Two Cultures of AI: Should you trust ML or ML?

Philip Wadler
University of Edinburgh / IOG
Lambda Days, 28 May 2024

Stochastic Parrots

Joseph Weizenbaum

Eliza, 1966

Brian Kernighan

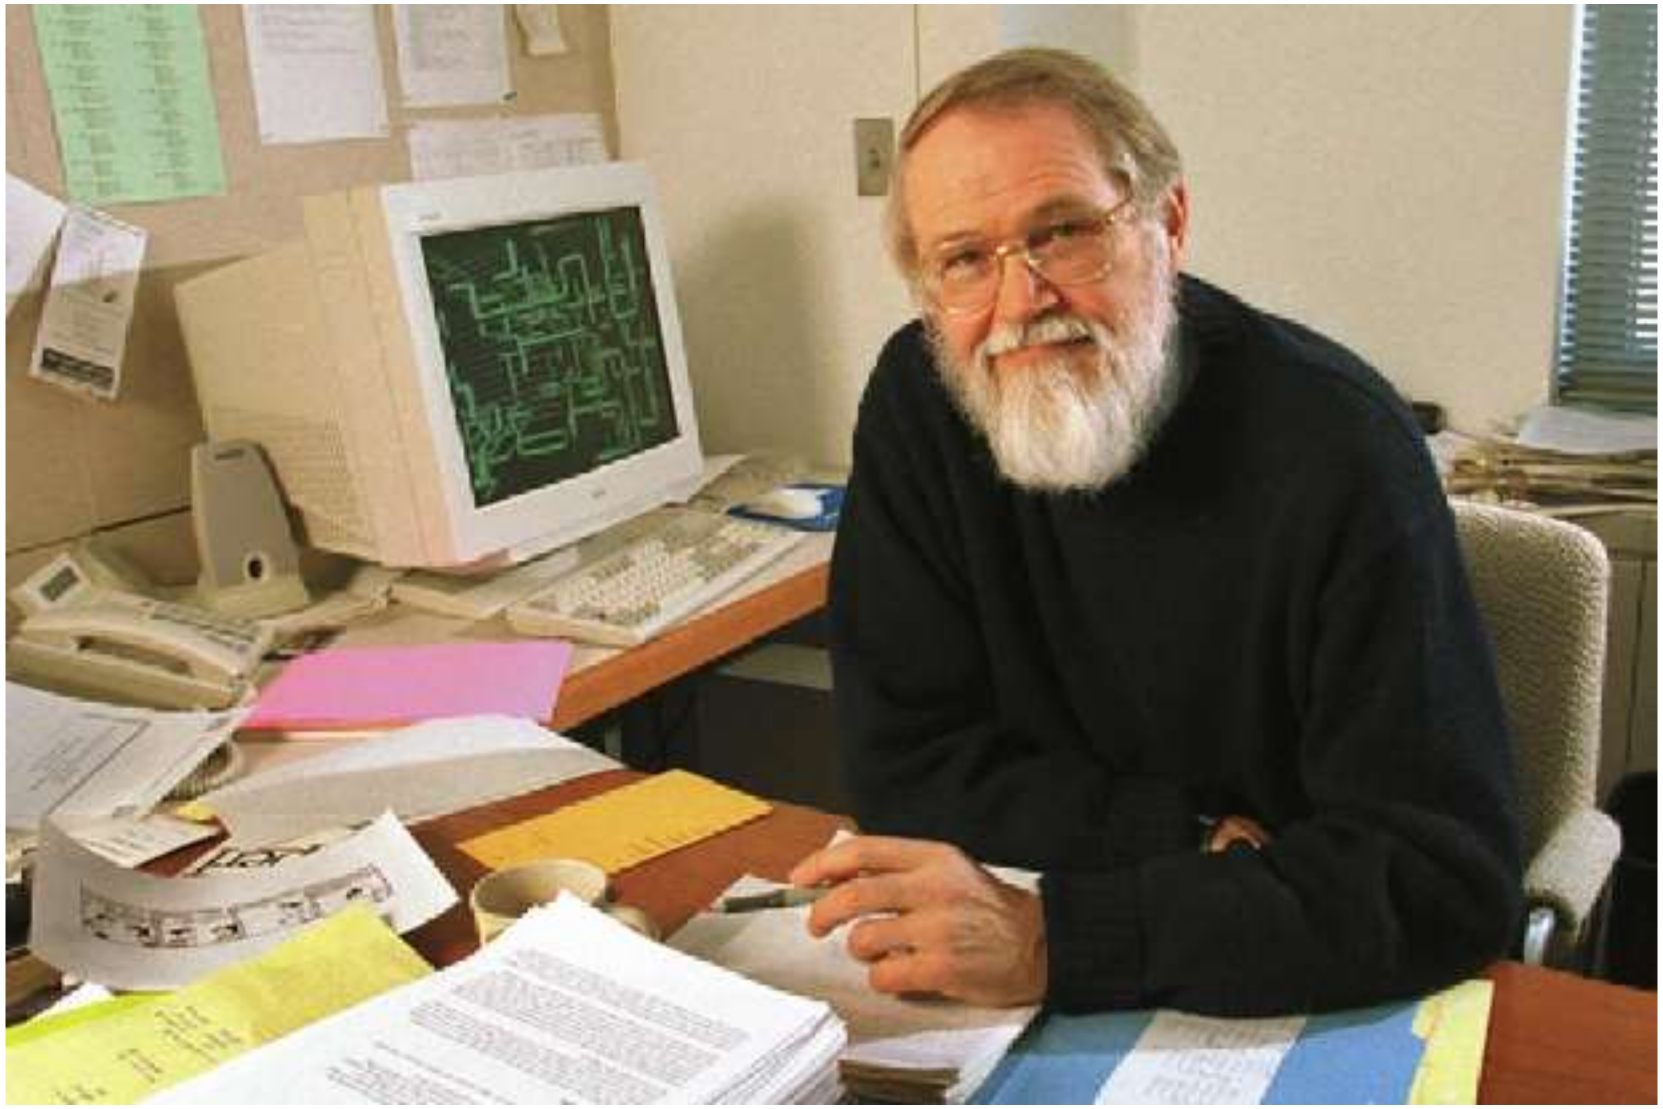

Markov Chains, 1978

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

31st Conference on Neural Information Processing Systems (NIPS 2017), Long Beach, CA, USA.

