Person: Pick up a big red block.
Computer: OK.
Person: Grasp the pyramid.
Computer: I DON'T UNDERSTAND WHICH PYRAMID YOU MEAN.
Person (changing their mind): Find a block which is taller than the one you are holding and put it into the box.
Computer: BY "IT", I ASSUME YOU MEAN THE BLOCK WHICH IS TALLER THAN THE ONE I AM HOLDING.
Computer: OK.
Person: What does the box contain?
Computer: THE BLUE PYRAMID AND THE BLUE BLOCK.
Person: What is the pyramid supported by?
Computer: THE BOX.
Person: How many blocks are not in the box?
Computer: FOUR OF THEM.
Person: Is at least one of them narrower than the one which i told you to pick up?
Computer: YES, THE RED CUBE.
Person: Is it supported?
Computer: YES, BY THE TABLE.
Person: Can the table pick up blocks?
Computer: NO.
Person: Can a pyramid be supported by a block?
Computer: YES.
Person: Can a pyramid support a pyramid?
Computer: I DON'T KNOW.
Person: Stack up two pyramids.
Computer: I CAN'T.
Person: Will you please stack up both of the red blocks and either a green cube or a pyramid?
Computer: OK.
Person: Which cube is sitting on the table?
Computer: THE LARGE GREEN ONE WHICH SUPPORTS THE RED PYRAMID.
:

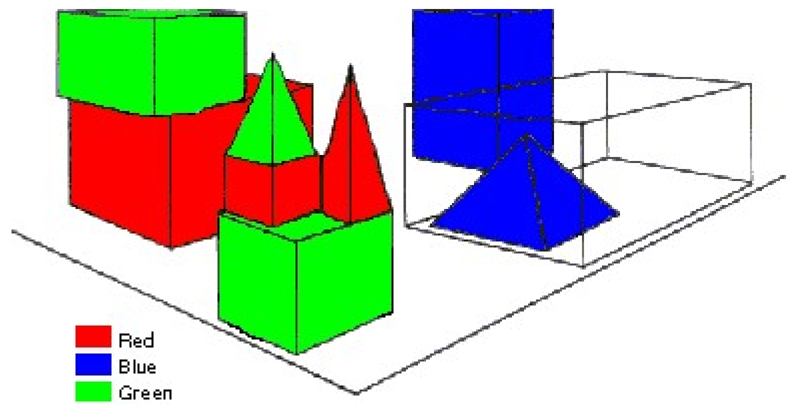

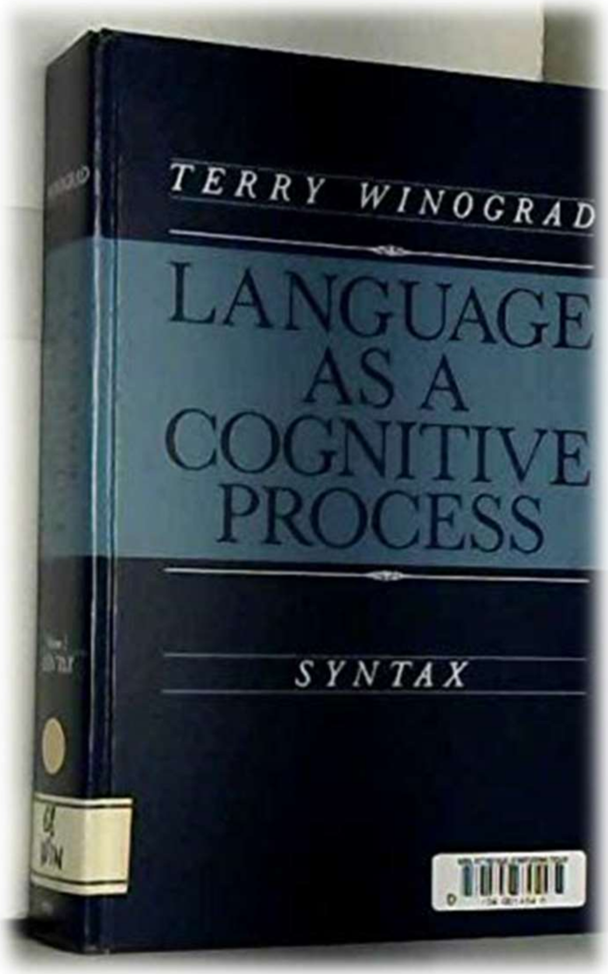

SHRDLU (1971)

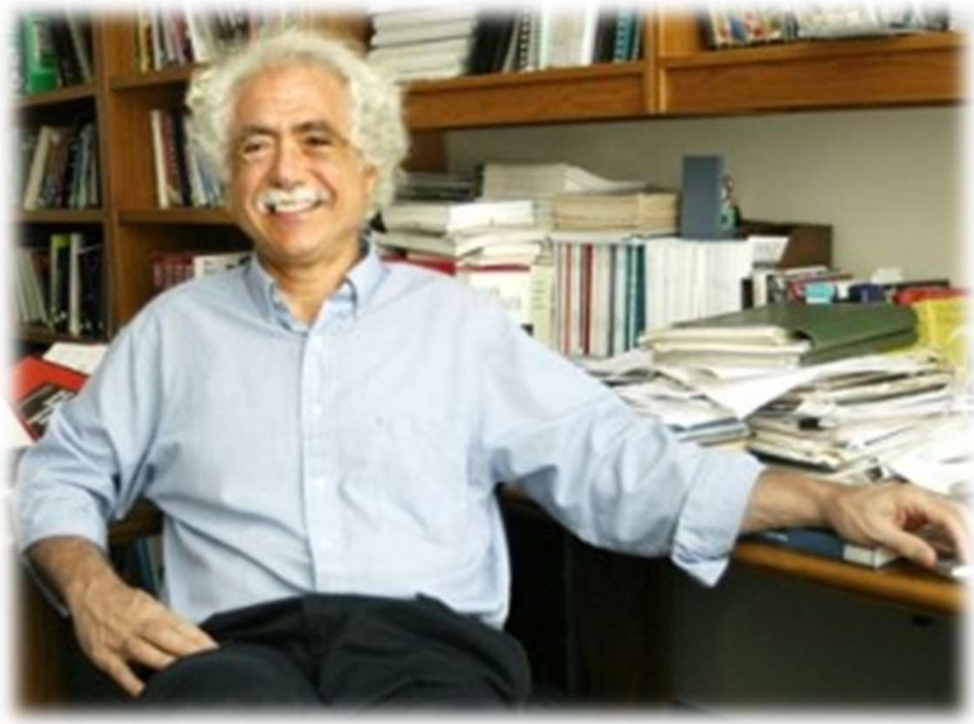

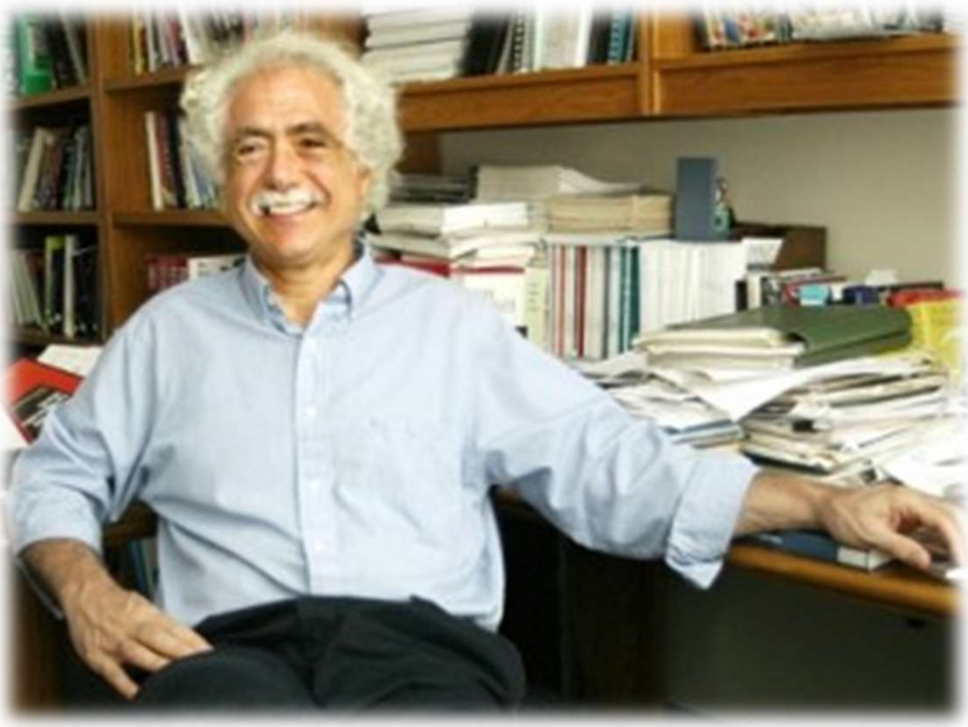

1997


Google

"… a fun analogy in trying to get some idea of what we're doing in trying to understand nature, is to imagine that the gods are playing some great game like chess… and you don't know the rules of the game, but you're allowed to look at the board, at least from time to time… and from these observations you try to figure out what the rules of the game are."

