Pain and Machine Learning

Shakir Mohamed, DeepMind
@shakir_za

With Daniel Ott
University of Cambridge
people.hps.cam.ac.uk/index/phd-students/ott

A Universal Feature of Human Experience

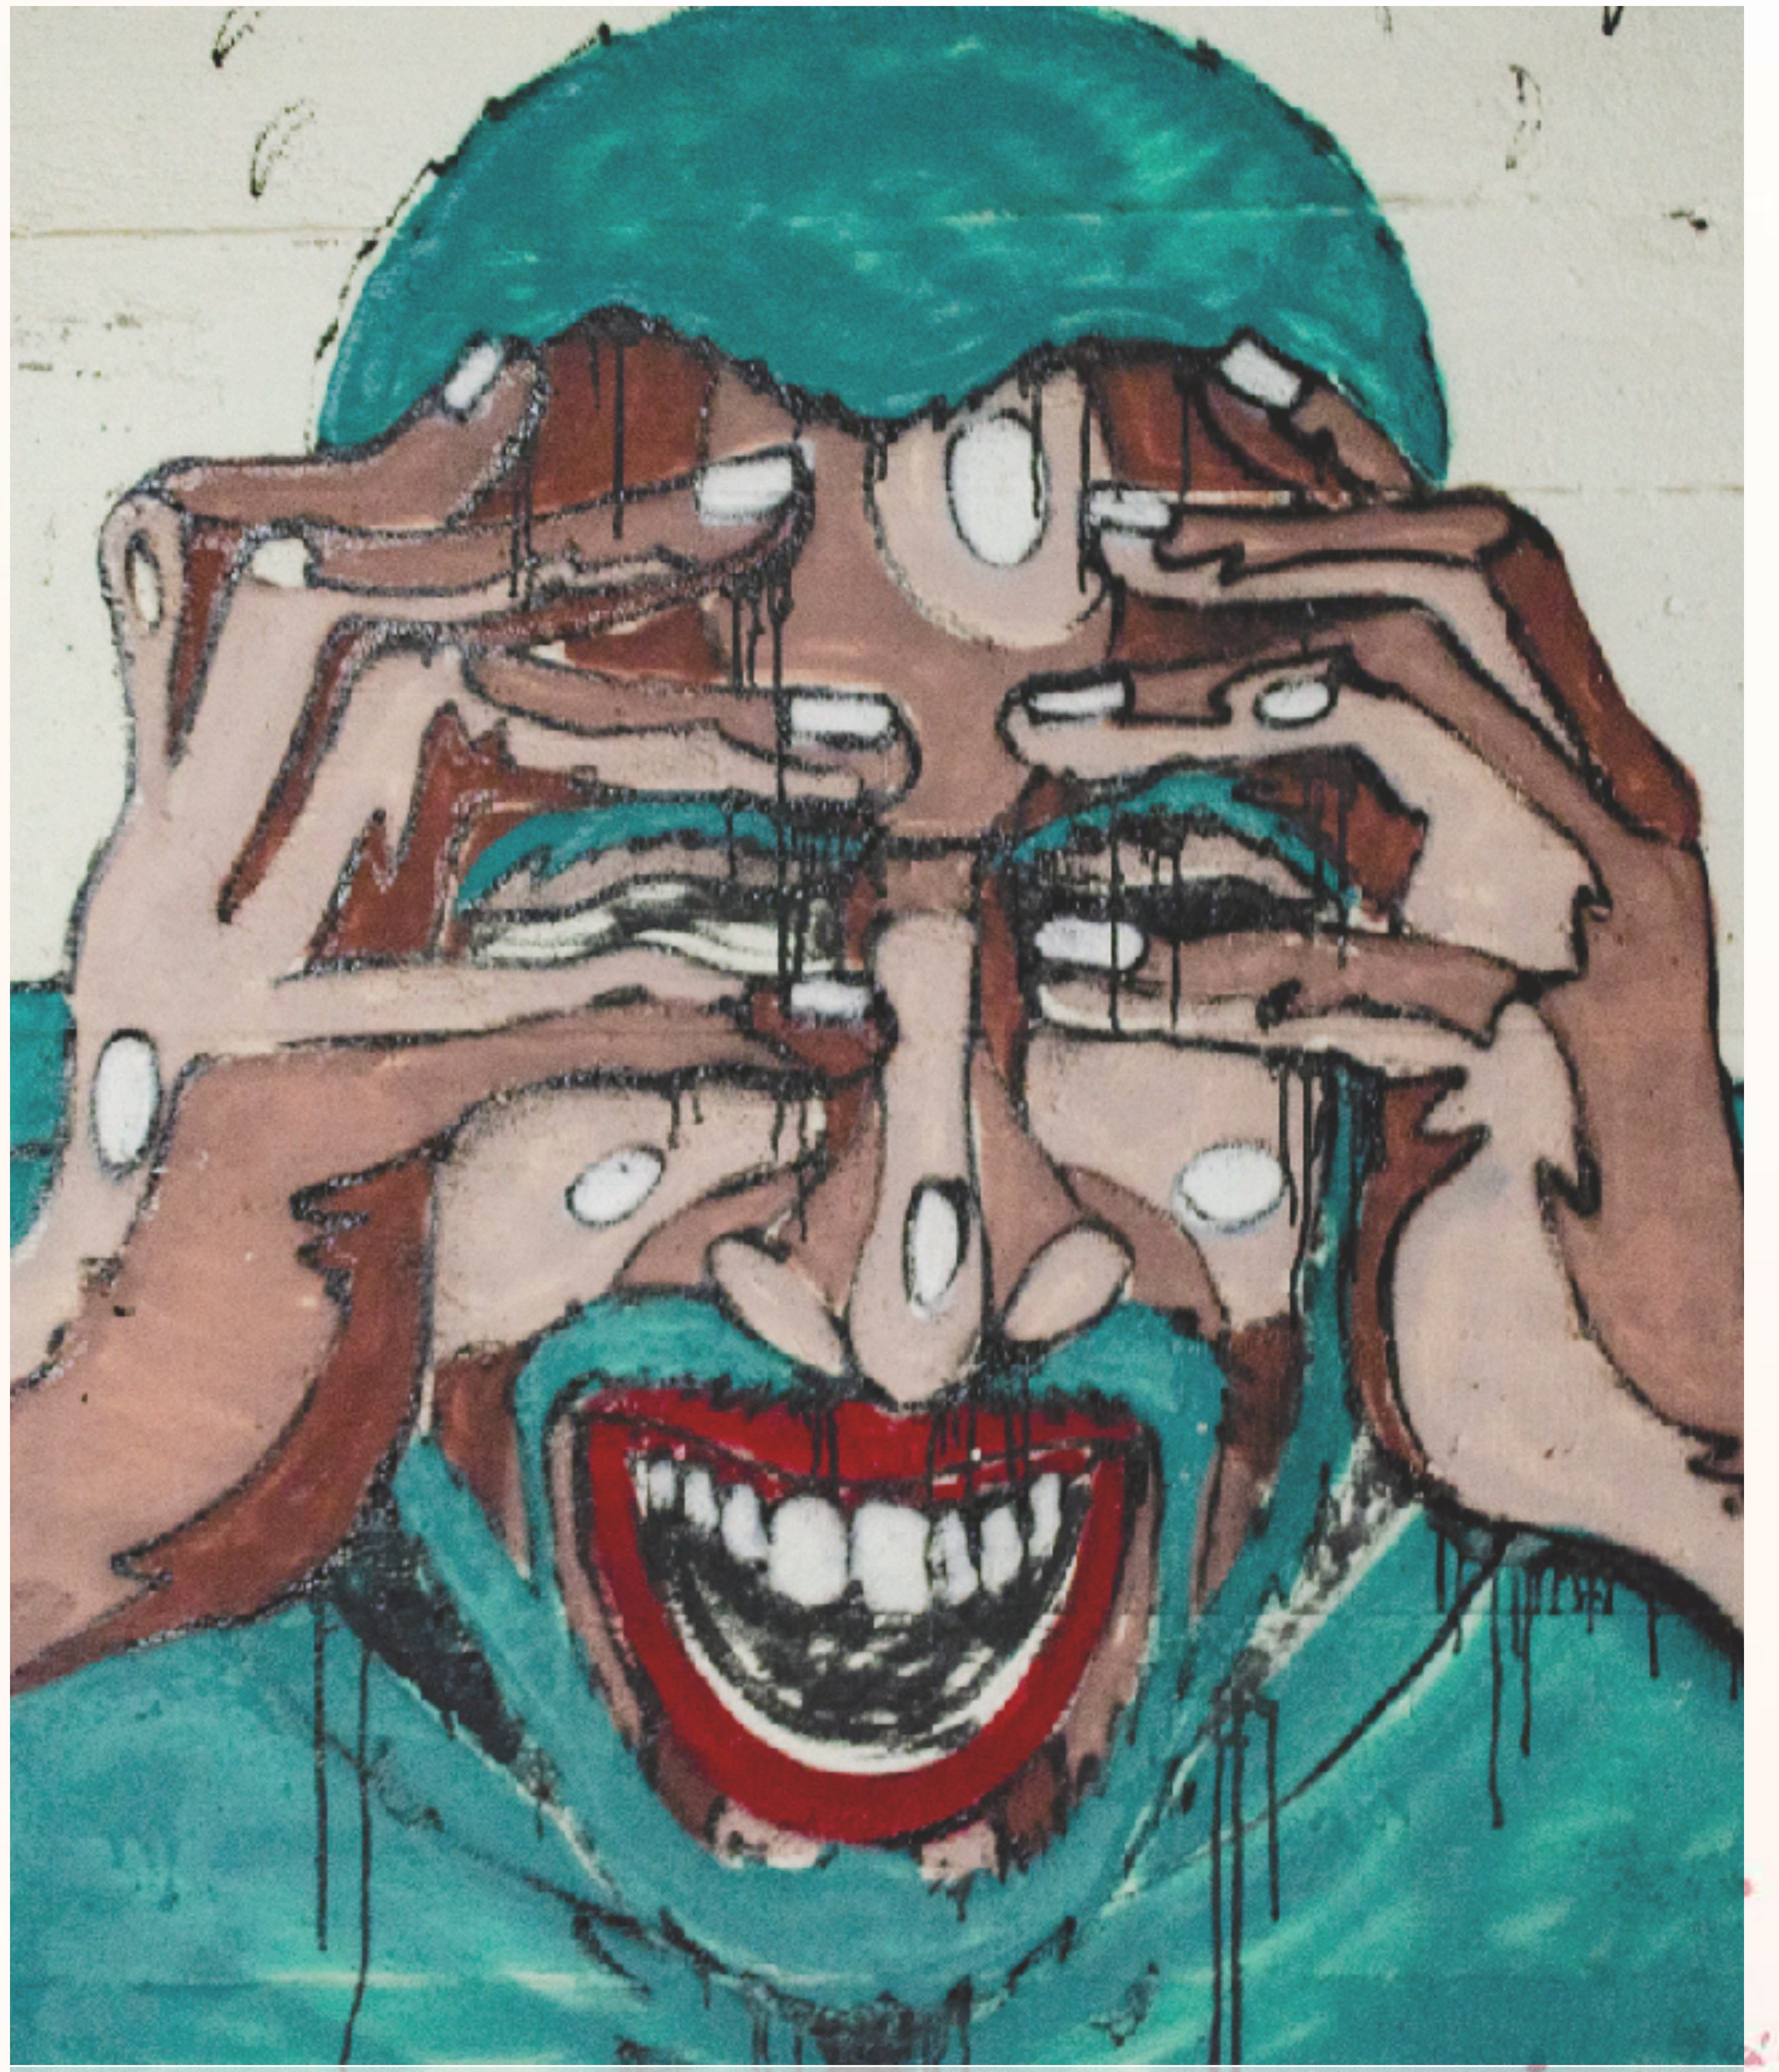

€200 billion
European Health systems

$635 billion
US Health systems

Scientific Pain

Pain studied scientifically, clinically and biologically

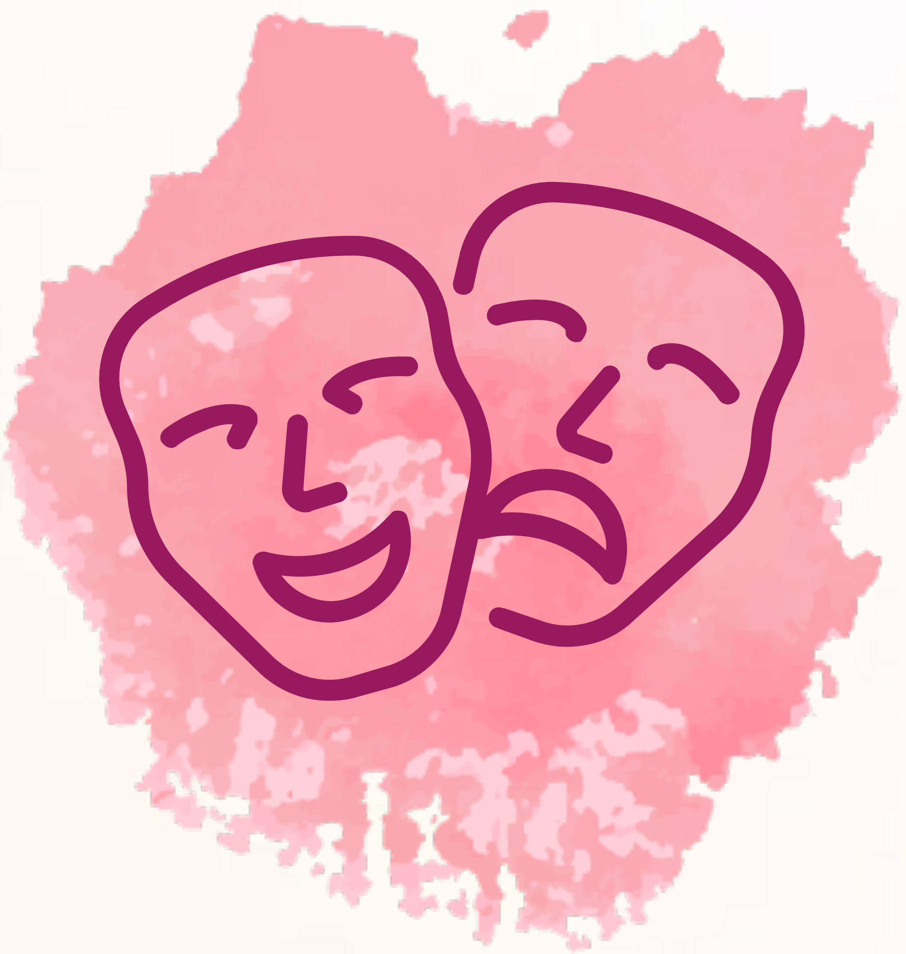

Folk Pain

Pain as understood amongst people in everyday experience

Pain is an unpleasant sensory and emotional experience associated with actual or potential tissue damage, or described in terms of such damage.

—The International Association for Study of Pain (IASP)

Pain is always subjective; It is unquestionably a sensation in a part or parts of the body, but it is also always unpleasant and therefore also an emotional experience;

Biologists recognize that those stimuli which cause pain are liable to damage tissue;

Many people report pain in the absence of tissue damage or any likely pathophysiological cause and that usually this happens for psychological reasons;

and there is usually no way to distinguish their experience from that due to tissue damage if we take the subjective report.

—The International Association for Study of Pain (IASP)

DANIEL C. DENNETT

WHY YOU CAN'T MAKE A COMPUTER THAT FEELS PAIN

It has seemed important to many people to claim that computers cannot *in principle* duplicate various human feats, activities, happenings. Such aprioristic claims, we have learned, have an embarrassing history of subsequent falsification. Contrary to recently held opinion, for instance, computers can play superb checkers and good chess, can produce novel and unexpected proofs