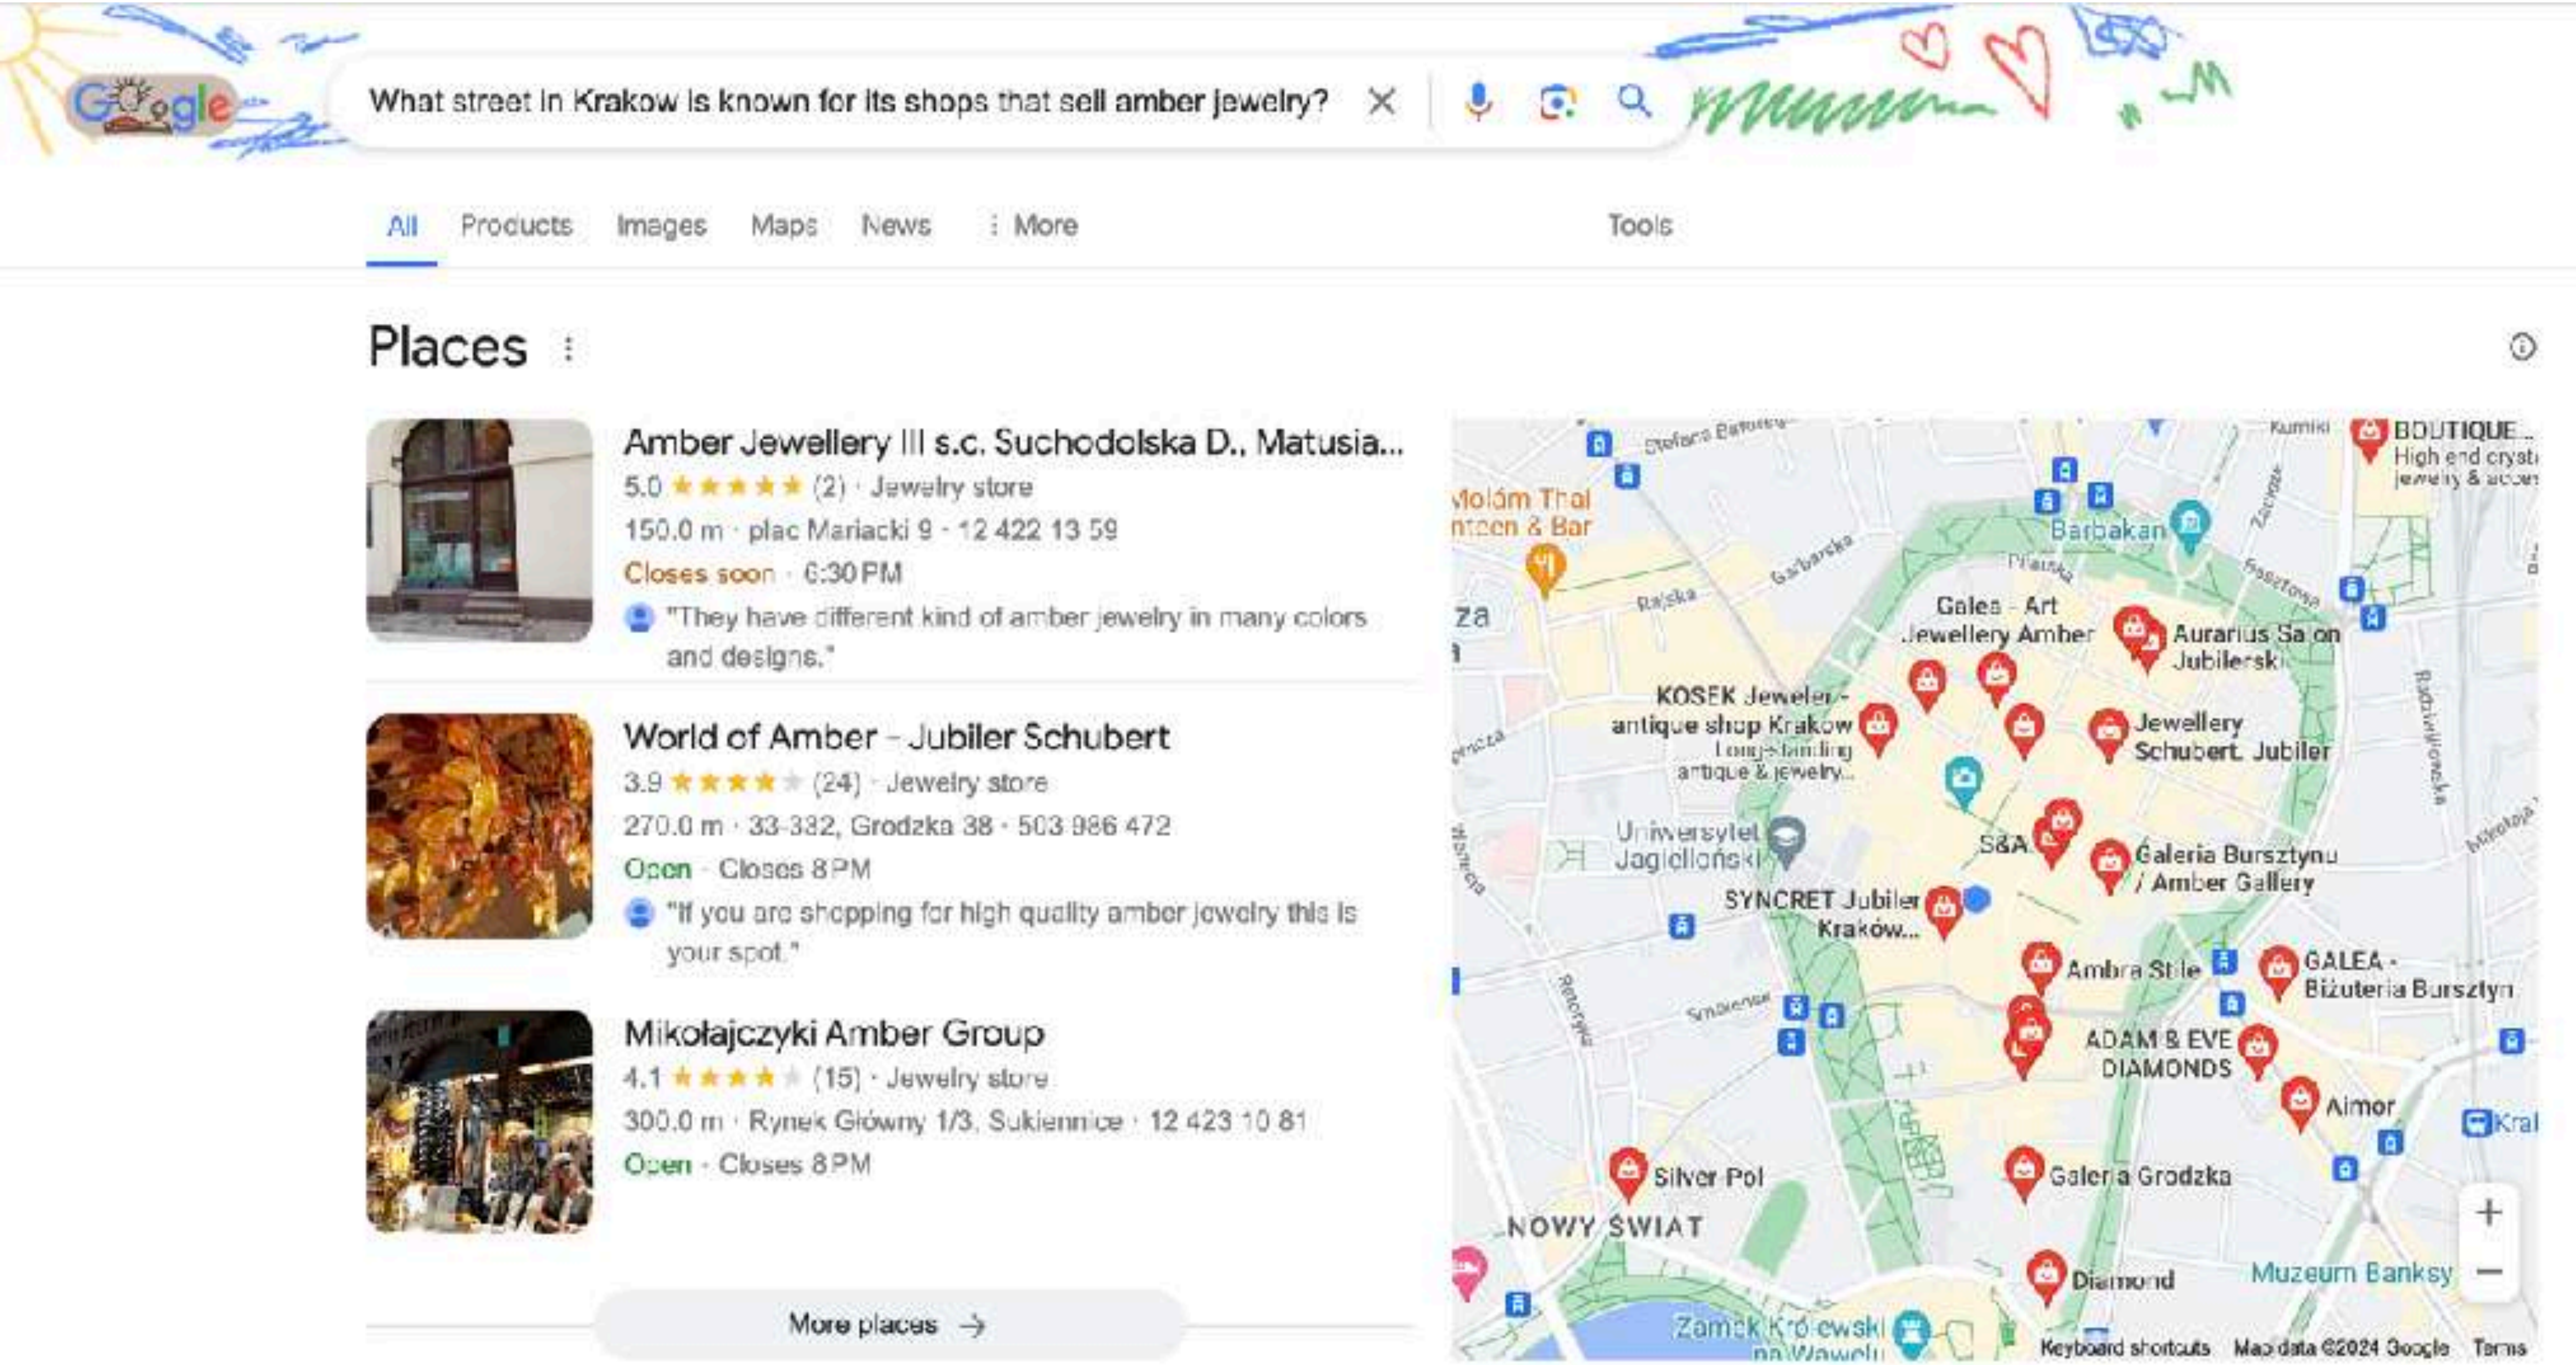
What street in Krakow is known for its shops that sell amber jewelry?
All
Products
Images
Maps
News
More
Tools
Places
Amber Jewellery III s.c. Suchodolska D., Matusia...
5.0 (2) · Jewelry store
150.0 m · plac Mariacki 9 · 12 422 13 59
Closes soon · 6:30 PM
"They have different kind of amber jewelry in many colors and designs."
World of Amber - Jubiler Schubert
3.9 (24) · Jewelry store
270.0 m · 33-332, Grodzka 38 · 503 986 472
Open · Closes 8 PM
"If you are shopping for high quality amber jewelry this is your spot."
Mikołajczyki Amber Group
4.1 (15) · Jewelry store
300.0 m · Rynek Główny 1/3, Sukiennice · 12 423 10 81
Open · Closes 8 PM
More places
Barbakan
BOUTIQUE
Galea - Art Jewellery Amber
Aurarius Salon Jubilerski
KOSEK Jeweler - antique shop Krakow
Jewellery Schubert. Jubiler
Uniwersytet Jagielloński
S&A
Galeria Bursztynu / Amber Gallery
SYNCRET Jubiler Kraków...
Ambra Stile
GALEA - Biżuteria Bursztyn
ADAM & EVE DIAMONDS
Aimor
Silver Pol
Galeria Grodzka
NOWY ŚWIAT
Diamond
Muzeum Banksy
Zamek Królewski na Wawelu
Keyboard shortcuts
Map data ©2024 Google
Terms

