufficiently accurate reflection of human psychology. It focuses on (generally) static, (usually) visible, entities, rather than the silent, invisible, cognitive process through which we humans experience our surroundings. Linhares and Freitas [91] have argued that cognitive scientists place too much emphasis on 'pattern recognition' — and scarce emphasis on *Experience Recognition*. The locus of Experience Recognition, in contrast, is a process through which the progression of the sensory influx of information is, at any point in time, mapped against a large store of experiences (or trajectories of information processing). While there are a number of similarities, the theory explored here differs in a number of aspects from Predictive State Representations (PSRs) and from other models which will be discussed.

Traditional (static) pattern recognition generally consists of training upon a data set which is presumed to be representative of all the

patterns to be seen, and subsequently classifying every encountered pattern based on what was 'learned'. One trains the system on a set of patterns, then one uses the system 'in real life', i. e. , on the data that is to be classified, based on the underlying premise that the training set was a representative sample of the space of possible patterns. While this can be useful in contained environments where there is a finite hypothesis space, the prevailing argument is that no sample can possibly represent the gamut of human experiences.

frame problem

This can be interpreted as a lightly adapted version of the *frame problem*, a well-established concept in cognitive science literature which was first stated by McCarthy and Hayes [95] and further explored by French and Anselme [45], who discuss in further detail in the same context as described here. They argue that "in any real-world situation, there are just too many representations to be considered".

A crude but useful analogy here would be that we show a child a dog, a bird and a spider, then ask them to classify all other living things with no further information. There is no training set that could completely span the breadth of variance displayed by living creatures. Humans are capable of inferring new knowledge from induction and composition based on prior experiences. Systems which can only work on a finite hypothesis space will be, by their very nature, unable to model the breadth of human experience, since there will always be something 'outside' the hypothesis space.

Finally, no mind has a 'learning/not-learning' switch. We are constantly aquiring, storing and crossing new information with already stored experiences. While there are systems which can constantly update themselves in light of new data [see 93], these still seem unable to model the influence of context on the cognitive aparatus. These systems have yet to show the flexibility to make different predictions given the same input at different times or circumstances. Moreover, rarely (if ever) are systems designed to use their own internal states as input, something which is clearly central to the human experience.

context

The term *context* is used in this work meaning "the circumstances that form the setting for a behavior, statement, or idea, and in terms of which it can be fully understood and assessed". In the realm of cognition, context may mean "the recent past that is yet exerting influence", be it a minute, an hour, or a day.

Dynamic (as opposed to static) pattern recognition, in contrast, achieves 'online' learning, that is, a system adds new data one-at-a-time, updating its internal state to make better predictions as the algorithm runs. In such models, every new pattern will update the model for the next iteration. However, this amounts to a constant retraining (although incremental) of the model at every time iteration, adding the pattern seen at time t_n to those used to evaluate the pattern at t_{n+1}; and every new pattern, every new experience (e. g. being

shot vs. having a nice cup of tea) holds equal value in updating the internal model.

This is a more flexible yet still lacking format, when the goal is a faithful model of human decision-making. Another crude but useful analogy may be to think of a person that would potentially re-evaluate their entire life at every new experience, relevant or not. Again, it intuitively proves to be a clearly implausible model of human cognition. We humans re-evaluate our experiences under highly contextual lenses, and most experiences fall by the proverbial wayside.

predictive state representations

Predictive State Representations (PSRs) by Littman, Sutton, and Singh [93], in turn, define a method for agents to constantly update their expectations of a future state of the environment given current information. The difference between prediction and reality is then used in tuning subsequent predictions. I posit that, while a useful mechanism, PSR does not encapsulate the interaction between the environment and the 'agent' (as a collection of working memory space, long-term memories and internal processing). Moreover, PSR fails to account for the internal processing as an input for further prediction, the internal state is not an input (see French and Anselme [45] for more on why this matters). In this sense, the agent itself and its own internal states must be accounted as input in prediction.

It is important to note that pattern recognition, as a field, is orthogonal to the pursuit of cognitive science: it is an immensely successful area of research with countless applied results, and does not hold the modeling of cognition as a principal (or even secondary, some might say) goal. We reinforce this notion by adding that the vast majority of pattern recognition models do not address fundamental issues of human decision-making, nor should they. Furthermore, we invoke pattern recognition here in order to address the fact that a number of research endeavours in cognitive science (see examples in [29, 55–57]) employ models and arguments that display characteristics and techniques of static pattern recognition.

Computational models based on (static or dynamic) pattern recognition usually are i) context-free (see the definition above); ii) temporal-sequence-independent, i. e. , the order in which patterns are presented does not affect the end result; iii) free of long-term memory biases, i. e. , there is no 'forgetting' and there are no early memories of encompassing influence; iv) do not utilize their own internal state as an input, and v) consider every input as having equal weight (which can be seen as a generalization of number iii).

In other words, the current pattern being processed is (typically) not affected by the previous one(s) seen; the temporal sequence of patterns the system has acquired is (usually) irrelevant; and there are no developed biases: one pattern is as good as any other. Past trajectories of information processing do not seem to affect future processing

to a large extent (other than, perhaps, incremental adjustments in preselected internal parameters). These characteristics seem to hold for the large majority of methods appearing in leading academic journals, such as *IEEE Transactions on Pattern Analysis and Machine Intelligence*. One notable exception (to numbers i, ii and iii above) are he already mentioned *predictive state representations* [93] which seem to be the closest approach to the one outlined here.

path-dependence

With these aspects in mind, I can postulate that the mind's pattern-recognizing has enormous *path-dependence*. In social science literature, path-dependence is defined as a self-reinforcing, positive feedback process, characterized by Pierson [111] as follows:

> "1. *Unpredictability.* Because early events have large effects and are partly random, many outcomes may be possible. We cannot predict ahead of time which of these possible end-states will be reached.
>
> 2. *Inflexibility.* The farther into the process we are, the harder it becomes to shift from one path to another. In applications to technology, a given subsidy to a particular technique will be more likely to shift the ultimate outcome if it occurs early rather than later. Sufficient movement down a particular path may eventually 'lock in' one solution. .
>
> 3. *Nonergodicity.* Accidental events early in a sequence do not cancel out. They cannot be treated (which is to say, ignored) as 'noise', because they feed back into future choices. Small events are remembered.
>
> 4. *Potential path inefficiency.* In the long run, the outcome that becomes established may generate lower payoffs than a foregone alternative would have. The process may be path inefficient." [111, p.18]

Its application as a mechanism of understanding processes has been shown in various fields, such as social science and politics [111], learning, and economics [9].

It is intuitive to view the mind's information-processing as highly path-dependent: someone who has experienced trauma is intrinsically different from someone who has never experienced it, yet is about to. And of course we are also strongly dependent on shared experiences: bound by taboos and customs, social norms, political institutions, formal and informal power (and status) structures, and other constructs which shape our interpretation of every context. [139].

Beyond path-dependence, in the scope of the individual, priming studies and studies from the heuristics and biases school show the enormous extent to which we are context-bound [1, 10, 131, 141, 142, 149]. Note that, in this sense, and as defined previously, the concept

of *context* is fundamentally different from that explored in standard vector-space models such as Latent Semantic Analysis [30]. The definitions are perhaps similar in that they both denote surrounding aspects: Deerwester et al. [30] use it in describing: "In different contexts or when used by different people the same term (e. g. 'chip') takes on varying referential significance". Again, in the present text, *context* implies the recent past which brings about an internal state in the cognitive apparatus, as we will explore further on.

This work postulates that the mind's pattern-recognizing is dynamically self-organizing, constantly creating and releasing constraints for further exploration: a process capable of molding itself. This flexibility and its potential neural bases will be further explored in Chapter 4.

AN ASIDE: THE COMPLEXITY OF CHUNKS HYPOTHESIS. Recently, the field of graph redescriptions in theoretical computer science has shown how all computation can be seen as transformations in graphs. Perhaps chunks, a basic 'unit of memory' [27, 96] (further defined in Section 3.3), if viewed in graph theoretical terms, are complex, with multiple hierarchical levels of description, and becoming increasingly difficult to cleanly tease apart from one another. It is possible that chunks have considerable overlap; encoding much more than the visible localist (i. e. atomic) superficial information (type, size, color, etc.). Linhares [85], Linhares [87], Linhares and Brum [88], Rehling [120] and Hofstadter and FARG [68] have proposed, for instance, that chunks encode abstract roles, and that chunks may also also encode intention. If chunks are sparse or easily clustered, like prisoners self-organized into different gangs or students into cliques, it might be easier to separate them for analysis. If they are dense, however, then attempts at a clean demarcation may be counterproductive.

The idea that chunks may contain highly abstract features and have a complex topology stems from the computational school in cognitive science that — as opposed to Fodor's claims of "modularity" — postulates that there can be no "perception" module. This school proposes that there is no vision, or any other sense for that matter, without interpretation — and the interpretation process is both culturally and contextually dependent, indeed intrinsically so [21, 42–44, 84], also [68, 69, 86].

Section 2.2 delineates the idea of process oriented modeling; and Section 2.3 discusses an example of a divergence between human information-processing and pattern-recognition algorithms. Section 2.4 specifies and proposes Experience Recognition as a new framework for cognitive modeling. Section 2.5 discusses limitations of the field of computational pattern recognition as currently defined in computer science, why these preclude their plausibility as applied cognitive

models, and closes providing a clear example of its misapplication to online marketing via contextual ads. The final sections summarize the contributions proposed in the text. The chapter closes with some final remarks on possible future applications of Experience Recognition and their expected outcome; and the direction of future research in XR.

2.2 PROCESS ORIENTED MODELING

information-- processing cone

One can informally define all the possible future information-processing pathways of an agent at any given point in time as an "information-processing cone". As a thought-experiment, let such a cone express all possible courses of interpretation and subsequent action by a decision-maker, and thus the entire space of causal reaction by that decision-maker. If one considers that each different possible interpretation within this cone has its own probability, the match between these probabilities and stored experiences may be seen as *relevance* within that information-processing cone. Within it, the most probable options are clearest when considered/undertaken. In this context, we can define understanding (used here as the apprehension of the significance of, causality of; or the possession of a meaningful anticipation in regards to) as *relevance within an information-processing cone*.

understanding

A choice-maker with a narrow cone is focused. They ignore and/or discard all alternative interpretations of a situation in favor of one (or a few), seen as preferable. Meanwhile, a choice-maker with a wide cone has many branching possibilities. In this latter, 'wide-cone' case, they are able to glimpse numerous pathways, that is, to consider a wider range of choice in reaction to a certain snapshot of stimuli. The range of possibilities (i. e. how wide the cone is) and their associated probabilities (i. e. how clearly we understand a scenario) reflects what is called *temperature* of the system at a given point in time; a system with high temperature has a wider range of possibilities than a system with a narrow range, and the probability of each scenario under consideration is spread throughout. This mechanism is present in models contained in Hofstadter and FARG [68], French [42], Marshall [94] and Mitchell [99].

temperature

Hofstadter and Mitchell describe how "as the system acquires more and more information, it starts creating a coherent viewpoint and focusing in on organizing themes. The more informed the system is, the more important it is that top-level decisions not be capriciously made" [68, p. 228]. It is worth noting that this definition diverges from that of a Bayesian hypothesis space, in that there is a constantly mutating hall of possibilities that may or may not be included in the cone, as it widens or narrows in response to stimuli.

Given a representational state at any point in time, the distribution of probabilities of the set of possible future states of a cogni-

tive system define its understanding, that is, its level of meaningful anticipation. This notion is close to the *memory-prediction framework* sketched by Hawkins and Blakeslee [59], and Littman, Sutton, and Singh [93]'s *Predictive State Representation*. Understanding is not a property of a particular state, of a particular representation, or of a particular chunk. *Understanding is a property of a process*.

Internal encoding should take precedence to raw external data (patterns) — this is the main argument of Linhares and Freitas [91], upon which this chapter expands. A war movie will bring different experiences to a veteran who has seen combat, than the experiences of other movie attendees: The same stimuli, not counting metrics of understanding, can lead to vastly different experiences in individuals.

2.3 XR: SIMAGIN'S BLUNDER AND INFORMATION-PROCESSING TRAJECTORIES

If the human mind is to be seen as a pattern-recognition device, what kind of pattern-recognition task is it accomplishing? What insights for future computational models can we gather from an experience recognition viewpoint? Modelers should strive for pattern-recognition that is i) highly path-dependent, ii) prioritizes internal encoding over external data, and iii) anticipates and utilizes the *adjacent possible*: the continuous cognitive process of anticipating multiple subconscious counterfactuals, and being surprised when the expectations are broken by an unanticipated state [75].

Concerning computational models, two additional ideas seem valuable: firstly, we should strive for theories of process. Secondly, we should focus analysis and comparisons on the information-processing pathways (not on final outputs that match/surpass human subjects). We expand on these concepts in this and the following sections.

If one's objective is to create a world chess champion computer — as it had been for the AI community for 50 years — how could we evaluate progress towards that goal? How could we know whether a new proposal was a scientific advance? How much can be gained by specific hardware, by a new pruning strategy, and so on? The historical answer has been to look at results obtained, in terms of quality of play. Hence, most research has been results-oriented. Yet, the criteria for measuring advances in psychologically-plausible, cognitive, models should be process-oriented, as opposed to output-oriented.

For an example from the domain of chess, consider the example in Figure 1. This denotes what is meant by information-processing trajectories: could a tree-searching computer fall for Simagin's blunder? Since this is a 'shallow' mate in three, there is no computer program of the tree-search paradigm that would fall for such mistake; as a small tree search leads to perfect play. But would its information-processing cone reflect Simagin's thought process?

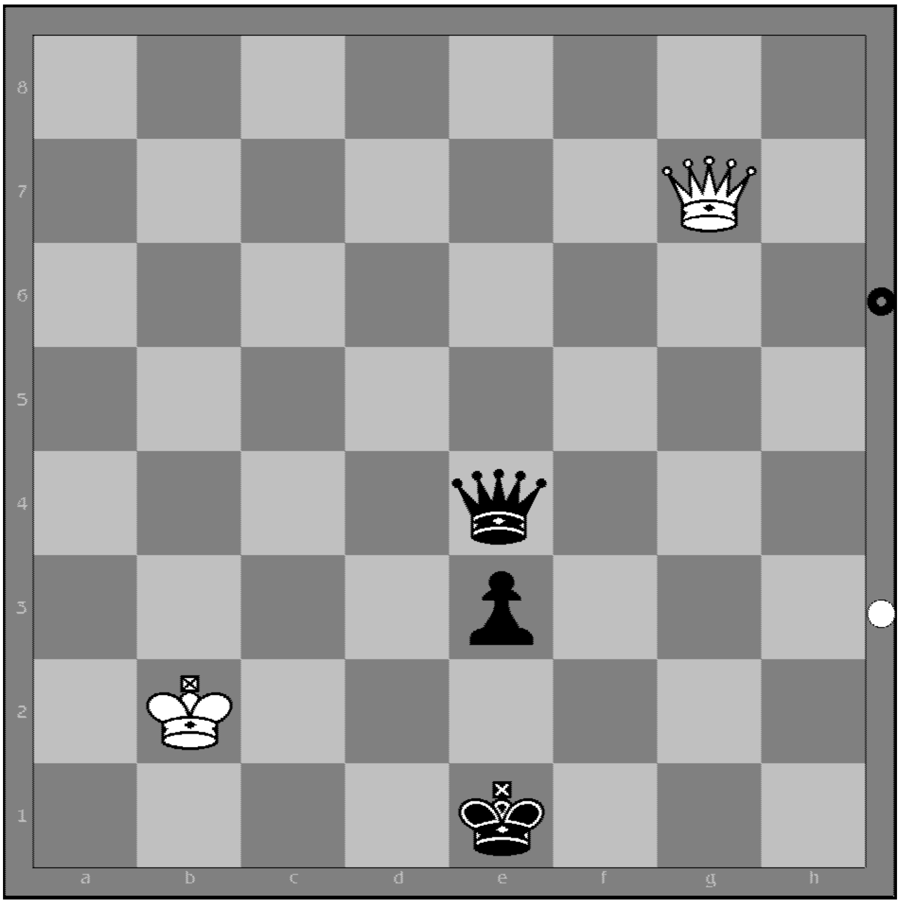

Figure 1: Batuyev versus Simagin, in Riga, 1954. Black to move. Simagin, a player of extraordinary skill, playing black, committed a blunder which no tree-searching machine ever would. With material advantage of a passed pawn, it seemed that Simagin could handle a winning endgame. However, in a rush to promote, he played pawn to e2. Since his advantage consisted on the passed pawn, one can only resonate with Simagin's attempt to promote. This led to a mate in three: white queen to g1 check, king escapes to d2, white queen to c1 check, king escapes to d3, queen to c3 checkmate.

Vladimir Simagin's clear advantage resided in having a passed pawn, that is, a pawn that no other pawn can stop from queening. Because of this, he was primed to promote it. The overconfident rush towards promotion shaped his information-processing cone into one much narrower than the computer's cone; that is, a cone that did not include (i. e. foresee) the opponent's imminent opportunity for a checkmate. Under the definition of understanding as *relevance within an information-processing cone*, we can say that a computer's 'understanding' of the combinatorial landscape involved in this scenario, however unnatural, would surpass Simagin's: there was a highly relevant adjacent counterfactual that Simagin overlooked and a machine wouldn't.

In this type of scenario, the computer resembles an automatic assembly line, diligently working on each required step, *ad nauseam*. In this case, the machine would consider each possible move and its ramifications with equal care. The chess player's process, on the

other hand, resembles craftsmanship; gradually converging into an interpretation of the situation or, in this case, hastily finishing a job, and seeing the piece ruined. And, as soon as the blunder is made, it is perceived. The realization of inevitable defeat, the response from an overlooking audience, and a host of other questions far beyond the combinatorics of chess, becomes inevitable. This shift of focus back to the real world cannot be accomplished by the machine, and while it is wired up in a form that understands the combinatorics required to win, it does not understand what 'to win' means, or what 'understanding the combinatorics' means—these are states outside of its adjacent possible, and are not encompassed in its information-processing cone. While this may sometimes prove an advantage (as shown) it does not accurately model the process or the results of cognition.

A static state in time can represent an anticipation of adjacent counterfactuals, but it is the process in which these anticipated counterfactuals change through time, the process of creation and removal of constraints of the information-processing itself, and the relevance of these adjacent counterfactuals, that should define whether or not a model is psychologically plausible. Cognition anticipates the adjacent possible using solely chunks, combinations of chunks, and knowledge of the different forms in which chunks may bind to each other: linear, hierarchical, 2-dimensional, 3-dimensional, circular, etc. In other words, the adjacent possible is wholly subjected to experience and inductive capacity. Path-dependence molds the information-processing cone by creating and releasing constraints into the next cycle of anticipation of the adjacent possible.

Path-dependence and the constant molding of the adjacent possible bring us to some questions we propose should be addressed in evaluating whether cognitive (computational) models have a semblance of Experience Recognition: contextual influence, influence of long-term experiences, domain-dependence, temporal-dependence and the ability of a model to readily self-criticize and dynamically self-organize (as one conjectures Simagin must have gone through).

2.4 METHODOLOGICAL DESIDERATA FOR XR

There are at least five questions to probe the path-dependence of a cognitive computational model's reflection of Experience Recognition:

i. *The information-processing cone must be "context-influenced".* A stimulus is never seen outside of a context. From anchoring to availability to framing, humans are incontrovertibly subject to strong influence from context. The ubiquity and enormous influence of priming is a noteworthy example: Williams and Bargh [149] conducted a striking experiment, in which subjects were asked to

interview a person and decide whether or not the person should be hired. The experimental setting had the subject meeting with a researcher, and both going up in an elevator to meet the interviewee. Inside the elevator, the subject was casually asked to hold a warm cup of coffee, or a cold cup of iced tea. The results imply that the physical sensation of warmth experienced during the brief moments in the elevator transfers to interpersonal warmth in evaluating a stranger: subjects that held a cold iced-tea cup were less likely to recommend hiring the interviewee. Context may radically alter the adjacent possible and each experience is interpreted in highly context-sensitive manner, be that context consciously acknowledged or not.

ii. *The information-processing cone must be influenced by long-term experiences.* Imagine a one-year old baby, laying down, and staring up. What would be in its field of view? Due to the availability heuristic, many readers are likely to imagine a room, toys, a cradle, walls and a ceiling, doors, windows, cabinets, and so forth. Yet, because the reader is very likely to come from a Western, Educated, Industrialized, Rich, and Democratic background—or, as Henrich, Heine, and Norenzayan [62] put it, "the WEIRDest people in the world" — caves, igloos, dense forests, ant trails, deserts, or houses made with straw, such as those from the Atr Tribes of Eritrea, are unlikely to be imagined. At this very moment, a large number of babies are most likely viewing things we cannot imagine — and such experiences will profoundly shape their psychology.