of nontrivial theorems, can conduct sophisticated conversations in ordinary if tightly circumscribed English. The materialist or computerphile who grounds an uncomplicated optimism in this ungraceful retreat of the skeptics, however, is in danger of installing conceptual confusion in the worst place, in the foundations of his own ascendant view of the mind. The triumphs of Artificial Intelligence have been balanced by failures and false starts. Some have asked if there is a pattern to be discerned here. Keith Gunderson has pointed out that the successes have been with task-oriented, *sapient* features of mentality, the failures and false starts with *sentient* features of mentality, and has developed a distinction between program-receptive and program-resistant features of mentality.[1] Gunderson's point is not what some have hoped. Some have hoped he had found a fall-back position for them: *viz.*, maybe machines can *think* but they can't *feel*. His point is rather that the task of getting a machine to feel is a very different task from getting it to think; in particular it is not a task that invites solution simply by sophisticated innovations in *programming*, but rather, if at all, by devising new sorts of *hardware*. This goes some way to explaining the recalcitrance of mental features like pain to computer simulation, but not far enough. Since most of the discredited aprioristic thinking about the limitations of computers can be seen in retrospect to have stumbled over details, I propose

Pain Eliminativism

The concept of pain fails to refer to anything empirical, so we would be better served by removing it from our usage and instead discovering other vocabularies to serve its place.

