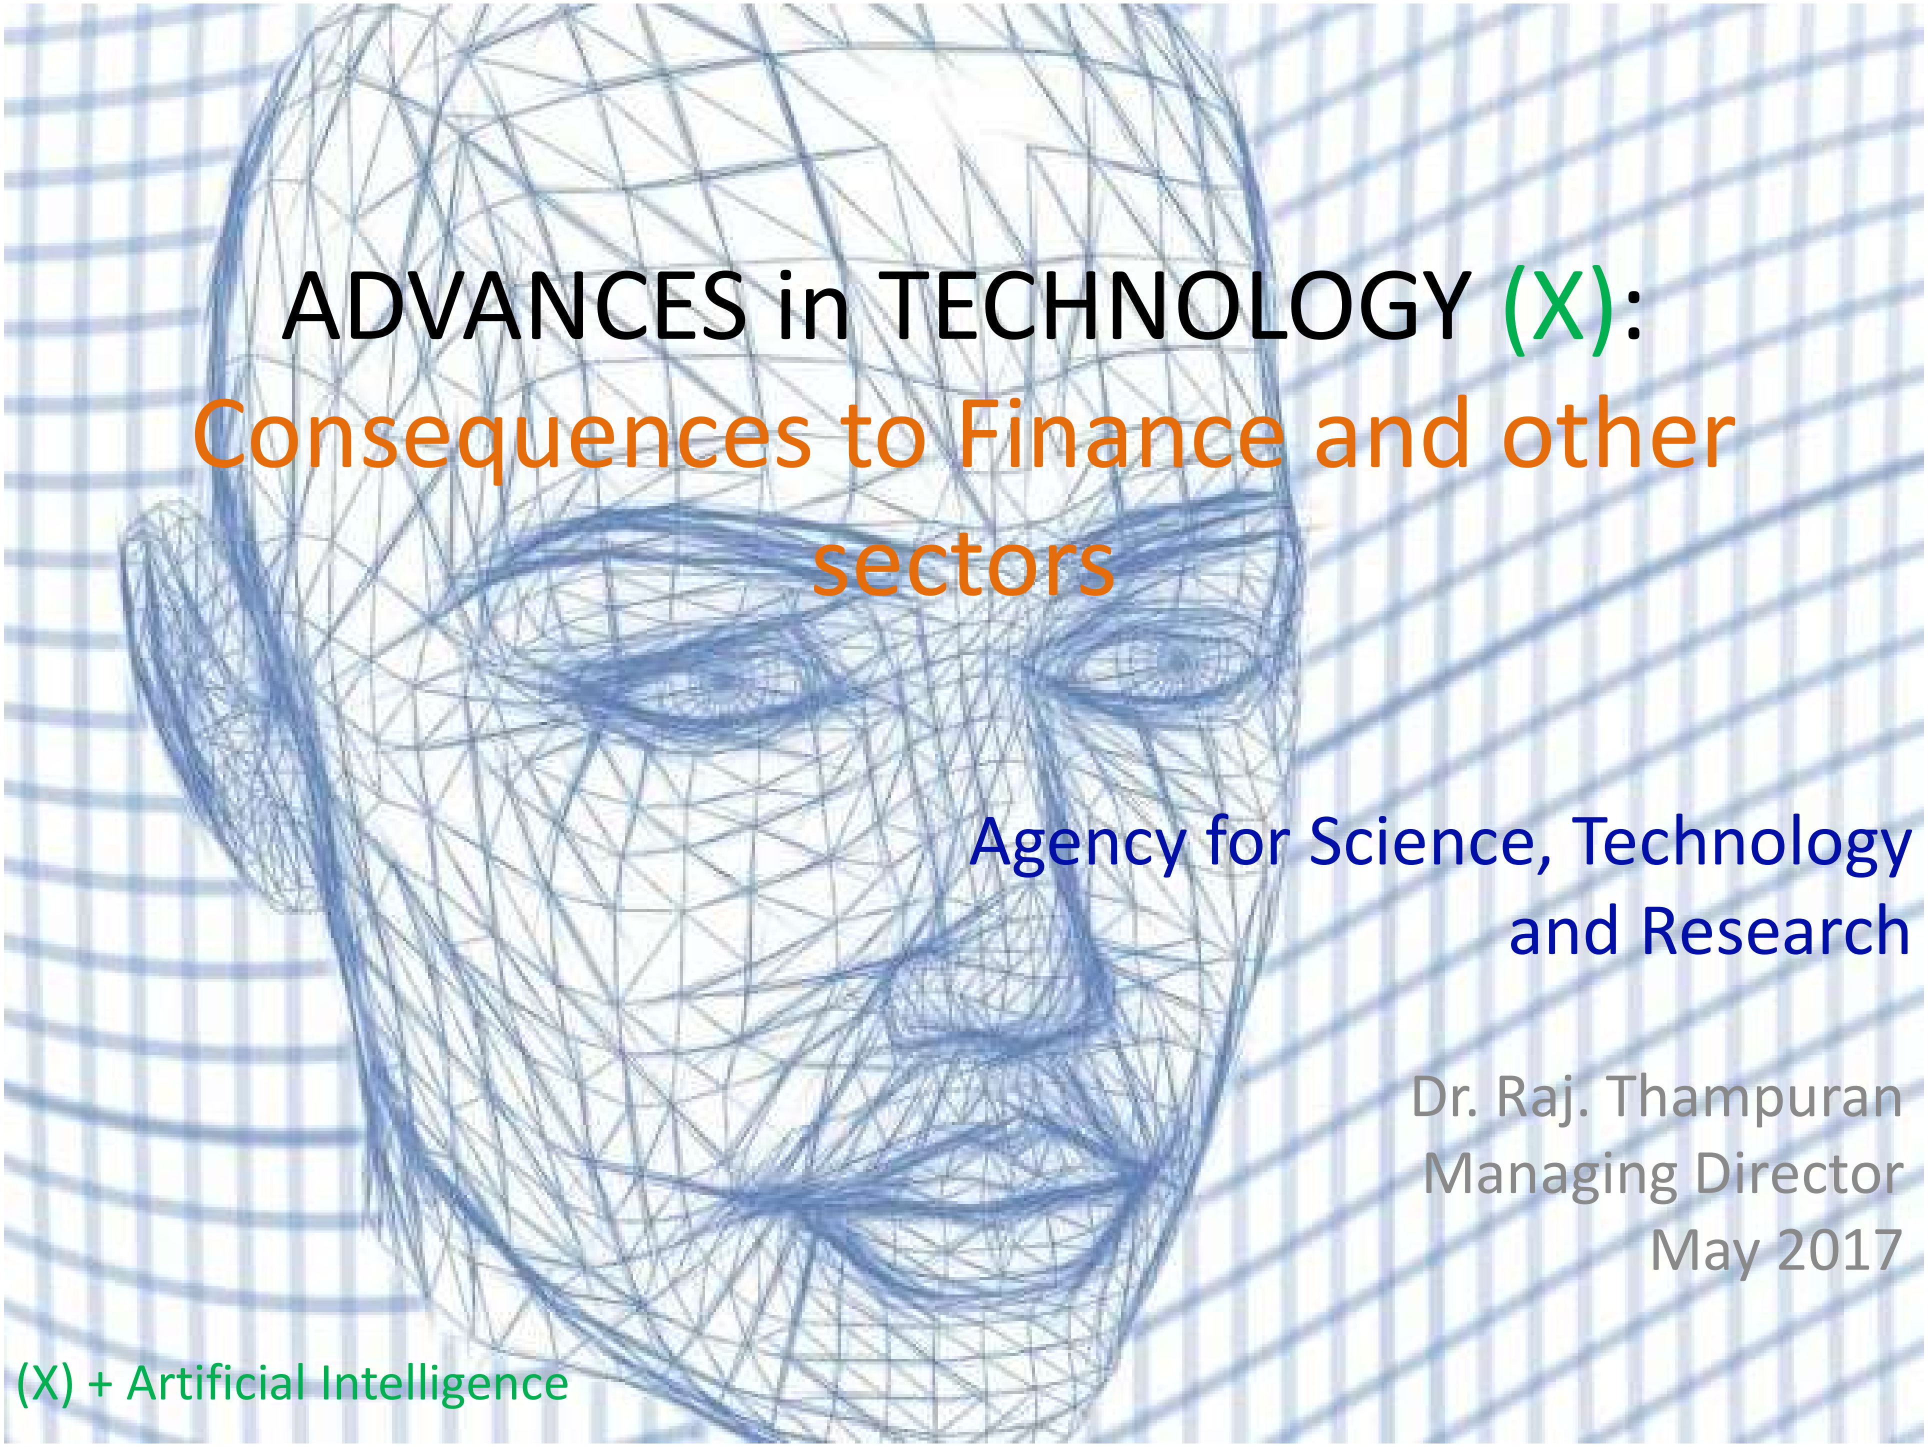

ADVANCES in TECHNOLOGY (X): Consequences to Finance and other sectors

Agency for Science, Technology and Research

Dr. Raj. Thampuran
Managing Director

May 2017

(X) + Artificial Intelligence

TRENDS

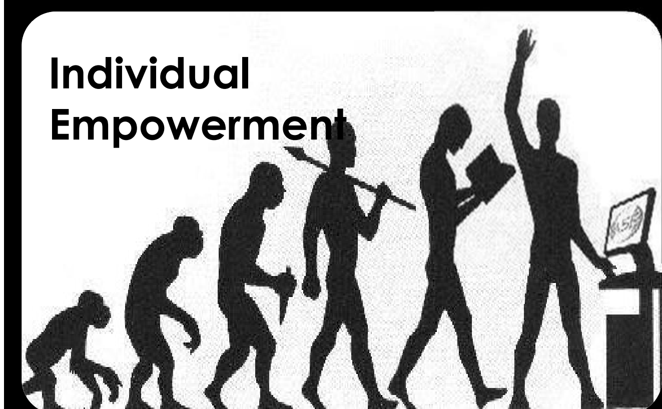

airbnb UBER

TECHNOLOGIES

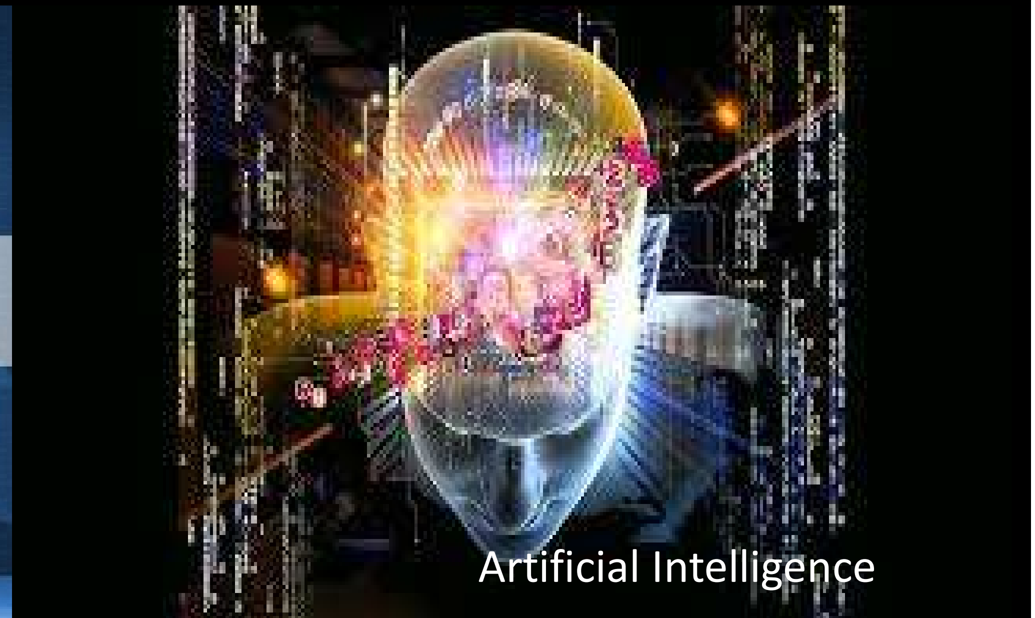
Artificial Intelligence


Quantum Cryptography

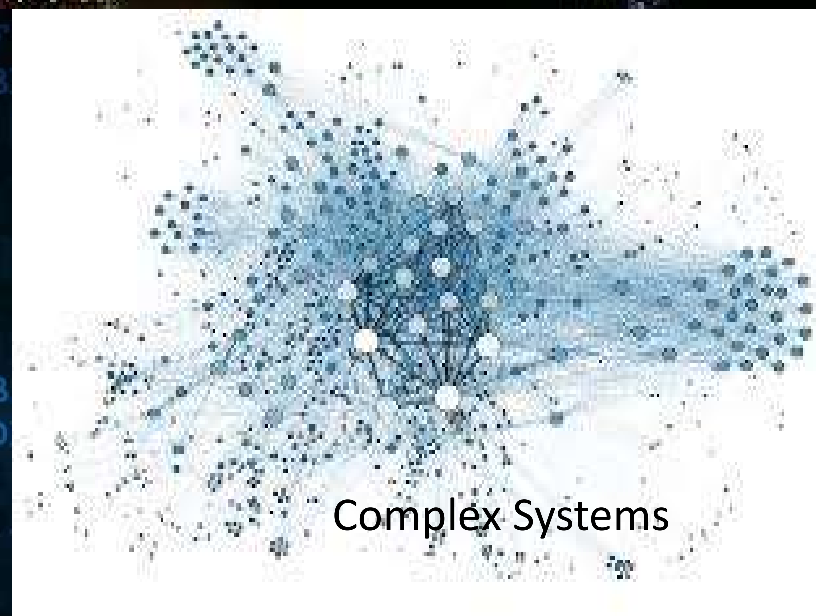
Complex Systems

DISRUPTORS

Exhibit E3

Data and analytics underpin six disruptive models, and certain characteristics make individual domains susceptible

Indicators of potential for disruption:

- Assets are underutilized due to inefficient signaling
- Supply/demand mismatch
- Dependence on large amounts of personalized data
- Data is siloed or fragmented
- Large value in combining data from multiple sources
- R&D is core to the business model
- Decision making is subject to human biases
- Speed of decision making limited by human constraints
- Large value associated with improving accuracy of prediction