In practice, consider visual illusions, such as those of depth perception. The Ponzo illusion superimposes two lines of exact same size over the background of a "train track" (see Figure 2), and, unsurprisingly, we tend to see the lines with different sizes: for our entire lives, we have lived in a carpentered world, with right angles at walls, doors, windows, cabinets, television sets, computer screens and keyboards, tables, roads and streets, buildings, books and articles, and so forth. Our experience has taught us to use these angles to see perspective, so that objects "farther away" in our projection must be larger if they have the same size in their light projection.

This is not universal, however. Those from extremely different backgrounds, who have lived most of their lives in a world that is not carpentered, do not fall prey to such illusions, and can point out that the Ponzo illusion lines do have similar size [127]. In fact, Frederick [41] shows how people with lower IQ are less susceptible to such illusions. Henrich, Heine, and Norenzayan [62] explain:

> "...these findings suggest that visual exposure during ontogeny to factors such as the "carpentered corners" of modern environments may favor certain optical calibrations and visual habits that create and perpetuate this illusion. That is, the visual system ontogenetically adapts to the presence of recurrent features in the local visual environment. Because elements such as carpentered corners are products of particular cultural evolutionary trajectories, and were not part of most environments for most of human history, the Müller-Lyer illusion is a kind of culturally evolved by-product." .

For some people, there is no illusion. The exact same pattern projected into different eyes can create vastly different representations. A converse version of this statement also holds: extremely different patterns can, through analogy, create strikingly similar psychological experiences. The authors in [42, 68, 69, 99] study this issue in depth.

iii. *Not all experiences are the same.* People can develop a choice set skewed towards some options (or against others). At the extreme, opiates (or the stimulation of the pleasure centers of the brain) rapidly become addictive, causing a shutdown of most alternative options of behavior that detract from further stimulation. At another extreme, people face strongly negative experiences, which can induce irrational overestimation of its chances and enormous effort in order to avoid a recurrence. At the very least, experiences should lie somewhere between undesirable, indifferent, and desirable. As we seem to be far from a information-processing theory of emotion (or of pain or pleasure), the issue here is not that traditional pattern-recognition systems cannot be strongly influenced by a new pattern. The issue is with claims that pattern-recognition systems be used to reflect human psychology. (Section 2.5 will show how some current pattern-recognition systems can generate extremely inappropriate results, from a human viewpoint.)

iv. *Knowledge and behavior are domain-dependent.* Consider Wason and Shapiro's [146, p. 64] selection task: given a set of four cards, with numbers on one side and letters on the other, and a rule, such as 'every card which has a vowel on one side has an even number on the other side', subjects find it difficult to select which cards should be turned to test the rule. If, however, the same logical test is applied in a setting in which subjects have experience, the task becomes 'trivial' (i. e. 'underage students cannot drink alcohol' leads to the correct test of those drinking alcohol and those underage). The exact same logic underlies both problems, yet, if presented within a setting in which we do not have experience,

less than 10% of subjects respond correctly. Lack of experience is a lack of chunks; or, as Wason and Shapiro [146, p. 70] write: “The abstract material has no unifying link: each card is distinct and separate rather than being parts of a whole.”

The university is one domain, while real life is another. Kahneman and Tversky have famously shown how even professional statisticians, when detached from the classroom, fall for basic probability biases. Frederick [41] found that numerous university-educated people prefer 100% chances of $500 to a 15% chance of a million dollars—breaking the ‘expected utility’ viewpoint by many orders of magnitude.

The implication for cognitive models seems clear: experience acquired in a contextual setting does not immediately transfer to other domains. It takes work to create the analogies and metaphors that will provide the necessary parallels between domains. This conclusion is anathema to many computational models, such as Newell, Shaw, and Simon’s ‘General Problem Solver’ system [102].

v. *The time dimension is necessary, but not sufficient.* While it is proposed that the temporal dimension is crucial for analysis, it is not, by itself, sufficient. There are numerous computational models that deal with patterns changing through time; for instance in pattern recognition for videos or sound waveforms, also the work of Littman, Sutton, and Singh [93]. However, the main concern here is how the internal processes, representations and mental states evolve in time; how they are affected by prior experiences and how they could have evolved. When we parse a phrase like:

"PROSTITUTES APPEAL TO POPE."

There is a temporal dimension to our processing in which we deal with the cognitive dissonance of such a phrase, and re-parse it. The first interpretation brought to mind is clearly untenable; the adjacent possible demands that we reject large amounts of previous knowledge. Here one is forced into an interpretational ’double-take’, the original input must be re-parsed, taking into account that the most likely pattern is to be explicitly ignored and rejected. Any temporal model will require input from long-term memory and context to process the above example. This cannot be achieved through temporal pattern-recognition alone. This temporal processing is different from the application of pattern-recognition methods to temporal data.

Furthermore, the temporal aspect must not be simply translated into adaptability within the model. Here one posits that the interpretation of the current pattern (say, at time t) must be most influenced by the most recent patterns experienced $(t-1, t-2 \ldots t-n)$. A quick visual of this process would be the old children’s gag,

where you ask the subject to say the word ‘fork’, then ask them to say it again and again, out loud, at least ten times, then you ask: ‘What do we use to eat soup?’. The idea is that the repetition of the word ’fork’ will activate the concept enough that it will cause the subject to automatically respond with ’fork’, although the appropriate answer to the final question is ’spoon’. Another version has the subject say the word ’toast’ ten times, then ask ’what do we put in a toaster?’ (the answer is ’bread’). While this is fairly straightforward to simulate in certain computational models (such as connectionist models), we believe it remains essential in any experience recognition model.

Some models have employed the concepts which partly fill the proposed characteristics. Recent models which address episodic memory and bias acquisition in cognitive systems potentially incorporate path-dependence, context-and-domain-dependence and the considerable influence of previous experience, though none all-at-once. For examples see [77, 78, 94, 117, 143]. The reader may note also that these points resemble Piaget’s [109, 110] notion of the assimilation–accommodation cycle. For more, see also the discussion in [71, p. 29].

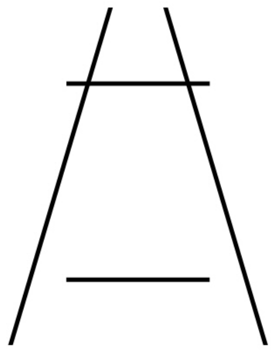

Figure 2: Ponzo Illusion – horizontal lines have same size.

2.4.1 *Affordances as a basis*

Gibson’s seminal work on vision [53] describes the theory of affordances as an interaction between properties (physical properties of things, properties of the environment, properties of events) and the observer. Affordances cannot be specified or measured strictly via the laws of physics, since they require a relation to an observer. The flat surface of the water in a still pond does not afford support to a human, but it does to a small spider. The author’s main focus is to define a theory for visual perception opposing the view that an interpretation based on visual glimpses can provide answers. In fact, Gibson proposes that “perception of the environment is not based on a sequence of snapshots but on invariant-extraction from a flux” [53, p. 304]. It is the flow of information between observer and observed, allowing the extraction of invariant properties. *affordances*

However, the author provides tools for the interpretation of all perception and interpretation: “The notion of invariants that are related at one extreme to the motives and needs of an observer and at the

other extreme to the substances and surfaces of a world provides a new approach to psychology".

Furthermore, he is careful to point out that an affordance does not change as the need of the observer changes, that is, the relation exists even if it is not realized, as opposed to the gestalt-based phenomenal relation. The ledge of an abyss affords the property of 'falling-off-ness' even if we do not experience it or realize it. A person may misperceive the affordance of a sheet of glass and walk into it. However, the author cautions that even at the highest level (speech, writing, pictures), no matter how powerful the affordances of human behavior become, they "still have to be perceived" [53, p. 137]. Here is where I would add they still have to be perceived 'as such'. Lastly, Gibson explains:

> "The theory of affordances rescues us from the philosophical muddle of assuming fixed classes of objects, each defined by its common features and then given a name. As Ludwig Wittgenstein knew, you *cannot*[1] specify the necessary and sufficient features of the class of things to which a name is given. They have only a "family resemblance." But this does not mean that you cannot learn how to use things anfd perceive their uses. You do not have to classify and label things in order to perceive what they afford."

Summarizing this chapter's contribution to the work initiated by Linhares and Freitas [91] and continued by Linhares and Chada [89]: If an affordance is a relation between the observer and the observed (environment, object, event, etc), then our experience is based on perceived affordances. In the example, Simagin's game afforded him defeat, though he was not able to see. Closing the loop, Experience Recognition is based on a path of perceived affordances.

The next section illustrates that it is possible to have high-quality pattern-recognition and low-quality experience recognition. This not only emphasizes the distinction between the two concepts in an applied scenario, but exemplifies a domain in which static pattern-recognition can be crucially orthogonal to experience recognition, and illuminates a field which would greatly benefit from the fruition of XR-based systems.

2.5 HIGH-QUALITY PATTERN-RECOGNITION AND LOW-QUALITY EXPERIENCE-RECOGNITION.

It is possible to have high-quality pattern recognition and low-quality experience-recognition. In other words, humans would be able to perceive a high-quality matching of patterns, simultaneous with a lack of understanding of at least one of the principles pointed out above

1 Emphasis present in the original work.

(context-sensitivity, desirability, domain-dependence, and the temporal dimension). A natural domain in which this arises is in the industry of contextual advertising.

Evans [34] shows a clear example of how network effects favor a winner-take-all industrial organization: Google, the market leader, has more pageviews, more advertisers (and hence a greater advertisement inventory). By comparing the results of a search for "Germany SIM cards", for instance, 8 out of 10 of Google advertisement inserts are directly relevant to the search, contrasted to only 2 out of 10 from Microsoft's system (as of this writing). One is then led to conclude, intuitively, that search-ad platforms with more advertisers will generally deliver more relevant advertisements to the searcher. This should raise the incentive for advertisers to stick with the market leader. Notice that this higher advertisement inventory effect is not only tied to web searches, but also to inserts placed on webpages. Google's massive access to web pages tends to increase their advertisement scope and relevance. The work mentions, in addition, that "like modern finance, online advertising relies heavily on advanced economic and statistical methods." Economic models (such as auctions) determine specific keyword advertisements pricing, while statistical methods are used for the automatic insertion of advertisements.

Targetted contextual advertising has been a topic of increasing scientific and technological importance [13, 22]. A growing number of studies and patents are dedicated specifically to ad-placement e.g. Henkin et al. [61], Agassi et al. [2], Bhatnagar and Papatla [13].

Figure 3: State-of-the-art pattern recognition can yield inappropriate experiences. Le Monde presents some of the first reports of the AF447 tragedy: an Air France flight from Rio de Janeiro to Paris collapsed into the Atlantic, in June 2009. Three selected advertisements from Korean Air are displayed, with soothing messages such as "le monde est ma destination" or "excellence in flight". Not only should the ads not be present next to this article, the message is wholly infelicitous.

In contextual advertisements, a webpage with baby-care content should not carry advertisements for sophisticated violins, pig medi-

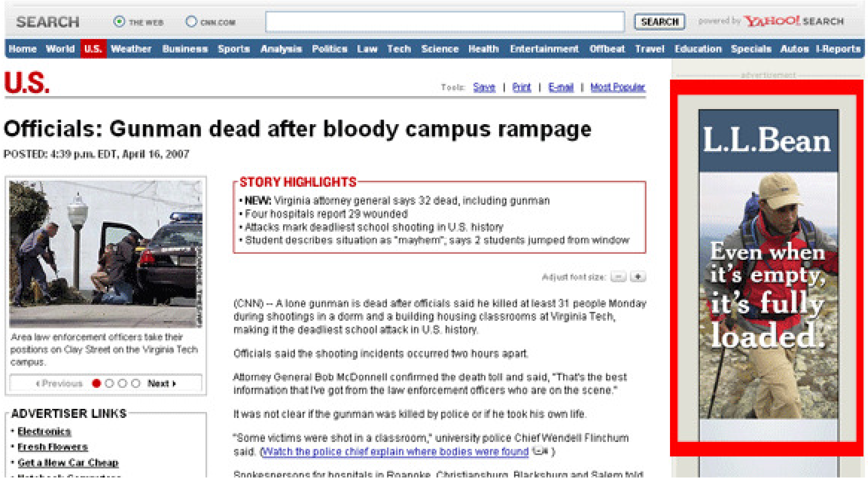

Figure 4: CNN reports the Virginia Tech Massacre. An advertisement on display reads "even when empty, it's fully loaded". Completely inappropriate in the context of a gun-violence tragedy.

cation, or SIM cards in Germany. There very little likelihood that the reader of such a webpage will have an a priori interest in violins. Even if that is the case, in the particular "baby-care" primed state in which readers may find themselves, their interest in violins will be, in all likelihood, considerably diminished. Contextual advertisement by pattern-recognition systems automatically select an advertisement that should, in principle, maximize revenue—by maximizing relevance for users. Yet, they can be extremely inappropriate. Despite being relatively rare, these inappropriate inserts might still occur in large absolute numbers due to sheer number of daily advertisements displayed on the web.

2.5.1 *Examples of inappropriate advertisements*

As examples of high-quality pattern recognition, and low-quality Experience Recognition, we refer readers to Figure 3, Figure 4, Figure 5, Figure 6, and Figure 7. These examples demonstrate advertisement systems that are able to produce high quality pattern-matching, while being oblivious to the inadequacy of the construed meaning. Billions of dollars are spent in these high-quality pattern-matching systems, yet they occasionally present incredibly inappropriate ads. Despite the name 'contextual' advertising, these messages are oblivious to context.

2.6 CONCLUSION

As a theory, Experience Recognition hopes to model the influence of an evolution of affordances, as defined by Gibson [53], on the realm of possibilities that an individual mind is able to process and consider,

washingtonpost.com Hello editor Change Preferences | Sign Out
The Washington Post Print Edition | Subscribe | PostPoints

Advertisement
Wherever you live, we'll find your best rates on auto insurance.
Enter your zip-code GO BillSavings

NEWS | POLITICS | OPINIONS | LOCAL | SPORTS | ARTS & LIVING | GOING OUT GUIDE | JOBS | CARS | REAL ESTATE | RENTALS | SHOPPING

SEARCH: go washingtonpost.com Web : Results by Google | Search Archives

washingtonpost.com > Technology

THIS STORY: READ + | Comments

Apple CEO Steve Jobs Takes Medical Leave of Absence

By Mike Musgrove and Peter Whoriskey
Washington Post Staff Writers
Wednesday, January 14, 2009; 6:41 PM

Apple chief executive Steve Jobs today announced that he will take a leave of absence as a result of health concerns.

THIS STORY
- Apple CEO Steve Jobs Takes Medical Leave of Absence
- Apple Media Advisory: Letter From Steve Jobs

"During the past week I have learned that my health-related issues are more complex than I originally thought," Jobs wrote in an e-mail sent to all Apple employees.

He said he intends to return to the company at the end of June.

Speculation about Jobs's health has swayed the company's share prices during the past year as his increasingly gaunt appearance left some investors worried. The iconic tech leader, who is often viewed as the main driving force for his company's recent successes, was treated for pancreatic cancer several years ago.

In this Oct. 14, 2008 file photo, Apple Inc. CEO Steve Jobs smiles during a product announcement at Apple headquarters in Cupertino, Calif. Jobs on Wednesday, Jan. 14, 2009 said he is taking a medical leave of absence until the end of June. Jobs told employees in an e-mail that his health issues are more complex than he thought. (AP Photo/Paul Sakuma, file) (Paul Sakuma - AP)

Enlarge Photo

TOOLBOX
Resize | Print | E-mail | Yahoo! Buzz | Save/Share +

Advertisement
JOBS BECOME OBSOLETE. TALENT DOESN'T.
Gain an all-important edge in an ever-changing job market.
GET A FREE LEARNING ASSESSMENT.
KAPLAN UNIVERSITY
Advertisement

Figure 5: The Washington Post reports that the late Steven Jobs, CEO of Apple, announces a leave of absence due to his cancer treatment. An advertisement starts with the message that "Jobs become obsolete. Talent doesn't." The intended meaning of the advertisement is clearly orthogonal to the story, yet the pattern-matching system is unable to perceived the inadequacy of the construed meaning, a connection which is inferred by humans in a matter of seconds.

be it a corporate decision-maker, child, or researcher . The scope of our thought is constantly and dynamically constrained and expanded by our affordances. Moreover, the capacities of the human cognitive apparatus are dependent on the path of affordances it has experienced. This summarizes the contribution of this chapter towards the Experience Recognition approach to cognitive modeling.

Experience Recognition seems a natural convergence of a number of previous ideas and models, including: Hofstadter's *analogy at the core of cognition* [67], Klein's *recognition-primed decision* framework [81], Hawkins and Blakeslee's *hierarchical temporal memory* [59], and, to a lesser extent, Saariluoma and Kalakoski's view of *apperception* [123], and closely related to the underlying concepts of *Predictive State Representations* [93]

This chapter has further defined the theory of Experience Recognition originated by Linhares and Freitas [91] and explored by Linhares and Chada [89]. It focused and further delineated its definition, elucidating the theory and presenting potential applications. Its relevance towards management and decision-science was discussed and defended.

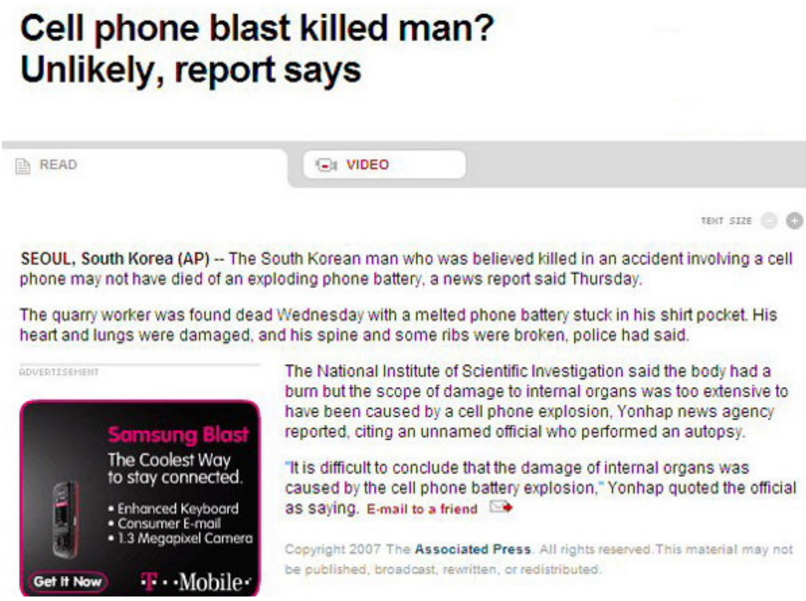

Cell phone blast killed man? Unlikely, report says

READ VIDEO

TEXT SIZE

SEOUL, South Korea (AP) -- The South Korean man who was believed killed in an accident involving a cell phone may not have died of an exploding phone battery, a news report said Thursday.

The quarry worker was found dead Wednesday with a melted phone battery stuck in his shirt pocket. His heart and lungs were damaged, and his spine and some ribs were broken, police had said.

The National Institute of Scientific Investigation said the body had a burn but the scope of damage to internal organs was too extensive to have been caused by a cell phone explosion, Yonhap news agency reported, citing an unnamed official who performed an autopsy.

"It is difficult to conclude that the damage of internal organs was caused by the cell phone battery explosion," Yonhap quoted the official as saying. E-mail to a friend

Copyright 2007 The Associated Press. All rights reserved. This material may not be published, broadcast, rewritten, or redistributed.

Figure 6: A story questions whether a cell phone battery blast killed a man in South Korea. The contextual advertisement on display is for a cell phone named "Samsung Blast". The keyword match is clear for the algorithm, but the unbecoming context is outside its capacities.