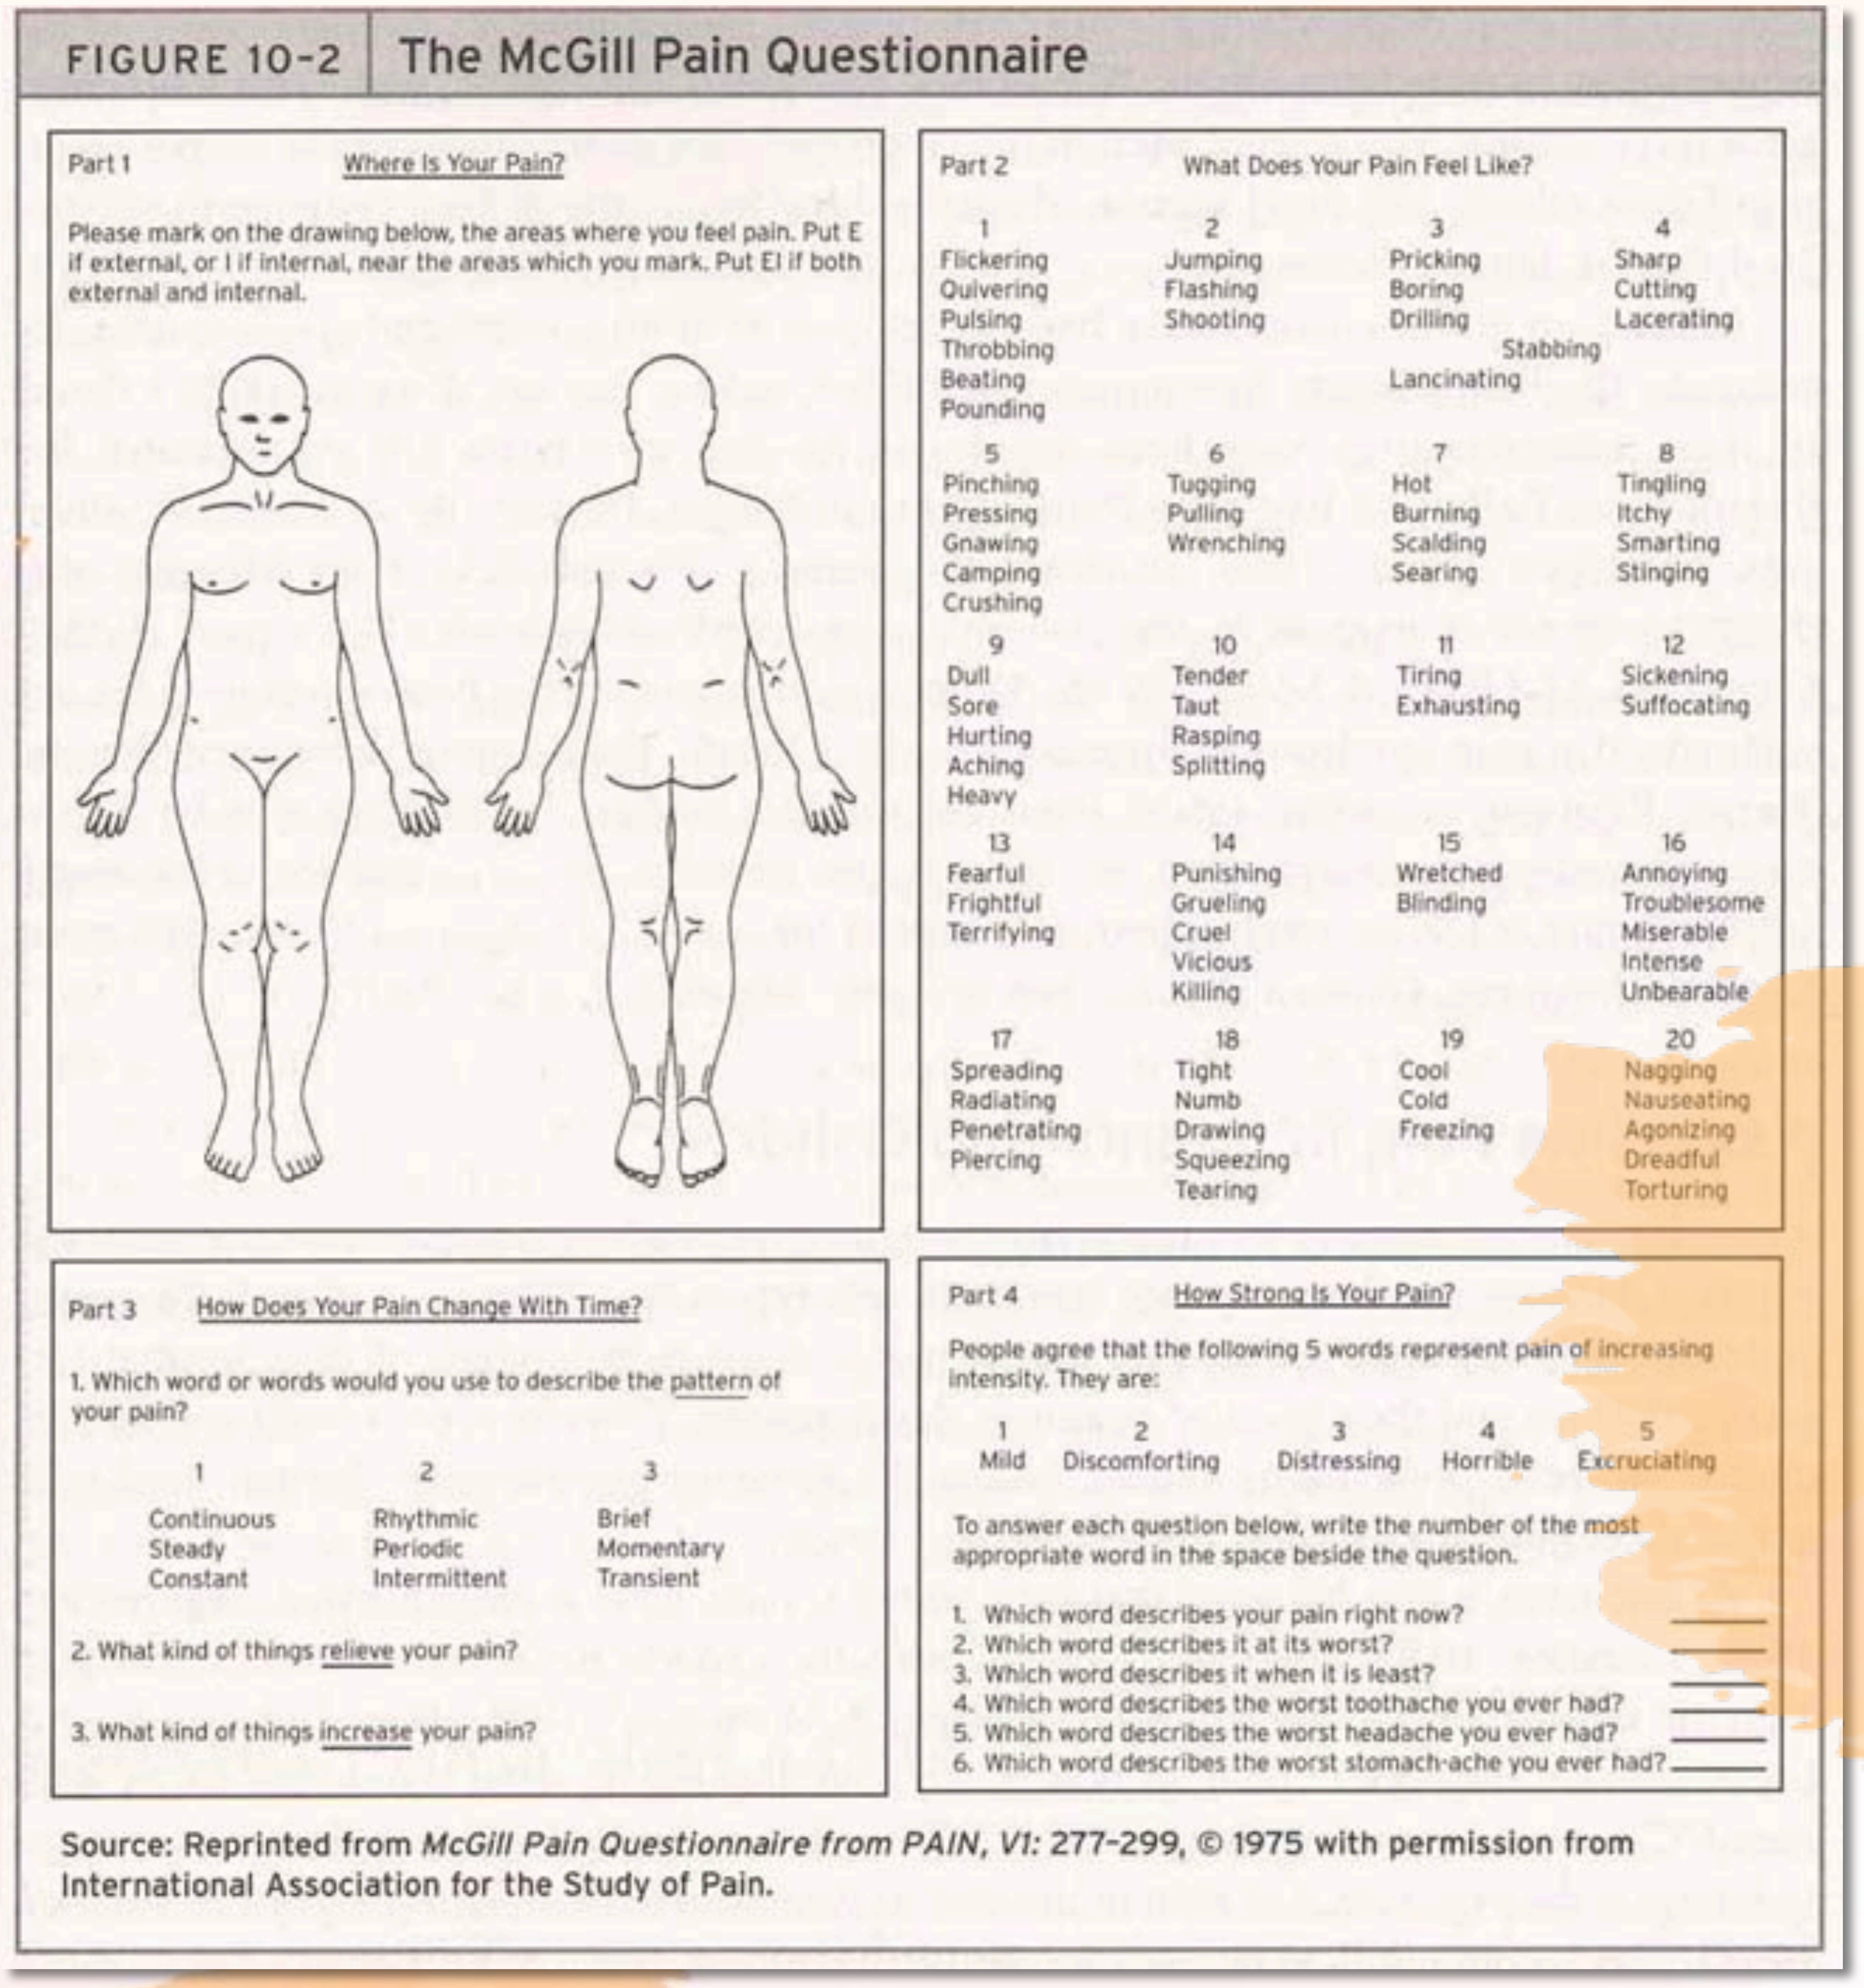

FIGURE 10-2 The McGill Pain Questionnaire

Part 1 Where Is Your Pain?

Please mark on the drawing below, the areas where you feel pain. Put E if external, or I if internal, near the areas which you mark. Put EI if both external and internal.

Part 2 What Does Your Pain Feel Like?

1	2	3	4
Flickering Quivering Pulsing Throbbing Beating Pounding	Jumping Flashing Shooting	Pricking Boring Drilling Stabbing Lancinating	Sharp Cutting Lacerating
5	**6**	**7**	**8**
Pinching Pressing Gnawing Camping Crushing	Tugging Pulling Wrenching	Hot Burning Scalding Searing	Tingling Itchy Smarting Stinging
9	**10**	**11**	**12**
Dull Sore Hurting Aching Heavy	Tender Taut Rasping Splitting	Tiring Exhausting	Sickening Suffocating
13	**14**	**15**	**16**
Fearful Frightful Terrifying	Punishing Grueling Cruel Vicious Killing	Wretched Blinding	Annoying Troublesome Miserable Intense Unbearable
17	**18**	**19**	**20**
Spreading Radiating Penetrating Piercing	Tight Numb Drawing Squeezing Tearing	Cool Cold Freezing	Nagging Nauseating Agonizing Dreadful Torturing

Part 3 How Does Your Pain Change With Time?

1. Which word or words would you use to describe the pattern of your pain?

1	2	3
Continuous Steady Constant	Rhythmic Periodic Intermittent	Brief Momentary Transient

2. What kind of things relieve your pain?

3. What kind of things increase your pain?

Part 4 How Strong Is Your Pain?

People agree that the following 5 words represent pain of increasing intensity. They are:

1	2	3	4	5
Mild	Discomforting	Distressing	Horrible	Excruciating

To answer each question below, write the number of the most appropriate word in the space beside the question.

1. Which word describes your pain right now? ____
2. Which word describes it at its worst? ____
3. Which word describes it when it is least? ____
4. Which word describes the worst toothache you ever had? ____
5. Which word describes the worst headache you ever had? ____
6. Which word describes the worst stomach-ache you ever had? ____

Source: Reprinted from *McGill Pain Questionnaire from PAIN, VI: 277-299*, © 1975 with permission from International Association for the Study of Pain.

Sensation Model

The painful sensation. Pain is not paired with any representation of a physical state or stimulus, but is seen to occur in correlation with them.