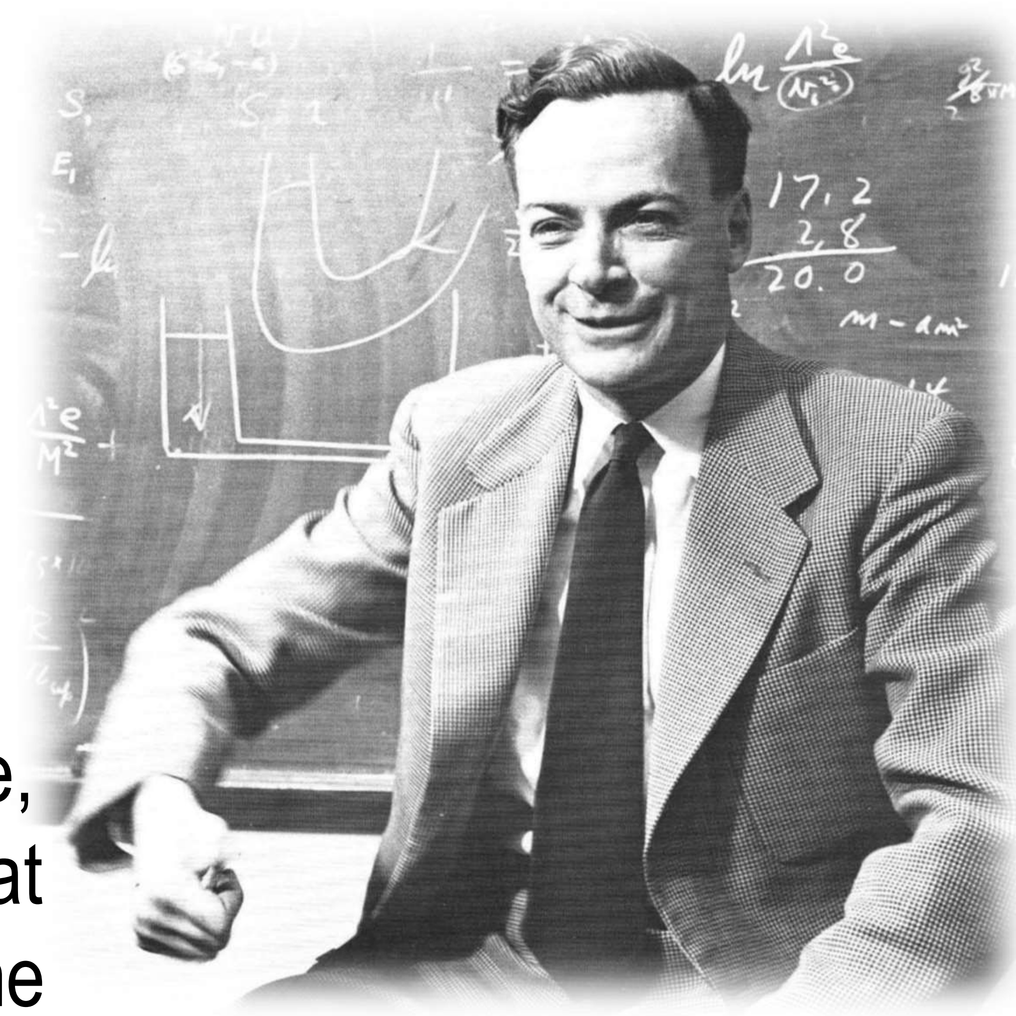

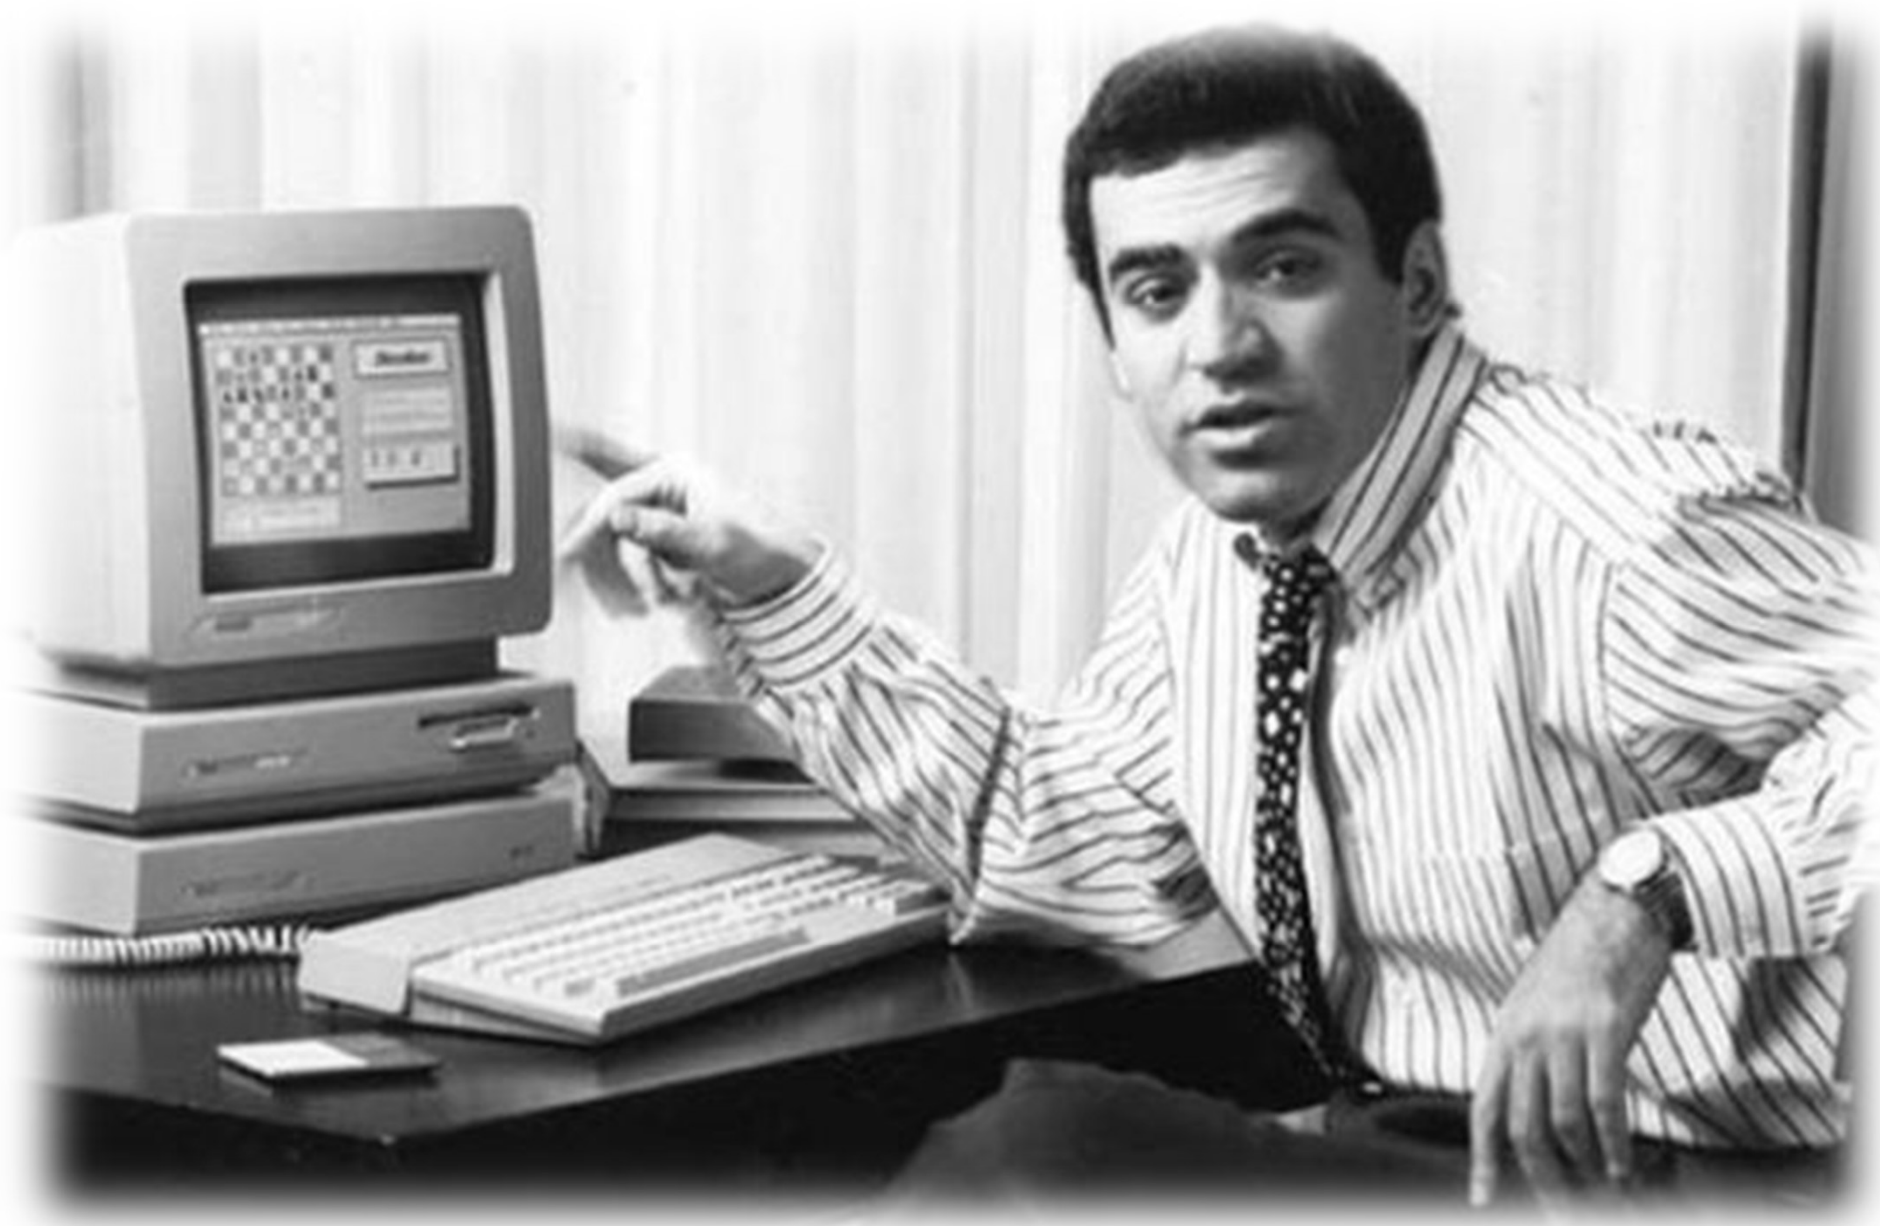

1992

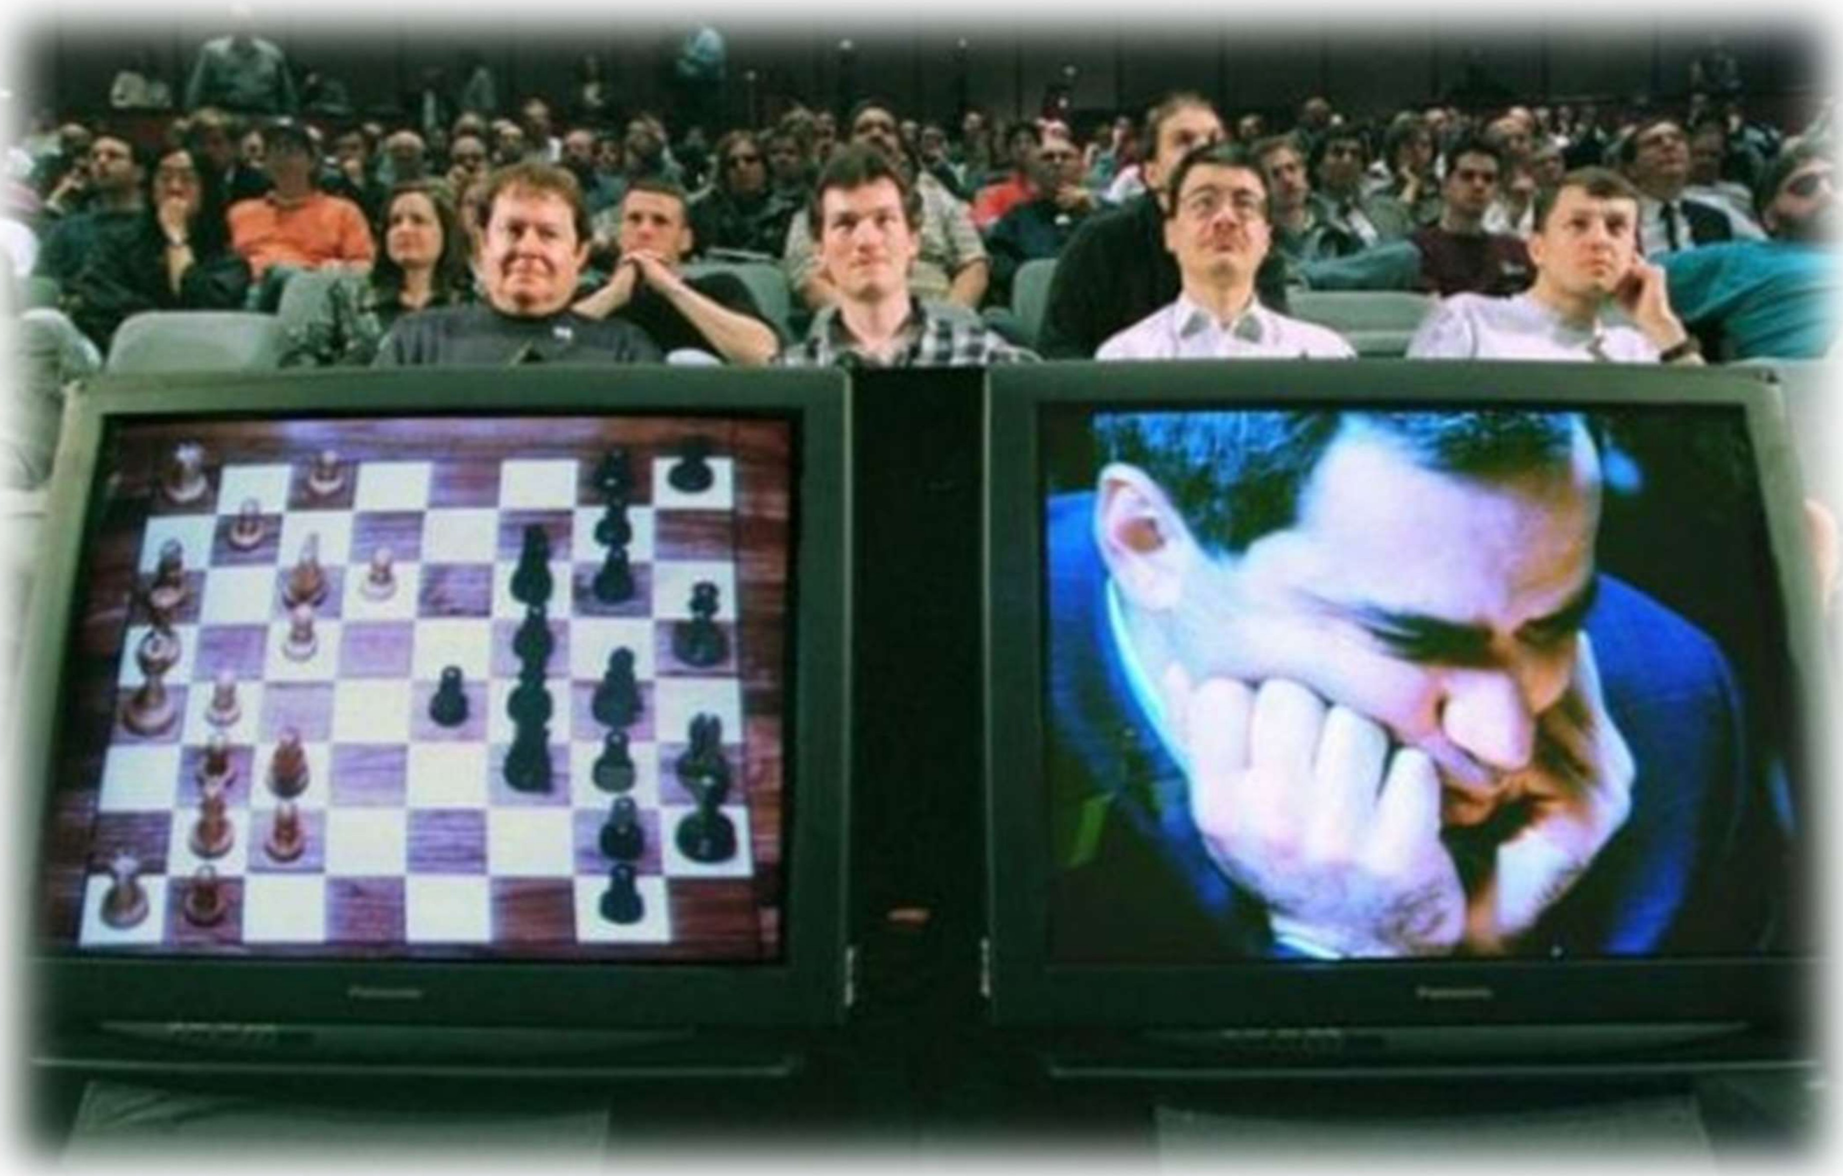

1996

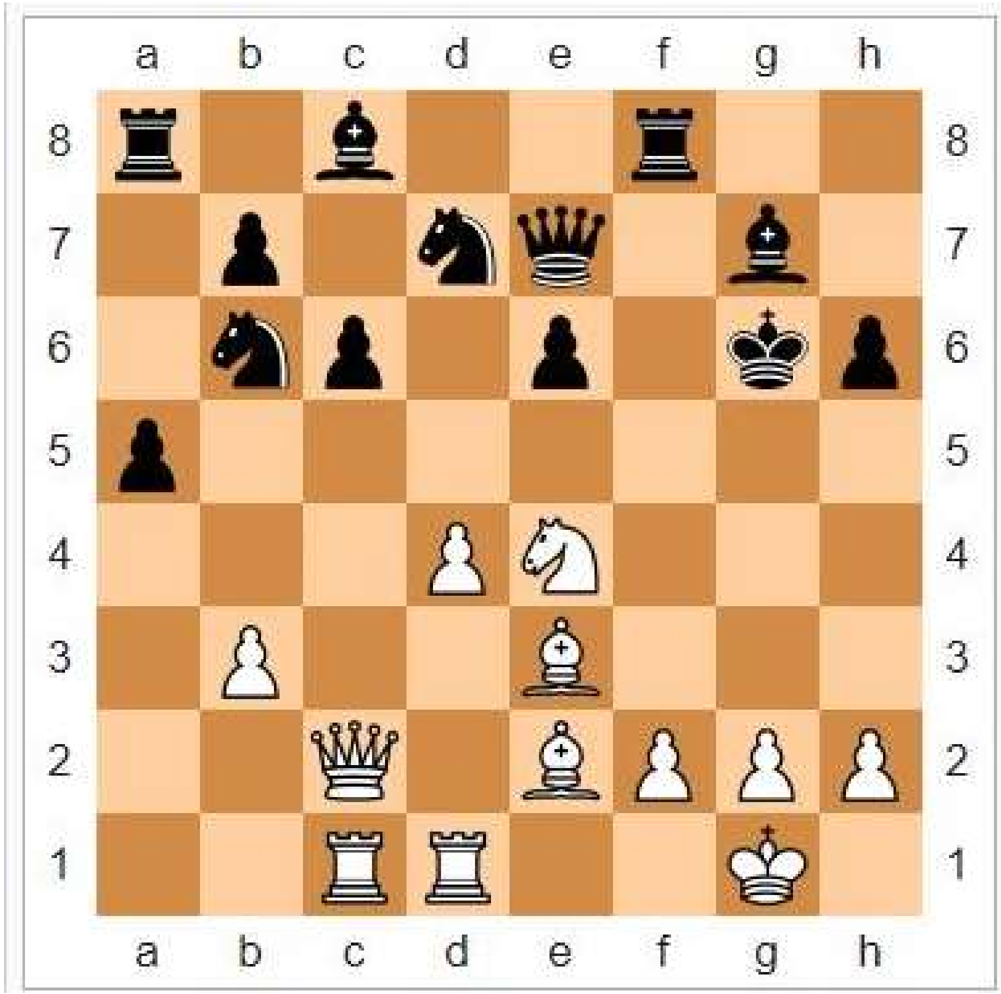
a b c d e f g h
8 7 6 5 4 3 2 1


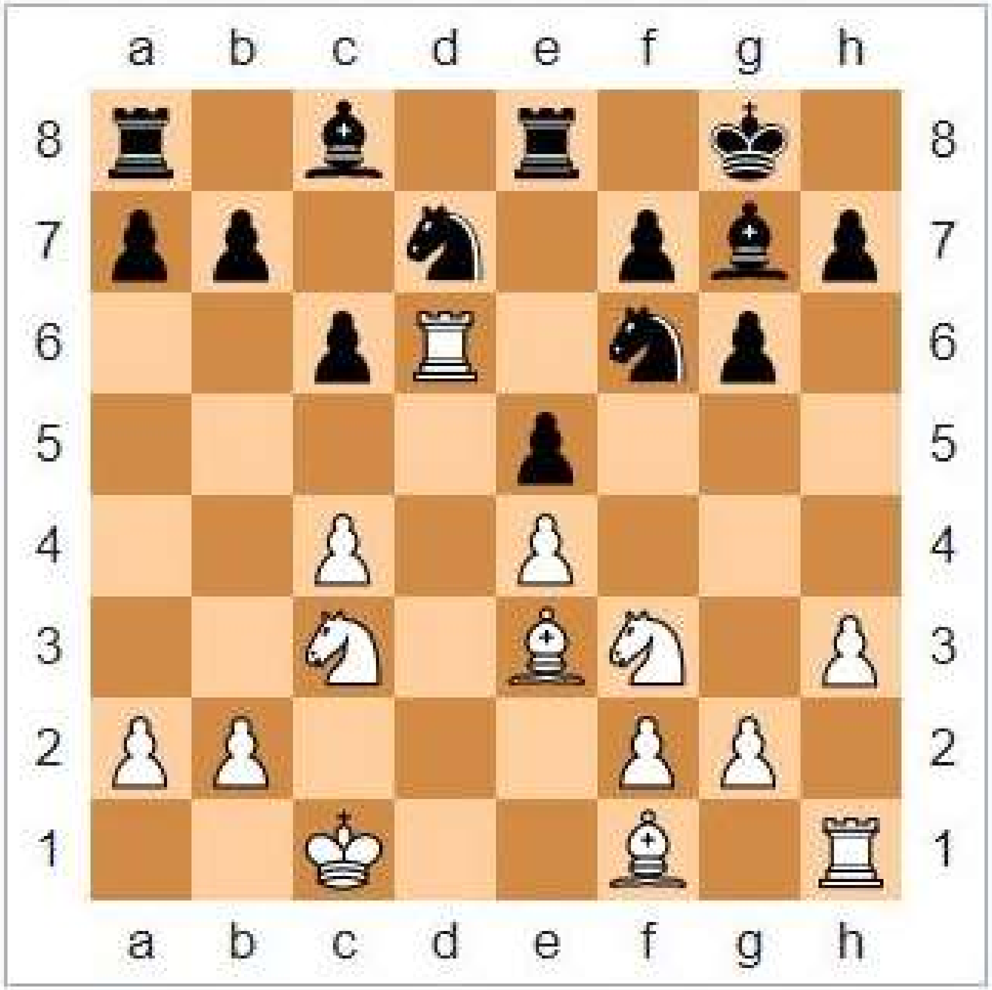
a b c d e f g h
8 7 6 5 4 3 2 1

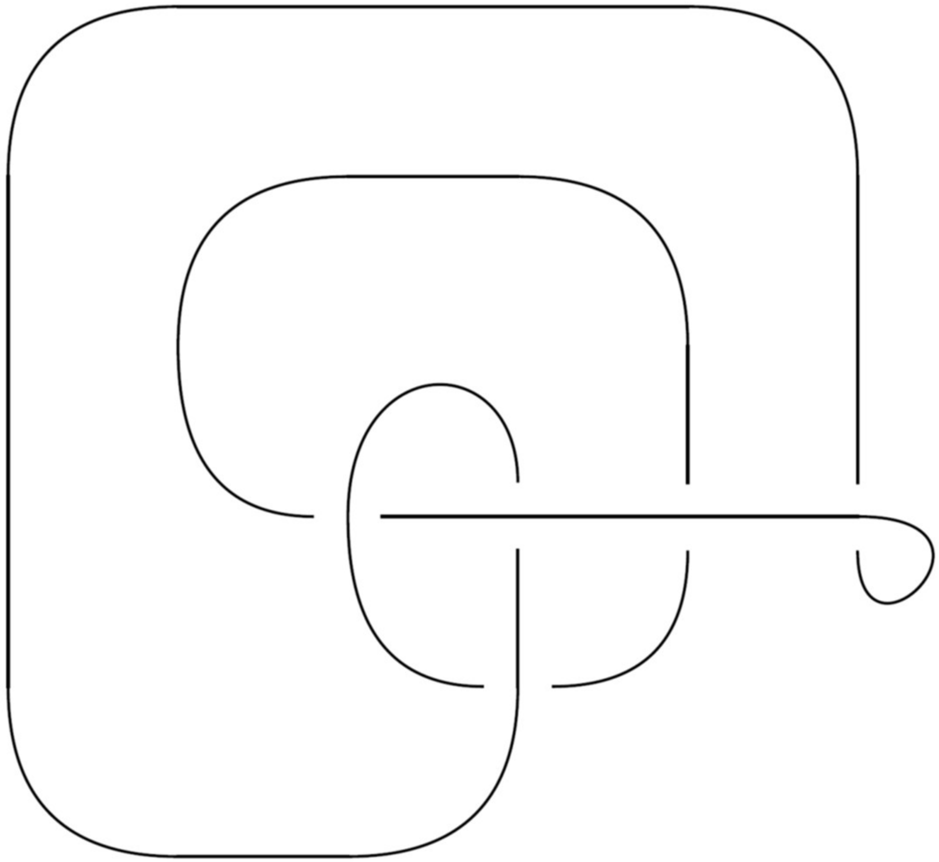

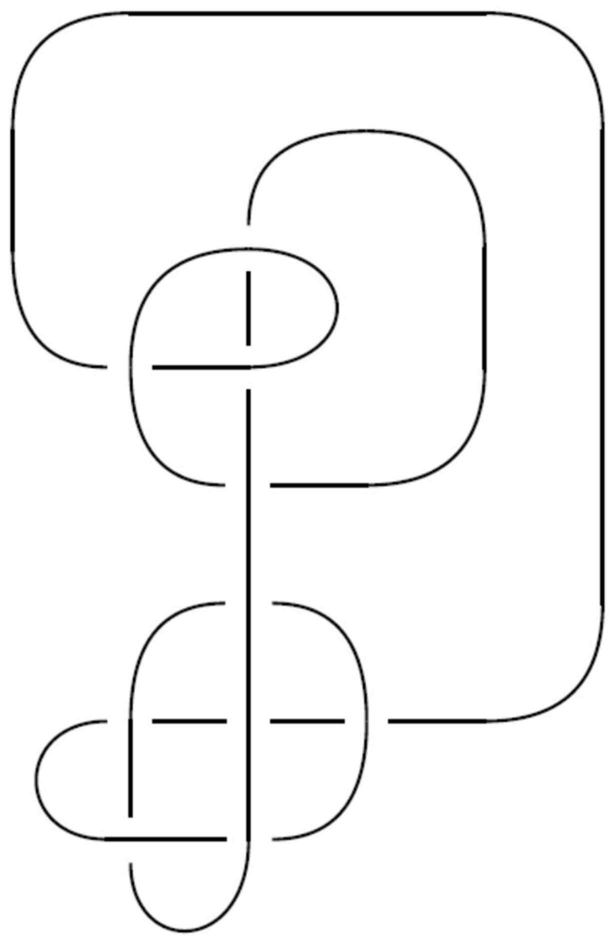

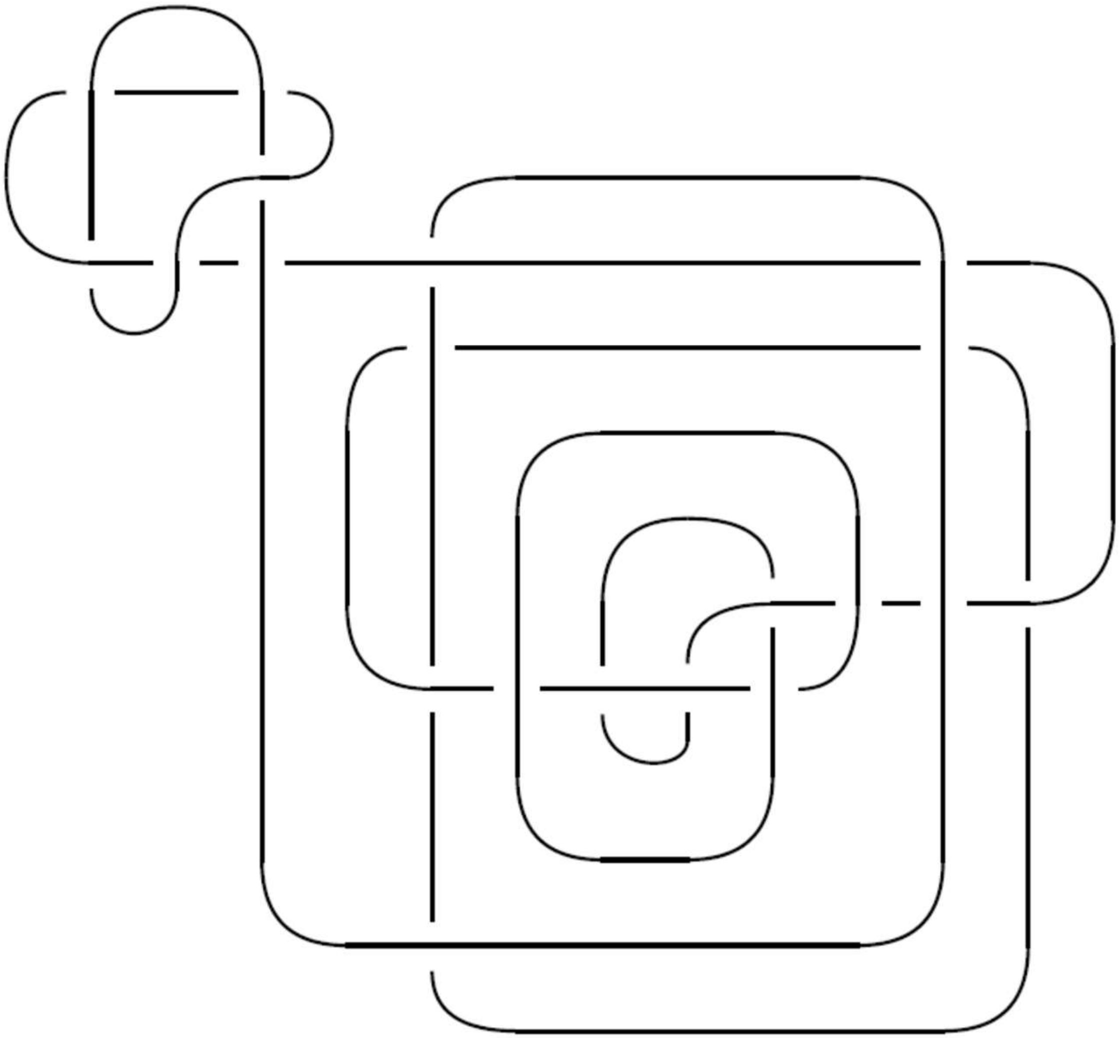