Archetype of disruption	Domains that could be disrupted
Business models enabled by orthogonal data	▪ Insurance ▪ Health care ▪ Human capital/talent
Hyperscale, real-time matching	▪ Transportation and logistics ▪ Automotive ▪ Smart cities and infrastructure
Radical personalization	▪ Health care ▪ Retail ▪ Media ▪ Education
Massive data integration capabilities	▪ Banking ▪ Insurance ▪ Public sector ▪ Human capital/talent
Data-driven discovery	▪ Life sciences and pharmaceuticals ▪ Material sciences ▪ Technology
Enhanced decision making	▪ Smart cities ▪ Health care ▪ Insurance ▪ Human capital/talent

SOURCE: McKinsey Global Institute analysis

MASSIVE INTEGRATION

Retail banks have opportunity to break their data silos, combining traditional and new data sources in data lakes

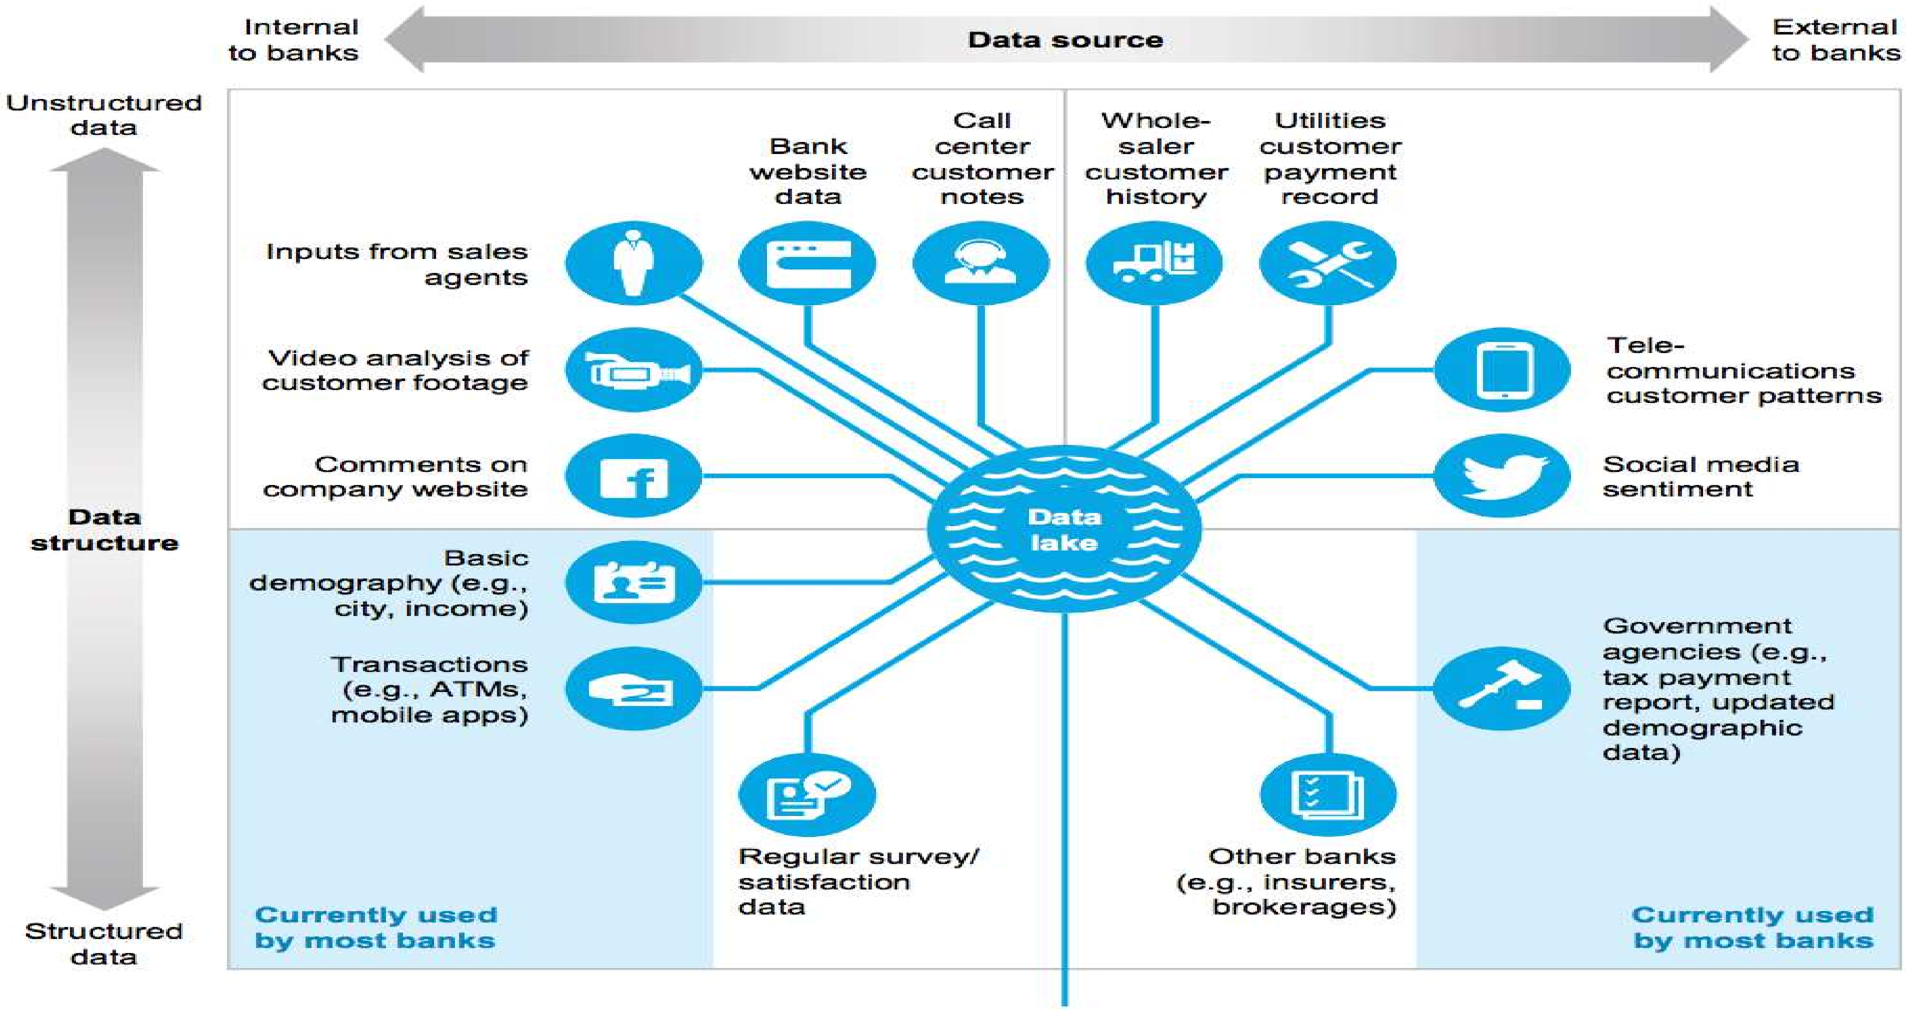

- Stores practically unlimited amounts of data of any format and type
- Silos minimized, and single source of truth accessible by whole organization
- Offers an improved platform to run analytics and data discovery
- Transformation to the data lakes environment can be done gradually