Representation Model

Pain is a representation, or abstraction, of a perceptual feature of one's environment or body.

Motivational Model

Pain is a request or command to protect a part of your body.

Theories of perception suffer from one fundamental flaw: they are theories of vision.

—**Ann-Sophie Barwich**

Does this resonate with the trajectory we have taken in Machine Learning?

Situational Assessment

Think of perception as always combining multiple sources of contextual information to form final perceptual states.

Allows for painful situations
Object-less account of sensory states.

Illustration by Freepik Stories

Pain as Inference

Pain is considered an active predictor of future bodily states, as well as an assessor of current afferent information, and Bayesian updating transforms multimodal prior experiences into future assessments.

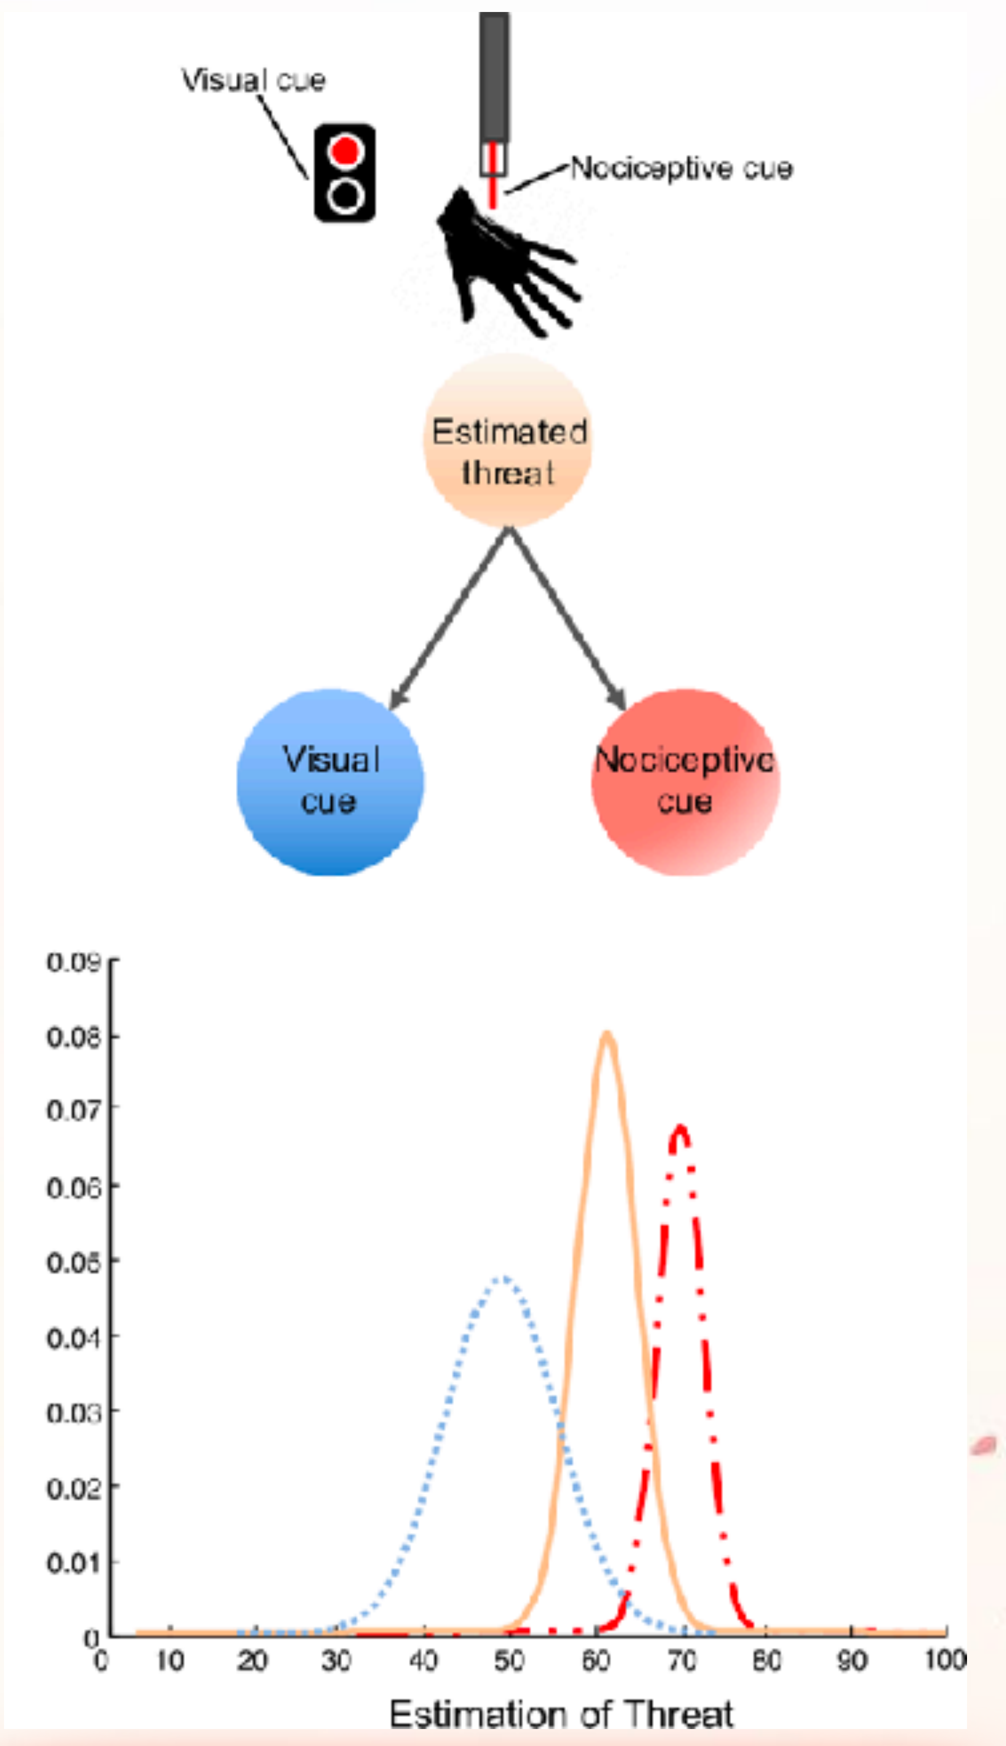

REVIEW

Pain: A Statistical Account

Abby Tabor[1], Michael A. Thacker[2,3], G. Lorimer Moseley[3,4], Konrad P. Körding[5]*

1 Centre for Pain Research, University of Bath, North East Somerset, United Kingdom, 2 Centre for Human and Aerospace Physiological Sciences/Pain Section, Neuroimaging, Institute of Psychiatry, Kings College London, London, United Kingdom, 3 Sansom Institute for Health Research, University of South Australia, Adelaide, South Australia, Australia, 4 Neuroscience Research Australia, Sydney, New South Wales, Australia, 5 Rehabilitation Institute of Chicago, Northwestern University, Chicago, Illinois, United States of America