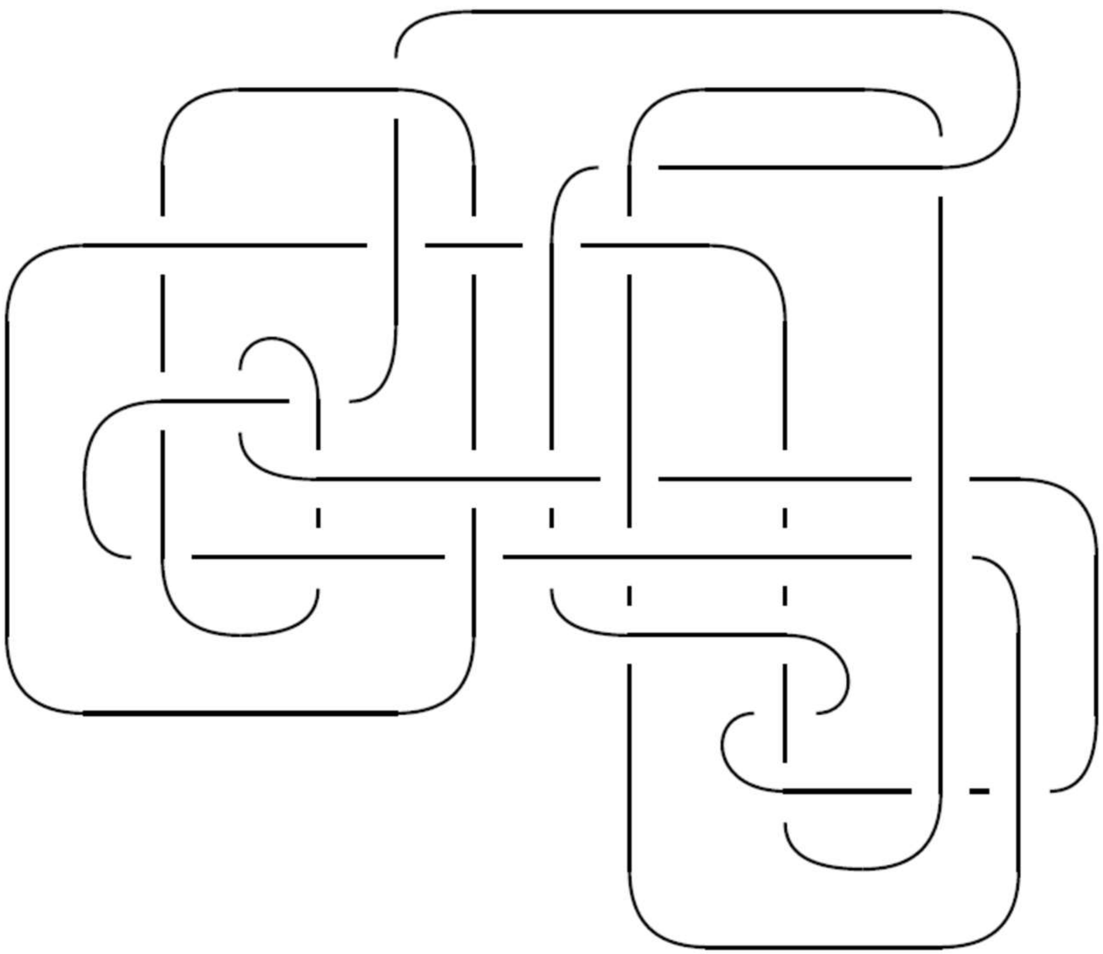

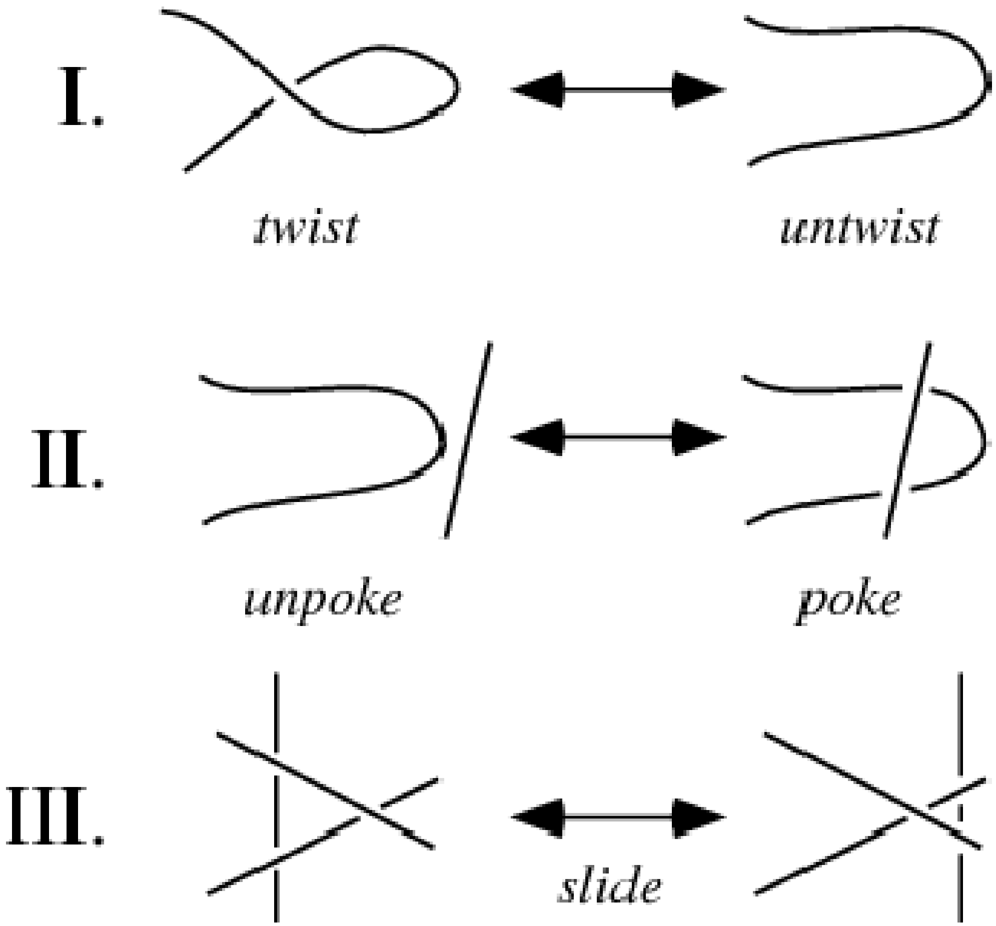

Kurt Reidemeister

Knotted ?

Knotted ?

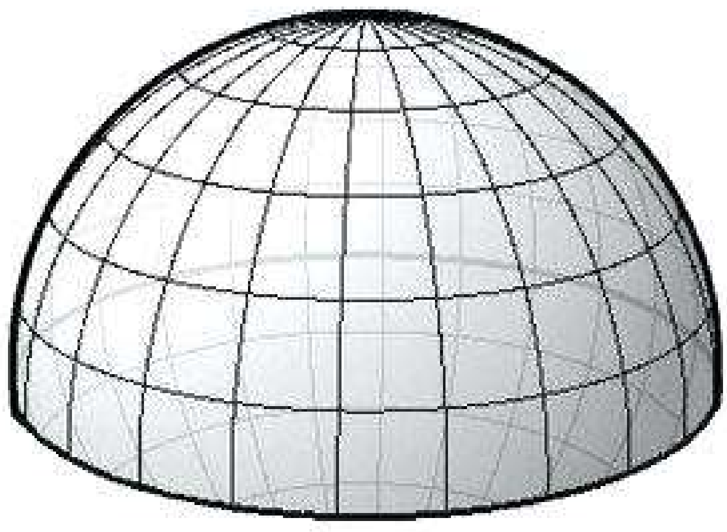

an exotic 4-ball has no smooth radius function with 3-sphere levels

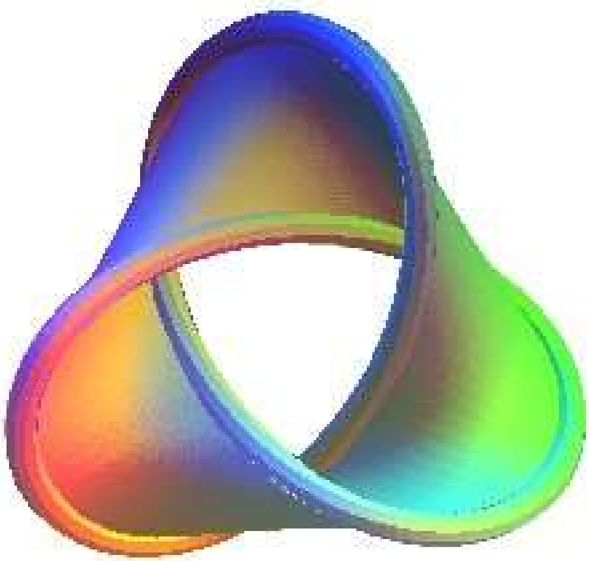
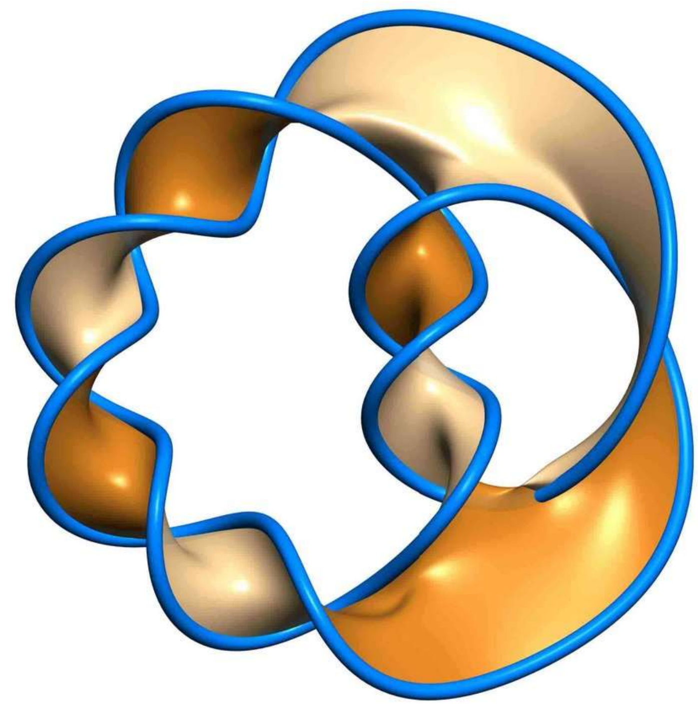

Man and machine thinking about the smooth 4-dimensional Poincaré conjecture.