SOURCE: McKinsey Global Institute analysis

Machine learning opportunities in finance

Highest-ranked use cases, based on survey responses	Use case type	Impact	Data richness
Personalize product offerings to target individual consumers based on multi-modal data (mobile, social media, location, etc.)	Radical personalization	1.2	1.7
Identify fraudulent activity using customer transactions and other relevant data	Discover new trends/ anomalies	1.0	1.3
Evaluate customer credit risk using application and other relevant data for less biased real-time underwriting decisions	Predictive analytics	0.9	1.0
Predict risk of churn for individual customers/clients and recommend renegotiation strategy	Predictive maintenance	0.7	0.7
Discover new complex interactions in the financial system to support better risk modeling and stress testing	Discover new trends/ anomalies	0.7	0.7
Predict risk of loan delinquency and recommend proactive maintenance strategies	Predictive analytics	0.5	1.0
Predict asset price movements based on greater quantities of data (e.g., social media, video feeds) to inform trading strategies	Forecasting	0.4	1.3
Optimize labor staffing and distribution to reduce operational costs in front and back office	Resource allocation	0.4	0.7
Route call-center cases based on multi-modal data (e.g., customer preferences, audio data) to increase customer satisfaction and reduce handling costs	Predictive analytics	0.1	1.7
Optimize branch/ATM network based on diverse signals of demand (e.g., social data, transactions)	Resource allocation	0.1	0.3

Capabilities in Informatics, Analytics & Computer Science

The Turing Test

1950: Alan Turing's "Computing Machinery and Intelligence"
(the "Turing Test")

A. M. Turing (1950) Computing Machinery and Intelligence. *Mind 49:* 433-460.

Can machines think?

COMPUTING MACHINERY AND INTELLIGENCE

By A. M. Turing

1. The Imitation Game

I propose to consider the question, "Can machines think?" This should begin with definitions of the meaning of the terms "machine" and "think." The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous, If the meaning of the words "machine" and "think" are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, "Can machines think?" is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words.

64

Chart 2.20
Average Annual Profit Growth of Selected Global Industries (2006–14)