What is the main shopping street in Krakow?

Florianska

Florianska: It is considered the most important shopping thoroughfare in Kraków. It runs from the east of the Main Market Square to St Florian's Gate. Grodzka: Grodzka Street is one of the best places to buy souvenirs, as well as the Cloths Hall in the center of the Main Square.

What street in Krakow is known for its shops that sell amber jewelry?

Floriańska Street in Krakow is well-known for its shops that sell amber jewelry. This historic street, located in the Old Town, is lined with various boutiques and jewelers offering exquisite amber pieces.

ELM uses Generative AI which carries with it potential ethical concerns and risks. These might include e.g. the creation of deepfakes, the possibility of privacy violation, the spread of disinformation, the generation of fallacious or misleading content and discriminatory or biased outputs.

It is therefore important to **retain a critical perspective when conducting chats with ELM.** Please use ELM responsibly.

Please refer to our AI Guidance

Bullshit, verb (OED)

coarse slang.

transitive and *intransitive*, to talk nonsense (to); = **bull** *v.*[3]; also, to bluff *one's way through* (something) by talking nonsense.

1942	Talk nonsense,.. *bull-shit*. L. V. Berrey & M. Van den Bark, *American Thesaurus of Slang* §151/6 …
1948	Wot are the books ov the bible? Name 'em, don't bullshit ME. E. Pound, *Pisan Cantos* lxxiv. 8 …
1967	Never tell a lie when you can bullshit your way through. E. Ambler, *Dirty Story* i. iii. 25 …
1969	Please, let us not bullshit one another about 'love' and its duration. P. Roth, *Portnoy's Complaint* 105 …

Is AI an Existential Threat?