* kk@northwestern.edu

Pain as Reward

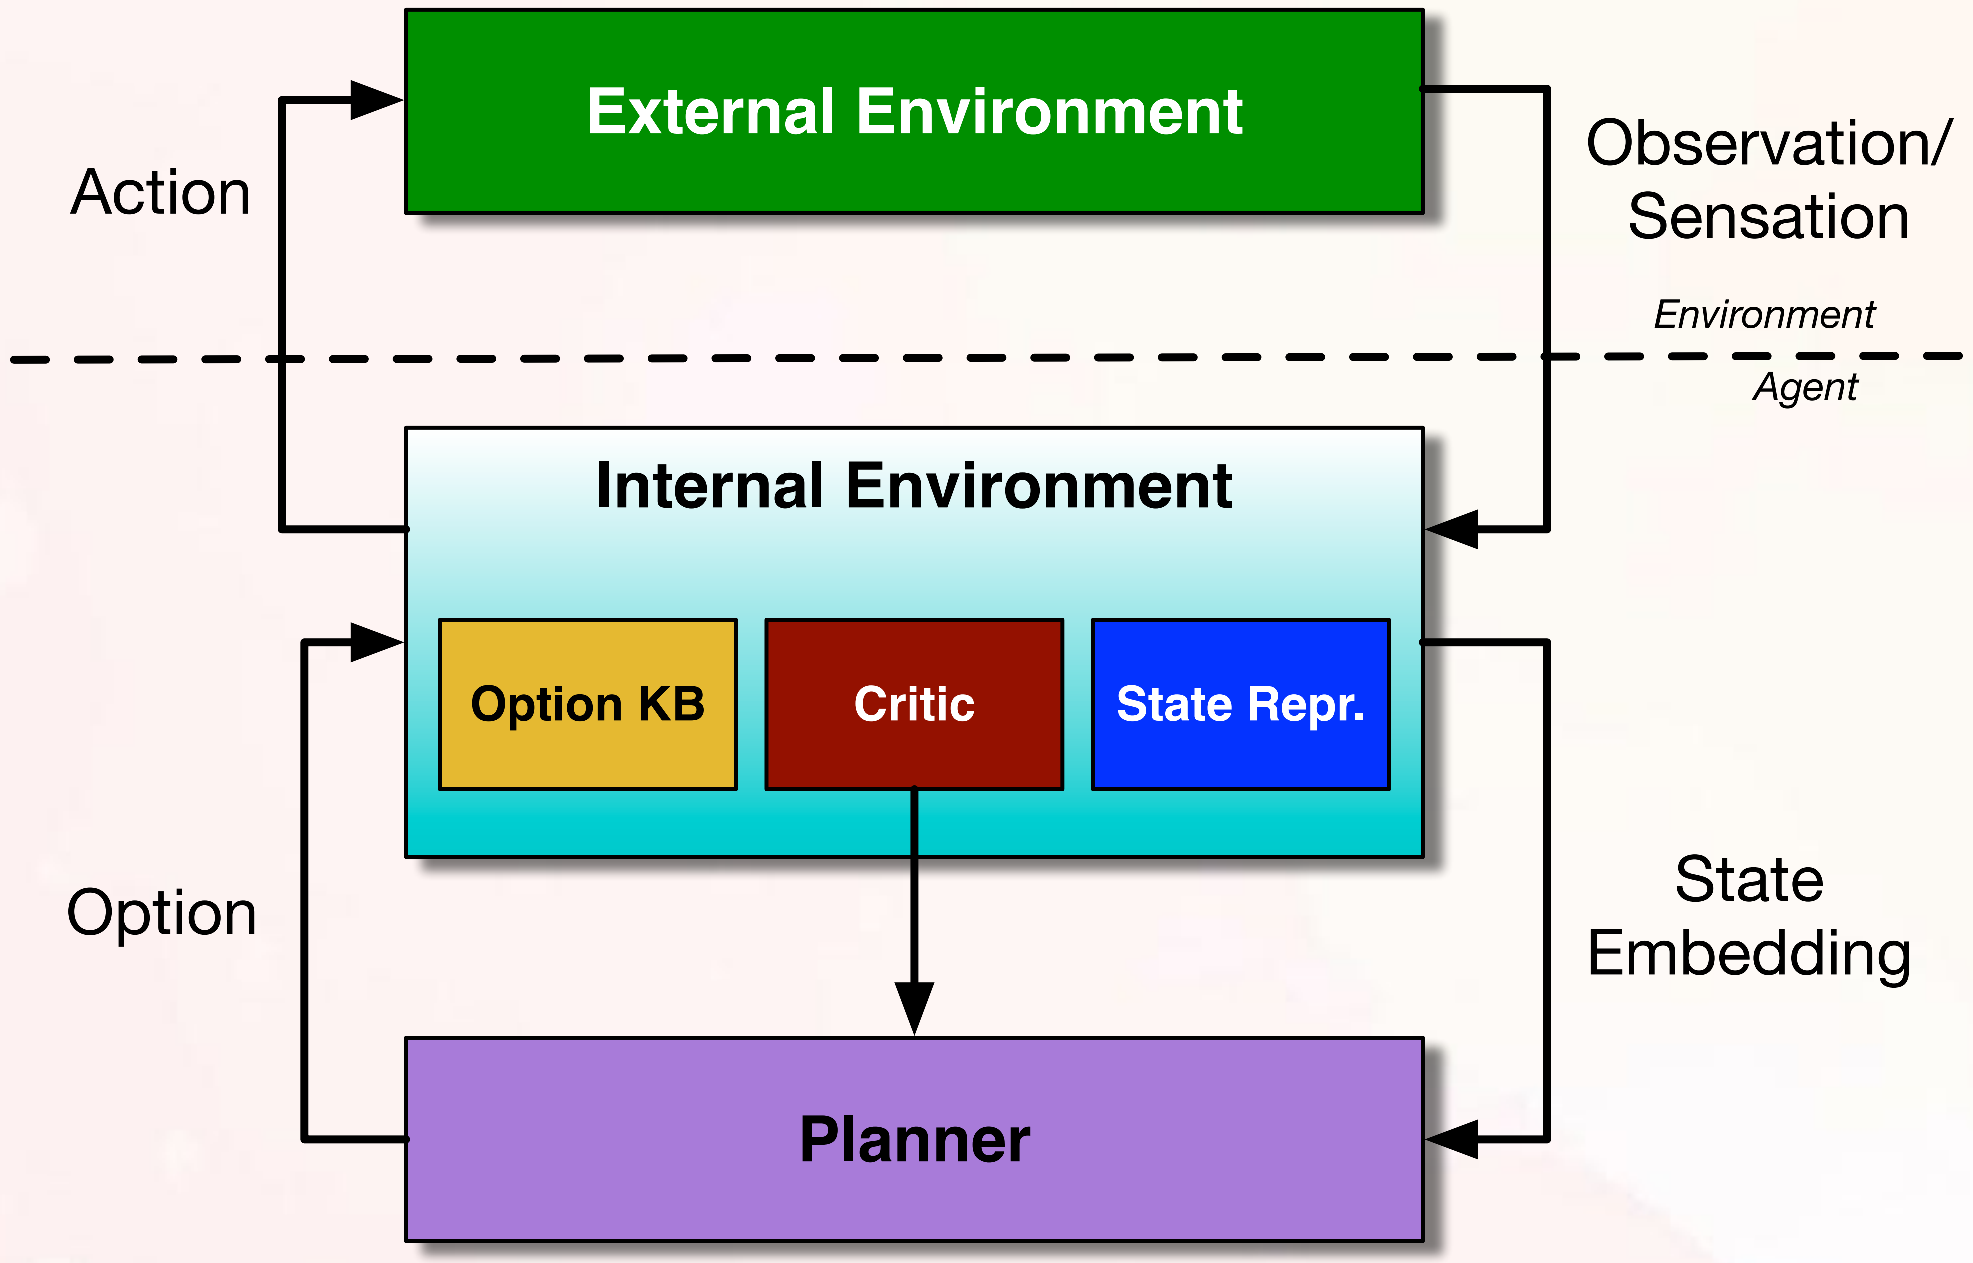

General type of reward formed by integrated sensory information – general type of cumulant – used within a risk-averse intrinsic motivation system.

Pain: A Precision Signal for Reinforcement Learning and Control

Ben Seymour[1,2,*]
[1]Center for Information and Neural Networks, National Institute of Information and Communications Technology, 1-4 Yamadaoka, Suita, Osaka 565-0871, Japan
[2]Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK
*Correspondence: bjs49@cam.ac.uk

https://doi.org/10.1016/j.neuron.2019.01.055

Since noxious stimulation usually leads to the perception of pain, pain has traditionally been considered sensory nociception. But its variability and sensitivity to a broad array of cognitive and motivational factors have meant it is commonly viewed as inherently imprecise and intangibly subjective. However, the core function of pain is motivational—to direct both short- and long-term behavior away from harm. Here, we illustrate that a reinforcement learning model of pain offers a mechanistic understanding of how the brain supports this, illustrating the underlying computational architecture of the pain system. Importantly, it explains why pain is tuned by multiple factors and necessarily supported by a distributed network of brain regions, recasting pain as a precise and objectifiable control signal.

Pain Learning Insights

Single-shot learning

Single exposure learning

Transfer learning

Generalisability to novel pain stimuli

Imitation learning

Social transfer of acquired pain knowledge

Computational

Algorithmic

Implementation

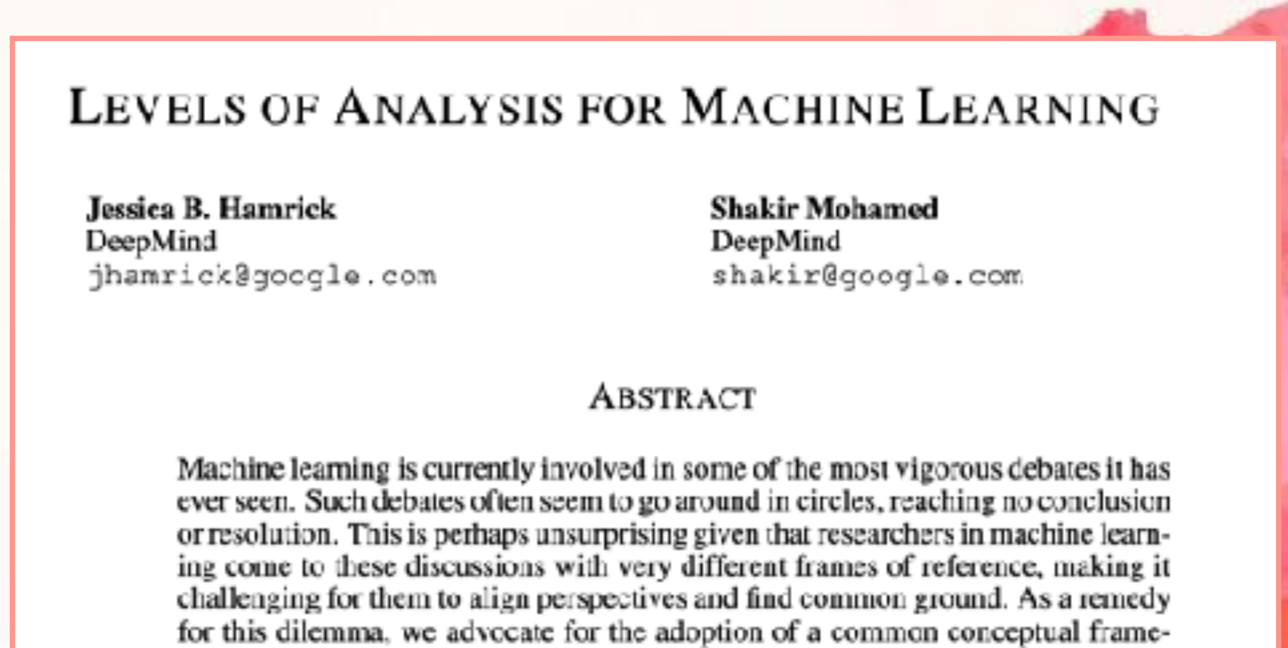

LEVELS OF ANALYSIS FOR MACHINE LEARNING

Jessica B. Hamrick
DeepMind
jhamrick@google.com

Shakir Mohamed
DeepMind
shakir@google.com

ABSTRACT

Machine learning is currently involved in some of the most vigorous debates it has ever seen. Such debates often seem to go around in circles, reaching no conclusion or resolution. This is perhaps unsurprising given that researchers in machine learning come to these discussions with very different frames of reference, making it challenging for them to align perspectives and find common ground. As a remedy for this dilemma, we advocate for the adoption of a common conceptual frame-