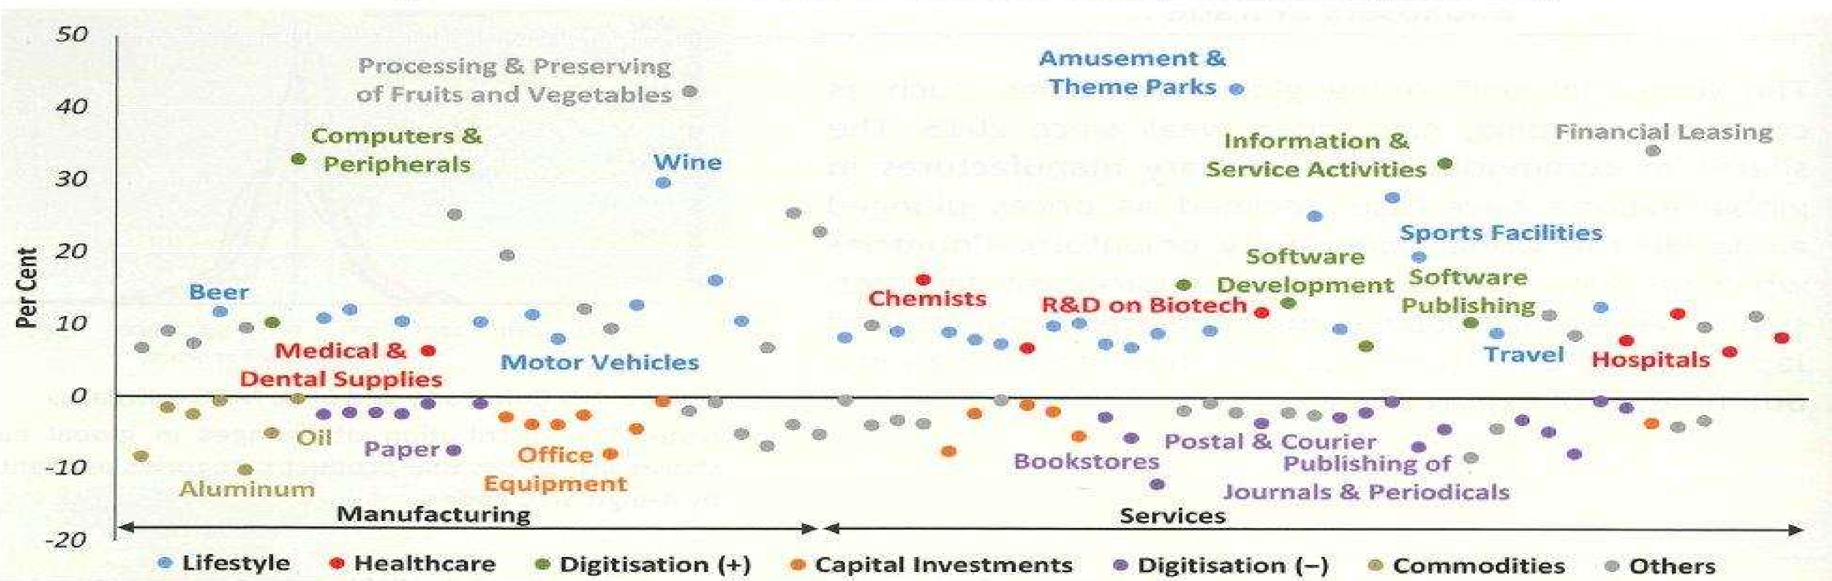

Source: Bureau Van Dijk and EPG, MAS estimates

Monetary Authority of Singapore

Economic Policy Group

Capabilities in Informatics, Analytics & Computer Science

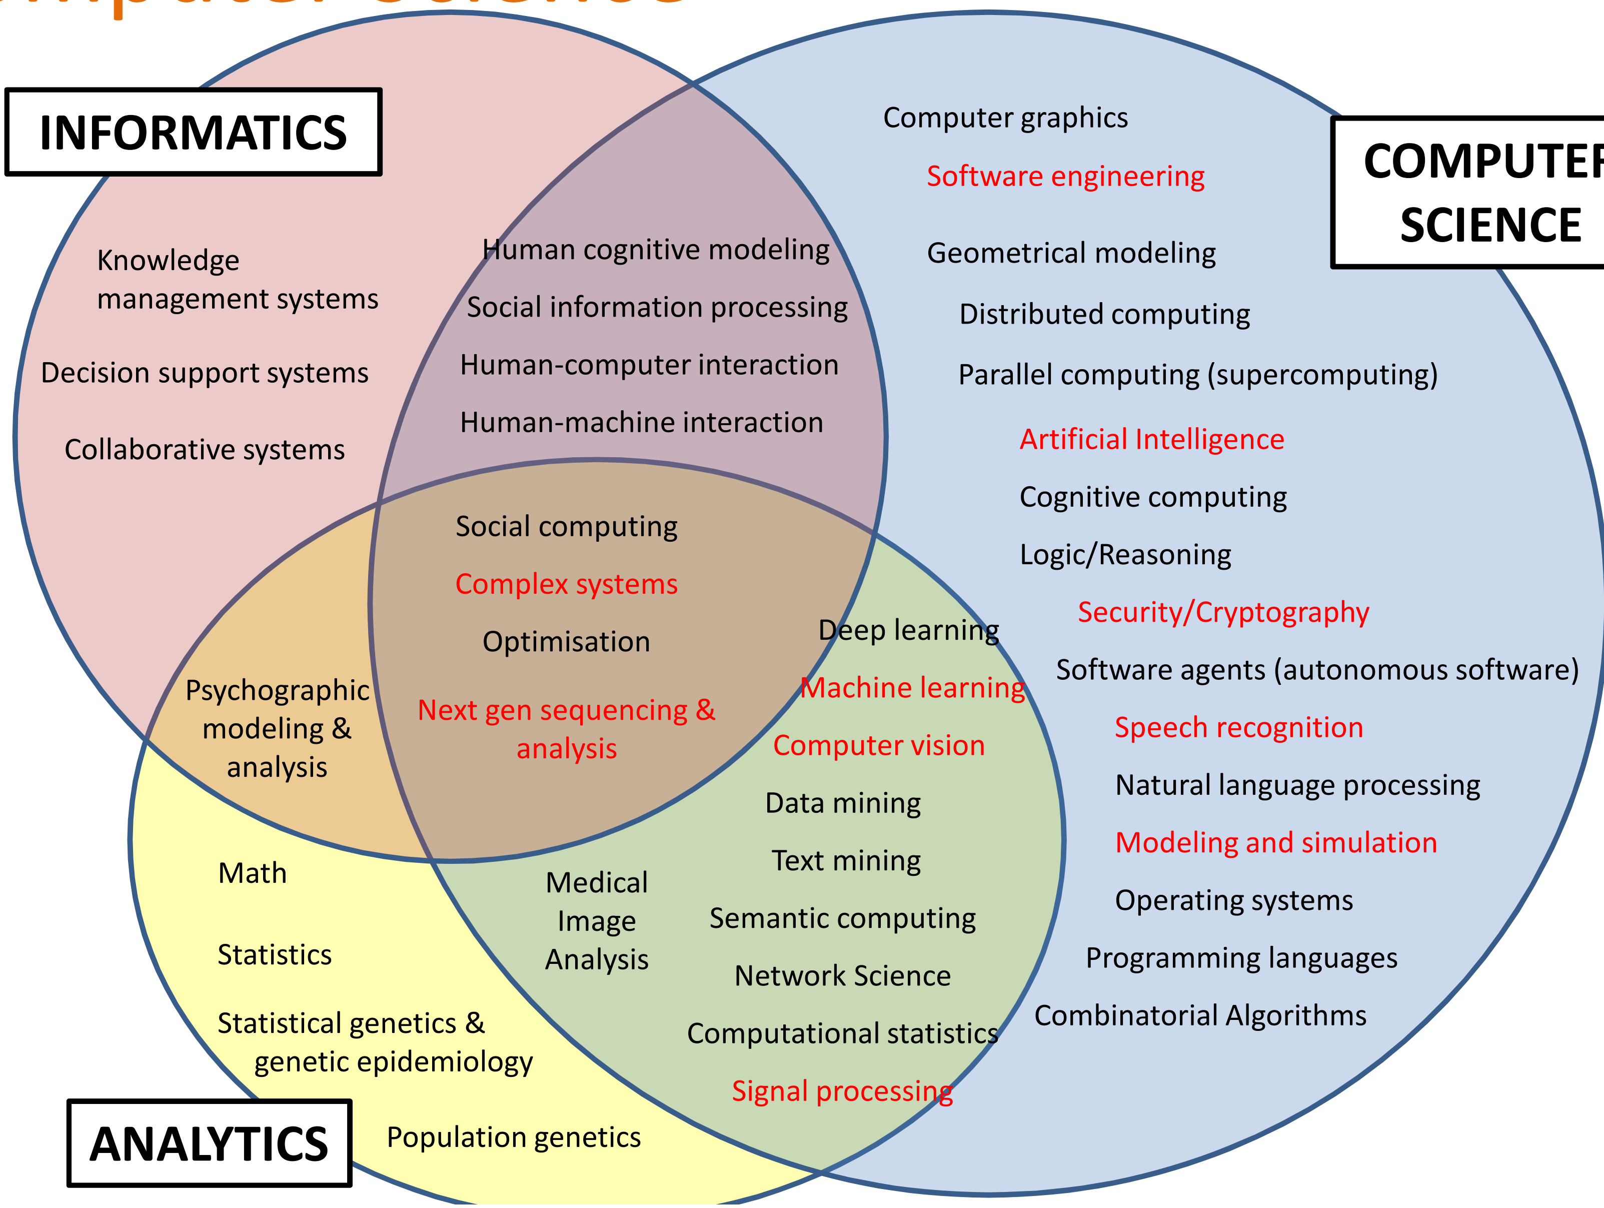

Data Analytics @ the Core

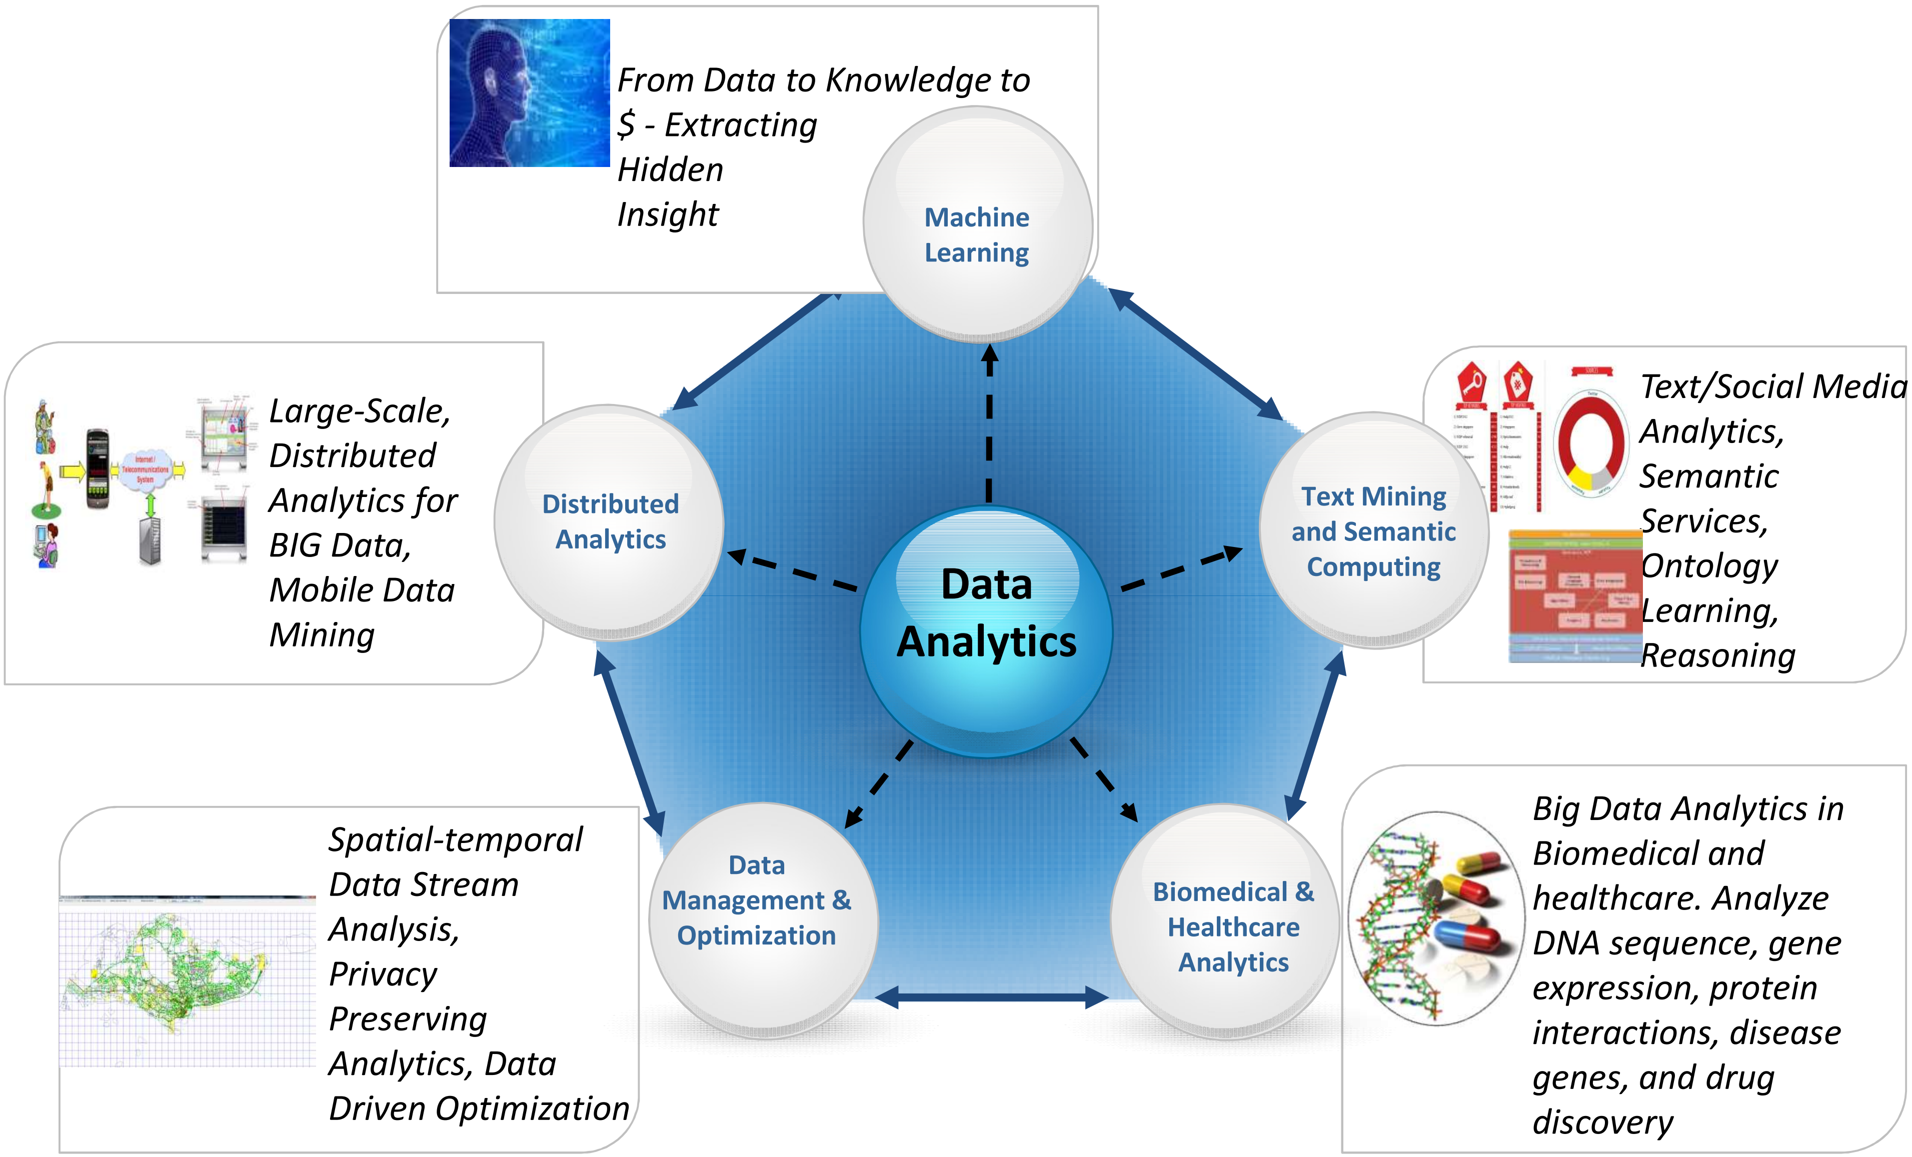

Easy to Use AI: Well within one's reach (not beyond one's grasp)

The Five Tribes of Machine Learning

Tribe	Origins	Master Algorithm
Symbolists	Logic, philosophy	Inverse deduction
Connectionists	Neuroscience	Backpropagation
Evolutionaries	Evolutionary biology	Genetic programming
Bayesians	Statistics	Probabilistic inference
Analogizers	Psychology	Kernel machines

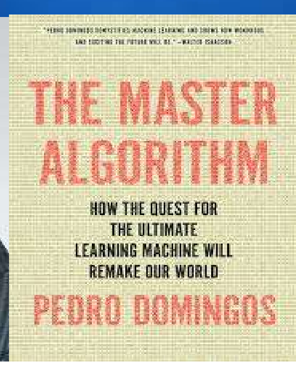

ARTIFICIAL INTELLIGENCE (Customer Intelligence) : some examples incorporating Social Intelligence and Psychological Knowledge (for Next-Generation Intelligent Systems)

- AI that negotiates win-win business contracts with humans **("Psychologically Intelligent Negotiation Agent")**
- AI that performs human-level text sentiment classification **("Fine-Grained Sentiment Analyzer – SentiMo")**
- AI that recognizes the personality and characteristics of humans **("Psychographic Profiling Engine")**