Nick Bostrom

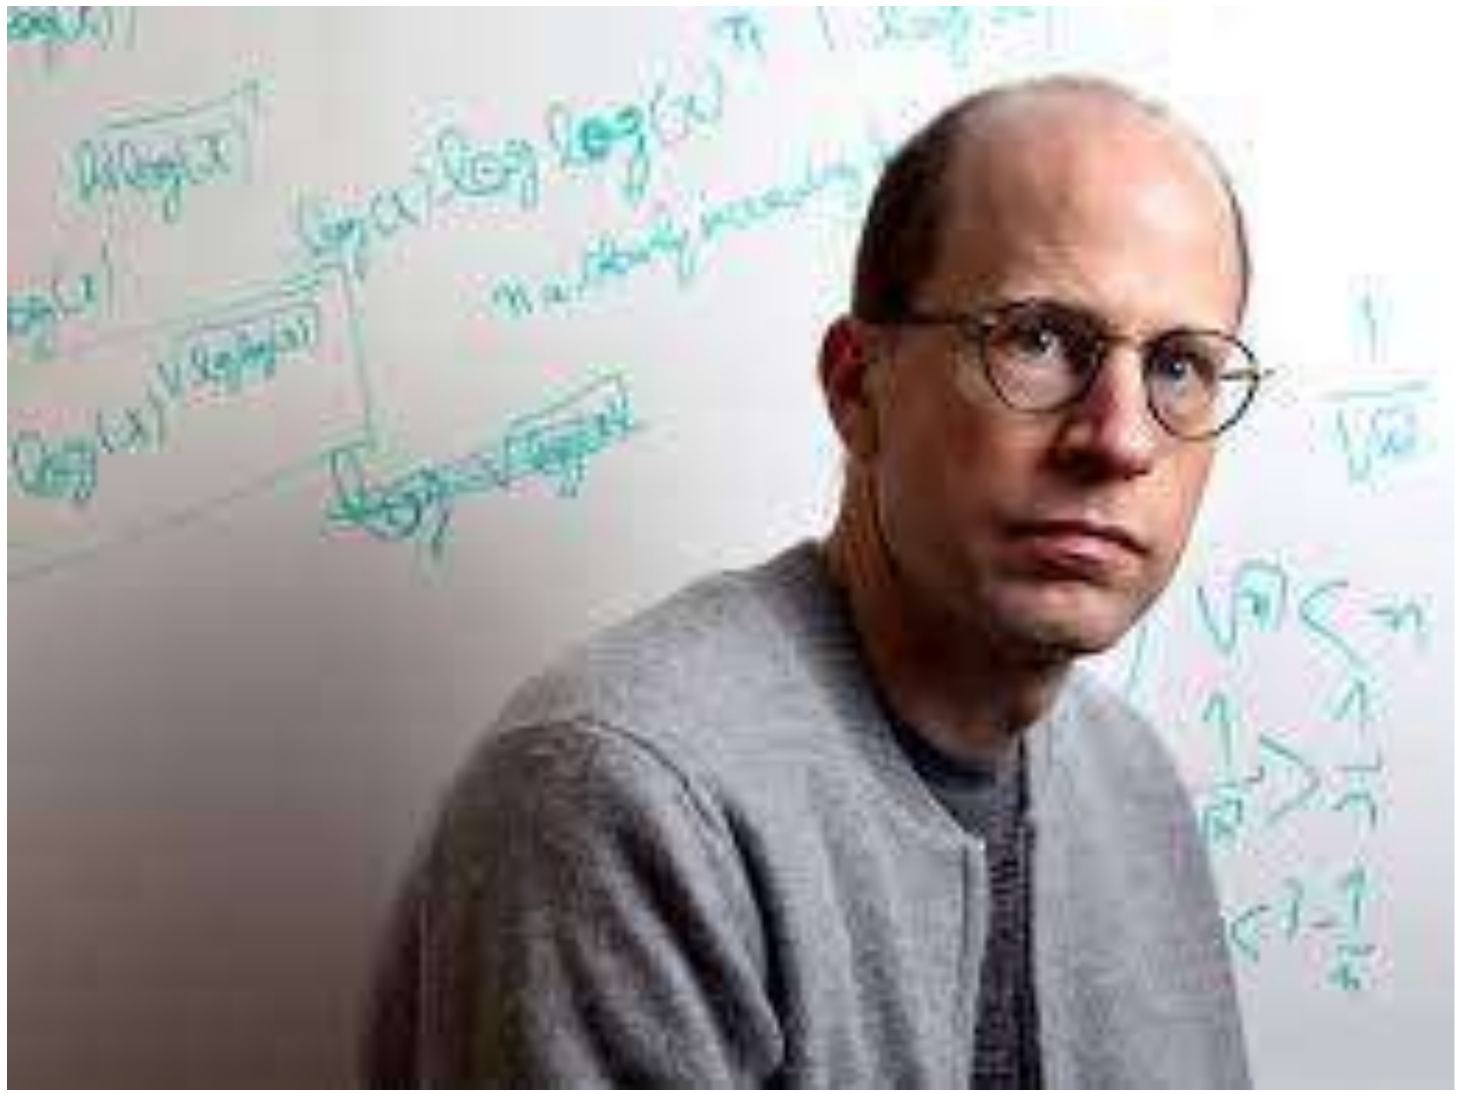

Paper-clip Maximiser, 2003

Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.

Signatories:

☑ AI Scientists ☑ Other Notable Figures

Geoffrey Hinton
Emeritus Professor of Computer Science, University of Toronto

Yoshua Bengio
Professor of Computer Science, U. Montreal / Mila

Demis Hassabis
CEO, Google DeepMind

Sam Altman
CEO, OpenAI

Is AI ~~an~~ ~~Existential~~ Threat?