Computational and Algorithmic Levels

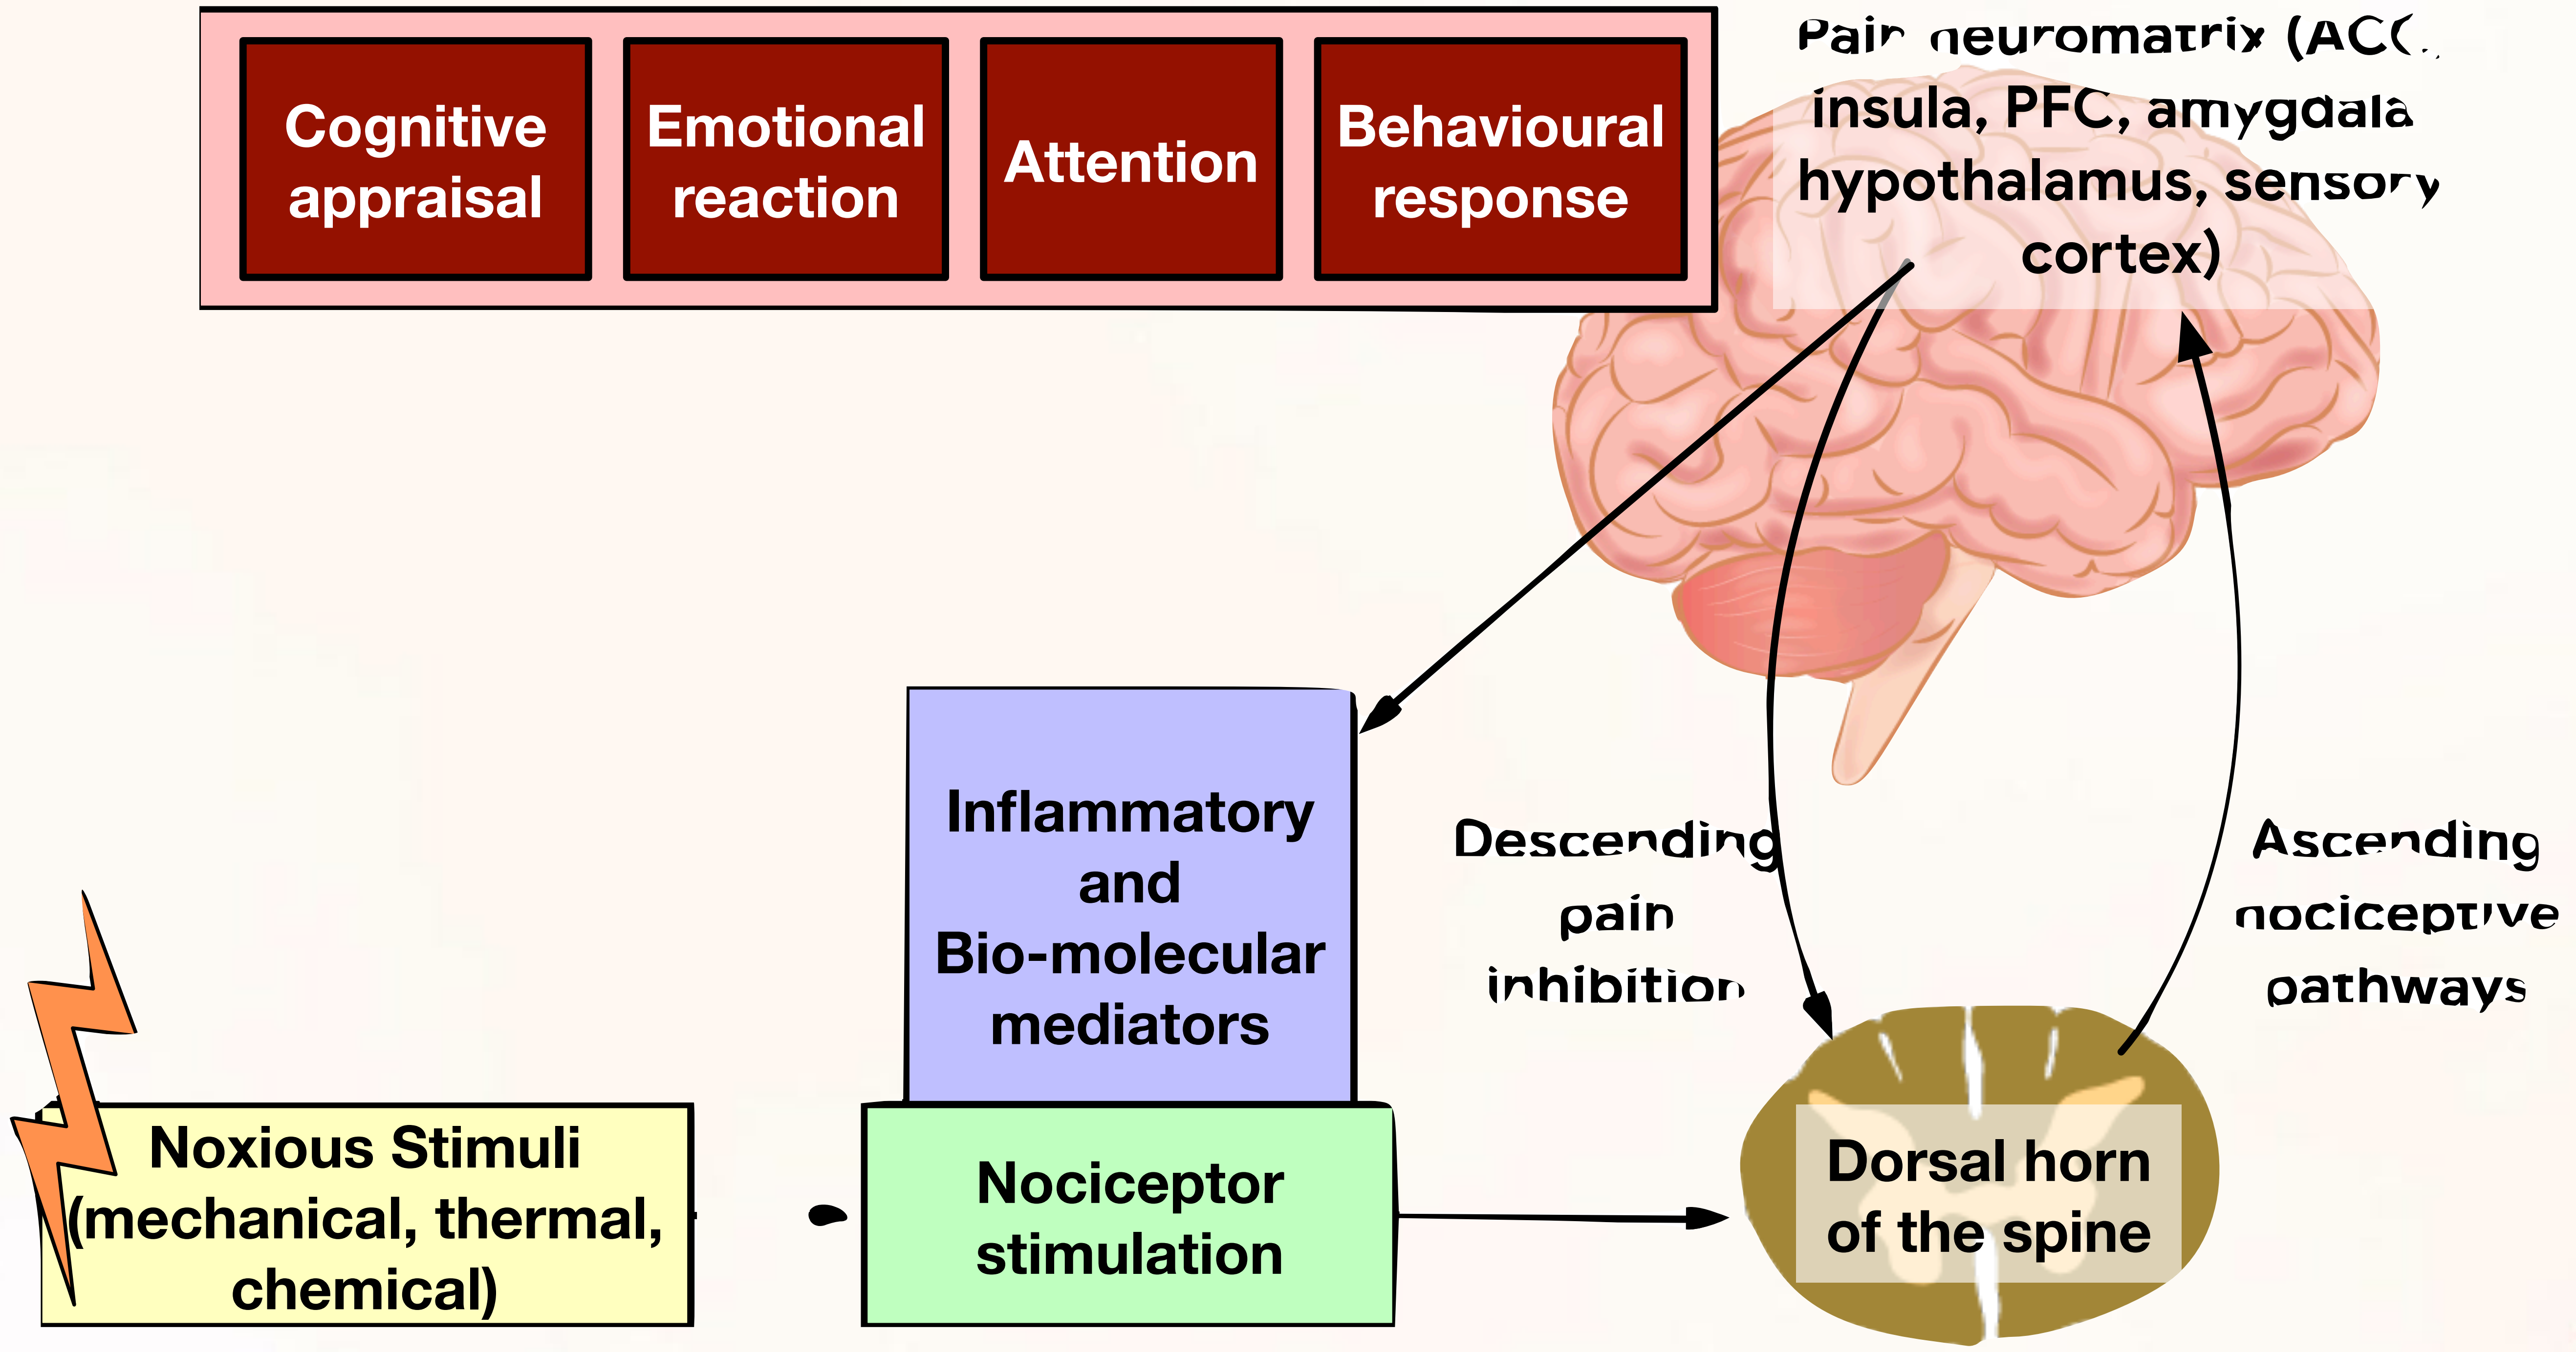

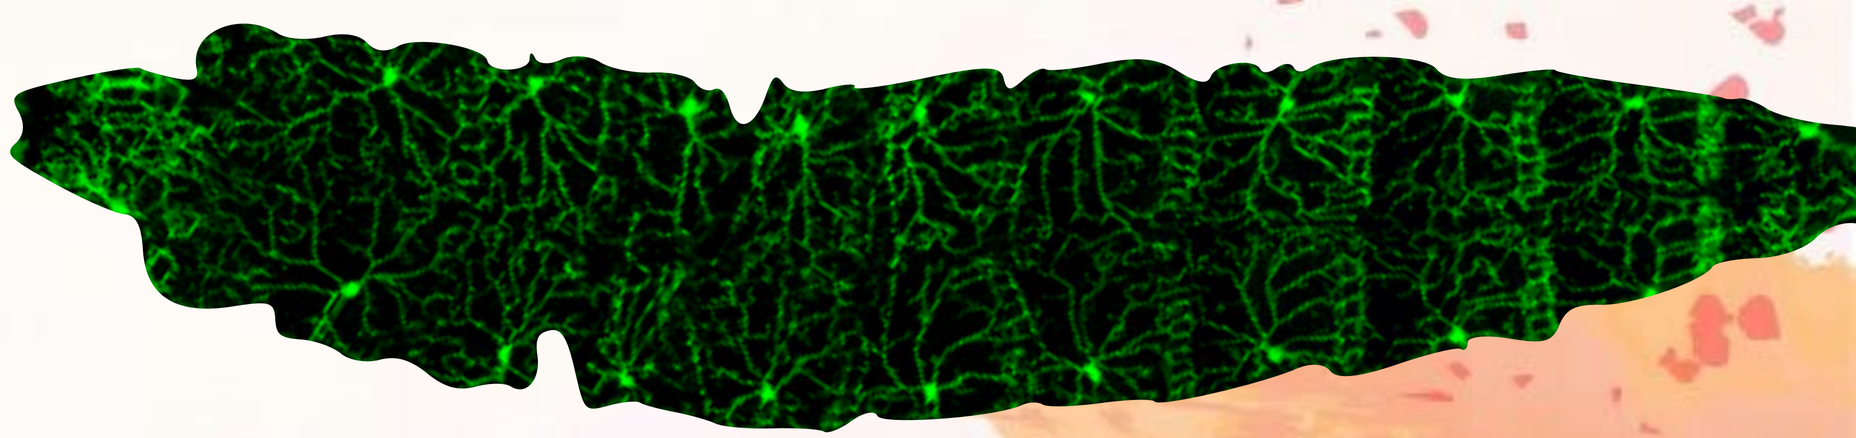

Look at Pain Research

View on the world of pain research and the many dimensions it takes, whether in technical, social, or the sociotechnical domain

Describe the Missing Level

To fill the missing algorithmic layer using situational assessment as a guide to developing new insights on learning.

On Pain

–Kahlil Gibran

And a woman spoke, saying, Tell us of Pain.

And he said:

Your pain is the breaking of the shell that encloses your understanding.

Even as the stone of the fruit must break, that its heart may stand in the sun, so must you know pain.

And could you keep your heart in wonder at the daily miracles of your life your pain would not seem less wondrous than your joy;

And you would accept the seasons of your heart, even as you have always accepted the seasons that pass over your fields.

And you would watch with serenity through the winters of your grief.

Much of your pain is self-chosen.

It is the bitter potion by which the physician within you heals your sick self.

Therefore trust the physician, and drink his remedy in silence and tranquility:

For his hand, though heavy and hard, is guided by the tender hand of the Unseen,

And the cup he brings, though it burn your lips, has been fashioned of the clay which the Potter has moistened with His own sacred tears.