Psychologically Intelligent Negotiation Agent

***Connectionists Systems:* AI that negotiates win-win business contracts with humans**

Incorporated Negotiation Psychology

Deployed system

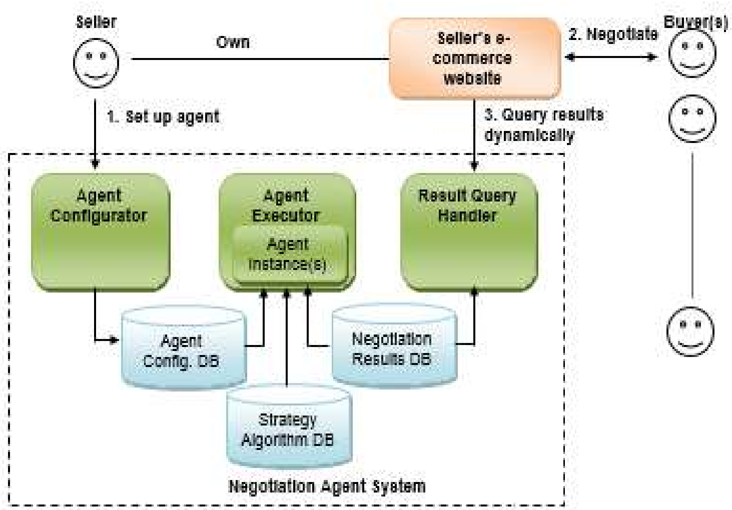

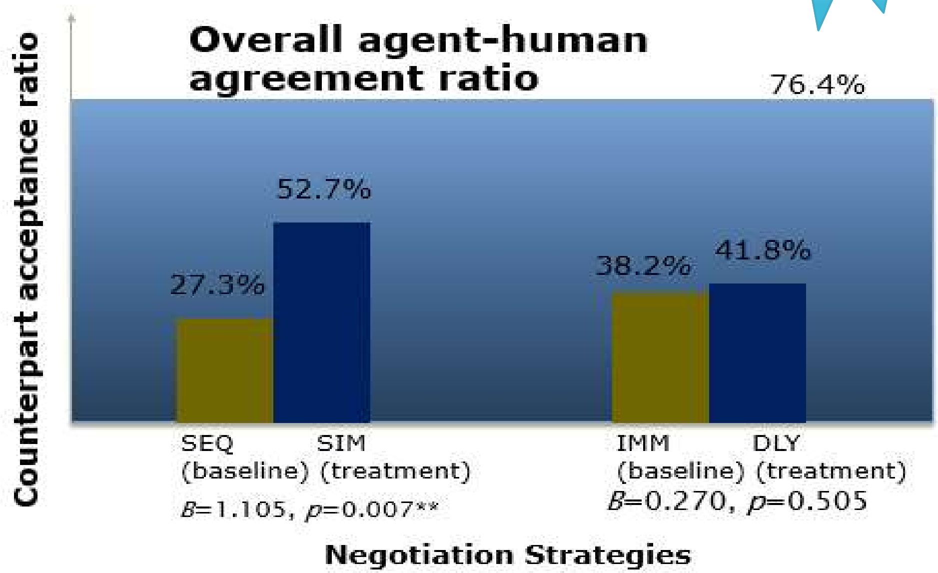

- Computational implementation of **psychologically grounded negotiation strategies (e.g., strategic delay, tactful information sharing)** which work without a priori data/knowledge about the counterpart
- Results showing **psychologically intelligent negotiation agent** can lead to better economic and social-psychological negotiation outcomes
- Deployed as a service in **HP's Mercado Cloud Computing Platform**

"System and method for negotiating a sale", US Patent Application 12/648,405 & "Alternate strategies for a win-win seeking agent in agent-human negotiations", JMIS & "Reducing Mistrust in Agent-Human Negotiations", IEEE Intelligent Systems

Fine-Grained Sentiment Analyzer (SentiMo)

*Connectionists Systems :**AI that performs human-level text sentiment classification***

Design Features and Novelty

Patent filed

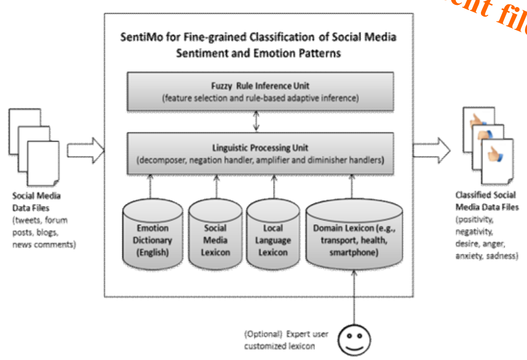

http://172.20.98.207:8080/sentimo-webportal/sentimo_api.html

- **Fine-grained multi-dimensional outputs** (positive, negative, neutral, mixed, sadness, anger, happiness, excitement…)
- **Comprehensive lexicons,** fully in-house developed (English, Internet slangs, local language and domain words collections)
- **Linguistic processing units** (decomposer, negation handler, amplifier handler…)

Performance

API & SDK released

Average of F1 - Score for Positivity, Negativity, Neutrality Recognition

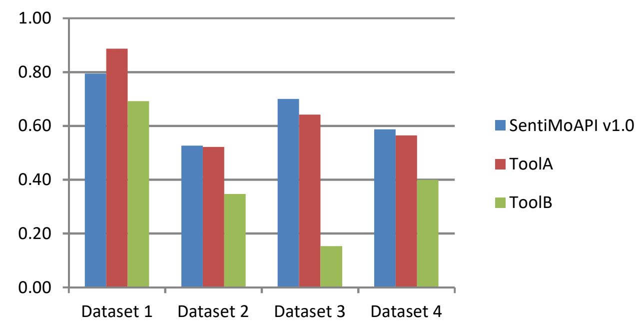

Our Real-World Sentiment Analysis Case Studies

Ground sensing of day-to-day commuter sentiments

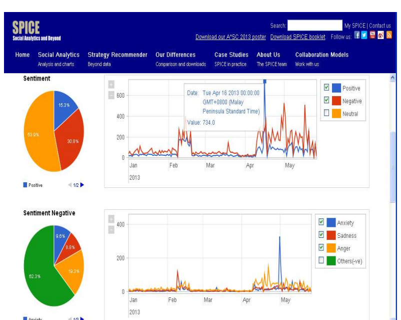

Discovering consumer preferences across products

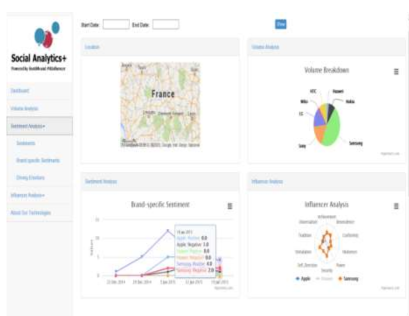

Understanding brand perceptions across cities

Quantifying positivity generated from public campaigns

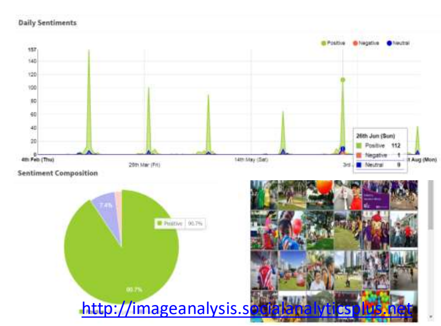

“A method and system for sentiment classification and emotion classification”, Patent Cooperation Treaty (PCT) Application PCT/SG2015/050469

Psychographic Profiling Engine

*Connectionists + Analogizer Systems: **AI that recognizes the personality and characteristics of humans***

Design Features and Novelty

Raw Individual Level Data (demographics, transactions, browser cookies, **questionnaire response**, **social network activities-text, images, likes**)

Psychographic Analysis Outputs (big-five personalities, personal values, **creditworthiness & sub-dimensions**, affective /cognitive-regulatory focus, Machiavellianism, interests, preferences)

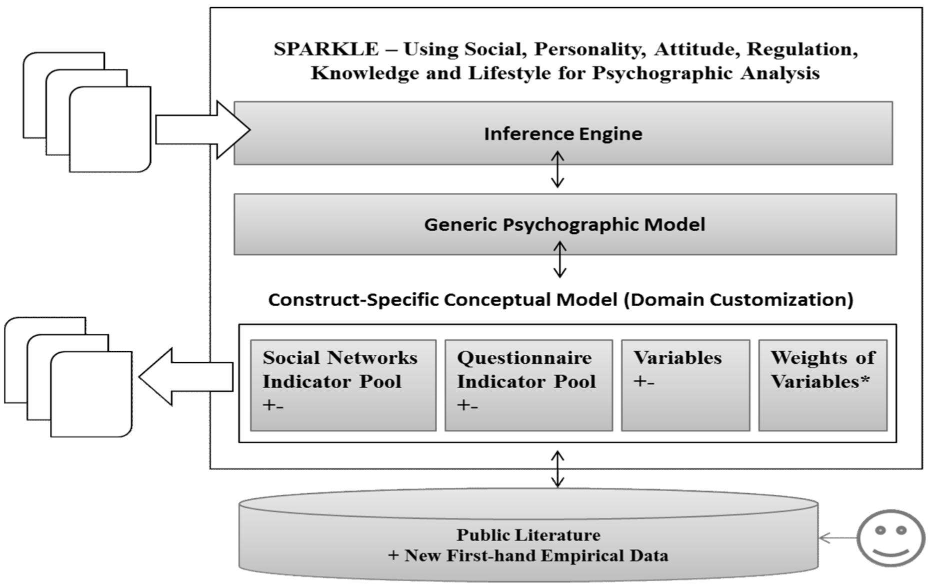

- This is an ongoing research where we train computational models to profile individuals based on **multiple sources of data** to describe their **personality traits**, **personal values**, or more context-specific characteristics such as creditworthiness and innovation propensity
- Leverage advanced data processing including **psycholinguistic/text analysis** and **image recognition**

Applications

SPARKLE Social Cognitive Creditworthiness Assessment

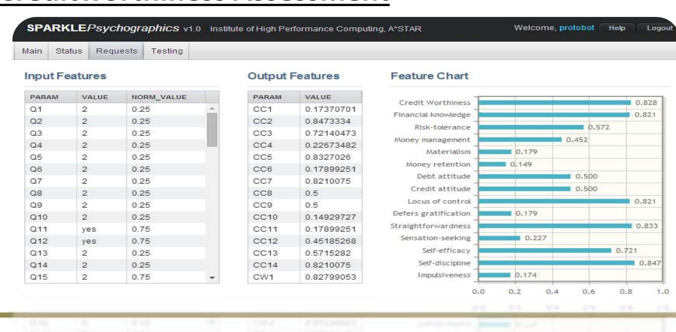

Influencer Analysis / Depicting Brand Persona

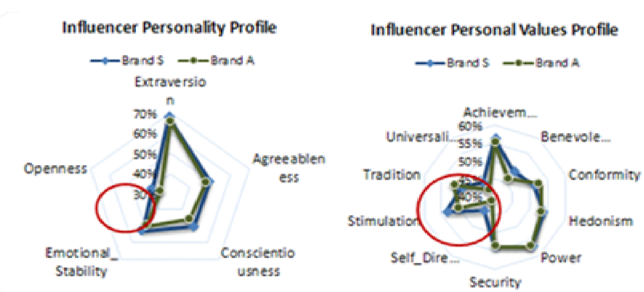

Rakuten-Viki Global TV Recommender Challenge

Connectionists Systems :Recommeder system for AI: A business opportunity

1st Prize

Rakuten viki

Motivation / Objectives

- To build a personalized TV Recommender system for world-wide Rakuten-Viki fans
- Recommend videos that a user is likely to watch (precision) and watch for long time (engagement)
- “Cold-Start” problem : 20+% users do not apprear in training data)
- Data sparsity problem : most users viewed <= 5 videos in training data