MICHAEL FREEDMAN
ROBERT GOMPF
SCOTT MORRISON
KEVIN WALKER

$$\frac{1}{2}|s(K)| \leq \mathrm{genus}_4(K)$$

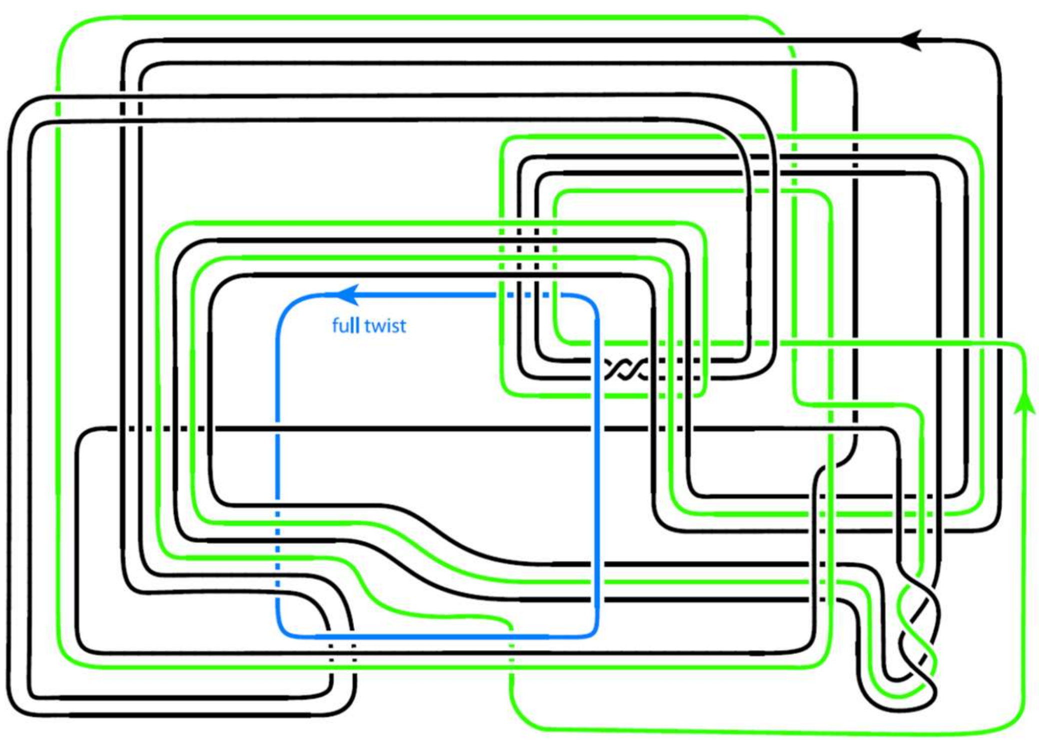

Knot type	$\text{genus}_4(K)$	Standard deviation
$11n_{34}$	0.496090381	0.303309781
$11n_{80}$	1.035585403	0.039595998
$12a_{153}$	1.069947992	0.081041807
$12a_{187}$	1.009274485	0.020159284
$12a_{230}$	1.019275038	0.029479206
$12a_{317}$	1.030372641	0.06259502
$12a_{450}$	1.009170415	0.013387401
$12a_{570}$	1.009285686	0.012420726
$12a_{624}$	1.027900245	0.042808674
$12a_{636}$	1.005604245	3.13×10^{-5}
$12a_{787}$	1.005587707	3.58×10^{-5}
$12a_{905}$	1.005527869	3.22×10^{-5}
$12a_{1189}$	1.07924728	0.092337155
$12a_{1208}$	1.023944944	0.028467024
$12n_{52}$	1.024184935	0.040119993
$12n_{63}$	1.010492662	0.066061094
$12n_{225}$	1.001823799	0.028759343
$12n_{239}$	1.016089445	0.025106408
$12n_{269}$	1.00551357	8.59×10^{-6}
$12n_{505}$	1.037368491	0.074247022
$12n_{512}$	1.027464205	0.049029925

Prediction for smooth slice genus

(2016)