Some Resources

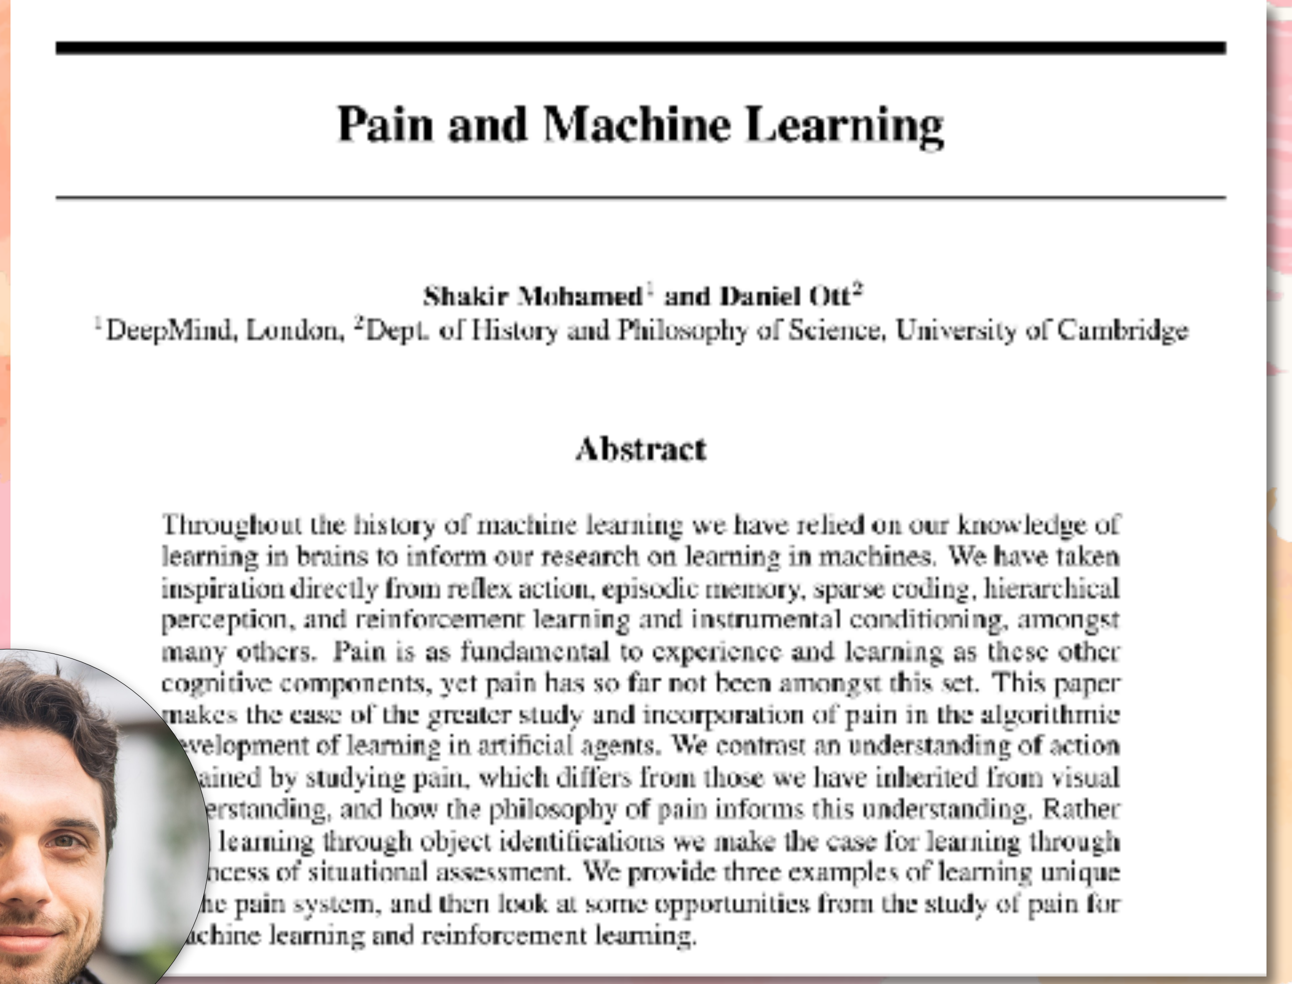

Pain and Machine Learning

Shakir Mohamed[1] and Daniel Ott[2]
[1]DeepMind, London, [2]Dept. of History and Philosophy of Science, University of Cambridge

Abstract

Throughout the history of machine learning we have relied on our knowledge of learning in brains to inform our research on learning in machines. We have taken inspiration directly from reflex action, episodic memory, sparse coding, hierarchical perception, and reinforcement learning and instrumental conditioning, amongst many others. Pain is as fundamental to experience and learning as these other cognitive components, yet pain has so far not been amongst this set. This paper makes the case for the greater study and incorporation of pain in the algorithmic evelopment of learning in artificial agents. We contrast an understanding of action ained by studying pain, which differs from those we have inherited from visual erstanding, and how the philosophy of pain informs this understanding. Rather learning through object identifications we make the case for learning through ncess of situational assessment. We provide three examples of learning unique he pain system, and then look at some opportunities from the study of pain for chine learning and reinforcement learning.