Approach

- Typical recommendation algorithms do not well here due to *sparsity* and *cold-start* problems
- Formulate as classification problem instead of a typical recommendation problem to predict the probability of a video that a user is likely to watch

Achievement / Impact / Value Capture

- 1st Prize Winner
- Overcome “cold-star” and data sparsity problems
- Robust and scalable approach for online recommendations
- Flexible to incorporate other general features

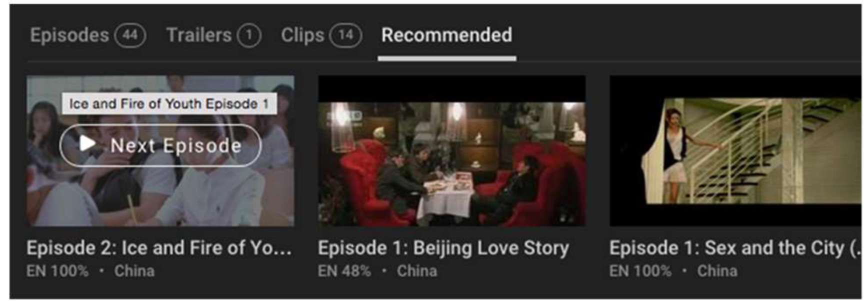

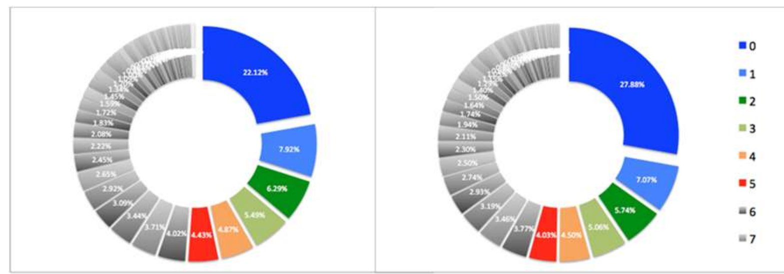

Distribution of No. of Videos viewed in the training data for Users tested in Feb 2915(Left) and Mar 2015 (Right)

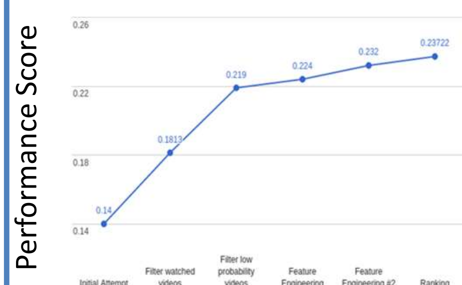

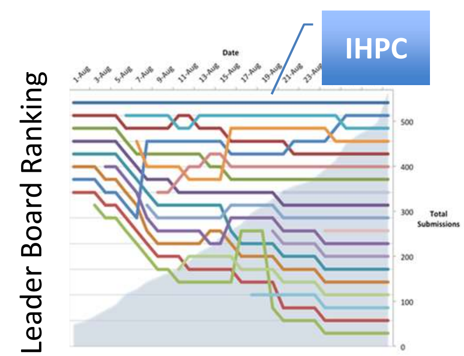

Deep Learning for E-Commerce Website Purchase Behaviour Prediction

Connectionists Systems :AI Example

- **Task** : Given a sequence of clicks,
- predict the next item that is likely to be clicked.
- Predict whether a user will buy something at the end of a clicking session.
- Predict what products a user will buy at the end of a clicking session.

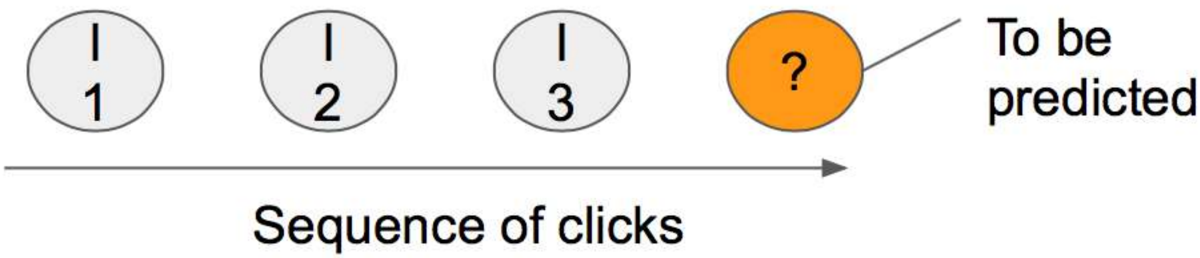

- Purchase Event?
- Which products will be purchased?

- **Approach**: Prediction using Deep Neural Network

- **Results** :
  - 70% top-20 accuracy for the next item prediction
  - ~10% improvement over reported results

Commonsense Knowledge Representation and Reasoning

*Symbolists Systems: **Human Centric AI***

Achievements:

. Codified a commonsense knowledge base (KB) using a semantic graph representation

- 3.4 million concepts involved in about 10 million relational assertions.
- From open source KBs such as ConceptNet, augmented by concepts from an 8-billion-word text corpus represented as word embedding in vector space using Word2Vec

- Applied KB in tasks such as topic categorisation, sentiment analysis and commonsense reasoning

Current Work:

- Implement structure for Commonsense KB based on noun and verb primitives to allow for inheritance of properties and attributes
- Develop representations for narrative knowledge (modelling temporally extended events)

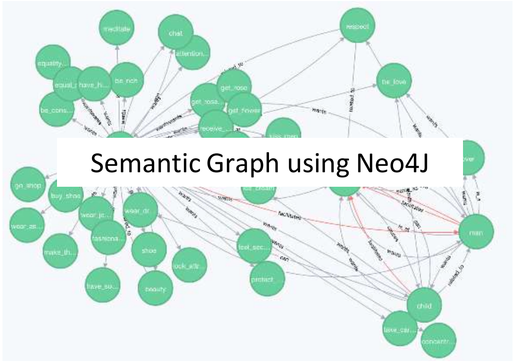

Semantic Graph using Neo4J

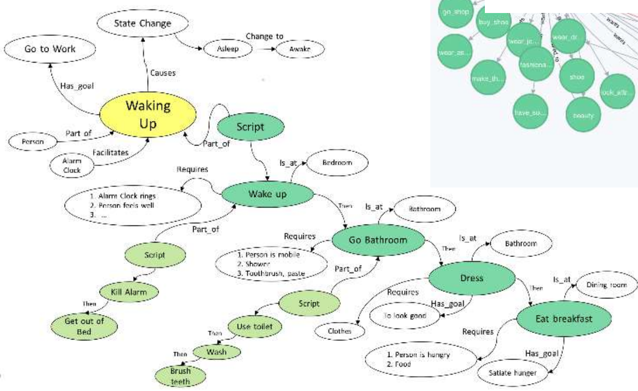

Representation for Events and Scripts

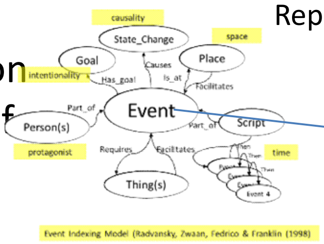

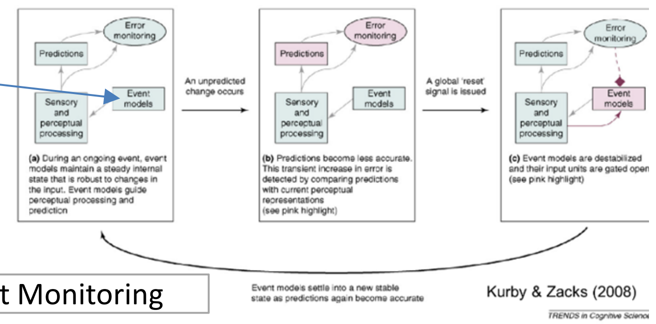

Event Models for Event Monitoring

OWL and RDF; Giant Global Graph; FOAF

Casuality & Inference: Rapid Causal Learning

Symbolists Systems: Human Centric AI

Ground Level Causal Learning (Temporal Correlation) → Knowledge Level Causal Learning

- Build **causal models** of the world that support explanation and understanding, rather than merely solving pattern recognition problems+
- **Causality from Temporal Correlation**

Strength(Cause(Event1, Event2)) =
*Prob(Cause(Event1, Event2)) – Wt * Uncert(Cause(Event1, Event2))*

- Inspired by *Contingency Model of Causal Learning* from Psychology

+ See also: *Building Machines that Learn and Think Like People* (Lake, Ullman, Tenenbaum, and Gershman, 2016) in [arXiv:1604.00289v3](https://arxiv.org/abs/1604.00289v3)

Experiment: Learning Causality from Experience

Relationship between Lightning and Thunder

Video

Lightning reliably predicts Thunder

Cause→Effect	Strength
H→H	0.847
L→T	**0.357**
H→L	0.273
L→G	0.245
G→H	0.239
G→L	0.212
B→L	0.137
H→G	0.082
H→B	0.052
B→H	0.040
G→G	-0.020

L = lightning, T = thunder, W = wind, S = vehicular sound, H = headlight

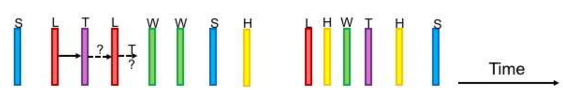

Similar Image Search using Privileged Information

Analogizer systems :
Multimodal AI

- **Motivation:**
  - How to use textual descriptions associated with the training images as privileged information for training hashing function