Mark Hughes

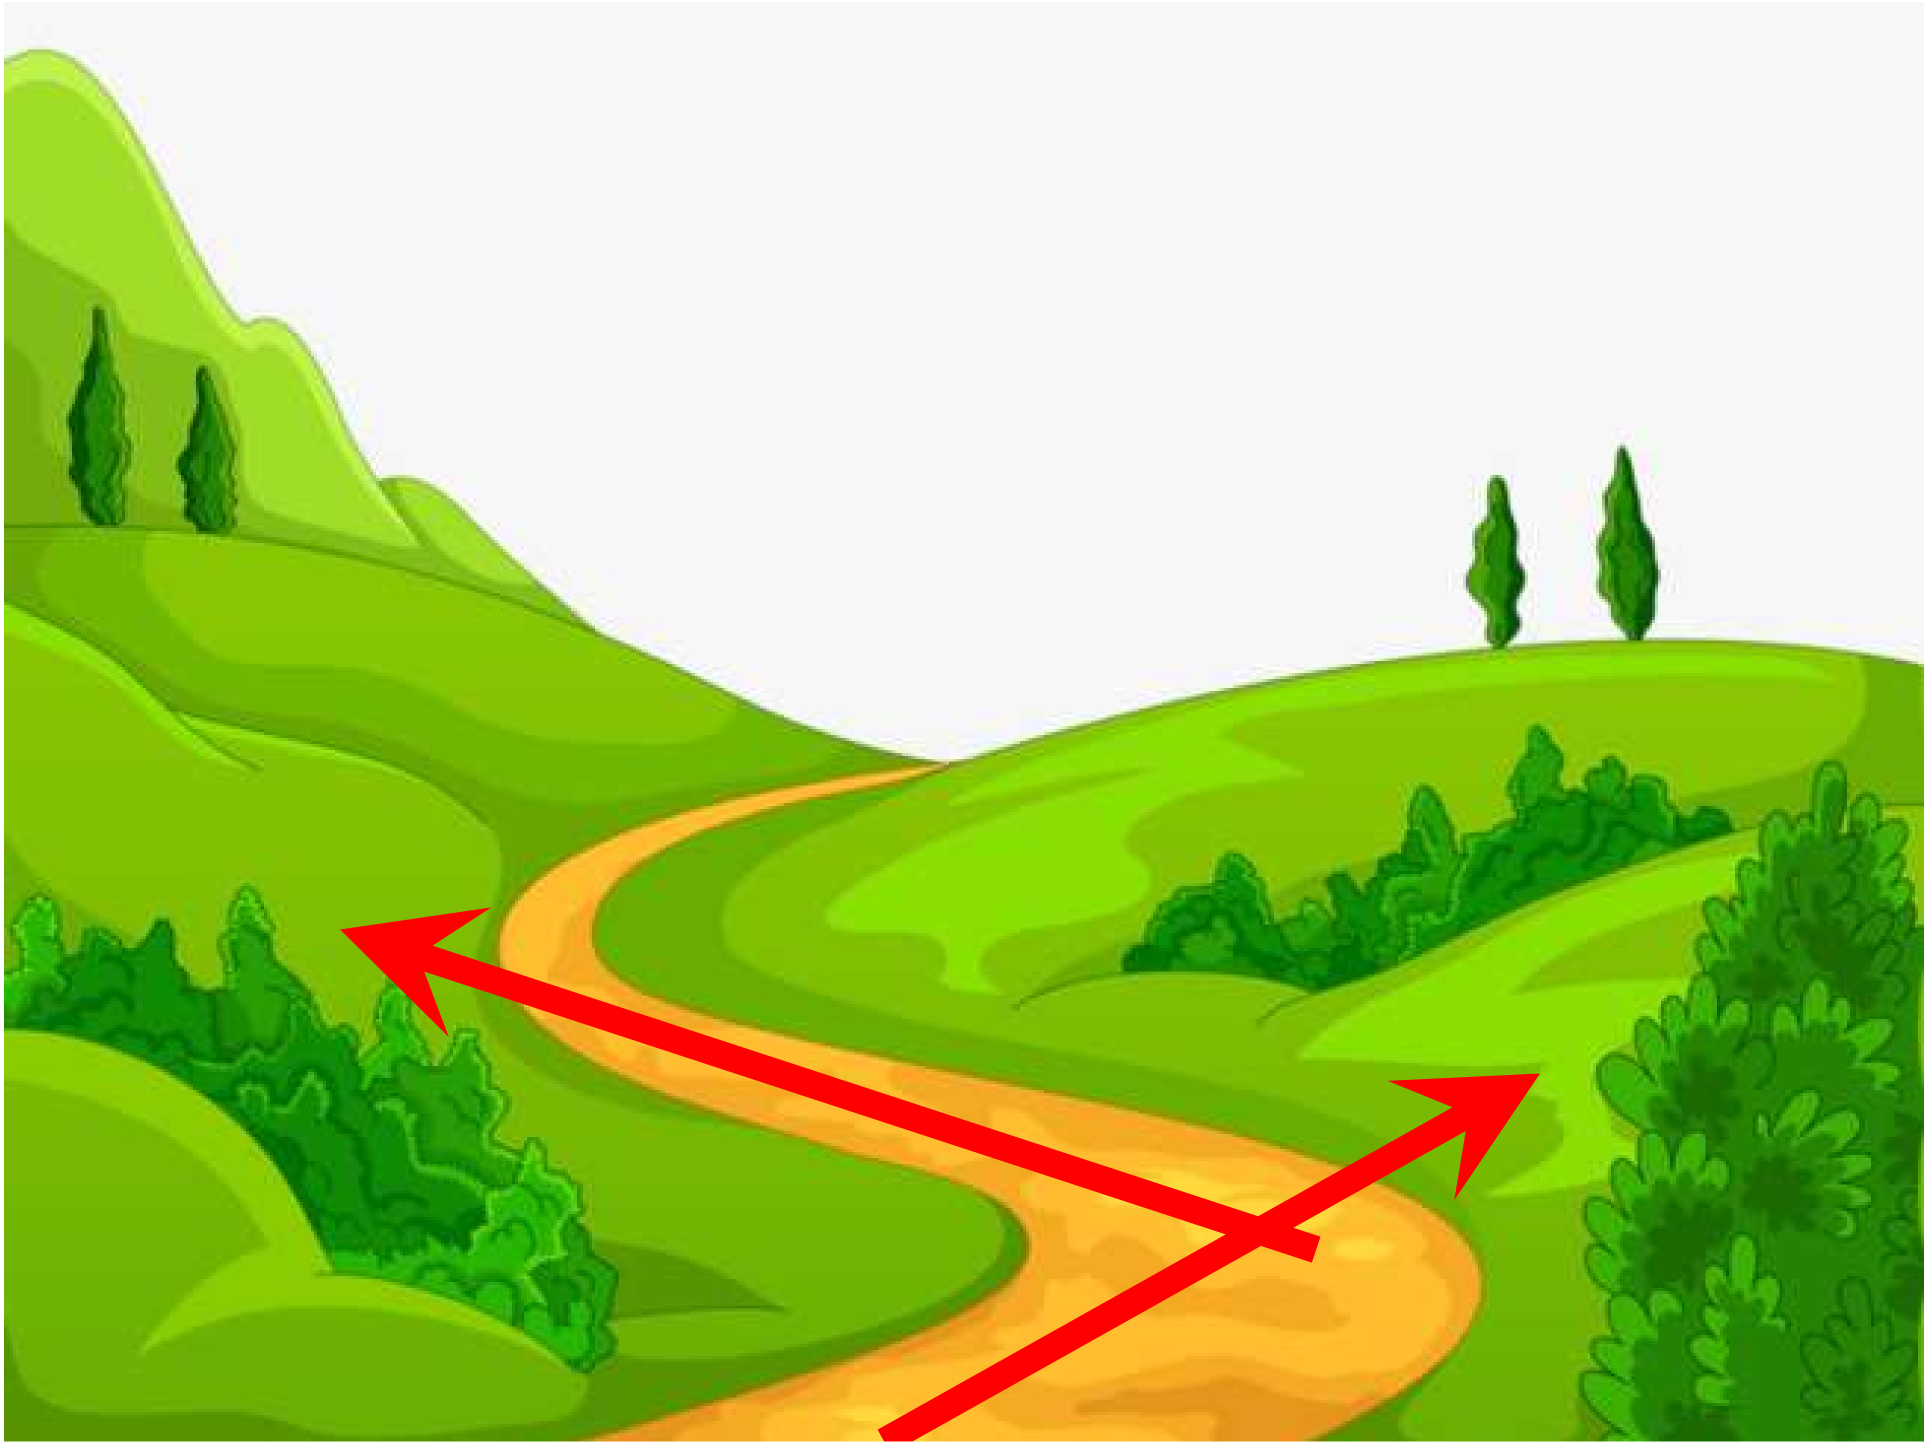

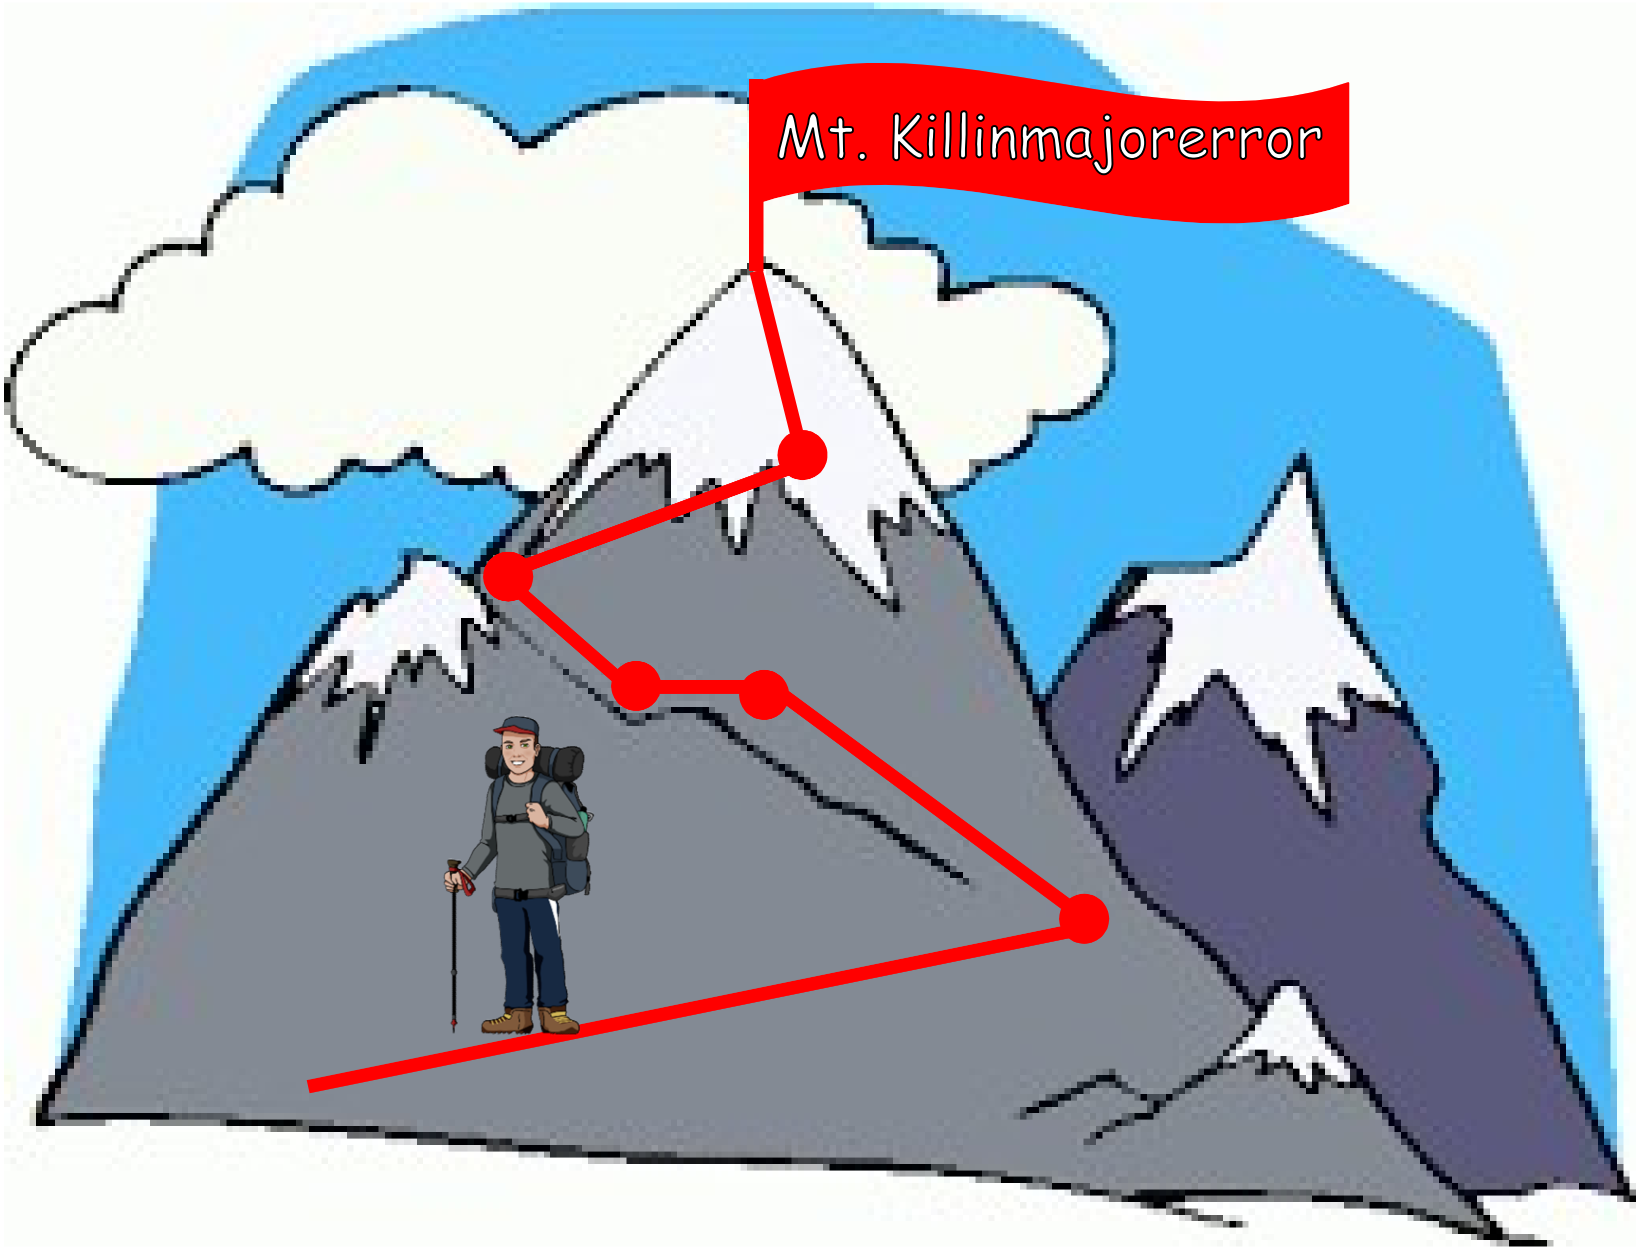
Mt. Killinmajorerror

“I grew up in France… I speak fluent ?__?”

"I grew up in France… I speak fluent ?__?"

Recurrent Neural Networks (RNNs)
Convolutional Neural Networks (CNNs)
Transformers
Reformers
⋮

Embedding: Vocabulary (Language) $\rightarrow \mathbb{R}^d$

$$\textit{Embedding: Vocabulary (Language)} \rightarrow \mathbb{R}^d$$

$$E\,(\text{king}) - E\,(\text{man}) + E\,(\text{woman}) \cong E\,(\text{queen})$$

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com
Noam Shazeer* Google Brain noam@google.com
Niki Parmar* Google Research nikip@google.com
Jakob Uszkoreit* Google Research usz@google.com
Llion Jones* Google Research llion@google.com
Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu
Łukasz Kaiser* Google Brain lukaszkaiser@google.com
Illia Polosukhin* ‡ illia.polosukhin@gmail.com

Abstract

The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

REFORMER: THE EFFICIENT TRANSFORMER

Nikita Kitaev* **Łukasz Kaiser*** **Anselm Levskaya**
U.C. Berkeley & Google Research Google Research Google Research
kitaev@cs.berkeley.edu {lukaszkaiser,levskaya}@google.com

ABSTRACT

Large Transformer models routinely achieve state-of-the-art results on a number of tasks but training these models can be prohibitively costly, especially on long sequences. We introduce two techniques to improve the efficiency of Transformers. For one, we replace dot-product attention by one that uses locality-sensitive hashing, changing its complexity from O(L^2) to O($L \log L$), where L is the length of the sequence. Furthermore, we use reversible residual layers instead of the standard residuals, which allows storing activations only once in the training process instead of N times, where N is the number of layers. The resulting model, the Reformer, performs on par with Transformer models while being much more memory-efficient and much faster on long sequences.

CALT-2020-046

CERN-TH-2020-179

Learning to Unknot

Sergei Gukov[1], James Halverson[2,3], Fabian Ruehle[4,5], Piotr Sułkowski[1,6]

Abstract

We introduce natural language processing into the study of knot theory, as made natural by the braid word representation of knots. We study the UNKNOT problem of determining whether or not a given knot is the unknot. After describing an algorithm to randomly generate N-crossing braids and their knot closures and discussing the induced prior on the distribution of knots, we apply binary classification to the UNKNOT decision problem. We find that the Reformer and shared-QK Transformer

Knots from Braids

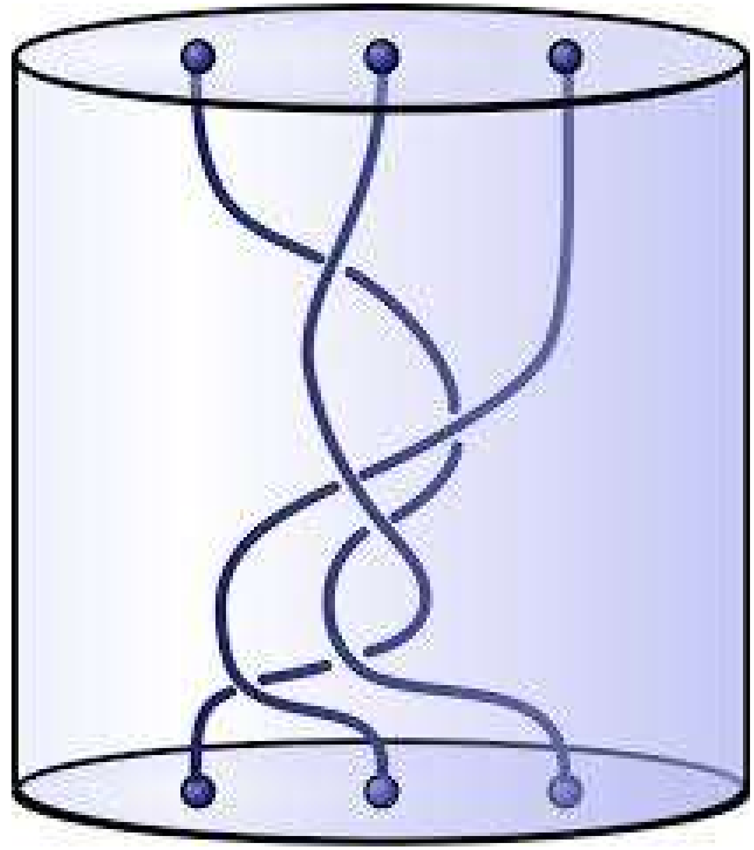

cf. “Dark Matter” of the Internet

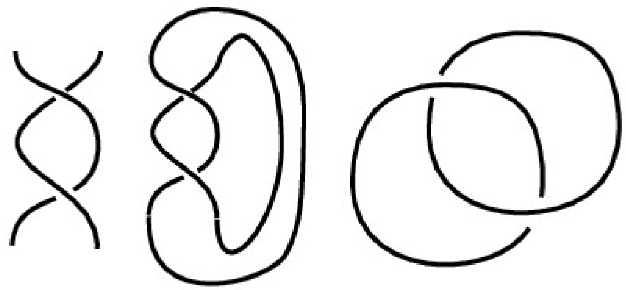

Hopf Link

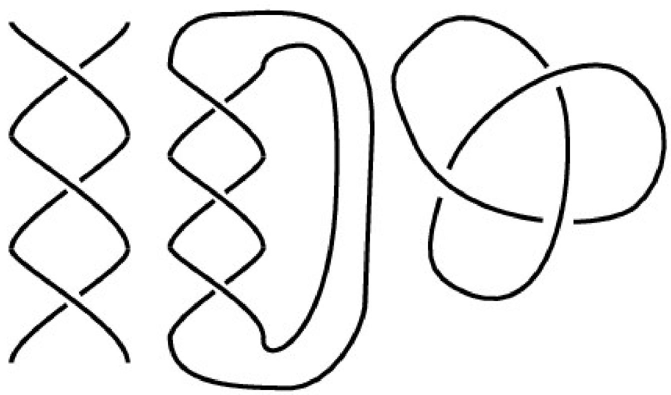

Trefoil Knot

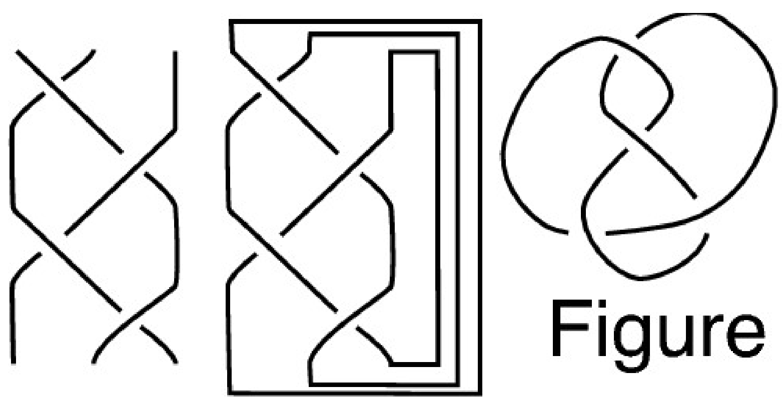

Figure Eight Knot

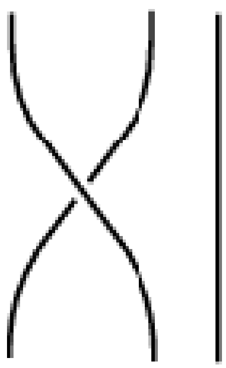

σ_1

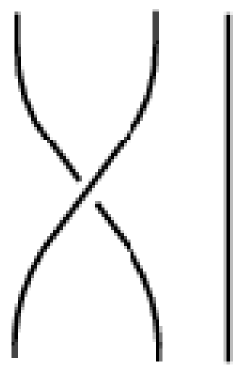

σ_1^{-1}

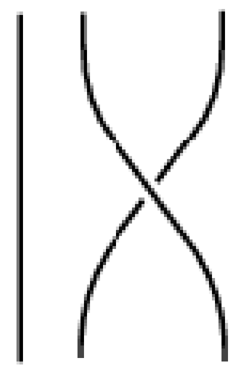

σ_2

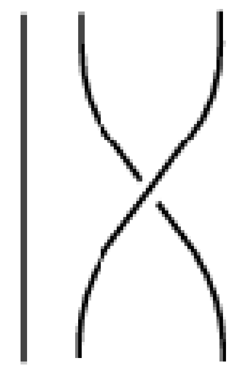

σ_2^{-1}

$$\sigma_i \sigma_j = \sigma_j \sigma_i \qquad \text{for } |i - j| > 1$$

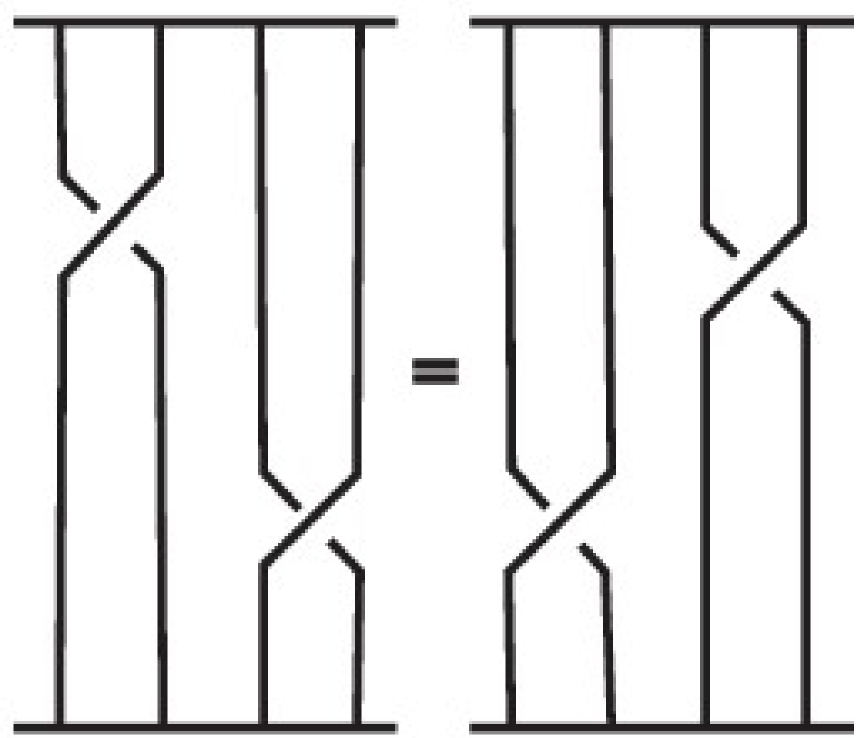

$$\sigma_i \sigma_{i+1} \sigma_i \;=\; \sigma_{i+1} \sigma_i \sigma_{i+1}$$

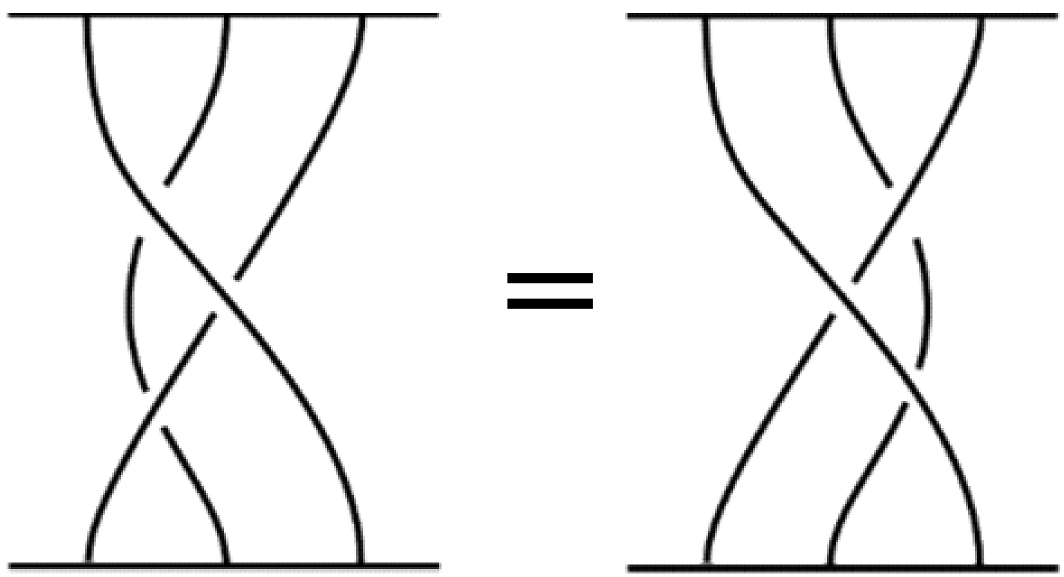

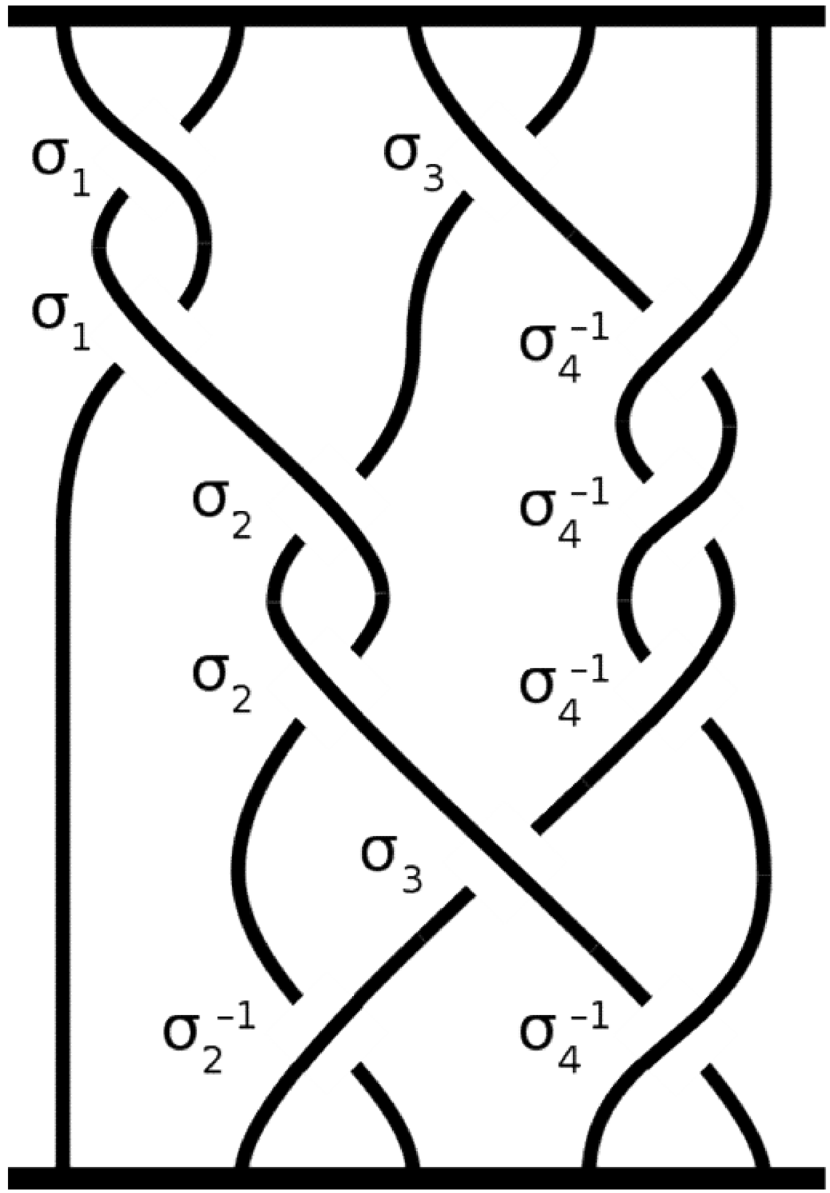
σ_1
σ_3
σ_1
σ_4^{-1}
σ_2
σ_4^{-1}
σ_2
σ_4^{-1}
σ_3
σ_2^{-1}
σ_4^{-1}