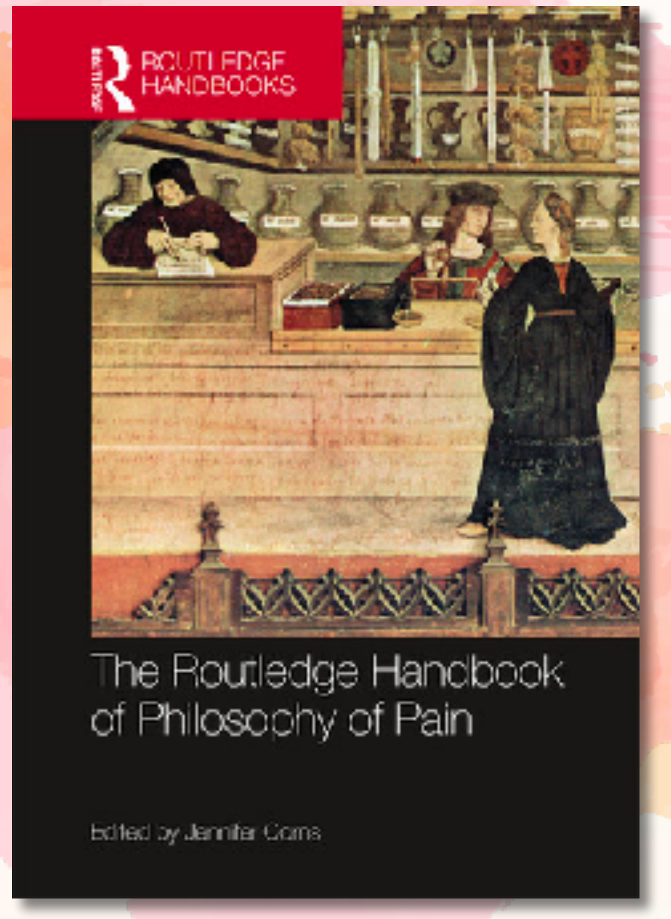

ROUTLEDGE HANDBOOKS

The Routledge Handbook of Philosophy of Pain

Edited by Jennifer Corns

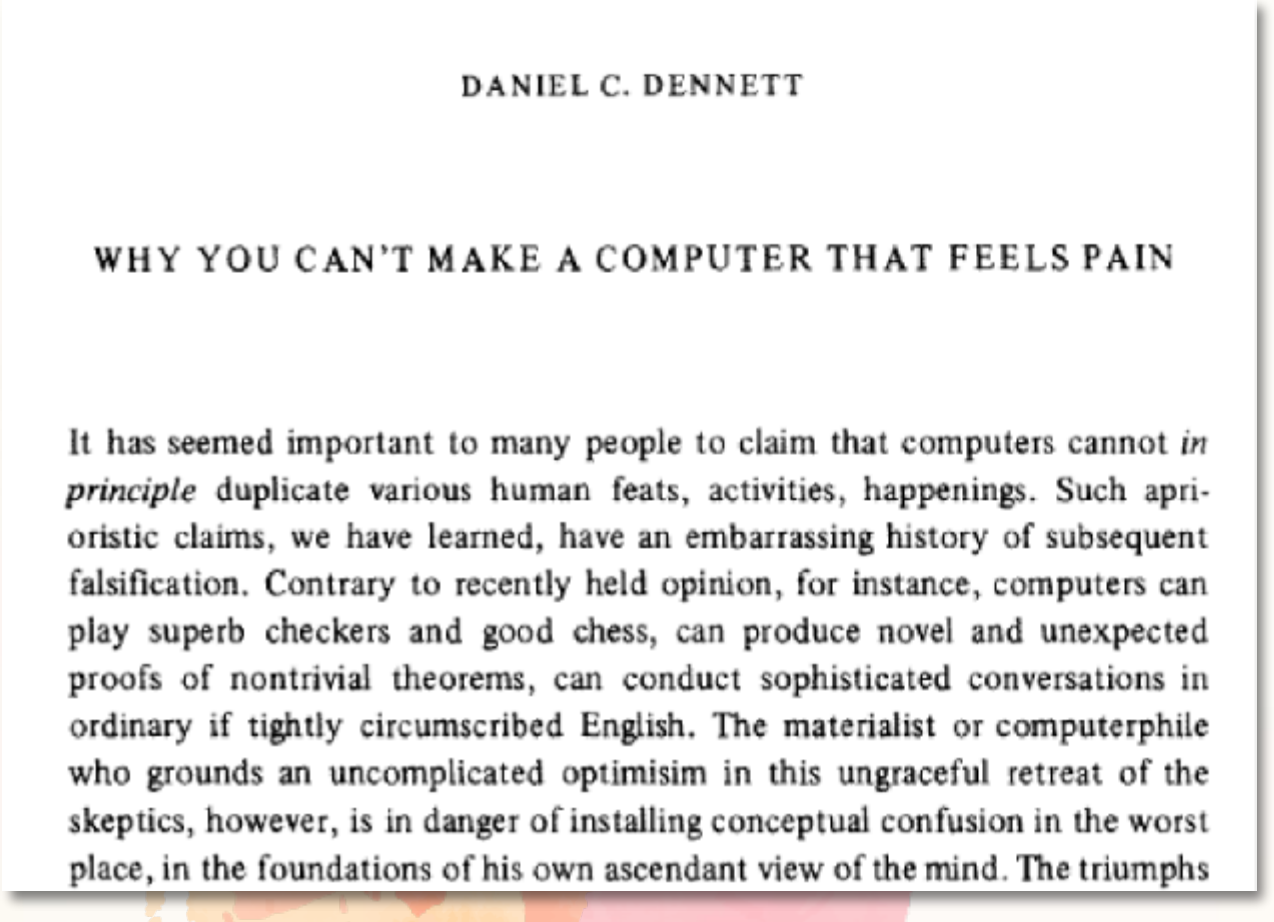

DANIEL C. DENNETT

WHY YOU CAN'T MAKE A COMPUTER THAT FEELS PAIN

It has seemed important to many people to claim that computers cannot *in principle* duplicate various human feats, activities, happenings. Such aprioristic claims, we have learned, have an embarrassing history of subsequent falsification. Contrary to recently held opinion, for instance, computers can play superb checkers and good chess, can produce novel and unexpected proofs of nontrivial theorems, can conduct sophisticated conversations in ordinary if tightly circumscribed English. The materialist or computerphile who grounds an uncomplicated optimism in this ungraceful retreat of the skeptics, however, is in danger of installing conceptual confusion in the worst place, in the foundations of his own ascendant view of the mind. The triumphs

Topical Review

PAIN

OPEN

Machine learning in pain research

Jörn Lötsch[a,b,*], Alfred Ultsch[c]

1. Introduction

Pain and pain chronification are incompletely understood and unresolved medical problems that continue to have a high prevalence.[16] It has been accepted that pain is a complex phenomenon.[2,33,72] Contemporary methods of computational science[21] can use complex clinical and experimental data to better understand the complexity of pain. Among data science techniques, machine learning is referred to as a set of methods (**Fig. 1**) that can automatically detect patterns in data and then use the uncovered patterns to predict or classify future data, to observe structures such as subgroups in the data, or to extract information from the data suitable to derive new knowledge.[11,49] Together with (bio)statistics, artificial intelligence and machine learning aim at learning from data.

Although statistics can be regarded as a branch of mathematics, artificial intelligence and machine learning have developed from computer science (Ref. 58; see also https://en.wikipedia.org/wiki/Artificial_intelligence). The initial definition of artificial intelligence originates from Alan Turing who proposed an experiment where 2 players, who can either be human or artificial, try to convince a human third player, that they are also humans.[68] The test of artificial intelligence is passed if the third player cannot tell who is the machine. Important steps in the development of machine learning were the first creation of the computer learning program, which was a checker game,[54] and the first neural network called the perceptron.[63] Statistics uses recognition, knowledge discovery, and data mining and share partly the same methods such as regression, which is used widely in statistics but is also considered as a classification method in machine learning (**Fig. 1**).

In the present research context, when provided with pain-related data, machine-learned methods are able to learn a mapping of complex features to a known class, that is, to predict a pain phenotype class from a complex pattern of acquired parameters. After the machine has learned the prediction of a pain-related phenotype, the algorithm can subsequently be used on new data from which the class membership of a novel yet unclassified subject can be identified. However, machine learning methods can also be used for pattern recognition in complex pain-related data to reveal traces of an underlying molecular background or for knowledge discovery in big data in a drug discovery or repurposing context. The increasing use of contemporary methods of computational science is reflected in the rising number of reports using machine learning for pain research (**Table 1**). This review is focused on machine-learned technologies applied to general pain research that allow one to analyze and predict pain phenotypes and to obtain knowledge from experimental and clinical pain-related data.

2. Pain research involving machine learning

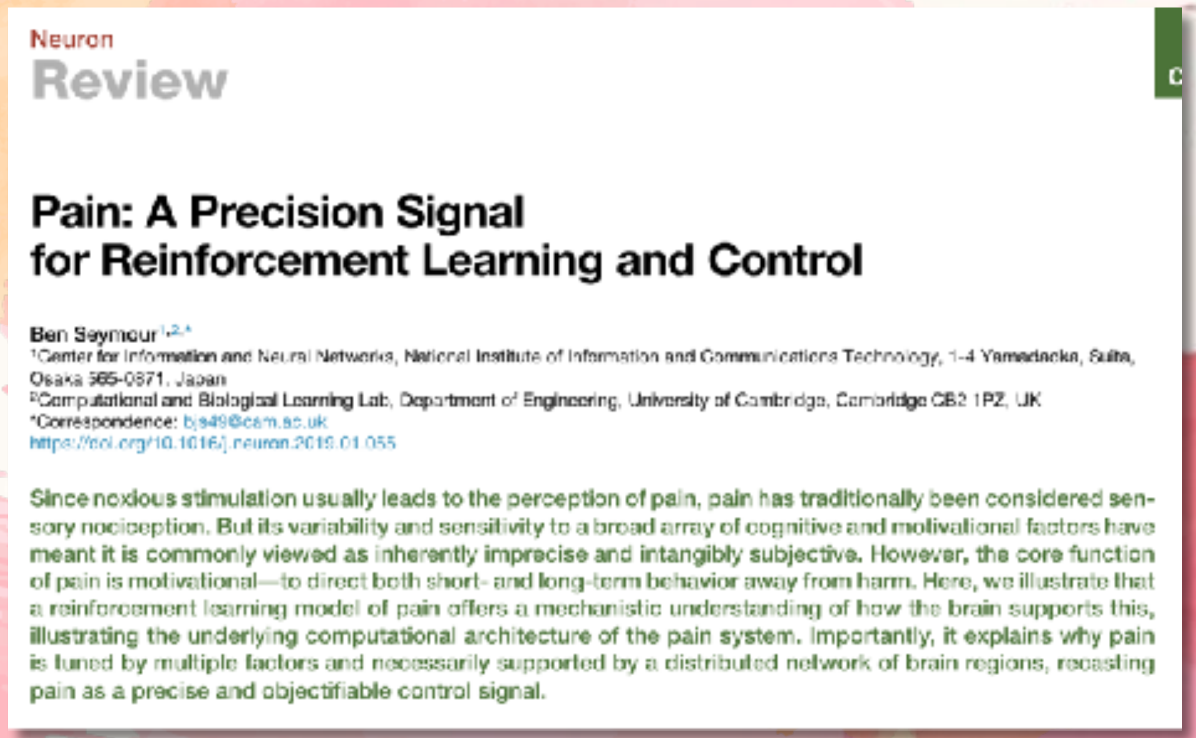

Neuron

Review

Pain: A Precision Signal for Reinforcement Learning and Control

Ben Seymour[1,2,*]
[1]Center for Information and Neural Networks, National Institute of Information and Communications Technology, 1-4 Yamadaoka, Suita, Osaka 565-0871, Japan
[2]Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK
*Correspondence: bjs49@cam.ac.uk

https://doi.org/10.1016/j.neuron.2019.01.055

Since noxious stimulation usually leads to the perception of pain, pain has traditionally been considered sensory nociception. But its variability and sensitivity to a broad array of cognitive and motivational factors have meant it is commonly viewed as inherently imprecise and intangibly subjective. However, the core function of pain is motivational—to direct both short- and long-term behavior away from harm. Here, we illustrate that a reinforcement learning model of pain offers a mechanistic understanding of how the brain supports this, illustrating the underlying computational architecture of the pain system. Importantly, it explains why pain is tuned by multiple factors and necessarily supported by a distributed network of brain regions, recasting pain as a precise and objectifiable control signal.