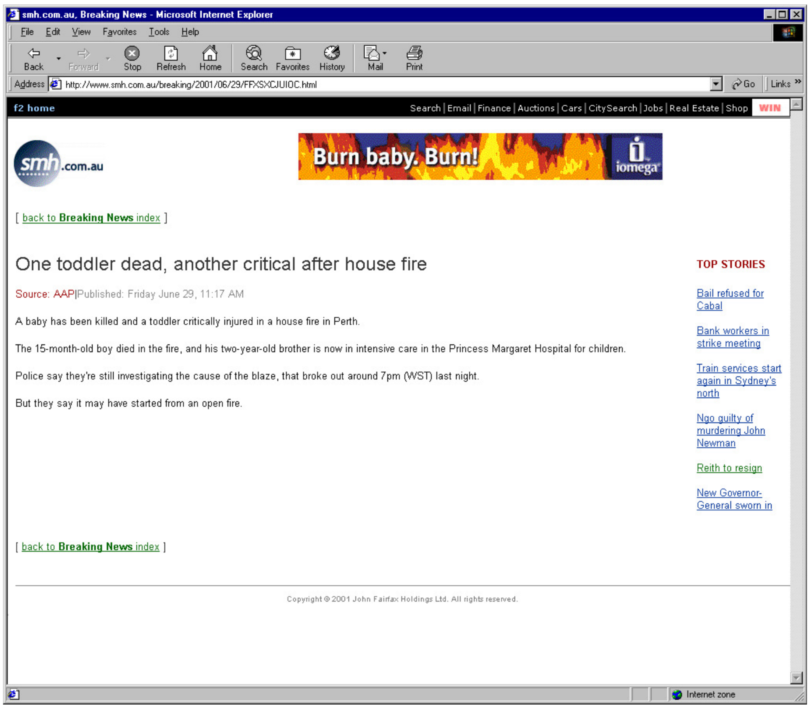

smh.com.au, Breaking News - Microsoft Internet Explorer

[back to Breaking News index]

One toddler dead, another critical after house fire

Source: AAP|Published: Friday June 29, 11:17 AM

A baby has been killed and a toddler critically injured in a house fire in Perth.

The 15-month-old boy died in the fire, and his two-year-old brother is now in intensive care in the Princess Margaret Hospital for children.

Police say they're still investigating the cause of the blaze, that broke out around 7pm (WST) last night.

But they say it may have started from an open fire.

[back to Breaking News index]

TOP STORIES
Bail refused for Cabal
Bank workers in strike meeting
Train services start again in Sydney's north
Ngo guilty of murdering John Newman
Reith to resign
New Governor-General sworn in

Copyright © 2001 John Fairfax Holdings Ltd. All rights reserved.

Figure 7: The news reports on the horrific event of a child's death in a house fire, while the algorith places the ad for a compact disc burning device with the atrocious (given the placement) line "Burn baby. Burn!".

2.7 chapter epilogue

With this theoretical framework defined, Part ii endeavours at more practical contributions:Chapter 3 introduces a cognitive framework called Rotational Sparse Distributed Memory (RSDM), based on Sparse Distributed Memory. It is briefly posited that RSDM may provide a neuro-psychologically model which is a path towards further confirming or denying the utility of Experience Recognition as a framework for cognitive modeling. However, the principal goal of the independent explorations of Chapter 3 and Chapter 4 is to establish and defend the relevance of RSDM towards the cognitive, decision-, and neuro- sciences.

Part II

ROTATIONAL SPARSE DISTRIBUTED MEMORY

"What lies behind you and
what lies in front of you,
pales in comparison
to what lies inside of you."

— Ralph Waldo Emerson (1803–1882)

3 SYMBOLIC ACTIVATION SPREADING IN A VECTOR SPACE MEMORY

THIS CHAPTER INTRODUCES the Rotational Sparse Distributed Memory (RSDM) framework, based on SDM. It describes a contribution to assuage the distributed/localist debate in cognitive science. It also contends that RSDM is a relevant architecture to the ends of cognitive science as a whole. Then, an applied example of RSDM is demonstrated instantiating semantic networks with spreading activation, a powerful cognitive architecture and a useful technique in recommendation systems for online marketing. The chapter closes with a discussion of RSDM's potential and future directions in research.

3.1 INTRODUCTION: SPARSE DISTRIBUTED MEMORY

Sparse Distributed Memory was developed by Kanerva [72]. It constitutes a memory model in which data is stored, in distributed fashion, in a vast and sparsely populated binary address space. As opposed to a traditional computer memory where each memory slot has an address, in SDM, a collection of instantiated slots act as address decoders, simulating a much larger memory.

SDM's address space is defined as a binary vector space $\{0, 1\}^N$. Thus, it allows 2^N possible locations, where N defines both the word length and the number of dimensions of the space (e. g. the memory holds binary vectors of length N). This way, the data stored is the same as the medium in which it is stored, in other words, the memory data are N-bit binary vectors being stored in N-dimensional binary addresses. We will differentiate between the two in that the instantiated address decoders are to be called hard location (HL), or neurons, or addresses as the context requires, while the full virtual space will be called memory items, items, elements, concepts, or bit strings (BStrs) — again according to context.

SDM uses Hamming distance [58] to gauge the distance between any two N-bit vectors. Hamming distance can be intuitively defined as the number of positions at which the corresponding bits on the compared bitstrings are different. The distance between 00001_b and 10101_b is two, since positions one and three (reading left-to-right) are different.

With the Hamming distance as a metric, the average distance between any two points in the space is given by the binomial distribution, and approximated by a normal curve with mean at $\frac{N}{2}$ and standard deviation $\frac{\sqrt{N}}{2}$. Given the Hamming distance and a large N,

the distance between any two points for the overwhelming majority of the space lies close to the mean. We can consider that a distance that is close to the mean signifies that almost any two bitstrings chosen at random will be orthogonal to each other, i. e. their distance will be $\frac{N}{2}$. This is the first important characteristic of the model since, unless explicitly connected by some other mechanism, any two concepts stored in memory are statistically completely orthogonal. As this text will show, there is a very simple mechanism which allows this connection to be established between any two concepts stored in the SDM, which is the main contribution of RSDM.

For large N (such as $N \geqslant 100$), it is impossible to physically instantiate every address of the space. It is here that Kanerva's insights concerning sparseness and distributed storage/retrieval come into play: 2^{20} — or a number around one million — physical (instantiated) memory locations, called hard locations, enable the representation of a large number of different bit strings. Items of a high-dimensional space with, say, 2^{1000} addresses can be stored in a mere $2^{20} \approx 10^6$ hard locations. As foreshadowing, the reader may draw the apt analogy of a finite number of neurons capable of storing and executing all human behavior and thought. For proportion, it is estimated that the adult human brain holds approximately 10^{11}, that is, approximately $2^{36.54121}$ neurons.

The disparity between the number of HLs and the size of the address space means that every single item is stored in several hard locations, and must, likewise, be retrieved in distributed fashion. Storage occurs by distributing the item in every HL within a certain threshold 'radius' given by the Hamming distance between the item's address and the associated hard locations. Different threshold values for different numbers of dimensions are used (the original work used $N = 100$; $1,000$ and $10,000$ as example dimensions). For $N = 1,000$, the distance from a random point of the space to its nearest (out of the one million) HLs will be approximately 424 bits [72, p. 56]. In this scenario, a threshold radius of 451 bits will define an N-dimensional *access sphere* containing around $1,000$ hard locations. In other words, from any point ζ of the $2^{1,000}$ space, approximately $1,000$ HLs (of the 2^{20} that are instantiated) lie within a 451-bit distance. All of these accessible HLs will be used in storing and retrieving whatever datum may be at virtual address ζ.

Here it is worth noting that SDM may be content addressable or not. In effect, the datum ζ may be stored at virtual location ζ, or it may be stored at some other location υ. In fact, the same SDM, may contain both kinds of data, having some nodes holding content-addressed data, and others holding different data (e. g. $\zeta@\zeta$ vs. $\zeta@\upsilon$). We will take advantage of this fact later on, in modeling RSDM.

A brief example of a storage and retrieval procedure in SDM is in order: to store an item x at a given (virtual) location ζ (in sparse

memory) one must access every HL within the access sphere of x and store the datum in each one. hard locations each carry N counters, or adders, one for each dimension. To store a bit string x at a HL ϕ, one must iterate through the adders of ϕ: If the i^{th} bit of x is 1, increment the i^{th} adder of ϕ, if it is 0, decrement it. Repeating this for all hard locations in x's access sphere will distribute the datum x throughout the hard locations.

Retrieval of data in SDM is also collective and distributed: to retrieve the contents of an activated (virtual) address, the adders from each hard location in range (i. e. within the access sphere) must be polled. One collects the values of each adder from all the HLs within range, assigning the i^{th} bit of the response bit string as a 1 or 0 if the sum of all the adders at the i^{th} position is positive or negative, respectively (a coin is flipped for a result of 0). Notice, however, that this information in itself is meaningless and may not correspond to any one specific datum previously stored.

In other words, to read from a location ζ in the $\{0, 1\}^N$ address space, one must activate all the hard locations in the access sphere of ζ and gather the sum of the adders at each bit position. The resulting datum will be the sum of all activated hard locations' adders'. If, for the i^{th} bit position, the sum of all bit adders is positive, the final read datum's i^{th} bit is set to 1, otherwise to 0. Thus, "SDM is distributed in that many hard locations participate in storing and retrieving each datum, and one hard location can be involved in the storage and retrieval of many data" [6]. In the case of a result of 0 that bit position is chosen at random (a 'coin flip').

All HLs within an access radius of ζ provide, collectively, the datum stored at ζ. Note Figure 8a, for an illustration of this concept. Note also that this process can be iterative. If the address obtained is somehow incomplete, corrupted, ambiguous (due to some bit positions having to be chosen at random) or has no information stored on it, it provides a new bitstring, which will activate a new access radius of hard locations to (possibly) converge to the desired original datum (see Figure 8b). For a deeper study of the iterative process and convergence in SDM, see [17, 80].

3.2 RELEVANT COGNITIVE CHARACTERISTICS OF SDM

As mentioned, SDM has the characteristic that nearly any two bitstrings at random are around $\frac{N}{2}$ bits distant, that is, they are orthogonal in the N-dimensional binary vector space. This makes it a propitious reflection of two facts about the organization of human cognition: orthogonality of random concepts, and close paths between random concepts:

ORTHOGONALITY OF RANDOM CONCEPTS: I posit that the vast majority of concepts in any human is orthogonal to all others. Con-

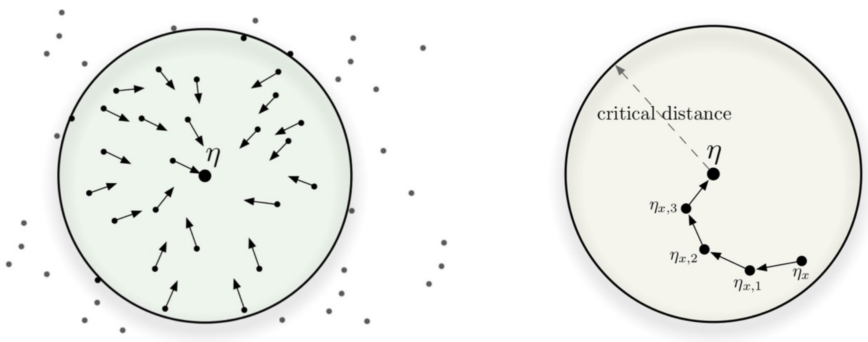

(a) hard locations pointing, approximately, to the target bit string.

(b) Iterative convergence towards a target bit string

Figure 8: Characteristics of SDM — used from Brogliato, Chada, and Linhares [17], with permission.

sider any random sample of concepts not related to cognitive science blurted out by a number of people during a cognitive science seminar, yet brought about by conscious thought: *birthdays, boots, dinosaurs, fever, executive order, x-rays*, and so on (we leave the reader to exercise their imagination and create a random sample with any number of people they see fit). Not only are the items unrelated to cognitive science, but they are also almost certainly unrelated to each other in anyone's mind.

CLOSE PATHS BETWEEN CONCEPTS: the organization of concepts in the mind seems to present a small-world topology, as defined by Watts and Strogatz [147] — for an empirical approach on words, for instance, Cancho and Sole [see 20]. To put it differently, for any two concepts, one can readily find a stream of thought relating two such items, e. g. 'Darwin gave *dinosaurs* the *boot*'; 'she ran a *fever* on her *birthday*'; 'isn't it time for the Supreme Court to *x-ray* that *executive order*?' . . . and so forth. So any two memory items can be easily and dinamically linked to form a connection (be it temporary or lasting). To emulate this, a connection between any two points in a cognitive model must be easily and dynamically created and destroyed. (Section 3.5 will show how RSDM accomplishes this.)

The iterative aspect of using the result of a SDM read as the input activation for another SDM read causes the iterative convergence towards a desired datum, or a simple divergence, where the returned results are random and orthogonal. However, if the requested address is far enough from the original item but not too far (approximately 209 bits if $N = 1,000$), repeated iterations of the process will not decrease (*or increase!*) the distance — and time to convergence goes to infinity. This is a particularly impressive emulation of humans' 'tip-of-the-tongue' phenomenon, in which one is certain about some features

of the desired memory item, yet has difficulty in retrieving it (sometimes being unable to do so — "what's the name of that american talk show host? The tall one ... with the goofy hair ... and the weird voice ... I can see his face!").

SDM is also robust against errors. This is so for at least two reasons: i) the contribution of any one hard location, in isolation, is negligible, and ii) the system can readily deal with incomplete information and still converge to a previously registered memory item. The model's sparse nature dictates that any point of the space may be used as a storage address, whether or not it corresponds to a hard location. By using about one million hard locations for $N = 1000$, the memory's distributed nature can 'virtualize' the large address space. The distributed aspect of the model allows such a virtualization.

Moreover, Kanerva [72] discusses the biological plausibility of the model, since the linear threshold function (see Haykin [60] for a comprehensive treatment) given by the access radius can be readily computed by neurons, and he suggests the interpretation of some particular types of neurons as address decoders. Chou [23] has explored the capacity limits of SDM, while Brogliato, Chada, and Linhares [17] explore the critical distance as a metric of expertise and recollection. Snaider and Franklin [135] propose an extended SDM where part of the bit string word is devoted to data and part towards addressing, effectively allowing both data and sequence information to be stored.

This robust convergence with repeated reading operations is seen as one of the principal advantages of SDM, but it is also what makes the model inflexible to vector *binding* and *bundling* (also called *composition*, or *chunking*, depending on the literature) via bit-operations such as bitwise exclusive-or (XOR) or convolution. The characteristics of these two operations will be discussed in the following section, along with other related models. The importance of vector composition will also be discussed, as the main contribution of this chapter is an alternative to this approach.

3.3 RELATED MODELS

vector symbolic architectures

Sparse Distributed Memory is usually compared to a larger class of connectionist models collectively called Vector Symbolic Architectures (VSAs). Levy and Gayler [82] provide a lucid summary of VSAs' differences from previous models:

> "VSA representations offer all of the desirable features of distributed (vector) representations (fast associative lookup, robustness to noise) while supporting systematic compositionality and rule-like behavior, and they do not rely on an innefficient or biologically implausible algorithm like back-propagation."

SDM, however, is unlike VSAs in that it does not directly implement the two principal operations that define a VSA: binding and bundling. These two operations serve as proxies for two psychological phenomena long discussed in cognitive science literature: binding and chunking. For a review on the binding problem in psychology and neuroscience, see Treisman [140]; for literature on chunking and its debate see Miller [96]; Cowan [27]. Here I will provide a brief (and summarized for the purpose of brevity and scope) description of both concepts

binding

Binding, as the pychological concept (not the vector operation) can take on a number of forms and be guided by plural mechanisms, but having one essence: it is the characteristic through which we 'attach' information pertaining to an object of concept, and which allows that object or concept to be distinguished from others.

chunking

Chunking was originally introduced by Miller [96] to account for humans' capacity to store a seemingly unlimited amount of information despite our short term memory and information processing limits. Accoring to the author, we store 'bits' of information in chunks (of limited size) which we can 'decompose' when we want to access the bits. Likewise, we can group and amalgamate chunks into other chunks with potentially unlimited recursion.

tensor product networks

VSAs were a development over Smolensky's Tensor Product Networks (TPNs) [133], which attempted to explore compositionality in filler/role bindings. TPNs proved unable to compactly represent vector composition of arbitrary complexity, given that the nature of the bundling operation multiplied the dimension of the vectors.

Kelly and West [76] define superposition and binding (the vector operation) in the following way:

> "Superposition is similarity preserving: the sum of two vectors is a vector that calls in the angle between them. Conversely, binding is similarity destroying: the circular convolution of two vectors is roughly orthogonal to the two original vectors. The purpose of superposition is to combine representations to create a new representation that is similar to all of the combined representations. The purpose of binding, on the other hand, is to create 'chunks': unique identifiers for combinations of representations."

Early VSA models, especially Holographic Reduced Representations (HRRs), [113] used vector addition for superposition, and circular convolution and vector correlation for binding and unbinding, respectively. As an example of binding, the circular convolution operation over two 3-dimensional vectors c and x, resulting in a vector t is as

follows: $t = c \circledast x$, whose elements $t_0, t_1 \ldots t_{n-1}$ can be expressed as (from Plate [113, p. 95]):

$$t_j = \sum_{k=0}^{n-1} c_k x_{j-k} \text{ for } j = 0 \text{ to } n-1$$

...making each element of t (subscripts are modulo n):

$$\begin{aligned} t_0 &= c_0x_0 + c_2x_1 + c_1x_2 \\ t_1 &= c_1x_0 + c_0x_1 + c_2x_2 \\ t_2 &= c_2x_0 + c_1x_1 + c_0x_2 \end{aligned}$$

binary spatter code

Another distributed architecture, the Binary Spatter Code (BSC) [73] makes use of random vectors with high dimensionality to achieve a level of composition, embedding sparse and dense vectors of different dimensionality into one another. In binary models, the XOR operation, which is the equivalent of computing the Hamming Distance, is generally used for superposition [74]. Both convolution and permutation have been shown to be plausible in encoding word order information [118, 124], an empirically important aspect of context in reading tasks. VSAs' functionality has been shown to be aplicable in a neurobiologically plausible architecture, called hierarchical temporal memory [105]. For further review of VSAs and applications, see Kelly and West [76] and Levy and Gayler [82].

While there is some variation in the implementation, these models unequivocally utilize methods combining vectors in order to simulate the psychological constructs of binding and chunking, either making the resulting vector an orthogonal composite or a proximate composite of its original argument vectors. This, in turn, implies the use of some vectorial metric of distance to gauge the semantic distance between concepts stored. In this aspect, the model proposed herein is an alternative account to this interpretation.

There is ample neuroscientific evidence (though it is beyond the scope of this chapter to review it) that there is specialization of function in populations of neuronal networks. At the basic level, physical areas of the neocortex and hippocampus can be identified as having a major role in this function or that function. However, it is not implied that conceptual semantics must bear the same characteristic. Indeed, there is literature to indicate the crucial influence of flexibility in representation for higher cognitive functions [21, 42, 43, 46, 66, 68, 99].

Hummel et al. [70] stake the strong claim that, as a research goal, "seeking to achieve role-filler binding, independence and integration simultaneously in a connectionist system is a powerful formula for simulating numerous aspects of human perception and cognition, and for realizing the promise of compositional connectionism". However, vector metrics imply that any representation that is a composition of other representations (any concept composed of other con-

cepts) is not a random vector. This characteristic may hinder the flexibility and psychological plausibility of models that make use of them: as we explored earlier, orthogonality of random concepts and close paths between concepts are central aspects in providing a representational framework with the flexibility required for cognitive tasks such as analogy-making, fluid representation, and path-dependence processing. In this context, despite the many problems they solve, architectures using vector operations become less psychologically-plausible.

One may argue that all concepts are made of other concepts, without there being a core atomic set of concepts upon which others are built. This is cited as one of the main rigidities in early symbolic models, which compositional connectionism (see Hummel et al. [70]) attempts to address: concepts were hard, atomic constructs. The vector metrics mechanism used to this end in nearly all connectionist architectures, however, entails some potentially insurmountable compromises.

With vector metrics, building new concepts as composite operations of other concepts forces a particular concept into a point in space. As a brief example let's say that, to me the concept of 'polar bear' is a combination of 'bear', 'white' and 'snow'. If I am exposed to images of a dancing polar bear via a popular commerical, the concept of dancing will now have to make-up my representation of 'polar bear', in order to be accessible when 'polar bear' comes to mind.

In this example, the system employing vector metrics composition will be forced to recalculate and wholly reposition 'polar bear' in the representational space in a way that it becomes a composition of four, rather than three constituent concepts. By consequence, any concept whose composition includes 'polar bear' will be required to be recalculated. As a consequence of the constant flow of stimuli and memory formation that occurs in humans, there would be a perpetual deluge of redescriptions, forcing a constant repositioning of concepts in order to maintain the appropriate representational compositions and distances. This remains true whether one uses filler/role bindings, or other forms of vector metrics.

Connectionist models, and VSAs have shown remarkable results in recent literature. Gayler and Levy [52] show that analogical mapping is achievable in a strictly distributed model, using a VSAs as an implementation, and graph isomorphisms. In related fashion, Chapter 4 addresses families of small (mostly 3 and 4 node) graphs as a building block of higher cognitive function. Other results have been achieved in analogical mapping using BSC [33] and SDM [32].