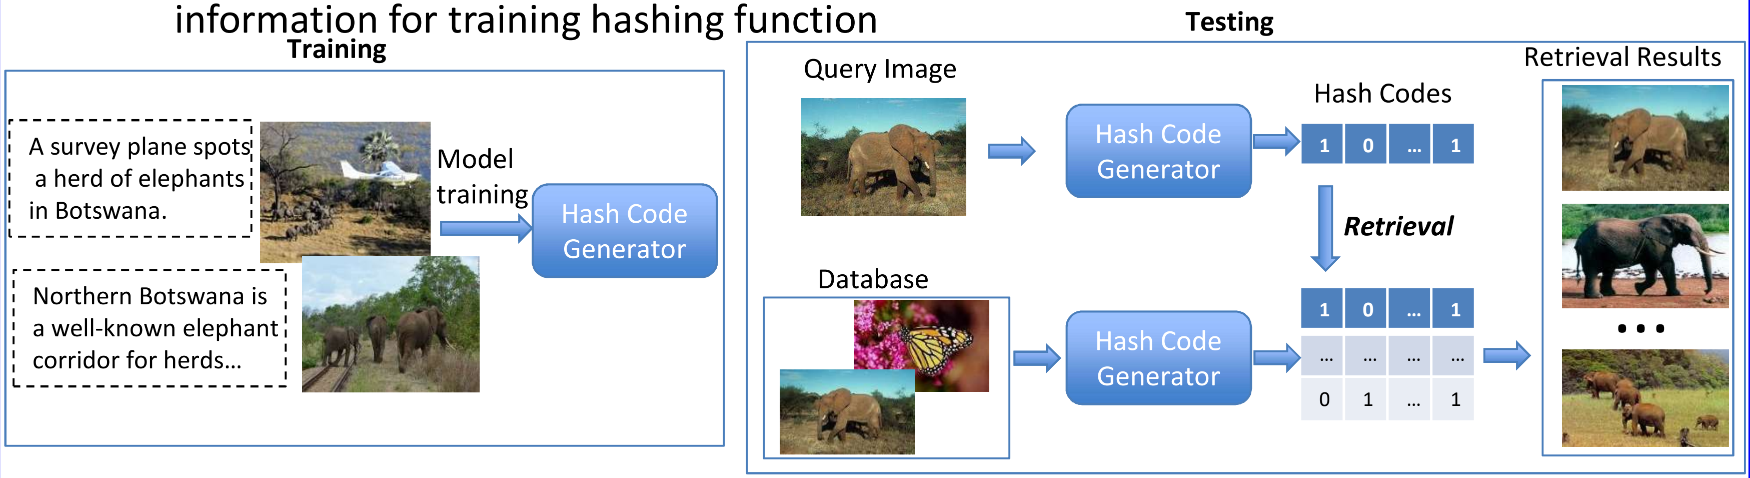

- **Our Contributions:**
  - Proposed an objective function that can utilize additional textual descriptions for hashing
  - Proposed an efficient solution to optimize the proposed objective function
  - The first work for hashing with privileged information

- **Results:**
  - Content based image retrieval using textual descriptions as privileged information
  - 45.40% in MAP on NUS-wide compared with 42.45 % for baseline without text description

Video Action and Event Recognition using Heterogeneous Sources

Analogizer systems:
Multimodal AI

❑ Motivation:

- ❑ Classify unlabeled consumer videos for action and event recognition
- ❑ A large number of freely available videos (e.g., from Flickr video search engine) and Web images (e.g., from Bing and Google image search engines) are available
- ❑ Additional textual descriptions are often available for both Web videos and images

❑ Our Contributions

- ❑ Proposed a new multi-domain adaptation approach to utilize the Web videos and images for training the classifier
- ❑ Proposed a new method to utilize the additional textual descriptions in the training set

❑ Results

- ❑ Applied the proposed methods without requiring any labeled samples from the target domain
- ❑ Applied the algorithm for the recognition of the video event (e.g, wedding, birthday, and sports) and action (e.g., eat, kiss and run)