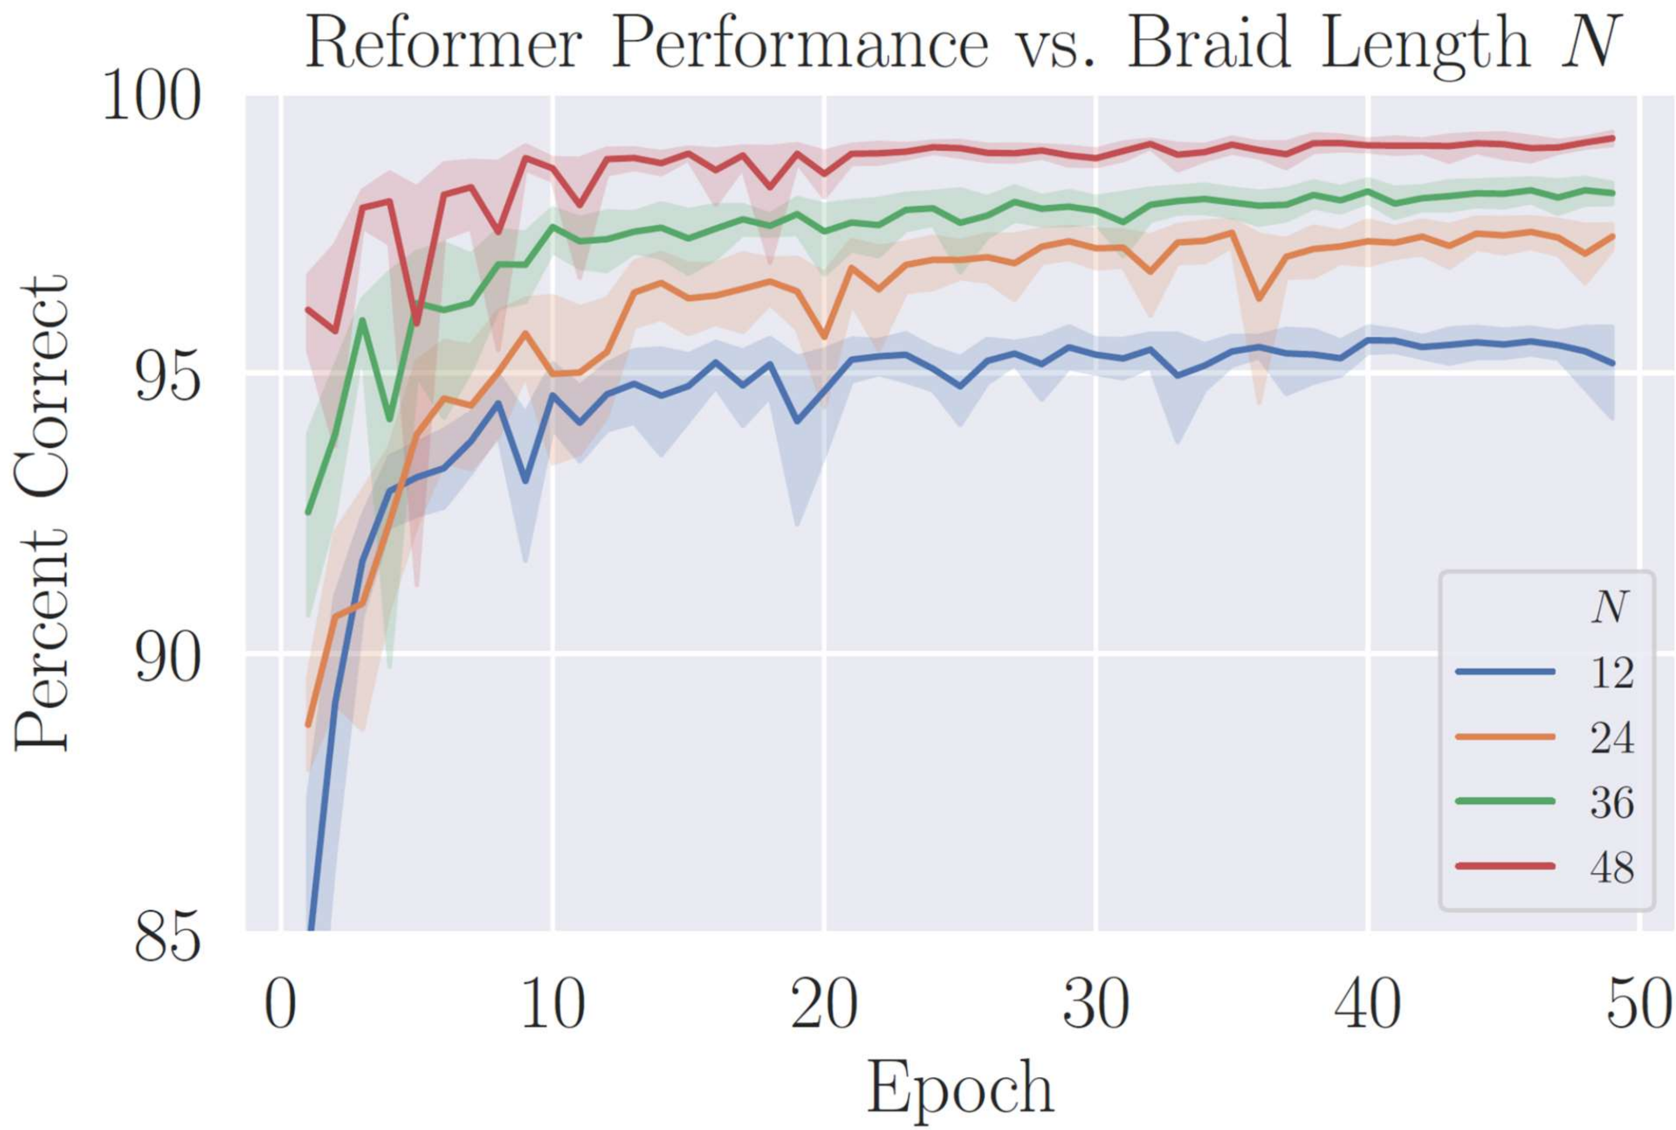

Reformer performance on UNKNOT as function of braid length. Performance increases with N.

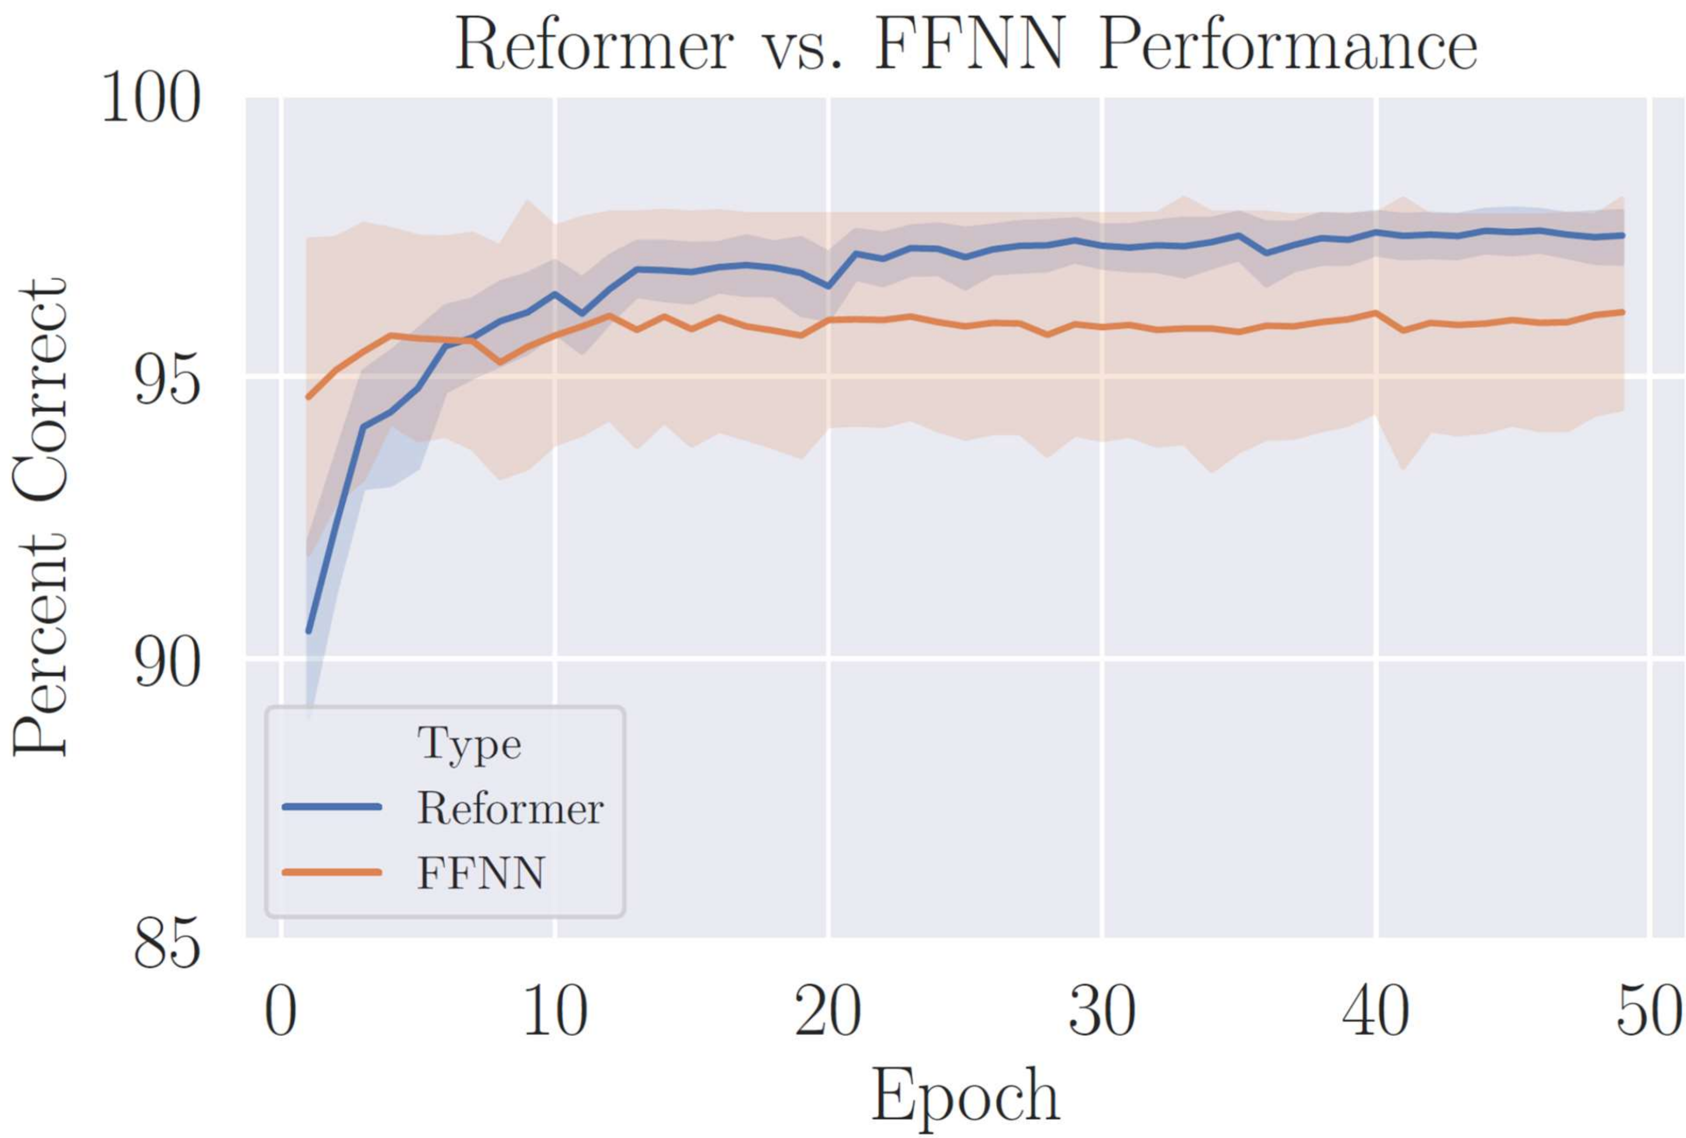

Performance comparison between reformer and feedforward network (FFNN).