One of the contributions of the RSDM model contained herein is to provide a means to account for spreading activation in a 'reductionist' (a strictly binary vector space) distributed model. If it holds true in RSDM, it may be possible to generalize towards other models, such as HRR, BSC and other Vector Symbolic Architectures.

Section 3.3.1 will address the connectionist vs. symbolic and the distributed vs. localist debate in cognitive science, and place RSDM in its context. Section 3.4 touches upon on spreading activation models and their significance within cognitive science literature. Section 3.5 defines the Rotational Sparse Distributed Memory model. Section 3.6 will illustrate how RSDM implements spreading activation. Section 3.7 discussed implications of this new approach, and future exploration to be undertaken.

3.3.1 *Connectionist/Distributed vs. Symbolic/Localist AI*

Connectionist architectures were initially developed to address implausibilities and rigidities present in the symbolic models of early artificial intelligence (AI). The latter models were mainly characterized by localist (i. e. stored in one place) and symbolic, distinct, atomic representations of concepts, which were resistant to compositional dynamics, and unable to process any form of noisy input. While both approaches have delivered significant results, both have shown some difficulty in replicating the strengths of 'the other side'. As Smolensky, Legendre, and Miyata [134] describe:

Central Paradox of Cognition

> "Formal theories of logical reasoning, grammar and other mental faculties compel us to think of the mind as a machine for rule-based manipulation of highly structured arrays of symbols. What we know of the brain compels us to think of human information processing in terms of manipulation of a large unstructured set of numbers, the activity levels of interconnected neurons. Finally, the full richness of human behavior, both in everyday environments and in the controlled environments of the psychological laboratory, seems to defy rule-based description, displaying strong sensitivity to suble statistical factors in experience, as well as to structural properties of information. To solve the *Central Paradox of Cognition*[1] is to resolve these contradictions with a unified theory of the organization of the mind, of the brain of behavior, and of the environment".

More recent literature shows that it may be possible to unite and amalgamate the approaches, thus strengthening the theory behind both and further corroborating experimental results. As shown, VSAs' results in the field of analogical mapping are one example of this potential hybridization. Inherent in the RSDM model is the potential for spreading activation to be implemented via a distributed model. This contribution, the reader may note, is a step in bridging the results in connectionism with other architectures that use localist (atomic) symbols.

1 emphasis present in the original

3.4 SPREADING ACTIVATION

spreading activation

Spreading activation was a model first proposed by Quillian [115, 116] to build human semantic structure, which he attempted to implement into a computer simulation of memory search. If human concepts are thought to be arranged in a network, i.e. a graph of nodes, signifying concepts, connected by edges, signifying relations, the activation of any one concept would propitiate the activation of the concepts connected to it. The activation of those would then cause the activation of other concepts connected to those, and so forth, propagating down the network.

Collins and Loftus [26] later expanded the theory, proposing it as a framework capable of accounting for a number of experimental results. Since then, it has been widely corroborated when applied to studying psychological phenomena, including Alzheimer's [39], child development [11], aphasia [128], Parkinson's [36], depression [38], memory recall [37] and others. It has also yielded a number of results in computational models of cognition [see 4, 42, 68, 94, 99]. Wolverton and Hayes-Roth [151] show how spreading activation can be used to retrieve even semantically distant analogies in a concept network.

conceptual slippage

Research into models of analogy-making (not based on "A is to B as C is to D" analogical mapping) shows that it is required that any one concept be allowed to 'slip' out of its central definition, and attempt a convergence to another. French [42] explains *conceptual slippage*, a term first coined by Hofstadter [65], as "the replacement of one concept in the description of some situation by another related one. The perceived similarities that engender slippage may range from the superficial to the most abstract". This slippage is implemented in a number of applied micro-domains of limited scope. However, the underlying architectures in a number of successful models [42, 68, 94, 99] rely on models of spreading activation (though using symbolic networks).

AN ASIDE ON GRAPH REPRESENTATIONS: Beyond semantic networks, graphs provide highly flexible formats of representation. Be it a distributed model [52], a spreading activation model [4, 42, 99] or a statistical Bayesian application [77, 78]. This latter approach shows how simple generative graph-grammars can generate classes of structural forms to account for a plurality of organizations of data. Graph representations are able to capture relational properties at any level of abstraction (see Figure 9).

The following section defines and describes the principal contribution of this work: Rotational Sparse Distributed Memory. Section 3.6

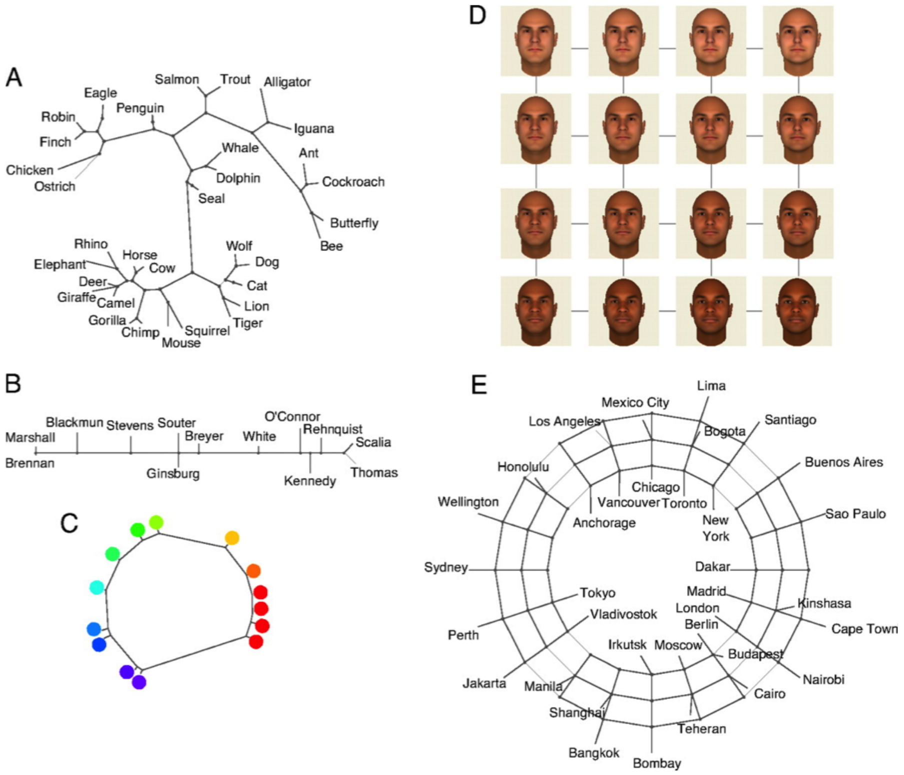

Figure 9: Structures learned from: A) biological features B) Supreme Court votes; C) judgments of the similarity between pure color wavelengths; D) Euclidean distances between faces represented as pixel vectors; E) distances between world cities — reproduced from Kemp and Tenenbaum [78].

then oulines its implementation, and presents the results in sample domains.

3.5 ROTATIONAL SPARSE DISTRIBUTED MEMORY

As we have seen, SDM shows several key-features whose neural-correlates grant the model a strong position in psychological-plausibility, all of which RSDM inherits. In other words, SDM is a model with several emergent characteristics which mirror characteristics of human psychology. However, SDM has also proven to have limitations that were addressed by either significantly altering the model (i. e. the composition of two bitstrings is no longer random but rather a calculation [74]), which would cause the same issues as in other VSAs; or building different, more complex mechanisms based on vector metrics, as in [32]. Furthermore, as described earlier, current research is investigating models that are able to provide the composition of one concept out of two or more others. In other words, the model must allow for any two stored elements to be brought together and composed, and this compostition must be another element that can be stored. Finally, the opposite operation must also hold.

However, this type of vector composition requires that concepts be composed of other concepts in a "turtles all the way down" fashion: every concept is made up of other concepts, and a concept is the compilation of those others that make it up, or 'point to it'. It may be argued that a non-hierarchical format of composition, closer to a symbolic network, may allow increased flexibility in representation and be more cognitively plausible, in the space of abstract concepts.

VSAs provide methods of composition and mapping via bit-wise operations like vector addition and convolution. These methods have achieved extensive results both in application and in furthering theory. However, the flexibility which is intrinsitc of symbolic models remains to be achieved.

In SDM, a random bit string in an N-dimensional binary space can be seen as a collection of N independent Bernoulli trials, having a binomial distribution with $p = 0.5$, having a mean of $\frac{N}{2}$ and variance $\frac{N}{4}$. Let us examine then, a rotation of the bits in any bit string α in $\{0, 1\}^N$ vector space.

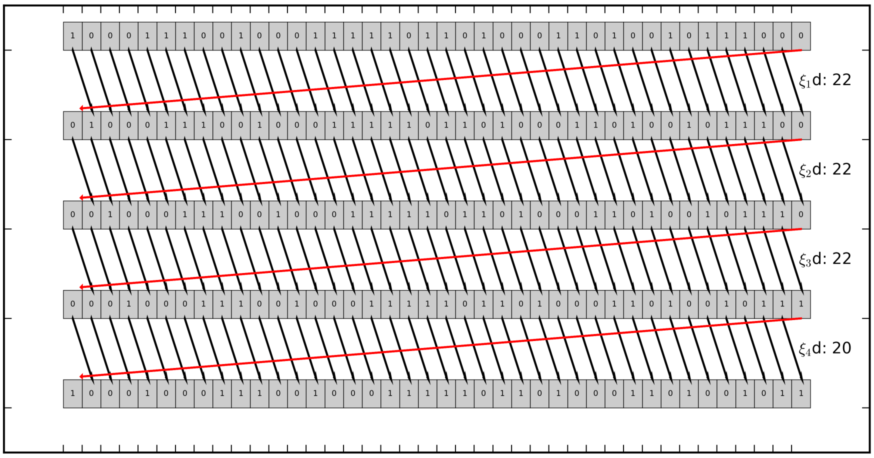

Figure 10: Randomly generated binary vector and rotations, showing the hamming distance between each rotation. This exemplifies the orthogonality of rotating a random binary vector.

Figure 10 depicts a random bit string and its rotations. It illustrates how a rotation of a bit string will yield another random bit string with a Binomial Distribution (therefore orthogonal). As N grows, so does the tendency towards orthogonality between any two vectors. So any rotation is likely to be $\frac{N}{2}$ bits away in hamming distance. Through this, rotating bits allows nearly any random BStr α in the vector space to lead to $N - 1$ other, completely different BStr in a non-stochastic fashion. This alone, however, does not allow any connection in the space to be made. To establish a connection between two previously stored locations in the vector space, any point β has to become reachable from any starting point α, not just α's $N - 1$ rotations.

Thus, one more step is required to connect any two points in the space. The vector space must be filled with both content-addressable

and pointer-style items. To store a connection between a pattern α (stored at bit string α) and pattern β (stored at β) one must choose a rotation of α, say α_{12} (that is, a 12 bit right-rotation of α), whereupon β will be stored. In this way, any two points having stored data can be connected by a non-content addressable BStr which is, in fact, orthogonal to both α and β.

Storing β at α_i is in line with Kanerva's original method of storing sequences in a Sparse Distributed Memory, where "the memory should allow the present situation to act as a retrieval cue, and, when presented with a situation that is *similar* to some previouly encountered situation, it should retrieve the *consequences* of that previous situation" [72, p.81, author's original emphases].

Storing the information of a different address β in the i^{th} rotation of an original address α makes β reachable from α, however, β has no information of how to reach any of the addresses which point to it. That is, β cannot reach α or the α_i which points directly to it. To reach α from β, α would have to be stored in some rotation of β: β_j. This makes the connection one-directional, unlike those applied by other models, such as VSAs.

Convolution and vector addition are both bi-directional in the sense that they are reversible. This means an operator which uses two vectors and results in a third can be *reversed*, allowing for the computation of the two original vectors from the third one. In other words, this would be akin to having the number 7 and being able to decompose that *this* particular 7 came from the sum of a 3 and a 4. This is a popular argument towards compositionality in distributed representations.

criterialities

A one-directional association, on the other hand, is an important potential advantage of the bit-rotation mechanism. This lies in the fact that concept association need not be bi-directional. While activating one concept may consistently activate another, the reverse need not be true. Intuitively, to borrow from Collins and Loftus's example of *criterialities*: "it might be highly criterial for the concept of a typewriter that it is a machine, and not very criterial for the concept of machine that one kind is a typewriter" [26, p. 408]. This is not to say a bi-directional association is not possible in RSDM, as indicated, it is a matter of creating two pointers: storing a pointer to β in α_i and one to α in β_j, for example.

bit rotation in RSDM

The bit-rotation mechanism in Rotational Sparse Distributed Memory employs both content-addressable (nodes) and non-content-addressable (edges) data to be stored, and allows the connection of stored patterns in a single memory space, while maintaining other neuropsychological properties that SDM provides, such as robust convergence, tip-of-the-tongue, and others. This differentiates the model described here from other mechanisms. Many VSA models based on convolution (including HRRs) propose utilizing a separate SDM as a

'clean up memory'. While this retains psychological and neurological plausibility, it garners further complexity to any model. Given that the results presented here can provide a simpler pathway, it may already prove a valid contribution.

In addition, the mechanism described herein provides a replicable model for other architectures, including VSAs and other connectionist models, so that they may further bridge the gap towards the results of higher level cognitive function frameworks, discussed in Section 3.3.1. The distance between localist and distributed representation may not be as wide as once established. Many models of cognition utilize networks to represent knowledge and many memory phenomena are explained in these models in terms of the network structures and interactions, while distribuetd representations provide robustness against noisy inputs and neuroscientific plausibility.

The results of the RSDM architecture provide a path for connectionist architectures such as VSAs to bridge this gap. RSDM, as it is implemented, brings a series of potential avenues of exploration towards decision models and cognitive modeling, and by extension towards behavioral and management research.

3.6 IMPLEMENTATION AND RESULTS IN SPREADING ACTIVATION NETWORKS

The Rotational Sparse Distributed Memory is implemented as a canonical instance of Kanerva's specification of SDM, and all models and tests are generated using the *sdm-framework project*, freely available in GitHub[2], as the underlying architecture.

The RSDM in all tests was initiated with 1000 random bit strings recorded into it. As explored in Brogliato, Chada, and Linhares [17], a network that is empty will exhibit the property of premature convergence towards any nearest input. The hamming distance for activation follows Kanerva's calculation [72], and is set at 451 bits. This establishes an access sphere for reading/writing of around 1,000 hard locations. For the purposes of the simulations shown here, the activation values are stored in an auxiliary hash-table data structure, external to the RSDM, keyed by the hexadecimal representation of each bit string.

3.6.1 *Spreading activation and activation decay*

A spreading activation network propagates any stimulus, from the starting point(s), through each connected node in the network. However, each 'hop', that is, each propagation of the initial input, must

2 The original code which generated the results is available at https://github.com/danielc2112/sdm-framework, the base sdm-framework is available at https://github.com/msbrogli/sdm-framework.

cause a decay in the activation strength, otherwise a signal would spread endlessly throughout the network. In a network with cycles or backlinks, it is trivial to see that, without some form of signal decay, any initial stimulus would eventually 'overtake' the network, causing endless activation cycles.

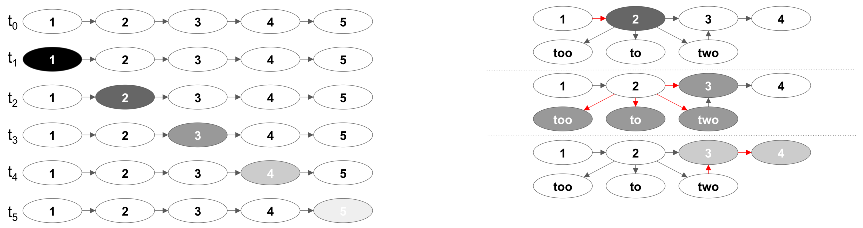

(a) Decaying spreading activation over 5 discrete time-steps.

(b) sample semantic network showing numerical and written (or phonetic) concepts, activation spreading over 3 time steps.

Figure 11: Spreading activation in semantic networks.

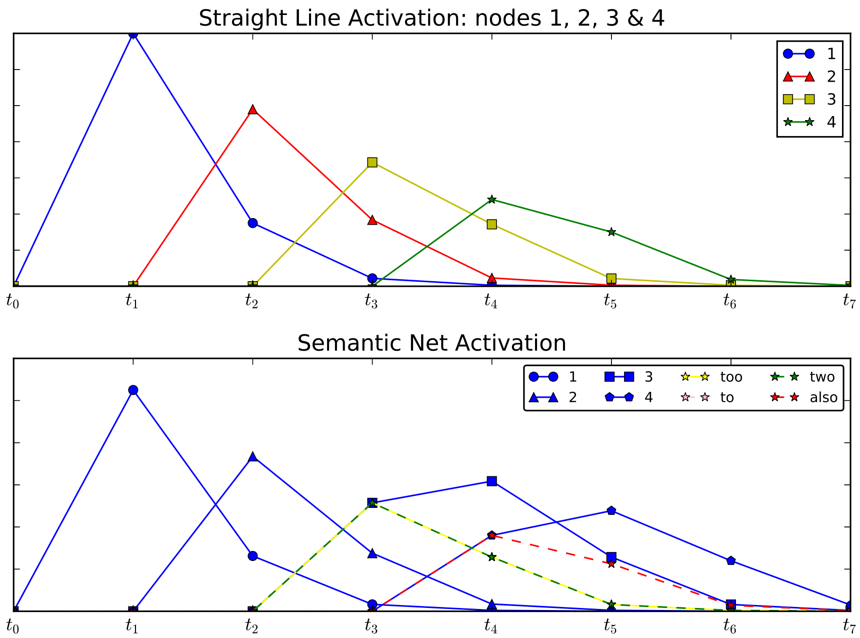

Figure 12: examples of activation in RSDM: [TOP] simulation of network activation from Figure 11a. Each note receives a decayed activation from the previous node, following Figure 11a. [BOTTOM] Semantic network activation values in RSDM for Figure 11b.

In order to provide a model of spreading activation, a metric of activation decay and thresholds of activation must be provided. Each bit string that is directly activated passes an amount of activation to each of its connected nodes (its rotations which contain data). The generic activation-decay equation is:

$$X_{n+1} = f(X_n); X_0 = \gamma \tag{1}$$

where n denotes the number of time-steps since the initial activation, $f()$ is some decay function, γ is an initial activation value. For the examples shown here γ was arbitrarily set to 20, but the input activation may be set as a function of the strength (or distance) of the activating stimulus. Additionally, a threshold value of 0.2 (see Equation 2) to consider a signal done, i.e. round it own to 0, was chosen merely for the sake of cleanliness and less calculation at each time-step. The system could allow for trace remains of a signal to approach zero on its own.

In Figure 11 we have two sample semantic networks — activation spreads from the stimulus of the concept to its connected nodes, losing strength at each hop. Figure 11a shows stimulus at '1' spreading over to '2', and so on to "5" ... we can trace the values of activation at each node retrieved from the RSDM and compare with values provided by $f()$. This is the way to corrobrate that the RSDM is functioning exactly as a semantic network with the same characteristics would. Figure 11b shows the activation at time steps 2–4 of a more complex semantic network. In this network, the concepts for numbers are interconnected with phonetically similar nodes, so the node for number '2' will activate the phonetic word 'two', and its homonyms 'to' and 'too'. From there, the concept for the *word* 'two' is also connected back to the number '3', and the concept for 'too' is connected to the concept for the word 'also'. A network entailing these same nodes can be far more interconnected, but this proved enough to demonstrate the functionality and corroborate the needed hypotheses.

For both examples, the function $f()$ modeling the decay of a node from time-step t to $t+1$ is defined within the RSDM by the following:

$$x_{t+1} = f(x_t) = \begin{cases} 0 & \text{, if } x_t \leqslant 0.2 \\ 0.125x & \text{, otherwise} \end{cases} \tag{2}$$

Spreading activation in RSDM works over an arbitrarily large sequence of nodes, given zero decay. Additionally, considering the orthogonality between one bit string rotation and another, the capacity of direct connections to be made at any arbitrary address of the memory is $N-1$ (in this specific case, $1,000-1=999$). For the examples shown, the strength of activation being passed *between* nodes, that is, the activation of one node at time t that passes to its connected nodes at $t+1$ is defined by the following:

$$\beta = g(\alpha) = \begin{cases} 0.85\alpha & \text{, if } \beta \in \Phi_\alpha \\ 0 & \text{, otherwise} \end{cases} \tag{3}$$

where Φ_α is the set of all rotations of bit string α.