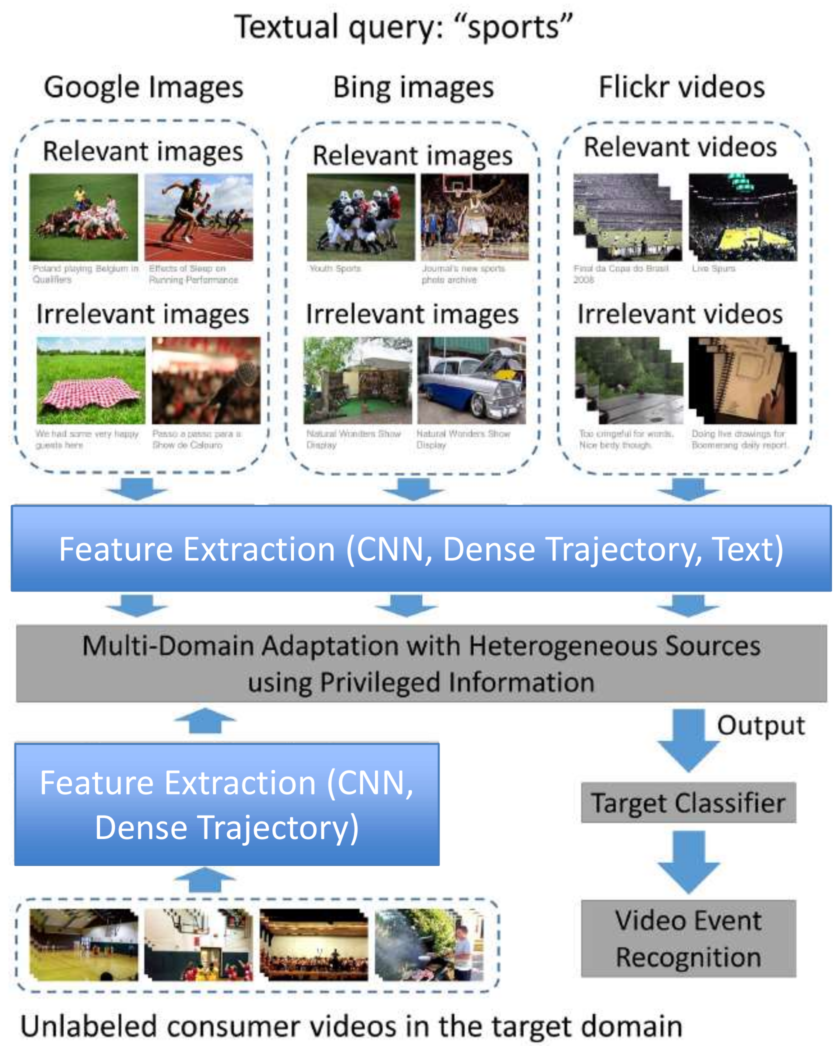

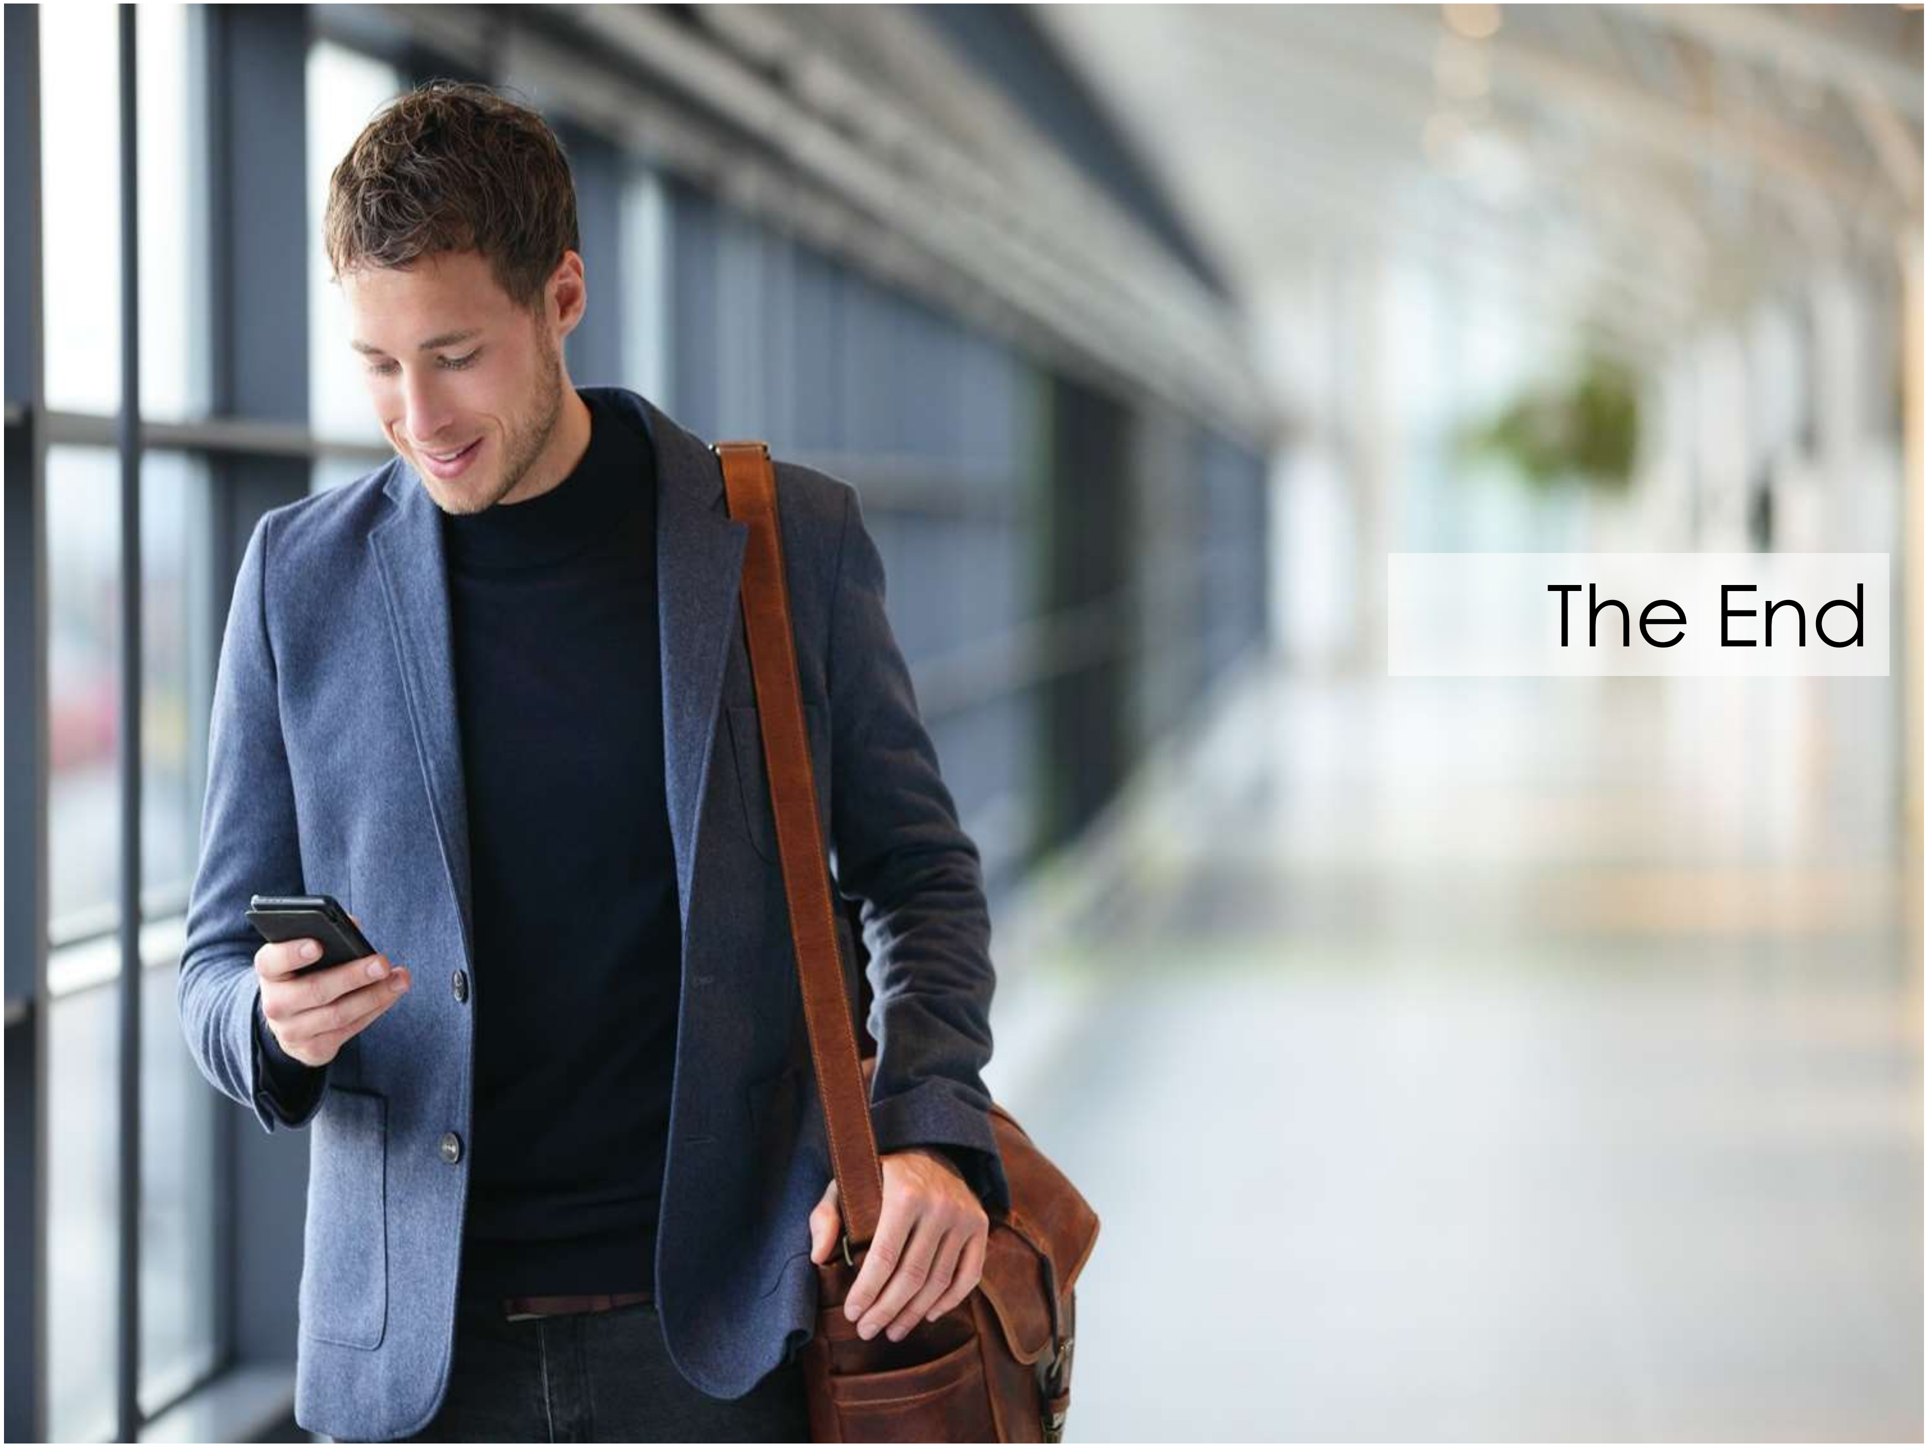
The End

AI FOR VISION

Human Centric AI

frontiers in Neuroscience

PERSPECTIVE
published: 13 October 2015
doi: 10.3389/fnins.2015.00374

Benchmarking neuromorphic vision: lessons learnt from computer vision

*Cheston Tan [1], Stephane Lallee [1] and Garrick Orchard [2,3]**

[1] Agency for Science, Technology, and Research (A*STAR), Institute for Infocomm Research, Singapore, Singapore, [2] Singapore Institute for Neurotechnology (SINAPSE), National University of Singapore, Singapore, Singapore, [3] Temasek Labs, National University of Singapore, Singapore, Singapore

Neuromorphic Vision sensors have improved greatly since the first silicon retina was presented almost three decades ago. They have recently matured to the point where they are commercially available and can be operated by laymen. However, despite improved availability of sensors, there remains a lack of good datasets, while algorithms for processing spike-based visual data are still in their infancy. On the other hand, frame-based computer vision algorithms are far more mature, thanks in part to widely accepted datasets which allow direct comparison between algorithms and encourage competition. We are presented with a unique opportunity to shape the development of Neuromorphic Vision benchmarks and challenges by leveraging what has been learnt from the use of datasets in frame-based computer vision. Taking advantage of this opportunity, in this paper we review the role that benchmarks and challenges have played in the advancement of frame-based computer vision, and suggest guidelines for the creation of Neuromorphic Vision benchmarks and challenges. We also discuss the unique challenges faced when benchmarking Neuromorphic Vision algorithms, particularly when attempting to provide direct comparison with frame-based computer vision.

Keywords: neuromorphic vision, computer vision, benchmarking, datasets, sensory processing

OPEN ACCESS

Edited by:
Michael Pfeiffer,
Swiss Federal Institute of Technology in Zurich, Switzerland

Reviewed by:
Teresa Serrano-Gotarredona,
Consejo Superior de Investigaciones Cientificas, Spain
Francisco Barranco,
University of Granada, Spain;
University of Maryland, USA

Correspondence:
Garrick Orchard
garrickorchard@nus.edu.sg

Specialty section:
This article was submitted to Neuromorphic Engineering, a section of the journal Frontiers in Neuroscience

Received: *25 July 2015*
Accepted: *28 September 2015*
Published: *13 October 2015*

Citation:
Tan C, Lallee S and Orchard G (2015) Benchmarking neuromorphic vision: lessons learnt from computer vision. Front. Neurosci. 9:374. doi: 10.3389/fnins.2015.00374

1. INTRODUCTION

Benchmarking using widely accepted datasets is important for algorithm development. Such benchmarking allows quantitative performance evaluation and comparison between algorithms, promoting competition and providing developers with tangible state-of-the-art targets to beat. Computer Vision (CV) is an obvious example where open access to good datasets has been integral in rapid development and maturation of the field (Kotsiantis et al., 2006).

We use the term "Computer Vision" (CV) to denote the conventional approach to visual sensing, which begins with acquisition of images (photographs), or sequences of images (video). Each image is a regular grid of pixels, each pixel having an intensity or color value. Such images are a widely accepted, and largely unquestioned first step in visual sensing.

However, the much younger field of Neuromorphic Vision (NV) takes a r
approach, doing away with images completely. The term "Neuromorphic Visio
approaches which rely on custom designed bio-inspired vision sensors which
non-frame-based manner. The most mature and common of these sensors are
asynchronous temporal contrast vision sensors. Other NV sensors have not ye
of maturity where they can be used to reliably capture datasets. Nevertheless,

PLOS | ONE

RESEARCH ARTICLE

Neural Tuning Size in a Model of Primate Visual Processing Accounts for Three Key Markers of Holistic Face Processing

Cheston Tan[1,2,3]*, Tomaso Poggio[1,2]*

1 McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America, 2 Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America, 3 Visual Computing Department, Institute for Infocomm Research, Singapore, Republic of Singapore

* cheston-tan@i2r.a-star.edu.sg (CT); tp@ai.mit.edu (TP)

OPEN ACCESS

Citation: Tan C, Poggio T (2016) Neural Tuning Size in a Model of Primate Visual Processing Accounts for Three Key Markers of Holistic Face Processing. PLoS ONE 11(3): e0150980. doi:10.1371/journal.pone.0150980

Editor: Nouchine Hadjikhani, Harvard Medical School, UNITED STATES

Received: January 9, 2016

Accepted: February 22, 2016

Published: March 17, 2016

Data Availability Statement: Minimal data set is available from the Figshare database (https://figshare.com/s/f5e1678fcde77fee36ca).

Funding: This work was funded by A*STAR (Agency for Science, Technology and Research), Joint Council Office grant #1335h00098, and National Science Foundation STC award CCF-1231216. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

Abstract

Faces are an important and unique class of visual stimuli, and have been of interest to neuroscientists for many years. Faces are known to elicit certain characteristic behavioral markers, collectively labeled "holistic processing", while non-face objects are not processed holistically. However, little is known about the underlying neural mechanisms. The main aim of this computational simulation work is to investigate the neural mechanisms that make face processing holistic. Using a model of primate visual processing, we show that a single key factor, "neural tuning size", is able to account for three important markers of holistic face processing: the Composite Face Effect (CFE), Face Inversion Effect (FIE) and Whole-Part Effect (WPE). Our proof-of-principle specifies the precise neurophysiological property that corresponds to the poorly-understood notion of holism, and shows that this one neural property controls three classic behavioral markers of holism. Our work is consistent with neurophysiological evidence, and makes further testable predictions. Overall, we provide a parsimonious account of holistic face processing, connecting computation, behavior and neurophysiology.

Introduction

Faces are an important class of visual stimuli with unique significance, and face processing is a longstanding topic of active study within neuroscience (e.g. [1–4]). Faces are ubiquitous throughout a person's life, and face recognition is important for daily social interaction. An important way in which visual processing of faces and non-face objects differs, is that faces have been found to elicit certain characteristic behavioral markers. These have been explained qualitatively through the loose notion of "holistic processing". However, the exact nature of holism is poorly understood, with multiple definitions, interpretations and putative mechanisms [5–7].

Importantly, little is known about the neural mechanisms underlying holistic face processing. For face processing in general by the primate and human visual systems, multiple neural