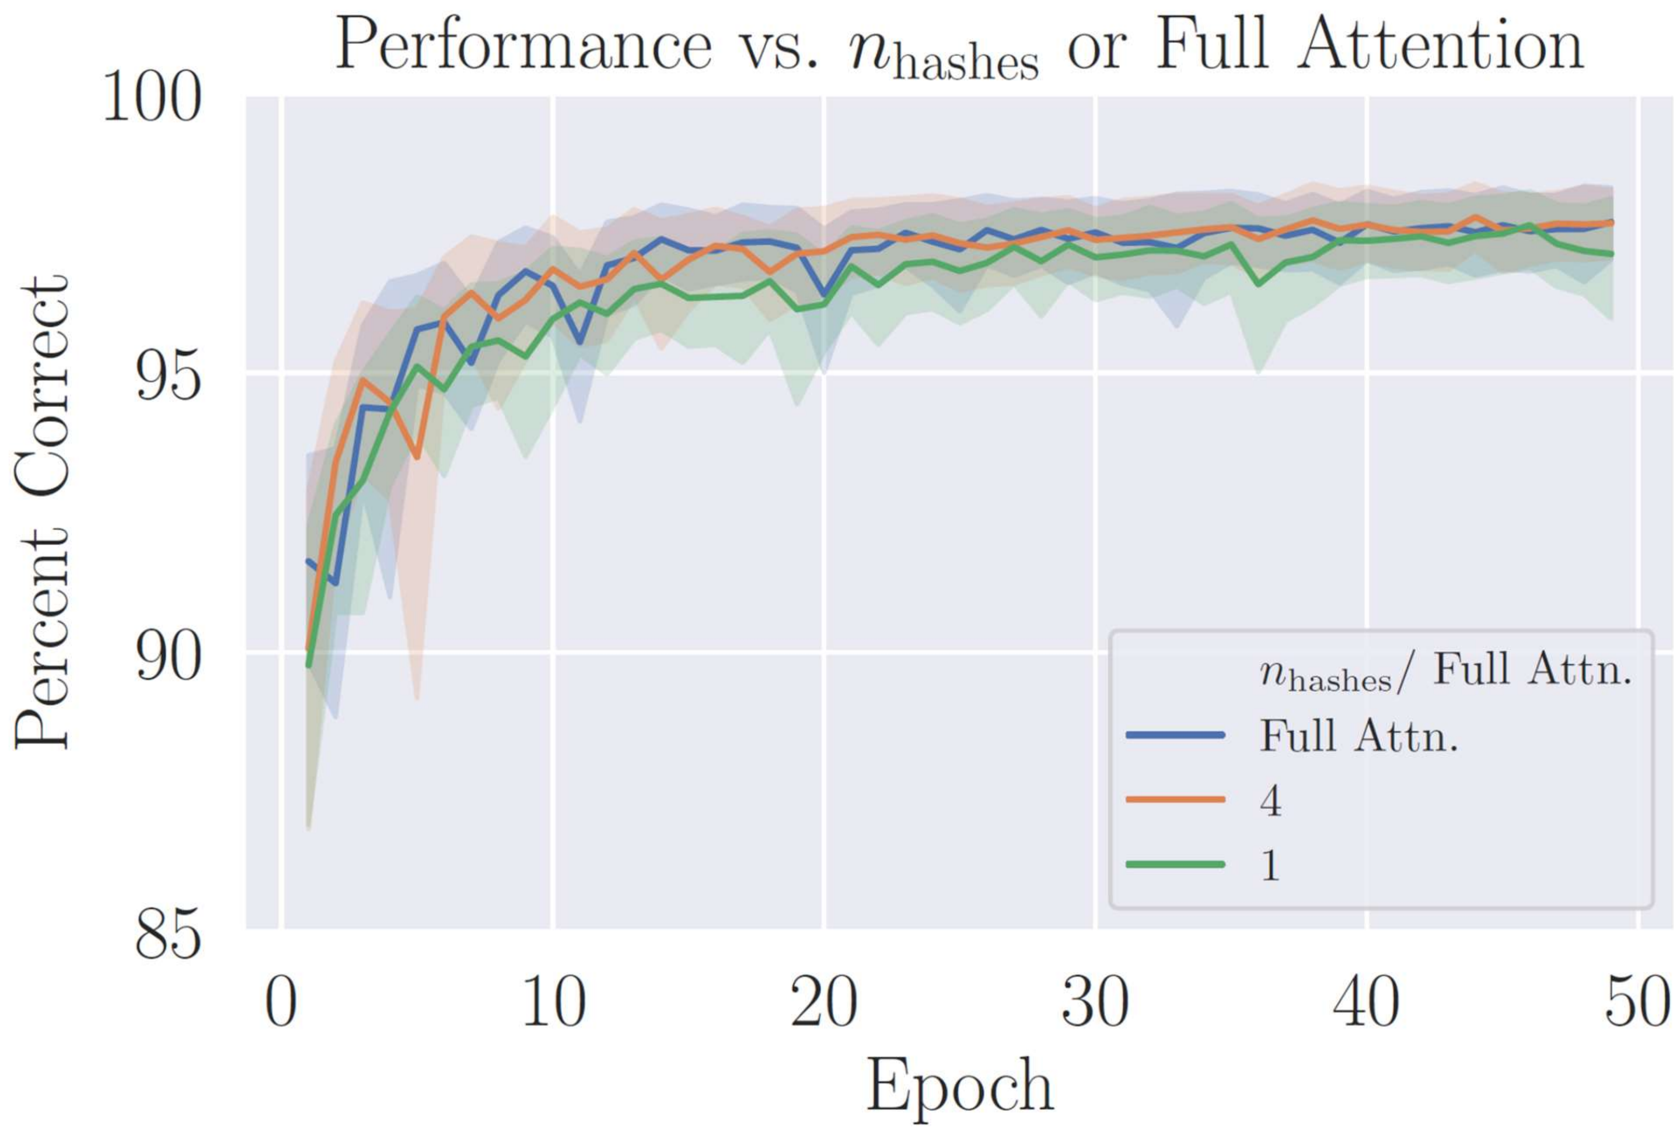

Performance dependence on the number of locality sensitive hashes (LSH).

The Computational Complexity of Knot and Link Problems

Joel Hass [*], Jeffrey C. Lagarias [†] and Nicholas Pippenger [‡]

February 1, 2008

Abstract

We consider the problem of deciding whether a polygonal knot in 3-dimensional Euclidean space is unknotted, capable of being continuously deformed without self-intersection so that it lies in a plane. We show that this problem, UNKNOTTING PROBLEM is in **NP**. We also consider the problem, UNKNOTTING PROBLEM of determining whether two or more such polygons can be split, or continuously deformed without self-intersection so that they occupy both sides of a plane without intersecting it. We show that it also is in **NP**. Finally, we show that the problem of determining the genus of a polygonal knot (a generalization of the problem of determining whether it is unknotted) is in **PSPACE**. We also give exponential worst-case running time bounds for deterministic algorithms to solve each of these problems. These algorithms are based on the use of normal surfaces and decision procedures due to W. Haken, with recent extensions by W. Jaco and J. L. Tollefson.

Knottedness is in NP, modulo GRH

Greg Kuperberg*
Department of Mathematics, University of California, Davis, CA 95616

Given a tame knot K presented in the form of a knot diagram, we show that the problem of determining whether K is knotted is in the complexity class NP, assuming the generalized Riemann hypothesis (GRH). In other words, there exists a polynomial-length certificate that can be verified in polynomial time to prove that K is non-trivial. GRH is not needed to believe the certificate, but only to find a short certificate. This result complements the result of Hass, Lagarias, and Pippenger that unknottedness is in NP. Our proof is a corollary of major results of others in algebraic geometry and geometric topology.

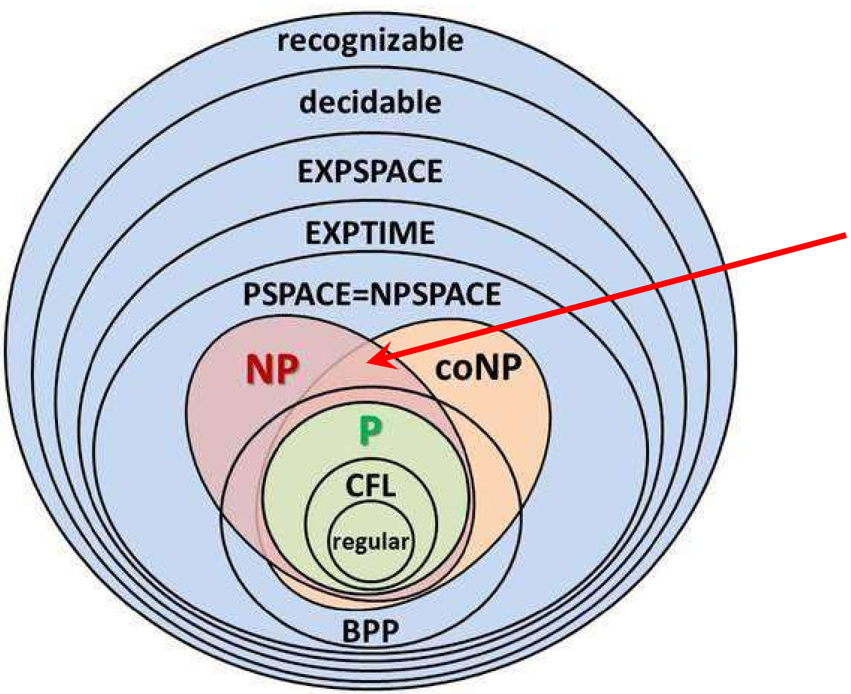

$$\text{Unknottedness} \in \mathsf{NP} \cap \mathsf{coNP}$$

integer $\stackrel{?}{=}$ product of two primes

:

THE EFFICIENT CERTIFICATION OF KNOTTEDNESS AND THURSTON NORM

MARC LACKENBY

1. Introduction

How difficult is it to determine whether a given knot is the unknot? The answer is not known. There might be a polynomial-time algorithm, but so far, this has remained elusive. The complexity of the unknot recognition problem was shown to be in NP by Hass, Lagarias and Pippenger [10]. The main aim of this article is to establish that it is in co-NP. This can be stated equivalently in terms of the Knottedness decision problem, which asks whether a given knot diagram represents a non-trivial knot.

Theorem 1.1. Knottedness *is in NP.*

In some sense, this result is not new. It was first announced by Agol [1] in 2002, but he has not provided a full published proof. In 2011, Kuperberg gave an alternative proof of Theorem 1.1, but under the extra assumption that the Generalised Riemann Hypothesis is true [19]. In this paper, we provide the first full proof of the unconditional result.

Combined with the theorem of Hass, Lagarias and Pippenger [10], Theorem 1.1 gives the following corollary.

Corollary 1.2. *If either of the decision problems* Unknot recognition *or* Knottedness *is NP-complete, then NP = co-NP.*

This is because if any decision problem in co-NP is NP-complete, then the complexity classes NP and co-NP must be equal. Since this is widely viewed not to be the case (see Section 2.4.3 in [7] for example), then it seems very unlikely that either of these decision problems is NP-complete.

The Unbearable Hardness of Unknotting*

Arnaud de Mesmay[1], Yo'av Rieck[2], Eric Sedgwick[3], and Martin Tancer[4]

Abstract

We prove that deciding if a diagram of the unknot can be untangled using at most k Riedemeister moves (where k is part of the input) is **NP**-hard. We also prove that several natural questions regarding links in the 3-sphere are **NP**-hard, including detecting whether a link contains a trivial sublink with n components, computing the unlinking number of a link, and computing a variety of link invariants related to four-dimensional topology (such as the 4-ball Euler characteristic, the slicing number, and the 4-dimensional clasp number).

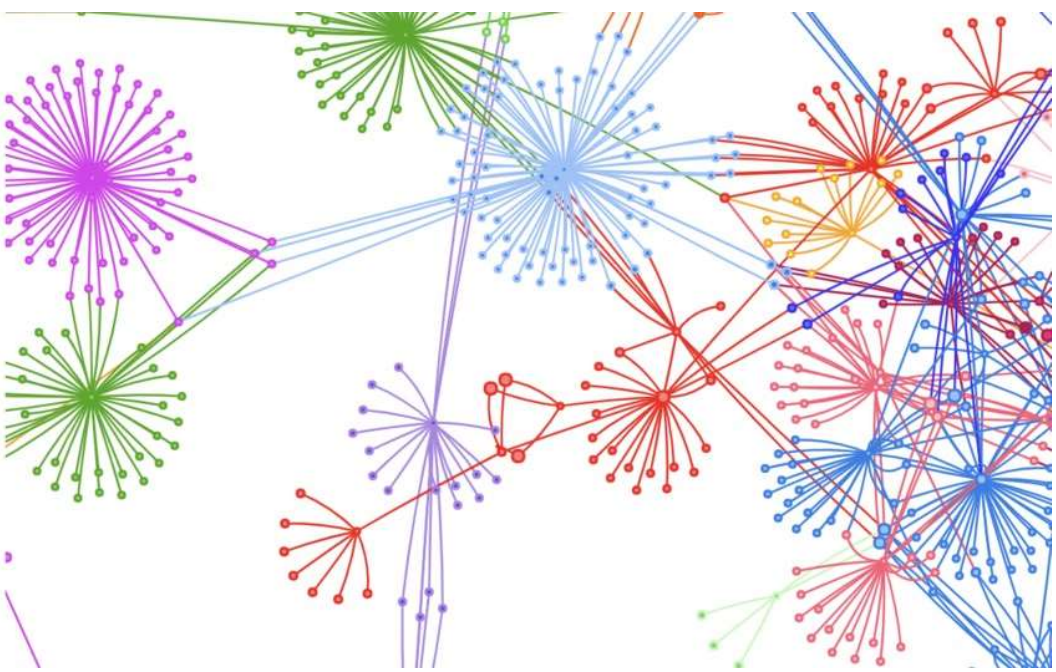

cf. connected components of a graph:

- Not finite
- Not explicitly presented

Coloring invariants of knots and links are often intractable

Greg Kuperberg*
University of California, Davis

Eric Samperton†
University of California, Santa Barbara

(Dated: July 16, 2019)

Let G be a nonabelian, simple group with a nontrivial conjugacy class $C \subseteq G$. Let K be a diagram of an oriented knot in S^3, thought of as computational input. We show that for each such G and C, the problem of counting homomorphisms $\pi_1(S^3 \setminus K) \to G$ that send meridians of K to C is almost parsimoniously #P-complete. This work is a sequel to a previous result by the authors that counting homomorphisms from fundamental groups of integer homology 3-spheres to G is almost parsimoniously #P-complete. Where we previously used mapping class groups actions on closed, unmarked surfaces, we now use braid group actions.

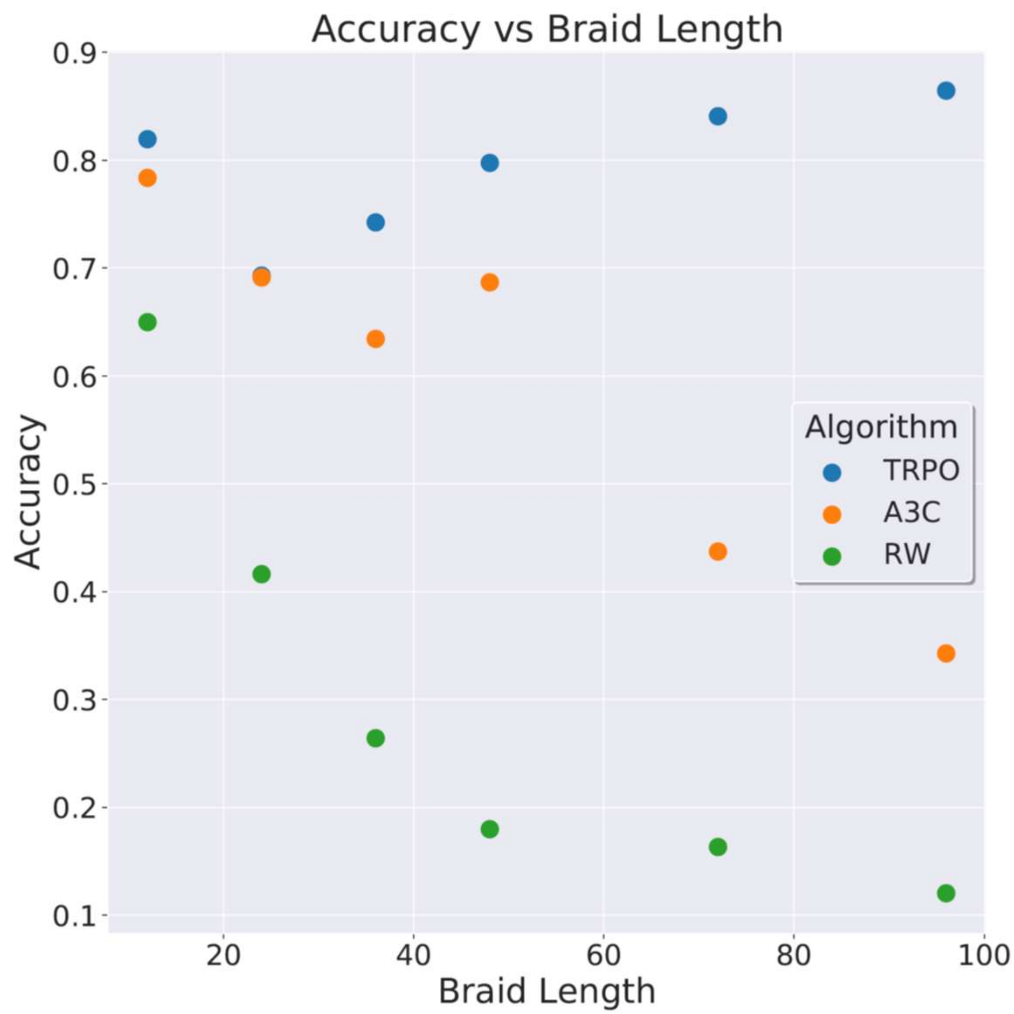

Fraction of unknots whose braid words could be reduced to the empty braid word as a function of initial braid word length.

 

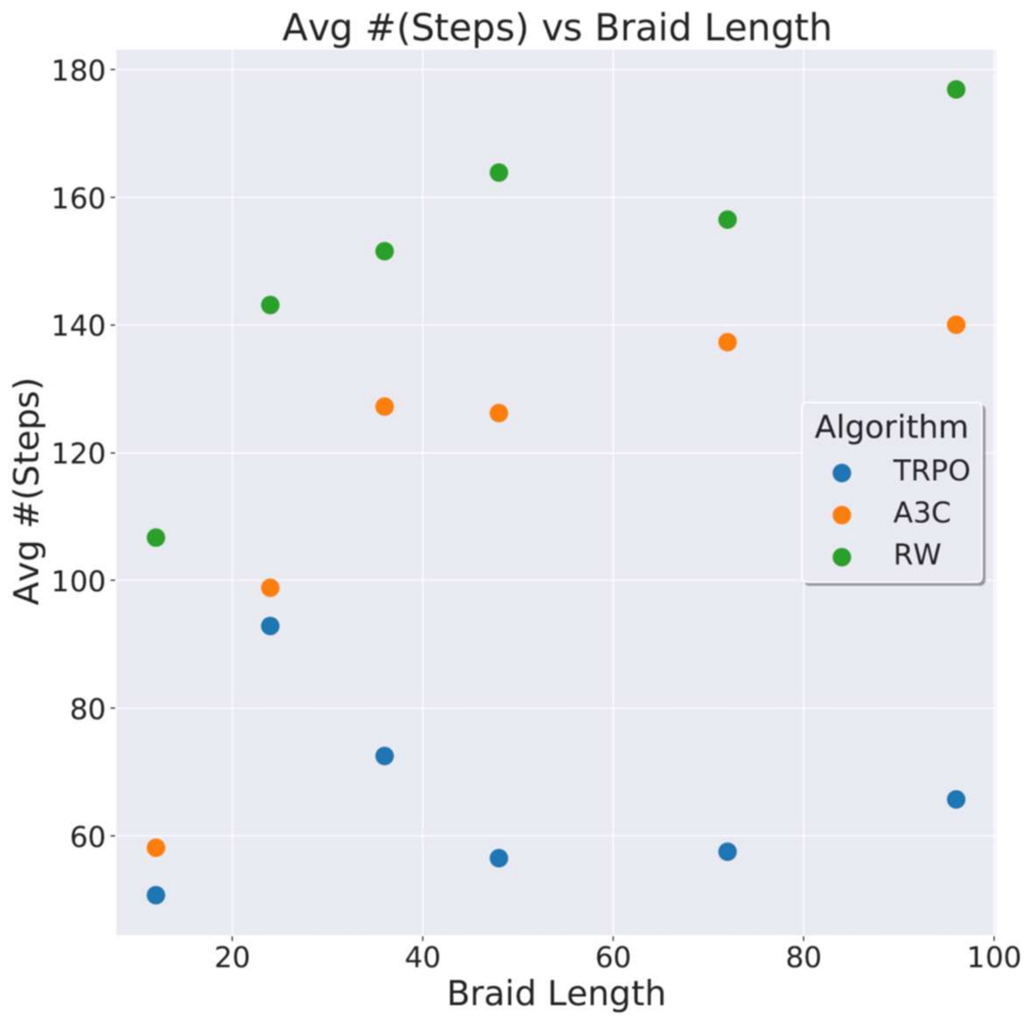

Average number of actions necessary to reduce the input braid word to the empty braid word as a function of initial braid word length.