It is noteworthy that, while in these examples the activation propagation characterstics (Equation 2 & Equation 3) are the same for every node in the entire memory, they need not be. a different $g()$ may be stored for different classes, or it may be node specific (though this may incur a prohibitively high storage cost). The final activation ξ_α of a node α at a given time-step is the sum of all activations ξ of the set of nodes which point to it.

$$\xi_\alpha = \sum_{\psi \in \Psi_\alpha} \xi_\psi \tag{4}$$

where Ψ_α is the set of all nodes which connect to node α via their bit rotations.

Figure 12 shows the associated activation values for each node in Figure 11. The top and bottom diagrams correspond to Figure 11a and Figure 11b, respectively. The values are activations propagated between nodes connected via bit-rotation in RSDM. Each line represents a node's activation values through each time-step of the simulation. In the top diagram, I opted to keep 5 out of the actiavtion graph, for clarity only. There are two notable aspects in the bottom graph of Figure 12. The first is the equivalence of certain nodes: the same activation path describes nodes 'too' (yellow star), 'to' (pink star) and 'two' (green star), which is perfectly in line with the hypothesized path. Second, we note that node '3' follows the same path as the above until t_3, but diverges in t_4 increasing its activation for one more time-step. This is also in agrees with the hypothesized diagram in that at t_2, node '3' receives the same activation as the other nodes connected to node '2', but at t_3 it *again* receives activation from node 'two'. Furthermore, the reader may note that node '4' receives two activations propagated from node '3', at t_4 and t_5, which matches the propagation being sent by node '3'. These instances quickly corroborate the necessary mechanisms for spreading activation semantic networks functioning in RSDM.

3.7 DISCUSSION AND FUTURE WORK

Spreading activation of concepts in human cognition is a highly context-based, path-dependent phenomenon. As a consequence, it is a good initial applied candidate for an architecture which aims at instantiating XR. There are many aspects which may play important roles in spreading activation, and are not modeled in this version of RSDM, among them rank the influence of context, and the acquisition of long-term biases.

Although the following is not within the scope of this thesis and has been left for future endeavours, RSDM *can* model contextual differences in activation and acquired long-term biases. Without much

effort, one may notice that a mechanism which 'propagates' the rotation which led to a node's activation can model context. More explicitly, if an activated node knows what length of rotation activated it, that is, if β was activated by α_{285}, β may activate only its rotations around β_{285}, say $\beta_{\{280,281,\ldots,290\}}$. In this way, only a certain rotation range would be propagated in any activation of a bit string, so this range-of-rotation is a proxy for a contextual frame-of-mind in high-level cognition, and contextual effects would be changes in the propagation of range-of-rotation.

With a bit more imagination, we may foresee, as consequence, a mechanism for short-term human memory. If only 7 ± 2 rotations are activated (as defined by Miller [96]) or 4 ± 1 rotations (as defined by Cowan [27]) are activated, we may conceive a dirty mechanism for high-level chunking. Linhares, Chada, and Aranha [90] explored a different mechanism for chunking, and acquired boundaries for the emergence of Miller [96]'s number in an SDM.

Again, as a rough thought experiment, the reader may envision that ranges of rotations within the $\{\alpha_1, \alpha_2, \ldots, \alpha_{n-1}\}$ domain may be used to record the same edge, that is, rotations $\alpha_{\{i,i+1,\ldots,j\}}$ pointing to the same node. This constitutes a simple model which can mirror two important characteristics: 'distance', widely used in distributed vectorial models, and 'closeness', or 'connection-strength' between nodes in localist/symbolic models. Hence, long-term acquired biases may be simulated by the length of a range of pointers. The strength of a connection from α to β is the length of the range of $\alpha_{\{i,i+1,\ldots,j\}}$. Relative connection strengths may be calculated as a fraction of the total $N-1$ possible outgoing connections. One outcome of this is a natural proxy for an overwhelming trigger to any stimulus. Any activation of the concept 'fire' may, in any context, trigger a panicked response from a person whose bit string for 'fire' has all its rotations pointing towards 'danger'.

Future work may also address the validity and robustness of the parameters used herein. An initial work may seek cognitive correlates to corroborate the definition of $f()$ and $g()$ in Equation 2 and Equation 3, respectively. Can RSDM model a network with non-linear interactions? What network characteristics can it robustly exhibit? Can it model a small-world network [147]? Can is corroborate the characteristics of an evolved network [98]? Can it instantiate rich-club organization [25, 145]? All these are open questions which could strengthen the architecture's relevance, and must be explored.

Aside from its cognitive correlates, instantiating spreading activation positions RSDM as an important potential tool for management science, including marketing and decision support. In 1997, Crestani [28] explored uses of spreading activation in information retrieval. Liang et al. [83] demonstrate semantic spreading activation used to instantiate a full recommendation system, whereas Blanco-Fernandéz

et al. [14] show how spreading activation over a semantic network can be used in tandem with content-based filtering for stronger recommendantion systems. These examples are cited for their relevance in applied and highly relevant fields of marketing and management, aside from currently being a billion-dollar industry in online marketing and sales. If RSDM can instantiate spreading activation, it can be a useful model in state-of-the-art recommender systems, lending its robustness to noise and a host of other useful chracteristics to this new domain of application.

This chapter has introduced Rotational Sparse Distributed Memory as an extended architecture based on Sparse Distributed Memory. It has argued the relevance of RSDM as a tool for expanding and testing theory in cognitive science and as an applied tool in management science. In the vein of Experience Recognition, it is my contention that RSDM is a prime candidate for exploring the desired characteristics of XR in a cognitive model. This initial study has only begun to explore the possibilities supplied by RSDM, and it is my hope its contributions in this and future explorations may prove relevant in the aim of modeling cognition.

4 VECTOR SPACES, COGNITION, AND NEURAL ARCHITECTURES

DECISION-MAKING BEGINS WITH CONNECTED, FIRING NEURONS. All high-level cognitive behavior can be seen as emergent phenomena based on the interactions of neural and physiological mechanisms in the brain. Recent research has begun exploring the properties and similarities between neuronal networks (i. e. networks of interconnected neurons . . . as opposed to the computer data sctructure, called *neural networks*), complex networks and their graph theoretic parallels. Papo et al. [106] provide an eloquent overview of this initiative.

This chapter provides results linking the characteristics of Rotational Sparse Distributed Memory (RSDM), defined previously, with several key results in neuroscience, network science and decision-making, pertaining to network motifs, flexible reconfiguration for dynamic processing and the neural basis of executive cognitive function (all defined below). Modeling different aspects of neuronal networks can be a good thing. Why, if decisions begin at the brain of the decision-maker, a model which can provide its basis on the fundamental architecture and functioning of the brain is a fruitful avenue for exploration of any aspect of cognitive modeling. A model that can, in theory, be applied towards real problems such as recommendation systems will be all the more resilient and sturdy if it can also present its bases in the neural underpinnings of decision making and cognition.

4.1 INTRODUCTION

In their review of new advances based on network interpretations of brain functions, Park and Friston [107] call for "a formal theory of hierarchical integration based on neuronally plausible models with coupling parameters and a well-defined computational objective". Furthermore, the authors hold that "computational modeling can serve as an interface between structure and function".

4.1.1 *Relevance Towards Cognitive Science and Decision-Making*

Once the sole province of behavioral research, decision-making has become the subject of neurobiologists, physiologists and neuroscience in general. The attractiveness of recent results signal a rise in multidisciplinary exploration of the neural bases of high level cognitive phenomena.

Giustiniani et al. [54] analyze the neural bases of strategic mechanisms under uncertainty using an Iowa Gambiling Task (a benchmark decision making assessment) experiment. Platt [114] uses the variations in the behavioral responses of monkeys given context, learned associations, and memory in order to argue that there are specific neural correlates to sensory judgements and to decisions derived from prior knowledge. These examples bespeak the encroaching importance of neuronal models of cognitive behavior in aiding research in the decision sciences. Smith and Ratcliff [132] review evidence relating recent results in behavioral psychology experiments with single-cell neuronal studies. Bogacz [15] presents evidence that optimal-decision theory can provide experimental predictions for both neurobiology and behavior research.

In the same vein, this chapter will provide correlates between RSDM and the network science interpretation of low-to-high level cognition. The contributions contained here are three-fold: i) the results show that RSDM can model the brain's principal network-motifs; ii) they also show that RSDM's architecture can provide a plausible analogue for the brain's hierarchical organization; and iii) it will be argued that RSDM's capacity for flexible reconfiguration makes it a suitable model to account and further test the structural-functional-effective mapping of brain processes.

4.2 REVIEW OF EXTANT LITERATURE

This section touches on relevant aspects of current research at the intersection of network-science, neuroscience, statistics and cognition, namely network motifs, hierarchical organization in brain networks, and dynamic, flexible reconfiguration of networks applied towards cognition. Here one hopes to champion the relevance of algorithms and systems which can model and simulate the phenomena that are presently at the forefront of these areas of investigation.

4.2.1 *Network Motifs*

network motifs

Research results indicate that real networks rely on "motifs", that is, small repeated patterns of connections between nodes that occur in statistically significantly greater numbers than they would if the network configuration was random. Network motifs are "recurring, significant patterns of interconnections" that is, "patterns for which the probability P of appearing in a randomized network an equal or greater number of times than in the real network is lower than a cut-off value", which is usually $P = 0.01$ [97]. Characteristic motifs have been found in networks such as food webs, gene regulation, electronic circuits, ecological food webs and the WWW. It is posited that theses motifs may play a part in the inherent and emergent characteristics

of the networks they constitute. Shen-Orr et al. [125], for example, conclude that much of *Escherichia Coli*'s (E. Coli) gene transcription network is dominated by three highly significant motifs.

Recurring patterns of connection, usually between three and four nodes, may provide insight into the underlying principles governing any particular network. Callaway et al. [19] prove that there are fundamental differences between evolved random networks versus static ones. Simply because the nodes in evolved networks have different lengths of existence, the networks which they form show a main interconnected component (called a giant component) which is far more robust to random edge deletion. They show that, in such a network, "twice as many edges would have to be removed from it to destroy its giant component".

Milo et al. [97] cite the two previous studies in arguing:

> "In information-processing networks, the motifs may have specific functions as elementary computational circuits [125]. More generally, they may be interpreted as structures that arise because of special constraints under which the network has evolved [19]. It is of value to detect and understand network motifs in order to gain insight into their dynamical behavior".

Indeed, early results indicate that possibly many information-passing networks may share common motifs. Furthermore, brain network motifs, small patterns of interconnected neurons (and interconnected populations of neurons) that repeat themselves throughout areas of the neocortex and hippocampus, have also been identified. These repeated patterns of neural circuitry, made of inhibitory and excitatory neurons and interconnecting synapses, are being explored as avenues towards gaining insights into the inner mechanisms of cognition. Recent work has begun studying the influence of network motifs on the structural and functional/effective levels of analysis of neuronal networks (discussed below).

Pernice et al. [108] explore how the structure of neuronal networks has strong influence over the activity dynamics of single neurons and populations. Local network properties have an effect on the dynamics of the system, and can influence the global behavior in emergent fashion.

Ferrante, Migliore, and Ascoli [35] use a realistic neuronal computer model to study feedforward inhibition circuits (shown in the diagram in Figure 13f). The authors describe how, despite the ubiquity of feedforward inhibition (FFI) in cortical regions and their functional involvement in a number of physiological mechanisms, "the specific role of FFI in shaping the neuronal I/O relationship is not yet thoroughly understood". They go on to argue:

> "Computational models can foster intuition about complex problems by requiring the precise and explicit definition of all variables and assumptions. Computer simulations cannot falsify a hypothesis, but may provide proof of concept of putative mechanisms as well as quantitative estimates of the parameter range of their aplicability."

A reasoning which emphasizes the apposite standing of the exploration in the present text. These results, and even explicit calls, bring relevance to any computational model which can exhibit the characteristics of network motifs, especially if that model will aggregate a number of other psycho-neurological behaviors which are inline with high-level cognitive phenomena (as seen in Chapter 3).

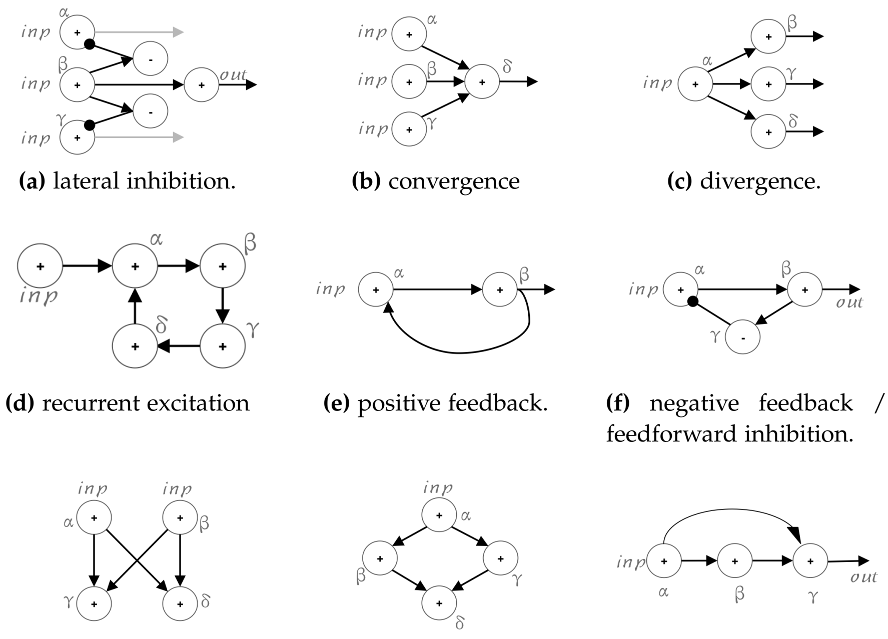

(a) lateral inhibition. **(b)** convergence **(c)** divergence.
(d) recurrent excitation **(e)** positive feedback. **(f)** negative feedback / feedforward inhibition.
(g) bi-fan. **(h)** bi-parallel. **(i)** feedforward loop.

Figure 13: neuronal motifs.

4.2.2 *Structural, Functional and Effective Brain Networks*

It is currently held that the brain can be interpreted via three hierarchical and overlapping notions: structural, functional and effective connectivity [16, 47, 107].

structural connectivity

Structural connectivity deals directly with the physical interconections between neurons and populations of neurons, i. e. the anatomy of interconnected cells. These connections are relatively stable in the short-term (e. g. minutes), but can be subject to significant changes over longer time-frames (e. g. hours to days).

FUNCTIONAL CONNECTIVITY deals with correlations in patterns of neuronal activity in remote areas of the brain. It "is defined as statistical dependencies among remote neurophysiological events" [47].

functional connectivity

EFFECTIVE CONNECTIVITY is an abstraction referring to information flow between areas of the brain. Friston [47], describes how "effective connectivity refers explicitly to the influence that one neural system exerts over another, either at a synaptic or population level". Park and Friston [107] define it "as the influence one node (neuronal population) exerts over another, under a particular network model of causal dynamics".

effective connectivity

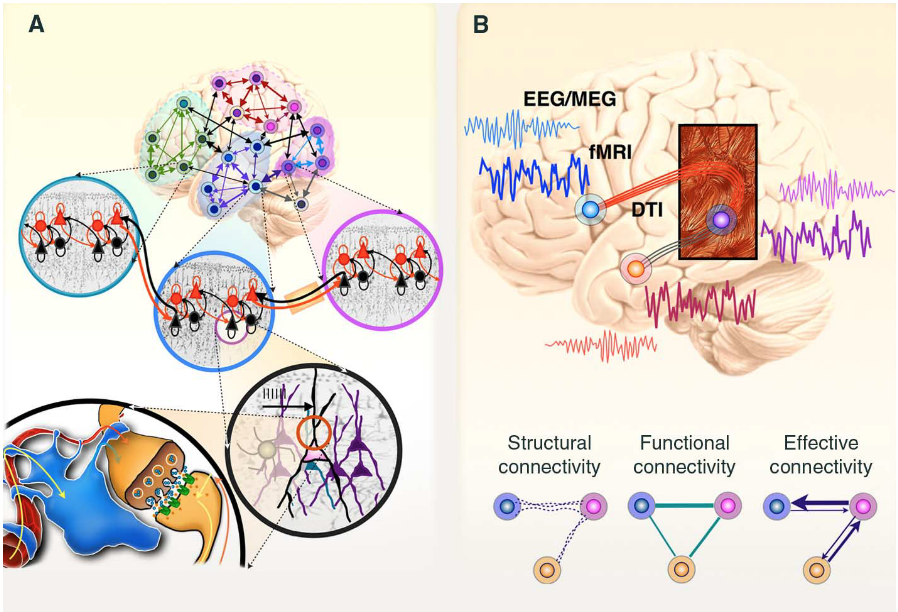

Figure 14: figure and text from: Park and Friston [107] — "*Node, edge, and organization in the brain network.* (A) Schematic of the multiscale hierarchical organization of brain networks: from neurons and macrocolumns to macroscopic brain areas. A network is composed of nodes and their links, called edges. A node, defined as an interacting unit of a network, is itself a network composed of smaller nodes interacting at a lower hierarchical level. (B) Depictions of "edges" in a brain network, as defined by three types of connectivity: structural, functional, and effective. Structural connectivity refers to anatomical connections and (macroscopically) is usually estimated by fiber tractography from diffusion tensor MRI (DTI). These connections are illustrated with broken lines in the bottom images. Functional and effective connectivity are generally inferred from the activity of remote nodes as measured by using BOLD-fMRI or EEG/MEG signals. Functional connectivity, defined by the correlation or coherence between nodes, does not provide directionality or causality and is therefore depicted without arrows. Because effective connectivity is estimated by using a model of neuronal interactions, it can evaluate directionality. This is illustrated by the one-sided arrows. Adjacency (or connectivity) matrices subserve graph theoretical analyses of brain systems and encode structural and functional connectivity between pairs of nodes".

Both functional and effective connectivity are statistical constructs borne from neuroimaging literature, but are currently central in interpreting high-level functions of the nervous system and cognitive behavior. In addition, both lessen the contribution of individual neurons or synapses to emphasize the collective functioning, where the unit

of analysis becomes a module, or population of neurons. While the unit for structural connectivity is actual synapses between neurons, functional and effective connections imply correlated links between populations of neurons (regardless of whether one specific synapse or another was activated). Effective connectivity, in particular, is able to illuminate directed flow of information and causality between functional modules of the brain. Figure 14, taken from Park and Friston [107] illustrates and describes the hiearachical nature of the concepts discussed.

It is also held that the structural connectivity (i. e. the physical interconnections between single neurons and populations) can constrain and shape the global behavior of a network [63, 126]. Meanwhile, functional and effective dynamics seem to remain flexible, exhibiting contextually incited behaviors and functions, even over the same structural network. In other words, while the structure stays the same, the function can (and does) vary given different stimuli and context.

These results suggest a divergence between the structural (physical) and functional/effective (correlation/causation level), in that one structural connectivity pattern allows for many possible patterns of behavior. This, in turn, leads to the next subject of review.

4.2.3 *Dynamic Reconfigurations and Functional/Effective Flexibility*

As an introduction to the subject, Park and Friston [107] again provide an apt description of the relevance of a dynamic reconfiguration approach:

> "Function may deviate from structure to exhibit dynamic and contextualized behavior. Such divergence of function from structure is perhaps the most intriguing property of the brain and invites further research. By studying the dynamics and self-organization of functional networks thereby enabled, we may gain insight into the true nature of the brain as the embodiment of the mind. The repertoire of functional networks will most likely emerge from the (hidden) structural architecture that enables the efficient global integration of local integrations."

Braun et al. [16] show that functional connectivity changes according to task (attention vs. memory) in a traditional "n-back" task. An n-back task is a benchmark performance task that is commonly used as an assessment in cognitive neuroscience to measure working memory and memory capacity. A subject is presented with a sequence of stimuli, and the task is indicating when the current stimulus matches the one from n steps before. As n grows, so does the difficulty of the task. Braun et al. [16] provide evidence that a reconfiguration of

large-scale distributed neural circuits occurs in order to adapt to cognitive demands. This is shown to be correlated to working memory performance and cognitive flexibility. Consequently, flexible reconfiguration (as opposed to strictly the activation of isolated brain areas) may be a behaviorally relevant aspect in human memory and information processing.

Bassett et al. [12] describe how 'network flexibility' in a given motor task during one training session is a predictor of the amount of learning in a subsequent session. In their own words: "'flexibility' is the number of times that each node changes module allegiance, normalized by the total number of changes. The flexibility of the network as a whole is then defined as the mean flexibility over all nodes". It can be described as the allegiance of a node to a specific functional module in the network. This indicates that the capacity for a particular 'node' (i. e. physical neuron or assembly of neurons) to participate in many 'modules' (i. e. functional groupings) is central to learning. Moreover, it signals that dynamic reconfiguration at the functional/ effective level of analysis is germane to human learning.

An experiment comprising a repeated association of simple drawings to their location on the screen took place in order to study *repetition supression*, that is, "a decrease of neural responses (that is, adaptation) to repeated stimulus presentations", which is "ubiquitous and a useful measure of learning-related changes intrinsic to unit or population responses". In this experiment, Büchel, Coull, and Friston [18] obtained results that, while corroborating the phenomenon, also show that "as responses within an area decrease with learning, effective connectivity between cortical areas increases. The correlation between the speed of learning and these changes in effective connectivity emphasize the relevance of plastic changes in functional integration for associative learning". A result which concludes that the adaptation of neural responses over time is material towards associations in human cognition.

To contextualize, a neuronal structure that allows functional flexibility plays a major role in every aspect of higher level cognition, from learning to memory to decision-making. One may conclude that a model which accurately simulates these dynamics has strong utility in the exploration of cognition and its epiphenomena. It is left to corroborate that RSDM can indeed model the required functional flexibility.

4.3 METHODOLOGY AND IMPLEMENTATION

Rotational Sparse Distributed Memory can be a basis for a number of applications. This chapter shows how neural micro-circuitry, including the desirable characteristics of motifs, dynamic flexibility and hierarchy can be instantiated via RSDM. Rather than re-state the de-

tails of RSDM itself (see Chapter 3), this section hopes to outline the procedures which permit the specific implementation to generate the results shown further on.

All circuits were instantiated on a $N = 1,000$ dimension RSDM, with an access radius of 451 and an activation threshold of 1.0, that is, if the activation associated with a bit string at any time t is $\geqslant 1$ that bit string propagates its signal to its connections (via its rotations).

Each write iteration was performed twice, in order to make a stronger imprint onto the HLs and prevent wrong convergence. As in Chapter 3, around $1,000$ random bit strings were written, to provide an initial population and prevent premature convergence in the RSDM. Each read was the end result of up to 6 steps of iteration for convergence.

Connecting nodes are given a valence ($+1$ or -1) which is multiplied to the exiting propagating signal, so an activated node with a negative valence propagates the negative of its activation value to all its connections. This was considered neurologically plausible since neurons are generally classified as excitatory or inhibitory, despite the same neuron being able to form more than one type of synaptic connection. So 'excitatory' nodes, when positively activated, propagate a positive signal, and 'inhibitory' nodes, likewise, propagate a negative signal.

At each time-step, all incoming signals of every bit string are added and activation is either propagated or not. Upon receipt, all signals are decayed, before being summed. The decay factor was chosen to be 0.95, and a single node's activation decay (not the propagated decay) from one time-step to the next was set at $t_{n+1} = 0.1 * t_n$. So if a note has activation 10 at t_5 and received no input, it will have activation 1 at t_6. For clarity in exposition, all activations were given a maximum/minimum of 10 and -10 respectively. Any activation above that was kept at the max/min value before decay and propagation.

For the results demonstrated in Section 4.4, all nodes were given these dynamics, effectively implementing standard artificial neurons with linear transfer functions (see Haykin [60] for a comprehensive overview).

4.4 RESULTS

The previous sections show that network motifs, hierarchy, and dynamic flexibility are central aspects of any cognitive mechanism. A computational architecture which is capable is proxying these characteristics may indeed be valuable in gaining insight into the innerworkings of human learning, behavior and decision-making.

Enter RSDM. The following results demonstrate how a high-dimensional binary vector space implemented in sparse, distributed fashion while applying rotational connectivity between its virtual nodes, can

model the characteristics of neuronal network motifs and dynamic reconfiguration of functional/effective connectivity.

It is shown here that RSDM can implement every major motif found in literature. In addition, it is demonstrated that RSDM can be dynamically (i. e. in execution-time) be reconfigured by realizing and cancelling directed connections between its nodes. Finally, I argue that the characteristics of the sparse and distributed implementation of RSDM coupled with its rotational connectivity is a natural model of the structural-functional-effective interpretation of brain hierarchy and modularity.

4.4.1 *Implementation of neuronal motifs in RSDM*

Results show that RSDM can model the most frequent neuronal motifs. The following figures show RSDM's implementation of the major neuronal motifs in current research.

convergence / divergence

Figure 15 shows the behavior of circuits from Figure 13c and Figure 13b, illustrating how RSDM can instantiate the most basic mechanisms of neuronal behavior: convergence and divergence. The addition of a neuronal signal causing strong activation in a node, and the propagation if that signal to n other nodes. Note that node 'mid' receives activation from 3 nodes, therefore its own activation is much higher (in fact capped). In turn, it propagates a much higher level of activation (at time t_3) to its target nodes μ, ν and η.

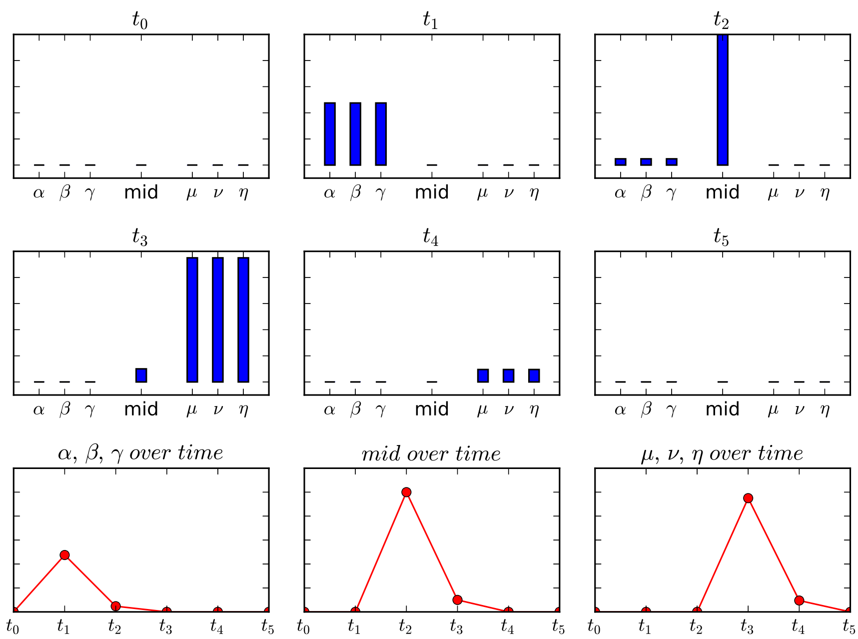

Figure 15: convergence / divergence — α, β and γ add their activation to *mid*, which in turn passes a higher activation to μ, ν and η. Diagrammed in Figure 13c and Figure 13b. Original *inp.* signal not shown, occurs only at t_0.

AN ASIDE ON FORMATTING: The activation diagrams contained in this section will all hold to a similar format. The initial panels will show the activation profile of the circuits' nodes at each pertinent time step, The first panel represents t_0, the second t_1 and so on, until the relevant behavior is fittingly exemplified. The bottom row of panels will show the evolution of one or more pertinent nodes, always compared to the original input signal received by the circuit, which is always displayed in a blue dashed line, for convenient identification. When the original input signal (always labeled 'inp' is not shown on a graph, the reader may presume it was a single input activation ocurring at t_0. Where aesthetics allow, the diagram of the motif in question will be repeated near to its activation results.

lateral inhibition

Figure 16 shows the application of *lateral inhibition*, illustrated in Figure 13a. This neuronal mechanism is known to be present in breathing and in attentional focus, dealing with execute control. The first 9 panels show how β maintains its activation while forcing α and γ to become deactivated, despite the continued incoming signal of the *input* node. The bottom 3 panels show the activation of each node, respectively, compared to the initial input (shown in the blue dashed line).

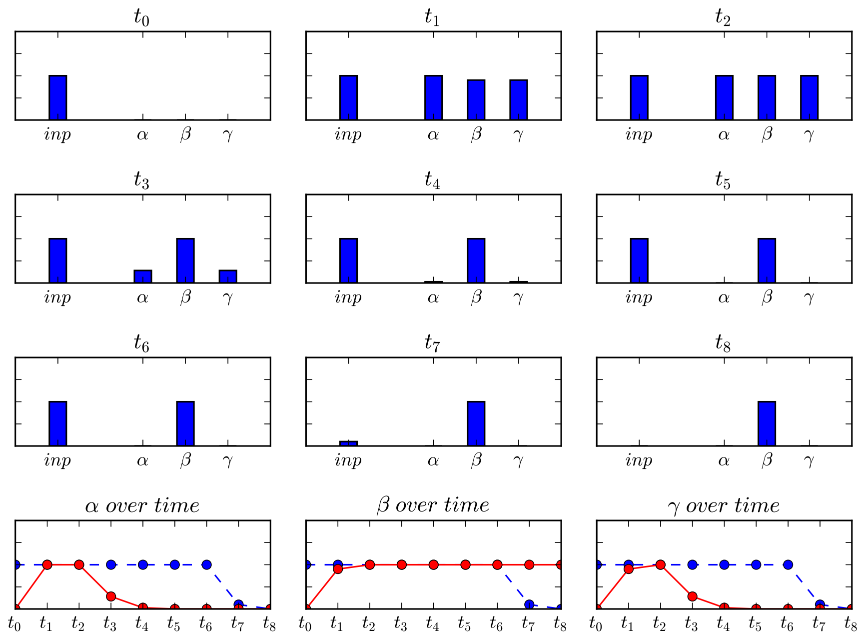

Figure 16: Lateral Inhibition — the first 9 panels show the activation behavior of nodes α, β and γ (diagram in Figure 13a), β inhibits α and γ despite sustained activation from the input. The bottom 3 panels show the activation of the 3 nodes as compared to the input stimulus shown in dashed blue ('inp' in the first 9 panels).

recurrent excitation

The *recurrent excitation* circuit (in Figure 13d) shows some interesting characteristics. In the base activation graphs, shown in Figure 17a, a single signal propagates across the circuit, and returns to α. If the input signal lasts for two time-steps, as in Figure 17b, a wave pat-

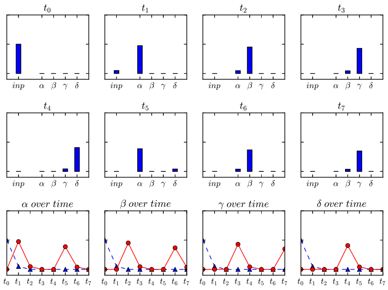

(a) single input burst at t_0 the single signal propagates forward and returns to α.

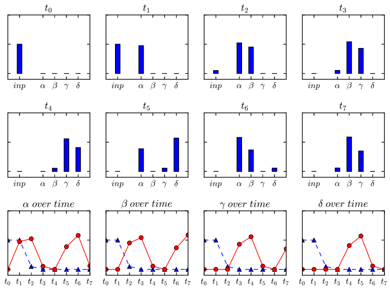

(b) double input burst at t_0 and t_1 creates a traveling wave of oscillation.

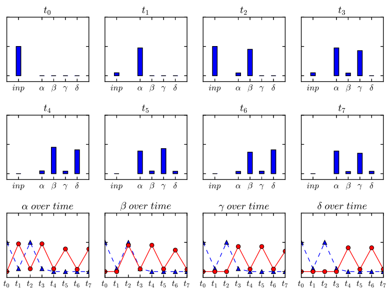

(c) intermittent double input burst at t_0 and t_2 creates a pattern of on-on, off-off activation, reminiscent of quadruped animal movement.

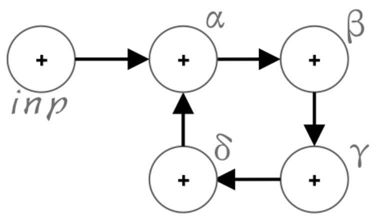

(d) motif diagram, as in Figure 13d, this same circuit exhibits the three adjacent patterns of behavior.

Figure 17: Recurrent Excitation — in the bottom line of panels in Figure 17a, Figure 17b and Figure 17c, inp. is in blue, dashed *triangles*, for clarity.

tern is formed, travelling in a circle through the circuit. This can be a proxy for wave generating neural mechanisms. In Figure 17c, the same circuit, upon receiving an intermittent burst of inputs (at t_0 and t_2) generates a "by-twos" oscillation, reminiscent of animal trotting (stepping with the front left and hind right at the same time, then the reverse). This shows traces how RSDM, in a very straightforward fashion, provides insight into how simple neural mechanisms can generate varied patterns of behavior.

positive feedback

The *positive feedback* mechanism, from the diagram in Figure 13e, shows the most basic feedback mechanism, where an activated node returns activation to its incoming node. Figure 18 shows how a single input generates a basic oscillation-style propagation in the circuit's ouput. This mechanism could be used, in turn, to generate the intermittent activation used as input to the circuit in Figure 13d generating the pattern in Figure 17c.

negative feedback

In Figure 19 the activations from circuit Figure 13f, which implement *negative feedback* are shown. Node α transmits its activation to β, which impedes the output from ever activating and, simultaneusly causes α to lose activation, despite sustained input. The reader

may note that this is a composite feedback mechanism (employing a 'middle' inhibitory node) as opposed to the direct feedback shown in Figure 13e

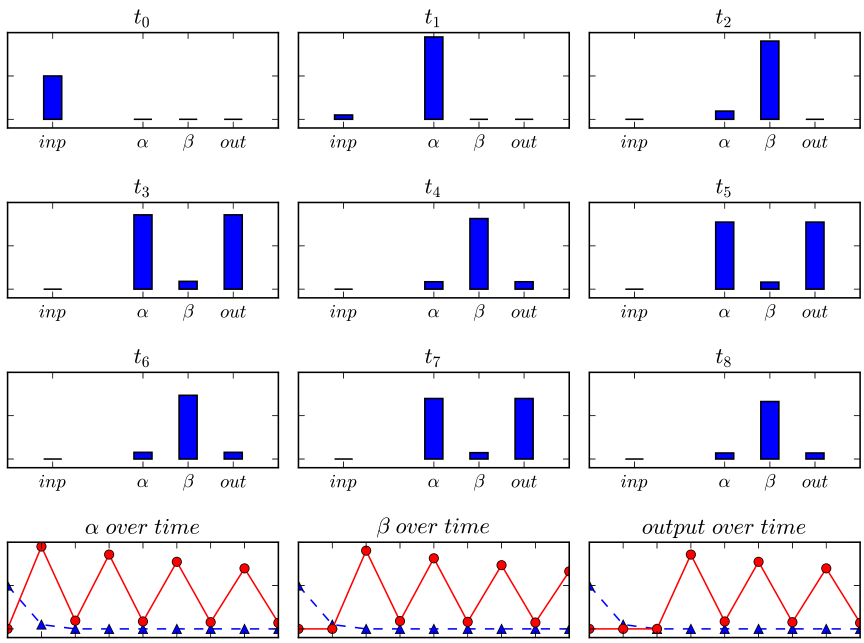

Figure 18: Positive Feedback — the activated β returns the signal to α.

feedforward loop / bi-fan / bi-parallel

The results in Figure 20a, Figure 20b, and Figure 20c show the instantiations of motifs Figure 13g, Figure 13h and Figure 13i, respectively. These are presented in unison, because they are modelled after the results in Milo et al. [97], showing that these motifs are present in statistically significant amounts in the neural architecture of *C. Elegans* — a benchmark of neuroscience research.

In Figure 20a, the input is passed to α and β, which simultaneously activate γ and δ. One may note that this could be applied as an 'OR' logical circuit, where both γ and δ are activated if α or β are activated. The characteristics of the *bi-parallel* circuit shown in Figure 20b show the amplification of the signal passed from α at δ.

4.4.2 *Flexible, dynamic reconfiguration in* RSDM

It remains to be shown that RSDM can flexibly reconfigure its connections. The lateral inhibition (Figure 13a) circuit was chosen to demonstrate this, due to its relative complexity, and simultaneous ease of demonstration. It is shown that lateral inhibition can be suppressed, then reinstated. In other words, the network instantiated in RSDM can be dynamically reconfigured. A node can begin by suppressing activation to its neighbors (β stops α and γ in Figure 16 at t_4), then the suppression mechanism, in the form of the negative connections, can be undone, allowing α and γ to grow active along with β, then the

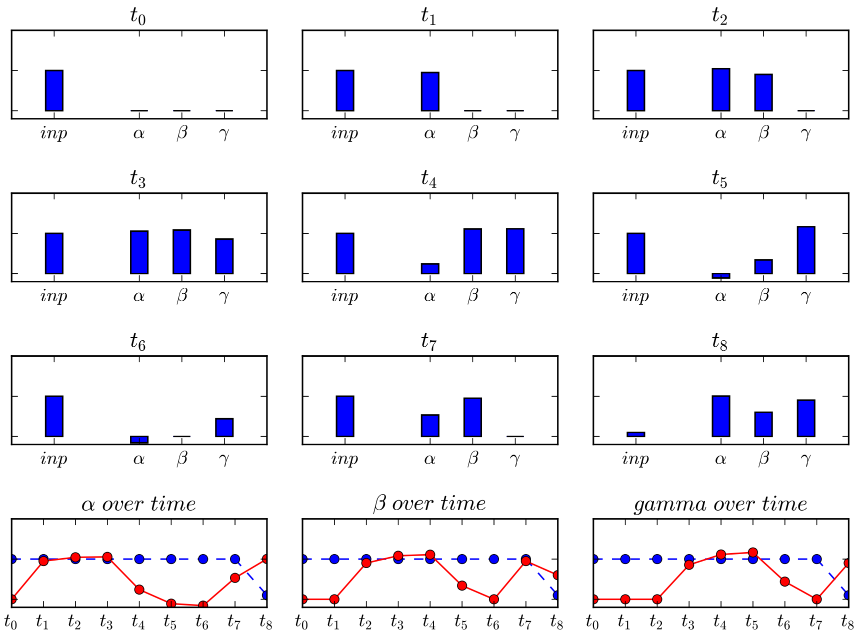

Figure 19: Negative Feedback — β, when activated, spreads negative activation towards α, cancelling the positive input.

suppression mechanism can be reinstated, all during the execution of the algorithm.

Figure 21 shows the activation profile in question: β, when activated, spreads negative activation via two inhibitory interlocutors towards α and γ, cancelling the positive signal from inp. At the end of t_5, the inhibitory connection between β and α/γ is undone, and the cancelling effect is negated, allowing them to express the activation received from inp. At t_{10} the connection is reestablished and the inhibition resumes.

So far, his section has been careful not to denote any of the constituents in any motif presented as 'neurons'. They may just as easily denote cell assemblies or entire modules of the brain. They may denote physically (i. e. at the structural level) proximate cells or relatively distant areas of the brain. In fact, I argue next that it is a more accurate description that the results shown by the RSDM be considered *effective*-level connections.

4.4.3 *Towards a computational model of the Structural-Functional-Effective*

Several research results presented in this chapter show that hierarchical interpretations of the brain's functioning is a fruitful avenue of exploration towards understanding cognition. Here I argue that there exists a straightforward correlation between the neuronal levels of analysis presented in Section 4.2.2 and RSDM. By analogy, we can interpret the RSDM architecture as holding a similar hierarchy

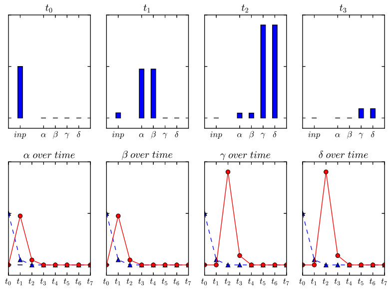

(a) Bi-Fan.

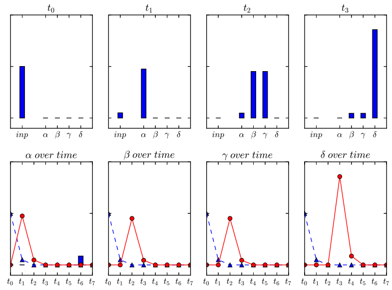

(b) Bi-Parallel.

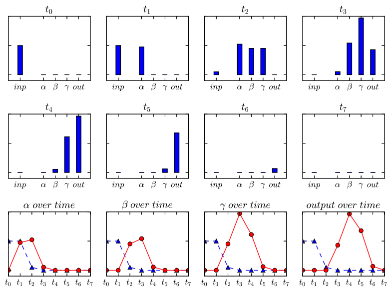

(c) Feedforward Loop.