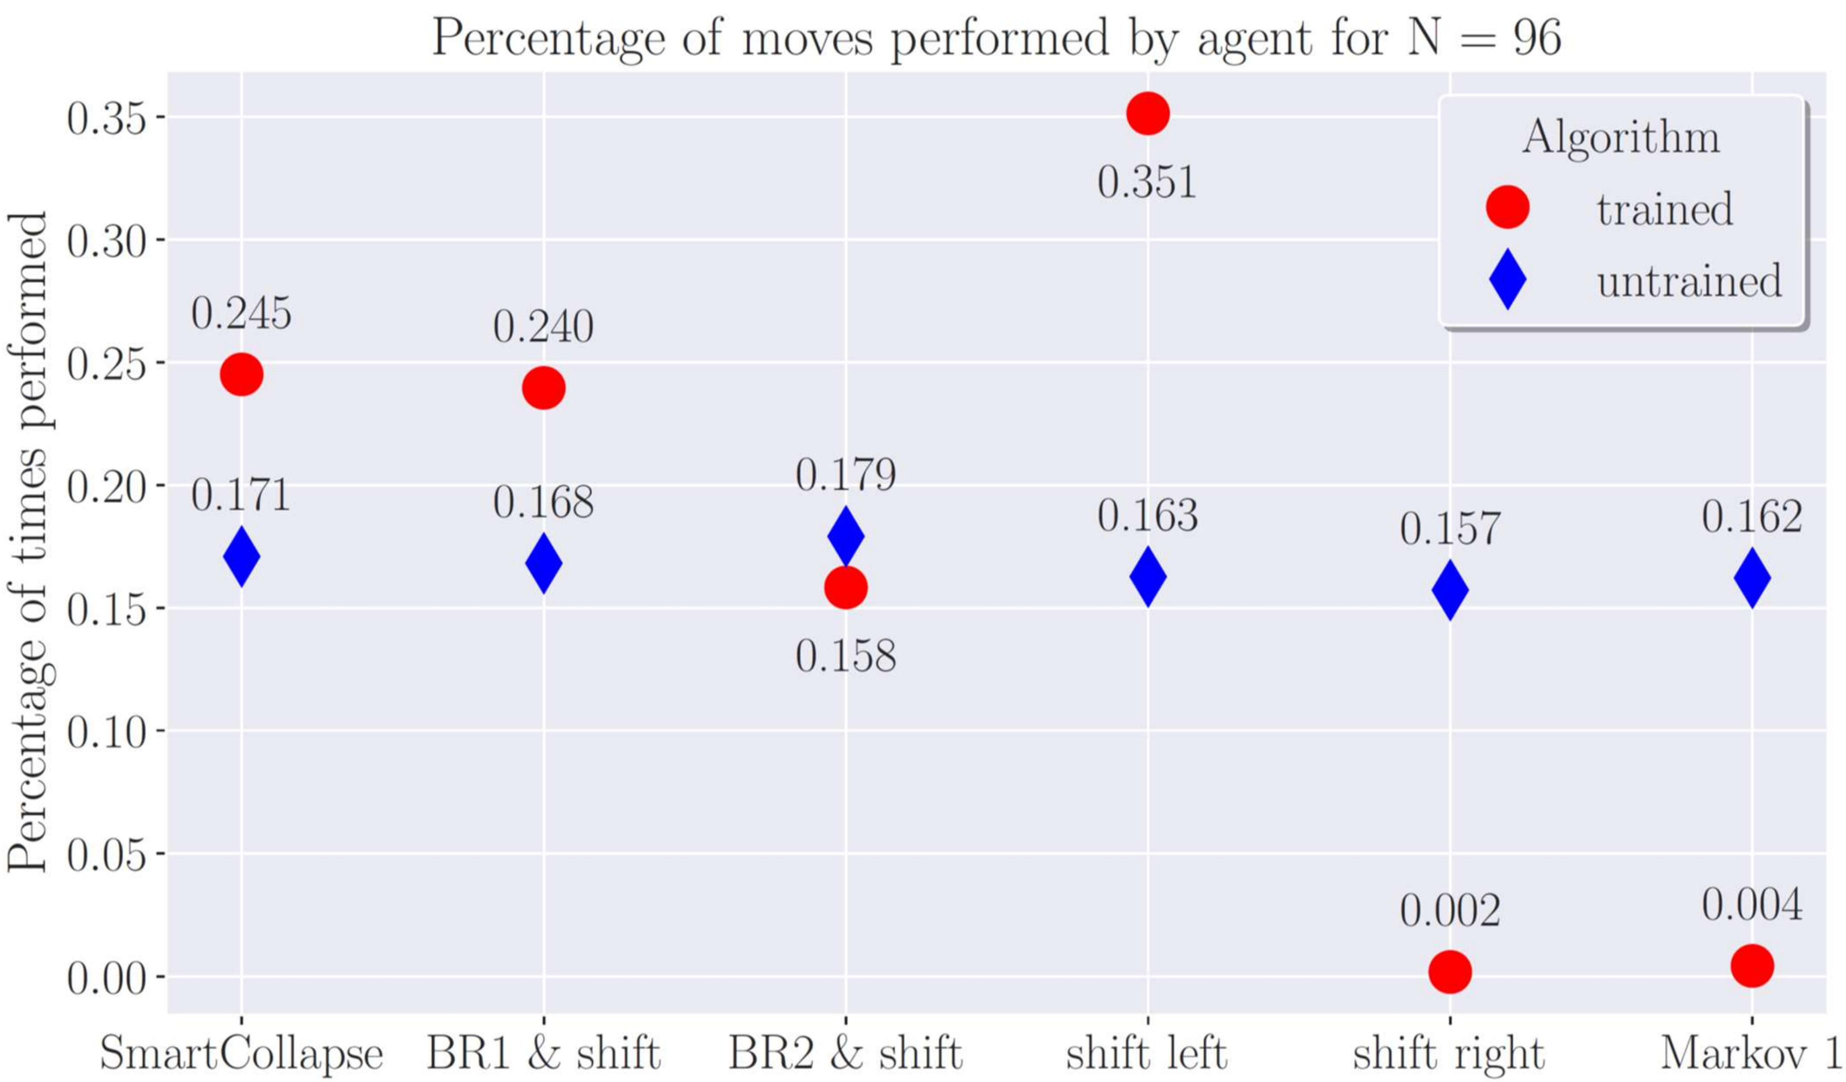
Percentage of moves performed by agent for N = 96
Percentage of times performed
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
Algorithm
trained
untrained
0.245
0.171
0.240
0.168
0.179
0.158
0.351
0.163
0.157
0.002
0.162
0.004
SmartCollapse
BR1 & shift
BR2 & shift
shift left
shift right
Markov 1

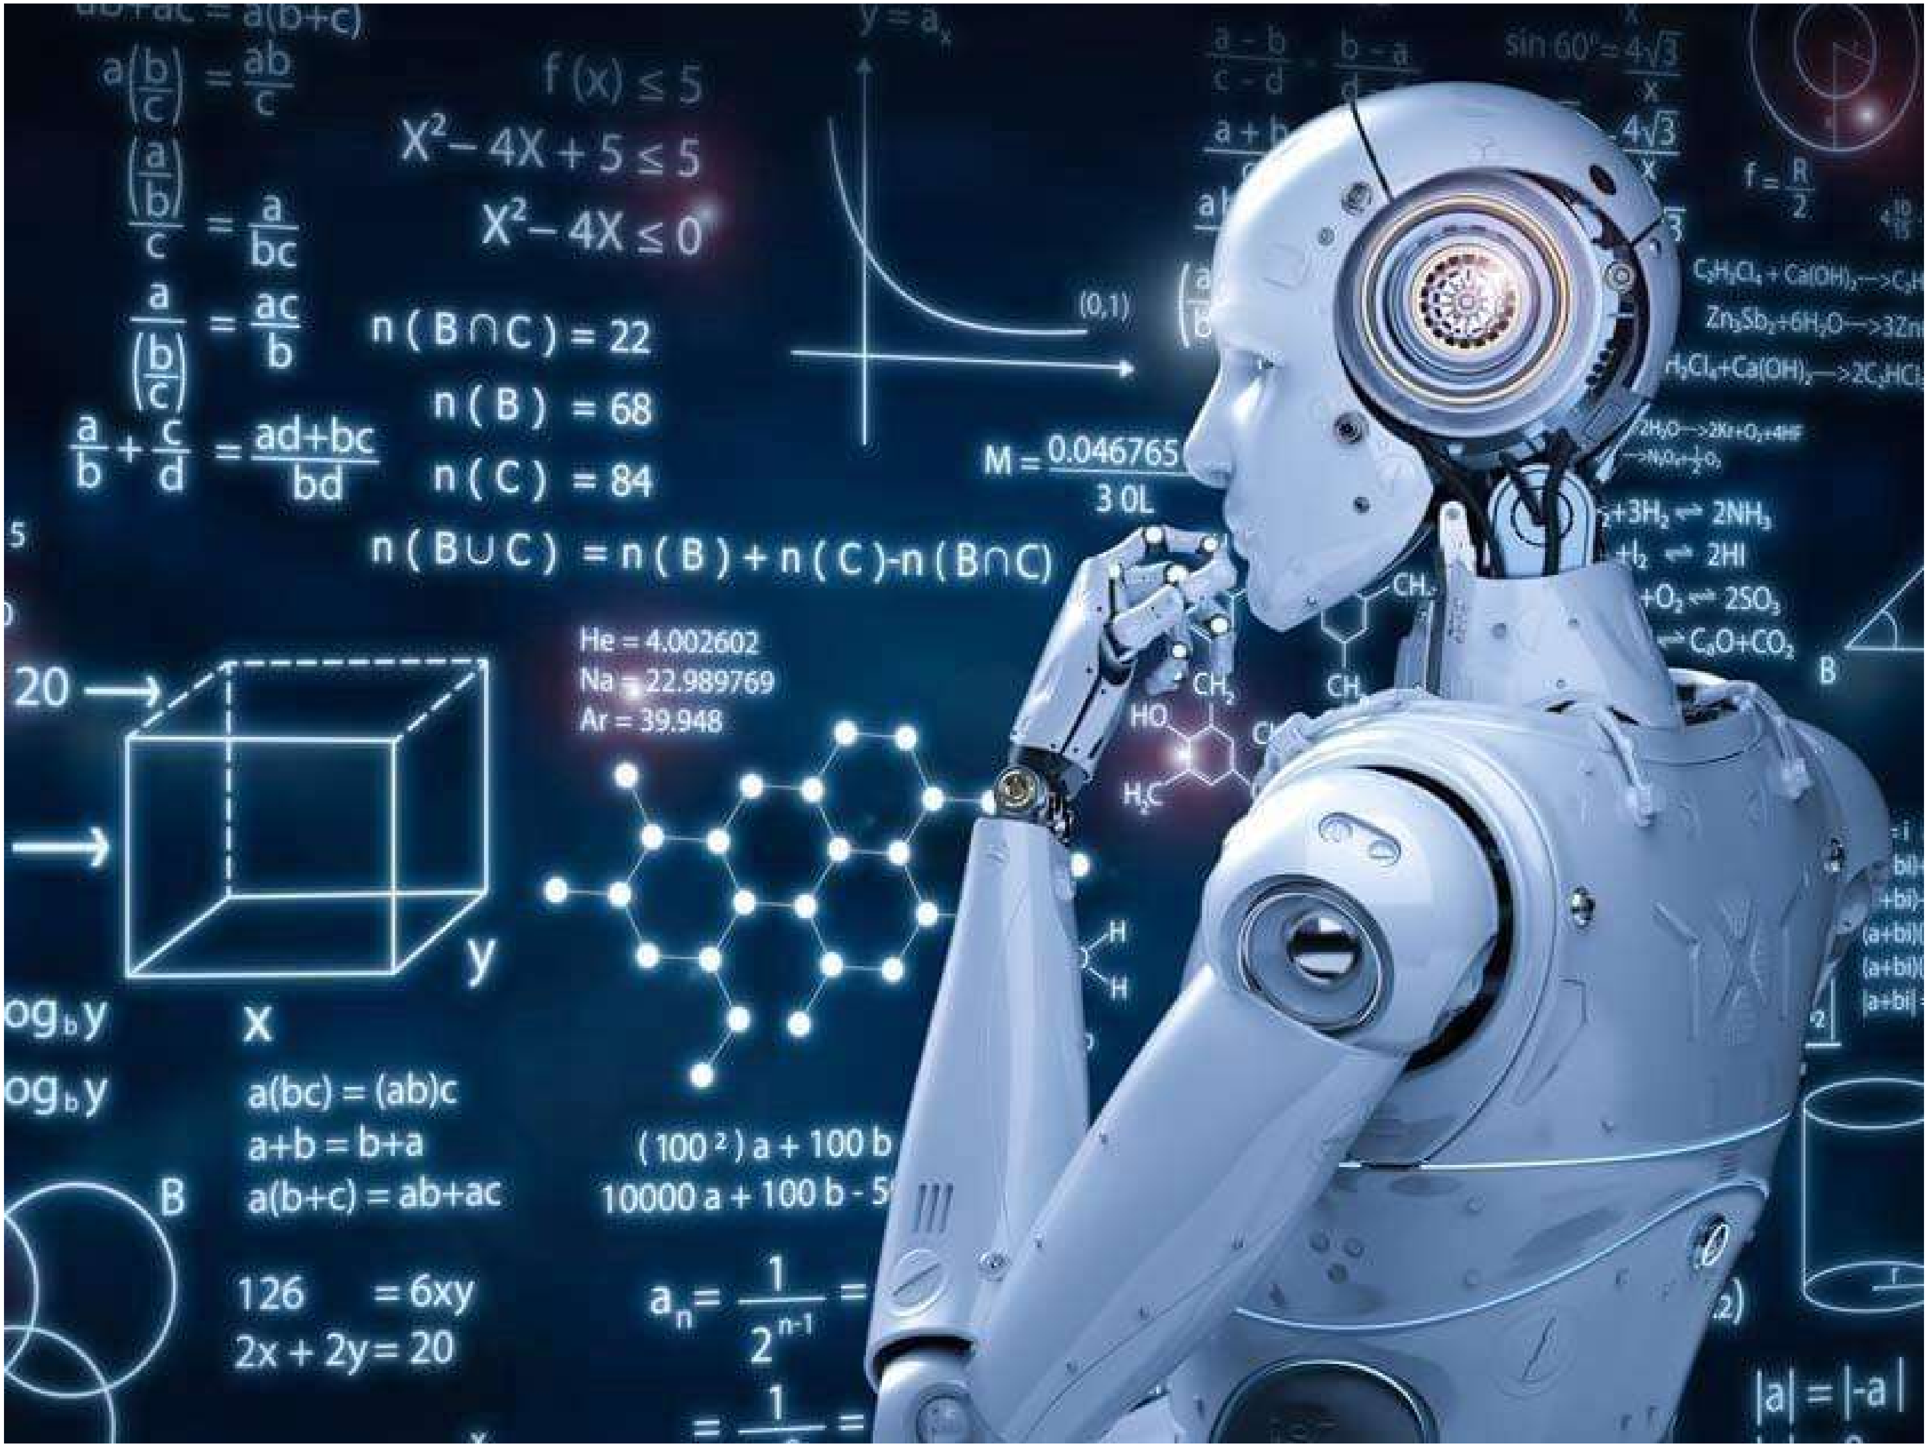

f(x) ≤ 5
X²– 4X + 5 ≤ 5
X²– 4X ≤ 0
n (B ∩ C) = 22
n (B) = 68
n (C) = 84
n (B ∪ C) = n (B) + n (C)-n (B∩C)
a/b + c/d = ad+bc/bd
M = 0.046765/3 0L
He = 4.002602
Na = 22.989769
Ar = 39.948
20 →
a(bc) = (ab)c
a+b = b+a
a(b+c) = ab+ac
126 = 6xy
2x + 2y = 20
(100 ²) a + 100 b
sin 60° = 4√3/x
f = R/2
|a| = |-a|