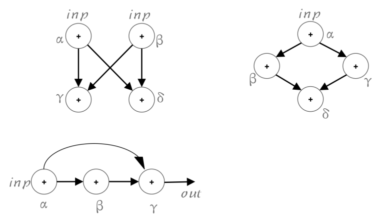

(d) diagrams

Figure 20: Significant Motifs in *C. Elegans* — from Milo et al. [97] — in the bottom line of panels in Figure 20a, Figure 20b and Figure 20c, inp. is in blue, dashed *triangles*, for clarity.

as the structural-funcional-effective intepretation of neuronal studies. This correlation should provide a solid basis for considering RSDM as a neuroscientifically plausible model of bottom-up cognitive exploration.

The initial parallel that must be drawn is at the structural level, where physical connections between neurons must be drawn. The physical distance in the brain is proxied by the *Hamming distance* in RSDM (as in SDM). The instantiated hard locations are the physical neurons, which is in line with Kanerva's original [72] interpretation. The activation of these given a certain input (in the form of a read operation at some binary address) is the activation of the brain's neurons given a certain stimulus, either from sensory information or memory itself.

The second layer of interpretation, which already deals in cell assemblies, is the functional layer. As described, here we no longer deal in single neurons, but in correlations in activation of different groupings of neurons. The equivalent in RSDM is the result of a read operation. Many 'neurons' (HLs) play a part in the activation, but any single neuron plays a minute part in the end result of the final activation. Conversely, a single neuron (hard location) may play a part several different responses to several diferent stimuli.

Lastly, and here is where this aplication of RSDM departs significantly from SDM, is the effective level, where populations of neurons

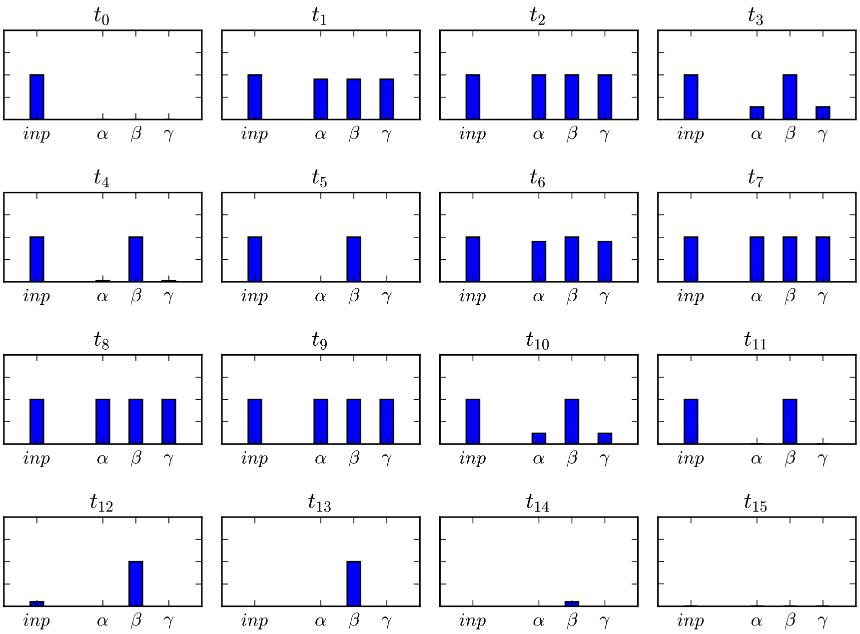

Figure 21: Reconfiguring a Lateral Inhibition in during execution — α and γ begin being inhibited by β, the inhibition is cancelled, then restored.

acting in unison become nodes, and flows of neuronal activity become edges through which information flows in the form of causally connected activation. Consequently, it becomes possible to view cognition as a network, with nodes performing different functions as individual units or as temporally amalgamated modules and clusterings. This is modeled in RSDM via the bit-rotation interconnection mechanism. It is through this mechanism that flexible, dynamic reconfiguration may be implemented and complex patterns of interconnection may arise, such as the motifs presented herein.

With these parallels drawn, one may view the RSDM architecture as an apt proxy of the human cognitive machinery. The following section will discuss some ramifications of the results (both practical and theoretical) shown here.

4.5 DISCUSSION AND FUTURE WORK

The results presented in this thesis indicate that RSDM is a prime candidate for modeling the mechanisms which generate some of the central aspects of cognitive function. As research shows, flexible reconfiguration and neuronal network motifs are important aspects of emergent higher level cognitive phenomena. Evidence accumulates towards their crucial role not only in physiological aspects such as breathing and muscular control but in executive function, including working memory, inhibitory control, attentional control and cognitive flexibility.

It is the contention of this thesis that valid models of higher-cognitive phenomena must take into account the Experience Recognition framework (see Part i), and that these mechanisms, applied via RSDM are a worthwhile path of exploration towards the end result of XR. Indeed, one may posit that to truly arrive at the first, the latter must be effectively implemented.

The subjects probed here are in the early stages of exploration in neuroscience, and computer models which implement these mechanisms for simulation or the corroboration of theories are scarce, while genuine real-world applications are even more scarce (at the time of this writing). Meanwhile, applied results in computer science and expert systems are yet dominated by traditional machine-learning algorithms or some variation of deep-learning, and models which apply concepts of the fundamental neuronal network mechanisms are almost wholly unexplored in decision-making literature.

The Rotational Sparse Distributed Memory architecture can provide an avenue for the exploration of applied expert systems, decision-support systems and simulations of behavior based on the central mechanisms of brain architecture, while maintaining utility as models of high-level cognitive phenomena, such as analogical-mapping, and semantic spreading activation.

There are a number of avenues for future work, at least three of which can be clearly contemplated. Firstly, there must be an attempt to attack some of the benchmark problems of decision-science: recommendation systems and other challenges undertaken by deep-learning architectures should provide fertile grounds. Secondly, further exploration of RSDM as a neuronal model must be undertaken — can oscillation and synchrony emerge? Can small-world and rich-club topologies be instantiated? And if so, what correlation in parameters can be made to actual brain networks? Lastly, what, if any, useful models can be drawn utilizing RSDM as a substrate for higher-level models of cognition — can RSDM be used as a network for an analogy-making system (such as those in [42, 68, 99])? If so, what advantages and drawbacks will this entail?

This study barely touches the surface of the connection between neuronal architectures and the decision sciences. There is yet much to learn about human decision-making, and about how to improve it. Inasmuch as many questions remain open, I hope to have provided a small contribution to this immensely rich and crucially important scientific endeavor. There remains a profusion of questions to be addressed, which makes for terrific research challenges and potentially ample results to be obtained. Let's get to work.

Part III

"..., COGNITION AND DECISION-MAKING"

"Everything that has a beginning has an end"

— The Oracle (Matrix Revolutions)

"And in the end, the love you take
is equal to the love you make"

— Paul McCartney (1942–)

5 FINAL REMARKS

WE CANNOT ESCAPE OUR EXPERIENCE. Every new experience we live is shaped by its predecessors. Yet, most extant cognitive models do not reflect the biases born out of accumulated experience, or the pervasive influence of context. This work postulates that the mind's pattern-recognizing process holds the following properties: it is a highly path-dependent process; it prioritizes internal encodings; it is a self-organizing process in constant change; and it constructs its future information-processing pathways by continuously recognizing the possibilities that lie within the adjacent possible. This entails a path-dependence of perceived affordances, are required traits of Experience Recognition.

Our vocabulary for the future was written in the past. The History of science and technology is filled with accidents of experience recognition; when one's personal experience is transferred to a new, poorly understood domain, and the words originally used to describe this new domain survive far into the future: meet the new Porsche, with enormous 'horsepower'. In Biology, the term 'cells' was given by Robert Hooke because, to him, their view under a microscope resembled rooms on a monastery. Later, parts of the cell were named as Chromosomes ('colored bodies') because dyes made them colorfully visible under a microscope. In Astronomy, Ptolemy rejected the idea of a moving Earth because that would imply unimaginably powerful winds, whereas the geostationary Earth was consistent with the lack of such winds. Ptolemy's personal experience prevented him from believing the Earth could be moving across a space without any such wind. As Indurkhya [71, p. 33] states:

> "When we consider changes in ontology, [an] issue to ponder is whether there are any criteria that identify certain ontology-changing moves rational and others irrational. There is obviously difficulty in positing any such criterion a priori. This is because in projecting a new theory onto a phenomenon, one cannot determine beforehand what the new ontology of the phenomenon would be and what consequences it might have".

As we attempt to model human information-processing, a productive avenue is, perhaps, to study how people match their surroundings to previous pathways of information-processing — *Experience Recognition*. If we focus on the experience, the path-dependence of affordances, we may be able to better understand how our computational models reflect, or fail to reflect, human cognition. Through

careful comparisons between human information-processing trajectories and trajectories generated by computational models, we may eventually be able to glance at advances in computational theories of intelligence, skill, and, perhaps, the mechanisms of inner human experience.

This work has provided three contributions, one towards further exploring the theory and viability of Experience Recognition as an avenue of exploration in cognitive models. The second presented Rotational Sparse Distributed Memory and established its validity as a model with psychological and neurological plausibility, in addition to showing its relevance as an applied tool. The final contribution applied RSDM towards modeling neuronal motifs, showing it can provide the flexibility and hierarchical organization inherent in current neuroscientific results.

There is much yet left to be explored, but I hope these small contributions may afford fruitful avenues for future research into the nature of human cognition.

BIBLIOGRAPHY

[1] J.M. Ackerman, C.C. Nocera, and J.A. Bargh. "Incidental haptic sensations influence social judgments." In: *Science* 328 (2010), pp. 1712–1715 (cit. on p. 10).

[2] S. Agassi et al. "Method and system for providing sponsored content based on previous provided content." Patent Application US 20070156519 (US). July 2007. URL: http://www.patentlens.net/patentlens/patent/US_20070156519/ (cit. on p. 21).

[3] D. Alba. *Uber Hits 2 Billion Rides As Growth in China Soars—— For Now*. Wired–https://www.wired.com/2016/07/uber-hits--2-billion-rides-growth-china-soars-now/. Jul. 2016 (last accessed Aug. 2016) (cit. on p. 4).

[4] J.R. Anderson. "A Spreading Activation Theory of Memory." In: *Journal of Verbal Learning and Verbal Behavior* 22.3 (1983), pp. 261–295 (cit. on p. 38).

[5] M. Andreessen. *Why Software Is Eating The World*. Wall Street Journal—http://www.wsj.com/articles/SB10001424053111903480904576512250915629460. 2104 (last accessed Aug. 2016) (cit. on p. 5).

[6] A. Anwar and S. Franklin. "Sparse Distributed Memory for 'conscious' software agents." In: *Cognitive Systems Research* 4.4 (2003), pp. 339–354 (cit. on p. 31).

[7] D. Ariely. *Predictably Irrational: The Hidden Forces That Shape Our Decisions*. New York: Harper, 2008 (cit. on p. 3).

[8] D. Ariely. *The Upside of Irrationality: The Unexpected Benefits of Defying Logic at Work and at Home*. New York: Harper, 2010 (cit. on p. 3).

[9] B.W. Arthur. *Increasing Returns and Path Dependence in the Economy*. Ann Arbor, MI: University of Michigan Press, 1994 (cit. on p. 10).

[10] J. A. Bargh and T.A. Chartrand. "The Unbearable Automaticity of Being." In: *American Psychologist* 54 (1999), pp. 462–479 (cit. on pp. 3, 10).

[11] R. Barr et al. "Age-Related Changes in Spreading Activation During Infancy." In: *Child Development* 85.2 (2014), pp. 549–563 (cit. on p. 38).

[12] D.S. Bassett et al. "Dynamic reconfiguration of human brain networks during learning." In: *Proceedings of the National Academy of Sciences* 108.18 (2011), pp. 7641–7646 (cit. on p. 55).

[13] A. Bhatnagar and P. Papatla. "Identifying Locations for Targeted Advertising on the Internet." In: *International Journal of Electronic Commerce* 5 (2001), pp. 23–44 (cit. on p. 21).

[14] Y. Blanco-Fernandéz et al. "Exploring synergies between content-based filtering and Spreading Activation techniques in knowledge-based recommender systems." In: *Information Sciences* 181 (2011), pp. 4823–4846 (cit. on pp. 5, 46).

[15] R. Bogacz. "Optimal decision-making theories: linking neurobiology with behaviour." In: *Trends in Cognitive Sciences* 11.3 (2007), pp. 118–125 (cit. on p. 50).

[16] U. Braun et al. "Dynamic reconfiguration of frontal brain networks during executive cognition in humans." In: *Proceedings of the National Academy of Sciences* 112.37 (2015), pp. 11678–11683 (cit. on pp. 52, 54).

[17] M.S. Brogliato, D.M. Chada, and A. Linhares. "Sparse Distributed Memory: understanding the speed and robustness of expert memory." In: *Frontiers in Human Neuroscience* 8 (2014), p. 222 (cit. on pp. 31–33, 42).

[18] C. Büchel, J. T. Coull, and K. J. Friston. "The Predictive Value of Changes in Effective Connectivity for Human Learning." In: *Science* 283.5407 (1999), pp. 1538–1541 (cit. on p. 55).

[19] D.S. Callaway et al. "Are randomly grown graphs really random?" In: *Physics Review E* 64.4 (2001), pp. 041902–1–041902–7 (cit. on p. 51).

[20] R.F. Cancho and R.V. Sole. "The Small World of Human Language." In: *Proceedings of the Royal Society* 268 (2001), pp. 6661–6665 (cit. on p. 32).

[21] D.J. Chalmers, R.M. French, and D.R. Hofstadter. "High level perception, representation, and analogy: a critique of artificial intelligence methodology." In: *Journal of Experimental and Theoretical Artificial Intelligence* 4.3 (1992), pp. 185–211 (cit. on pp. 11, 35).

[22] D. M. Chickering and D. Heckerman. "Targeted Advertising on the Web with Inventory Management." In: *Interfaces* 33 (2003), pp. 71–77 (cit. on p. 21).

[23] P.A. Chou. "Capacity of the Kanerva Associative Memory." In: *IEEE Transactions on Information Theory* 35.2 (1989), pp. 281–298 (cit. on p. 33).

[24] D. Clark. *Apple Buys Machine-Learning Startup Turi*. Wall Street Journal–`http://www.wsj.com/articles/apple-buys-machine-learning-startup-turi-1470448718`. 2016 (last accessed Aug. 2016) (cit. on p. 5).

[25] V. Colizza et al. "Detecting rich-club ordering in complex networks." In: *Nature Physics* 2.2 (2006), pp. 110–115 (cit. on p. 46).

[26] A.M. Collins and E.E. Loftus. "A Spreading Activation Theory of Semantic Processing." In: *Psychological Review* 82.6 (1975), pp. 407–428 (cit. on pp. 38, 41).

[27] N. Cowan. "The magical number 4 in short-term memory: A reconsideration of mental storage capacity." In: *Behavioral and Brain Sciences* 24.1 (2001), pp. 87–114 (cit. on pp. 11, 34, 46).

[28] F. Crestani. "Application of Spreading Activation Techniques in Information Retrieval." In: *Artificial Intelligence Review* 11.6 (1997), pp. 456–482 (cit. on p. 46).

[29] A. De Groot, F. Gobet, and R. W. Jongman. "Perception and Memory in Chess." In: (1996) (cit. on p. 9).

[30] S. Deerwester et al. "Indexing by Latent Semantic Analysis." In: *Journal Of The American Society For Information Science* 41.6 (1990), pp. 391–407 (cit. on p. 11).

[31] M.A.H. Dempster and V. Leemans. "An automated {FX} trading system using adaptive reinforcement learning." In: *Expert Systems with Applications* 30.3 (2006), pp. 543–552 (cit. on p. 5).

[32] B. Emruli and F. Sandin. "Analogical Mapping with Sparse Distributed Memory: A Simple Model that Learns to Generalize from Examples." In: *Cognitive Computation* 6.1 (2014), pp. 74–88 (cit. on pp. 36, 39).

[33] B. Emruli et al. "Analogical mapping and inference with binary spatter codes and sparse distributed memory." In: *International Joint Conference on Neural Networks (IJCNN)*. 2013, pp. 1–8 (cit. on p. 36).

[34] David S. Evans. "The Online Advertising Industry: Economics, Evolution, and Privacy." In: *The Journal of Economic Perspectives* 23.3 (2009), pp. 37–60 (cit. on p. 21).

[35] M. Ferrante, M. Migliore, and G.A. Ascoli. "Feed-forward inhibition as a buffer of the neuronal input-output relation." In: *Proceedings of the National Academy of Sciences* 106.42 (2009), pp. 18004–18009 (cit. on p. 51).

[36] P. Foster et al. "Spreading activation of lexical–semantic networks in Parkinson's disease." In: *Neuropsychologica* 46.7 (2008), pp. 1908–1914 (cit. on p. 38).

[37] P. Foster et al. "Recall of word lists is enhanced with increased spreading activation." In: *Aging Neuropsychology and Cognition* 20.5 (2013), pp. 553–566 (cit. on p. 38).

[38] P.S. Foster et al. "Increased spreading activation in depression." In: *Brain and Cognition* 77.2 (2011), pp. 265–270 (cit. on p. 38).

[39] P.S. Foster et al. "Differential Lexical and Semantic Spreading Activation in Alzheimer's Disease." In: *American Journal of Alzheimer's Disease and Other Dementias* 28.5 (2013), pp. 501–507 (cit. on p. 38).

[40] G.A. Fowler. *Thereś an Uber for Everything Now*. Wall Street Journal–`http://www.wsj.com/articles/theres-an-uber-for--everything-now-1430845789`. May 2015 (last accessed Aug. 2016) (cit. on p. 4).

[41] S. Frederick. "Cognitive Reflection and Decision Making." In: *Journal of Economic Perspectives* 19 (2005), pp. 25–42 (cit. on pp. 16, 18).

[42] R.M. French. *The Subtlety of Sameness*. Cambridge, MA: MIT Press, 1995 (cit. on pp. 11, 12, 17, 35, 38, 64).

[43] R.M. French. "When Coffee Cups Are Like Old Elephants, or Why Representation Modules Dont Make Sense." In: *Proceedings of the 1997 International Conference on New Trends in Cognitive Science*. Ed. by A. Riegler and M. Peschl. Austrian Society for Cognitive Science. 1997, pp. 158–163 (cit. on pp. 11, 35).

[44] R.M French. "Relational priming is to analogy-making as one-ball juggling is to seven-ball juggling." In: *Behavioral and Brain Sciences* 31 (2008), pp. 386–387 (cit. on p. 11).

[45] R.M. French and P. Anselme. "Interactively Converging on Context-sensitive Representations: a Solution to the Frame Problem." In: *Revue Internationale de Philosophie* 53.209 (3) (1999), pp. 365–385 (cit. on pp. 8, 9).

[46] R.M. French and C. Labiouse. "Why co-occurrence information alone is not sufficient to answer subcognitive questions." In: *Journal of Theoretical and Experimental Artificial Intelligence* 13 (2001), pp. 419–429 (cit. on p. 35).

[47] K.J. Friston. "Functional and Effective Connectivity: A Review." In: *Brain Connectivity* 1.1 (2011), pp. 13–36 (cit. on pp. 52, 53).

[48] S. Funahashi. "Neuronal Mechanisms of Executive Control by the Prefrontal Cortex." In: *Neuroscience Research* 39 (2001), pp. 147–165 (cit. on p. 6).

[49] G. Gavetti, D.A. Levinthal, and J.W. Rivkin. "Strategy Making in Novel and Complex Worlds: The Power of Analogy." In: *Strategic Management Journal* 26.8 (2005), pp. 691–712 (cit. on p. 3).

[50] G. Gavetti and J.W. Rivkin. "On the Origin of Strategy: Action and Cognition over Time." In: *Organization Science* 18 (2007), pp. 420–439 (cit. on p. 3).

[51] G. Gavetti and M. Warglien. “Recognizing the New: A Multi-Agent Model of Analogy in Strategic Decision-Making.” Strategy Unit Working Paper. 2007 (cit. on p. 3).

[52] R.W. Gayler and S.D. Levy. “A Distributed Basis for Analogical Mapping.” In: *Second International Conference on Analogy*. Sofia: New Bulgarian University Press, 2009, pp. 165–174 (cit. on pp. 36, 38).

[53] J.J. Gibson. *The Ecological Approach to Visual Perception*. New York: Psychology Press, 1986 (cit. on pp. 19, 20, 22).

[54] J. Giustiniani et al. “Neural Correlates of Successful and Unsuccessful Strategical Mechanisms Involved in Uncertain Decision-Making.” In: *PLoS ONE* 10.6 (2015), pp. 1–16 (cit. on p. 50).

[55] F. Gobet. “Expert memory: A comparison of four theories.” In: *Cognition* 66 (1998), pp. 115–152 (cit. on p. 9).

[56] F. Gobet and S. Jackson. “In search of templates.” In: *Cognitive Systems Research* 3.1 (2002), pp. 35–44 (cit. on p. 9).

[57] F. Gobet and H.A. Simon. “Five seconds or sixty? Presentation time in expert memory.” In: *Cognitive Science* 24 (2000), pp. 651–682 (cit. on p. 9).

[58] R.W. Hamming. “Error Detecting and Error Correcting Codes.” In: *Bell Systems Technical Journal* 29.2 (1950), pp. 147–160 (cit. on p. 29).

[59] J. Hawkins and S. Blakeslee. *On Intelligence*. New York, NY: Times Books, 2004 (cit. on pp. 13, 23).

[60] S. Haykin. *Neural Networks: A Comprehensive Foundation*. 2nd. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1998 (cit. on pp. 33, 56).

[61] A. Henkin et al. “System and method for real-time web page context analysis for the real-time insertion of textual markup objects and dynamic content.” Patent Application US 20050149395 (US). July 2005. URL: `http://www.patentlens.net/patentlens/patent/US_20050149395/` (cit. on p. 21).

[62] J. Henrich, S.J. Heine, and A. Norenzayan. “The weirdest people in the world?” In: *Behavioral and Brain Sciences* 33 (2010), pp. 61–83 (cit. on p. 16).

[63] A.M. Hermundstad et al. “Structural foundations of resting-state and task-based functional connectivity in the human brain.” In: *Proceedings of the National Academy of Sciences* 110.15 (2013), pp. 6169–6174 (cit. on p. 54).

[64] D. Hernandez. *Is Google Cornering the Market on Deep Learning?* Wired—http://www.wired.com/2014/08/deep-learning-yann-lecun/. 2104 (last accessed Aug. 2016) (cit. on p. 5).

[65] D.R. Hofstadter. *Gödel, Escher, Bach: An eternal golden braid*. New York: Basic Books, 1979 (cit. on p. 38).

[66] D.R. Hofstadter. "On Seeing A's and Seeing As." In: *Stanford Humanities Review* 4.2 (1995), pp. 109–121 (cit. on p. 35).

[67] D.R. Hofstadter. "Epilogue: analogy as the core of cognition." In: *The Analogical Mind: Perspectives from Cognitive Science*. Ed. by D. Gentner, K. J. Holyoak, and B. N. Kokinov. Cambridge: MIT Press, 2001, pp. 499–538 (cit. on p. 23).

[68] D.R. Hofstadter and FARG. *Fluid Concept and Creative Analogies: Computer Models of The Fundamental Mechanisms of Thought*. New York, NY: Basic Books, 1995 (cit. on pp. 11, 12, 17, 35, 38, 64).

[69] D.R. Hofstadter and E. Sander. *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking*. Basic Books, 2013 (cit. on pp. 3, 11, 17).

[70] J.E. Hummel et al. "A Solution to the Binding Problem for Compositional Connectionism." In: *Compositional Connectionism in Cognitive Science: Papers from the AAAI Fall Symposium*. Ed. by S.D. Levy and R. Gayler. AAAI Press, 2004, pp. 31–34 (cit. on pp. 35, 36).

[71] B. Indurkhya. "Rationality and reasoning with metaphors." In: *New Ideas in Psychology* 25 (2007), pp. 16–36 (cit. on pp. 19, 67).

[72] P. Kanerva. *Sparse Distributed Memory*. MIT Press, 1988 (cit. on pp. 6, 29, 30, 33, 41, 42, 62).

[73] P. Kanerva. "Binary spatter-coding of ordered K-tuples." In: *Artificial Neural Networks—ICANN 96: 1996 International Conference Bochum, Germany, July 16–19, 1996 Proceedings*. Ed. by C. von der Malsburg et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 1996, pp. 869–873 (cit. on p. 35).

[74] P. Kanerva. "What We Mean when We Say "What's the Dollar of Mexico?": Prototypes and Mapping in Concept Space." In: *AAAI Fall Symposium: Quantum Informatics for Cognitive, Social, and Semantic Processes*. 2010 (cit. on pp. 35, 39).

[75] S. Kaufmann. *Investigations*. New York, NY: Oxford University Press, 2000 (cit. on p. 13).

[76] M.A. Kelly and R.L. West. "From Vectors to Symbols to Cognition: The Symbolic and Sub-Symbolic Aspects of Vector-Symbolic Cognitive Models." In: *Proceedings of the 34th Annual Conference of the Cognitive Science Society (CogSci'12)*. Ed. by N. Miyake, D. Peebles, and R.P. Cooper. Cognitive Science Society. Austin, TX, 2012, pp. 1768–1773 (cit. on pp. 34, 35).

[77] C. Kemp. “The acquisition of inductive constraints.” PhD thesis. Massachusetts Institute of Technology, 2007 (cit. on pp. 19, 38).

[78] C. Kemp and J.B. Tenenbaum. “The Discovery of Structural Form.” In: *Proceedings of the National Academy of Sciences* 105.31 (2008), pp. 10687–10692 (cit. on pp. 19, 38, 39).

[79] Y. Kim et al. “Customer Targeting: A Neural Network Approach Guided by Genetic Algorithms.” In: *Management Science* 51.2 (2005), pp. 264–276 (cit. on p. 4).

[80] J. Kistoferson. *Best Probability of Activation and Performance Comparisons for Several Designs of Sparse Distributed Memory*. Tech. rep. 1995 (cit. on p. 31).

[81] G. Klein. *Sources of power: how people make decisions*. Cambridge: MIT Press, 1999 (cit. on p. 23).

[82] S. Levy and R.W. Gayler. “Vector Symbolic Architectures: A New Building Material for Artificial General Intelligence.” In: *Proceedings of the First Conference on Artificial General Intelligence (AGI'08)*. Amsterdam, The Netherlands: IOS Press, 2008, pp. 414–418 (cit. on pp. 33, 35).

[83] T. Liang et al. “A semantic-expansion approach to personalized knowledge recommendation.” In: *Decision Support Systems* 45 (2008), pp. 401–412 (cit. on p. 46).

[84] A. Linhares. “A glimpse at the metaphysics of Bongard problems.” In: *Artificial Intelligence* 121 (2000), pp. 251–270 (cit. on p. 11).

[85] A. Linhares. “An active symbols theory of chess intuition.” In: *Minds and Machines* 15 (2005), pp. 131–181 (cit. on p. 11).

[86] A. Linhares. “Dynamic sets of potentially interchangeable connotations: a theory of mental objects.” In: *Behavioral and Brain Sciences* 31 (2008), pp. 389–390 (cit. on p. 11).

[87] A. Linhares. “The emergence of choice: Decision-making and strategic thinking through analogies.” In: *Information Sciences* 259 (2014), pp. 36–56 (cit. on p. 11).

[88] A. Linhares and P.R.S. Brum. “Understanding our Understanding of Strategic Scenarios: What Role Do Chunks Play?” In: *Cognitive Science* 31 (2007), pp. 989–1007 (cit. on pp. 3, 11).

[89] A. Linhares and D.M. Chada. “What is the nature of the mind's pattern-recognition process?” In: *New Ideas in Psychology* 31.2 (2013), pp. 108 –121 (cit. on pp. 6, 7, 20, 23).

[90] A. Linhares, D.M. Chada, and C.N. Aranha. “The Emergence of Miller's Magic Number on a Sparse Distributed Memory.” In: *PLoS ONE* 6.1 (2011), pp. 1–6 (cit. on p. 46).

[91] A. Linhares and A.E.T.A. Freitas. “Questioning Chase and Simon’s (1973) "Perception in Chess": The "Experience Recognition" Hypothesis.” In: *New Ideas in Psychology* 28 (2010), pp. 64–78 (cit. on pp. xi, 6, 7, 13, 20, 23).

[92] A. Linhares et al. “Entanglement of Perception and Reasoning in The Combinatorial Game of Chess: Differential Errors of Strategic Reconstruction.” In: *Cognitive Systems Research* 13.1 (2012), pp. 1–15 (cit. on p. 3).

[93] Michael L. Littman, Richard S. Sutton, and Satinder Singh. “Predictive Representations of State.” In: *In Advances In Neural Information Processing Systems 14*. MIT Press, 2001, pp. 1555–1561 (cit. on pp. 8–10, 13, 18, 23).

[94] J.B. Marshall. “Metacat: a self-watching cognitive architecture for analogy-making and high-level perception.” PhD thesis. Indiana University, Bloomington, 1999 (cit. on pp. 12, 19, 38).

[95] J. McCarthy and P. Hayes. “Some Philosophical Problems From the Standpoint of Artificial Intelligence.” In: *Machine Intelligence 4*. Ed. by B. Meltzer and D. Michie. Edinburgh University Press, 1969, pp. 463–502 (cit. on p. 8).

[96] G. Miller. “The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information.” In: *Psychological Review* 63.2 (1956), pp. 81–97 (cit. on pp. 11, 34, 46).

[97] R. Milo et al. “Network Motifs: Simple Building Blocks of Complex Networks.” In: *Science* 298.5594 (2002), pp. 824–827 (cit. on pp. 50, 51, 60, 62).

[98] Ron Milo et al. “Superfamilies of Evolved and Designed Networks.” In: *Science* 303.5663 (2004), pp. 1538–1542 (cit. on p. 46).

[99] M. Mitchell. *Analogy-Making as Perception*. Cambridge, MA: MIT Press, 1993 (cit. on pp. 12, 17, 35, 38, 64).

[100] M. Natter et al. “The Effect of Incentive Schemes and Organizational Arrangements on the New Product Development Process.” In: *Management Science* 47.8 (2001), pp. 1029–1045 (cit. on p. 4).

[101] R. Neelamegham and P. Chintagunta. “A Bayesian Model to Forecast New Product Performance in Domestic and International Markets.” In: *Marketing Science* 18.2 (1999), pp. 115–136 (cit. on p. 5).

[102] A. Newell, J.C. Shaw, and H.A. Simon. “Report on a general problem-solving program.” In: *Proceedings of the International Conference on Information Processing*. 1959, pp. 256–264 (cit. on p. 18).

[103] J. Novet. *Is Google Cornering the Market on Deep Learning?* Venture Beat—`http://venturebeat.com/2015/04/13/ibm-acquires-health-tech-startups-explorys-and-phytel/`. 2015 (last accessed Aug. 2016) (cit. on p. 5).

[104] J. O et al. "Adaptive stock trading with dynamic asset allocation using reinforcement learning." In: *Information Sciences* 176.15 (2006), pp. 2121–2147 (cit. on p. 5).

[105] D.E. Padilla and M.D. McDonnel. "A Neurobiologically Plausible Vector Symbolic Architecture." In: *IEEE International Conference on Semantic Computing (ICSC)*. Newport Beach, CA: IEEE, 2014, pp. 242–245 (cit. on p. 35).

[106] D. Papo et al. "Complex network theory and the brain." In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1653 (2014) (cit. on p. 49).

[107] H. Park and K. Friston. "Structural and Functional Brain Networks: From Connections to Cognition." In: *Science* 342.6158 (2013) (cit. on pp. 49, 52–54).

[108] V. Pernice et al. "The relevance of network micro-structure for neural dynamics." In: *Frontiers in Computational Neuroscience* 7 (2013), p. 72 (cit. on p. 51).

[109] J. Piaget. "Play, dreams and imitation in childhood." In: ed. by C. Gattegno and F.M. Hodgson. W.W. Norton, 1962 (original work published 1945) (cit. on p. 19).

[110] J. Piaget. "The origin of intelligence in the child." In: ed. by M. Cook. New York: Penguin, 1977 (original work published 1936) (cit. on p. 19).

[111] P. Pierson. *Politics in Time*. Princeton, NJ: Princeton University Press, 2004 (cit. on p. 10).

[112] S. Piramuthu, H. Ragavan, and M.J. Shaw. "Using Feature Construction to Improve the Performance of Neural Networks." In: *Management Science* 44.3 (1998), pp. 416–430 (cit. on p. 4).

[113] T.A. Plate. *Holographic Reduced Representation: Distributed Representation for Cognitive Structures*. Stanford, CA: CSLI Publications, 2003 (cit. on pp. 34, 35).

[114] M.L. Platt. "Neural correlates of decisions." In: *Current Opinion in Neurobiology* 12.2 (2002), pp. 141–148 (cit. on p. 50).

[115] R. Quillian. "A Revised Design for an Understanding Machine." In: *Mechanical Translation* 7.1 (1962), pp. 17–29 (cit. on p. 38).

[116] R. Quillian. "Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities." In: *Behavioral Science* 12.1 (1967), pp. 410–430 (cit. on p. 38).

[117] U. Ramamurthy, S.K. D'Mello, and S. Franklin. "Realizing Forgetting in a Modified Sparse Distributed Memory System." In: *The 28th Annual Conference of the Cognitive Science Society*. Vancouver, Canada, 2006, pp. 1992–1997 (cit. on p. 19).

[118] G. Recchia et al. "Encoding Sequential Information in Vector Space Models of Semantics: Comparing Holographic Reduced Representation and Random Permutation." In: *Computational Intelligence and Neuroscience* 2015 (2015), pp. 865–870 (cit. on p. 35).

[119] A. Regalado. *Is Google Cornering the Market on Deep Learning?* Technology Review—`http://www.technologyreview.com/news/524026/is-google-cornering-the-market-on-deep-learning/`. 2014 (last accessed Aug. 2016) (cit. on p. 5).

[120] J. Rehling. "Letter Spirit (part two): Modeling Creativity in a Visual Domain." PhD thesis. Indiana University, Bloomington, 2000 (cit. on p. 11).

[121] A. Ries and J. Trout. *The 22 Immutable Laws of Marketing*. New York: HarperCollins, 1993 (cit. on p. 4).

[122] G.A. Ruz, S. Varas, and M. Villena. "Policy making for broadband adoption and usage in Chile through machine learning." In: *Expert Systems with Applications* 40.17 (2013), pp. 6728–6734 (cit. on p. 5).

[123] P. Saariluoma and V. Kalakoski. "Apperception and imagery in blindfold chess." In: *Memory* 6.1 (1998), pp. 67–90 (cit. on p. 23).

[124] M. Sahlgren, A. Holst, and P. Kanerva. "Permutations as a Means to Encode Order in Word Space." In: *Proceedings of the Thirtieth Annual Meeting of the Cognitive Science Society (CogSci'08)*. Cognitive Science Society. Washington, D.C., 2008 (cit. on p. 35).

[125] Shai S. Shen-Orr et al. "Network motifs in the transcriptional regulation network of Escherichia coli." In: *Nature Genetics* 31.1 (2002), pp. 64–68 (cit. on p. 51).

[126] K. Shen et al. "Network Structure Shapes Spontaneous Functional Connectivity Dynamics." In: *The Journal of Neuroscience* 35.14 (2015), pp. 5579–5588 (cit. on p. 54).

[127] E. Shiraev and D. Levy. *Cross-Cultural Psychology*. 3rd ed. New York, NY: Pearson Education, 2007 (cit. on p. 16).

[128] J. Silkes and M. Rogers. "Masked Priming Effects in Aphasia: Evidence of Altered Automatic Spreading Activation." In: *Journal of Speech Language and Hearing Research* 55.6 (2012), pp. 1613–1625 (cit. on p. 38).

[129] T. Simonite. *Facebook Launches Advanced AI Effort to Find Meaning in Your Posts*. Tech. Review–http://www.technologyreview.com/news/519411/facebook-launches-advanced-ai-effort-to-find-meaning-in-your-posts/. 2013 (last accessed Aug. 2016) (cit. on p. 5).

[130] J.E. Smith and D. von Winterfeldt. “Decision Analysis in Management Science.” In: *Management Science* 50 (2004), pp. 561–574 (cit. on p. 3).

[131] P.K. Smith and J.A. Bargh. “Nonconscious Effects of Power on Basic Approach and Avoidance Tendencies.” In: *Social Cognition* 26.1 (2008), pp. 1–24 (cit. on p. 10).

[132] P.L. Smith and R. Ratcliff. “Psychology and neurobiology of simple decisions.” In: *Trends in Neurosciences* 27.3 (2004), pp. 161 –168 (cit. on p. 50).

[133] P. Smolensky. “Tensor product variable binding and the representation of symbolic structures in connectionist systems.” In: *Artificial Intelligence* 46.1-2 (1990), pp. 159–216 (cit. on p. 34).

[134] P. Smolensky, G. Legendre, and Y. Miyata. *Principles for an Integrated Connectionist/Symbolic Theory of Higher Cognition*. Tech. rep. 1992 (cit. on p. 37).

[135] J. Snaider and S. Franklin. “Extended Sparse Distributed Memory and Sequence Storage.” In: *Cognitive Computation* 4.2 (2012), pp. 172–180 (cit. on p. 33).

[136] B. Solomon. *Leaked: Uber's Financials Show Huge Growth, Even Bigger Losses*. Forbes–http://www.forbes.com/sites/briansolomon/2016/01/12/leaked-ubers-financials-show-huge-growth-even-bigger-losses/. Jan. 2016 (last accessed Aug. 2016) (cit. on p. 4).

[137] L. Sun and A. Erath. “A Bayesian network approach for population synthesis.” In: *Transportation Research Part C: Emerging Technologies* 61 (2015), pp. 49–62 (cit. on p. 5).

[138] N.N. Taleb. *Antifragile: Things that gain from disorder*. New York: Random House, 2012 (cit. on p. 7).

[139] M. Tomasello et al. “Understanding and sharing intentions: The origins of cultural cognition.” In: *Behavioral And Brain Sciences* 28 (2005), pp. 675–735 (cit. on p. 10).

[140] A. Treisman. “The Binding Problem.” In: *Current Opinions in Neurobiology* 6.2 (1996), pp. 171–178 (cit. on p. 34).

[141] A. Tversky and D. Kahneman. “Judgment under Uncertainty: Heuristics and Biases.” In: *Science* 185.4157 (1974), pp. 1124–1131 (cit. on pp. 3, 10).

[142] A. Tversky and D. Kahneman. "The Framing of Decisions and the Psychology of Choice." In: *Science* 211.4481 (1981), pp. 453–458 (cit. on p. 10).

[143] U. U. Ramamurthy, S.K. D'Mello, and S. Franklin. "Modified Sparse Distributed Memory as Transient Episodic Memory for Cognitive Software Agents." In: *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*. The Hague, 2004 (cit. on p. 19).

[144] S. Usborne. *An Uber for everything: the surge in on-demand apps*. The Guardian–https://www.theguardian.com/technology/2016-/may/03/an-uber-for-everything-the-surge-in-on-demand-apps. May 2016 (last accessed Aug. 2016) (cit. on p. 4).

[145] P.E. Vértes, A. Alexander-Bloch, and E.T. Bullmore. "Generative models of rich clubs in Hebbian neuronal networks and large-scale human brain networks." In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 369.1653 (2014) (cit. on p. 46).

[146] P. C. Wason and D. Shapiro. "Natural and contrived experience in a reasoning problem." In: *Quarterly Journal of Experimental Psychology* 23 (1971), pp. 63–71 (cit. on pp. 17, 18).

[147] Duncan J. Watts and Steven H. Strogatz. "Collective dynamics of 'small-world' networks." In: *Nature* 393.6684 (1998), pp. 440–442 (cit. on pp. 32, 46).

[148] M. Weinberger. *The secret to Uberś success: Be everywhere all the time*. Business Insider–http://www.businessinsider.com/uber--apps-are-everywhere-2015-12. Dec. 2015 (last accessed Aug. 2016) (cit. on p. 4).

[149] L. E. Williams and J. A. Bargh. "Experiencing Physical Warmth Promotes Interpersonal Warmth." In: *Science* 322 (2008), pp. 606–607 (cit. on pp. 10, 15).

[150] W. Woertman et al. "Evidence based policy decisions through a Bayesian approach: The case of a statin appraisal in the Netherlands." In: *Health Policy* 112.3 (2013), pp. 234–240 (cit. on p. 5).

[151] M. Wolverton and B. Hayes-Roth. "Retrieving Semantically Distant Analogies with Knowledge-Directed Spreading Activation." In: *Proceedings pf the National Conference on Artificial Intelligence*. Ed. by AAAI. Seattle, 1994, pp. 56–61 (cit. on p. 38).

DECLARATION

This thesis is a presentation of my original research work. Wherever contributions of others are involved, every effort is made to indicate this clearly, with due reference to the literature, and acknowledgement of collaborative research and discussions.

The work was done under the guidance of Professor Alexandre Linhares, at the Escola Brasileira de Administração Pública e de Empresas — EBAPE , Fundação Getulio Vargas — FGV.

Rio de Janeiro, October 2016

Daniel de Magalhães Chada