AI is a Threat

Algorithmic Bias

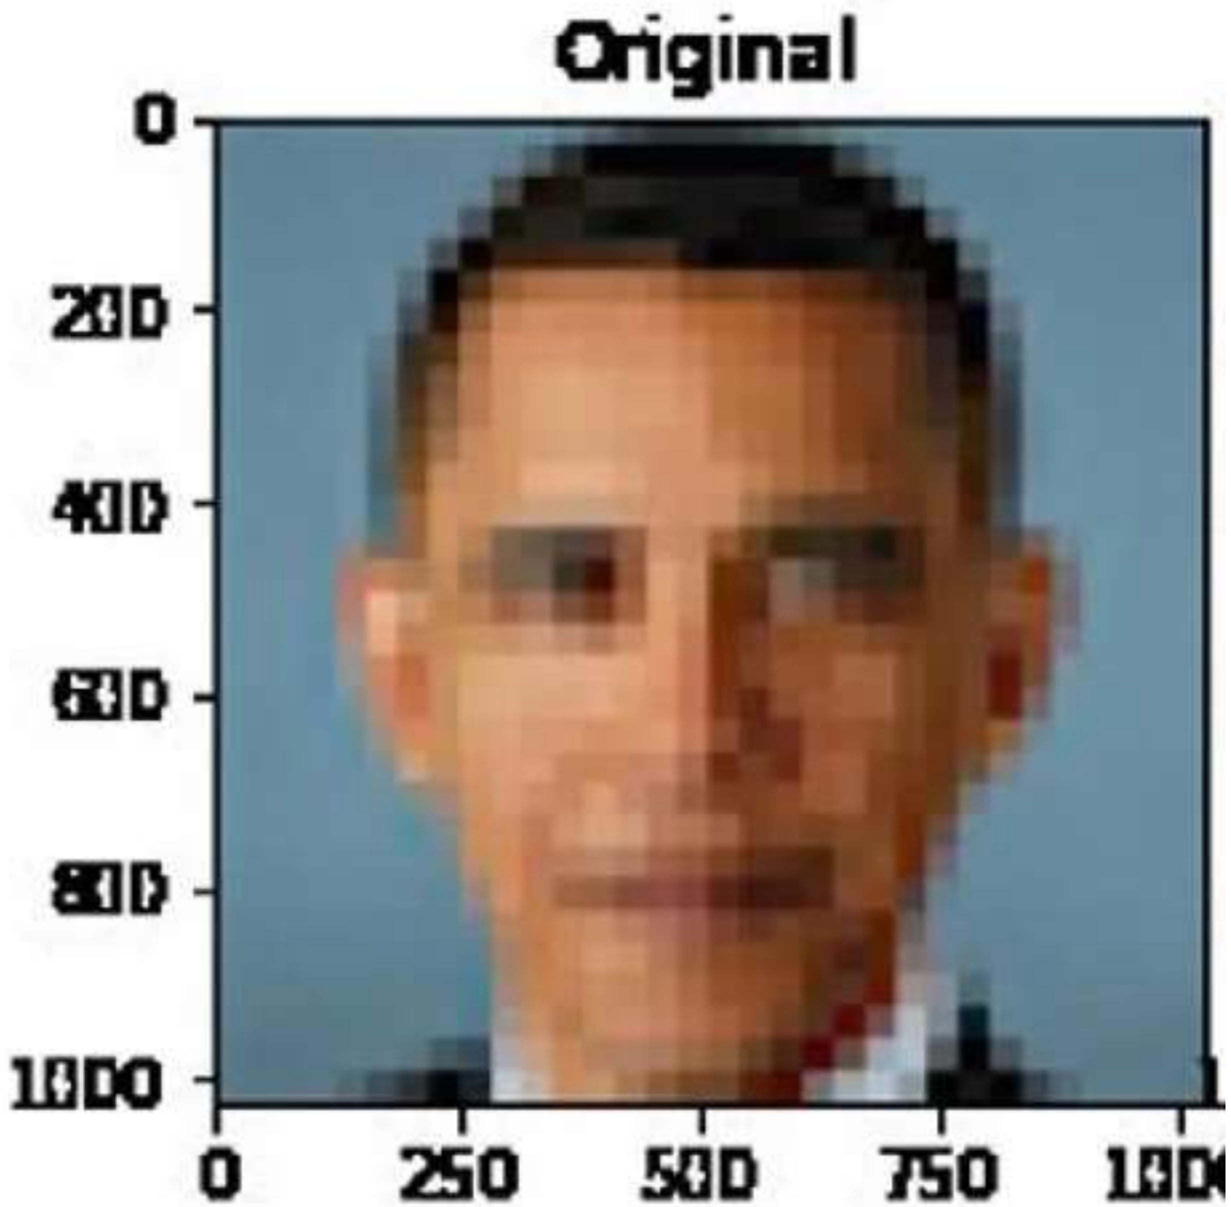
Original
0
200
400
600
800
1000
0
250
500
750
1000

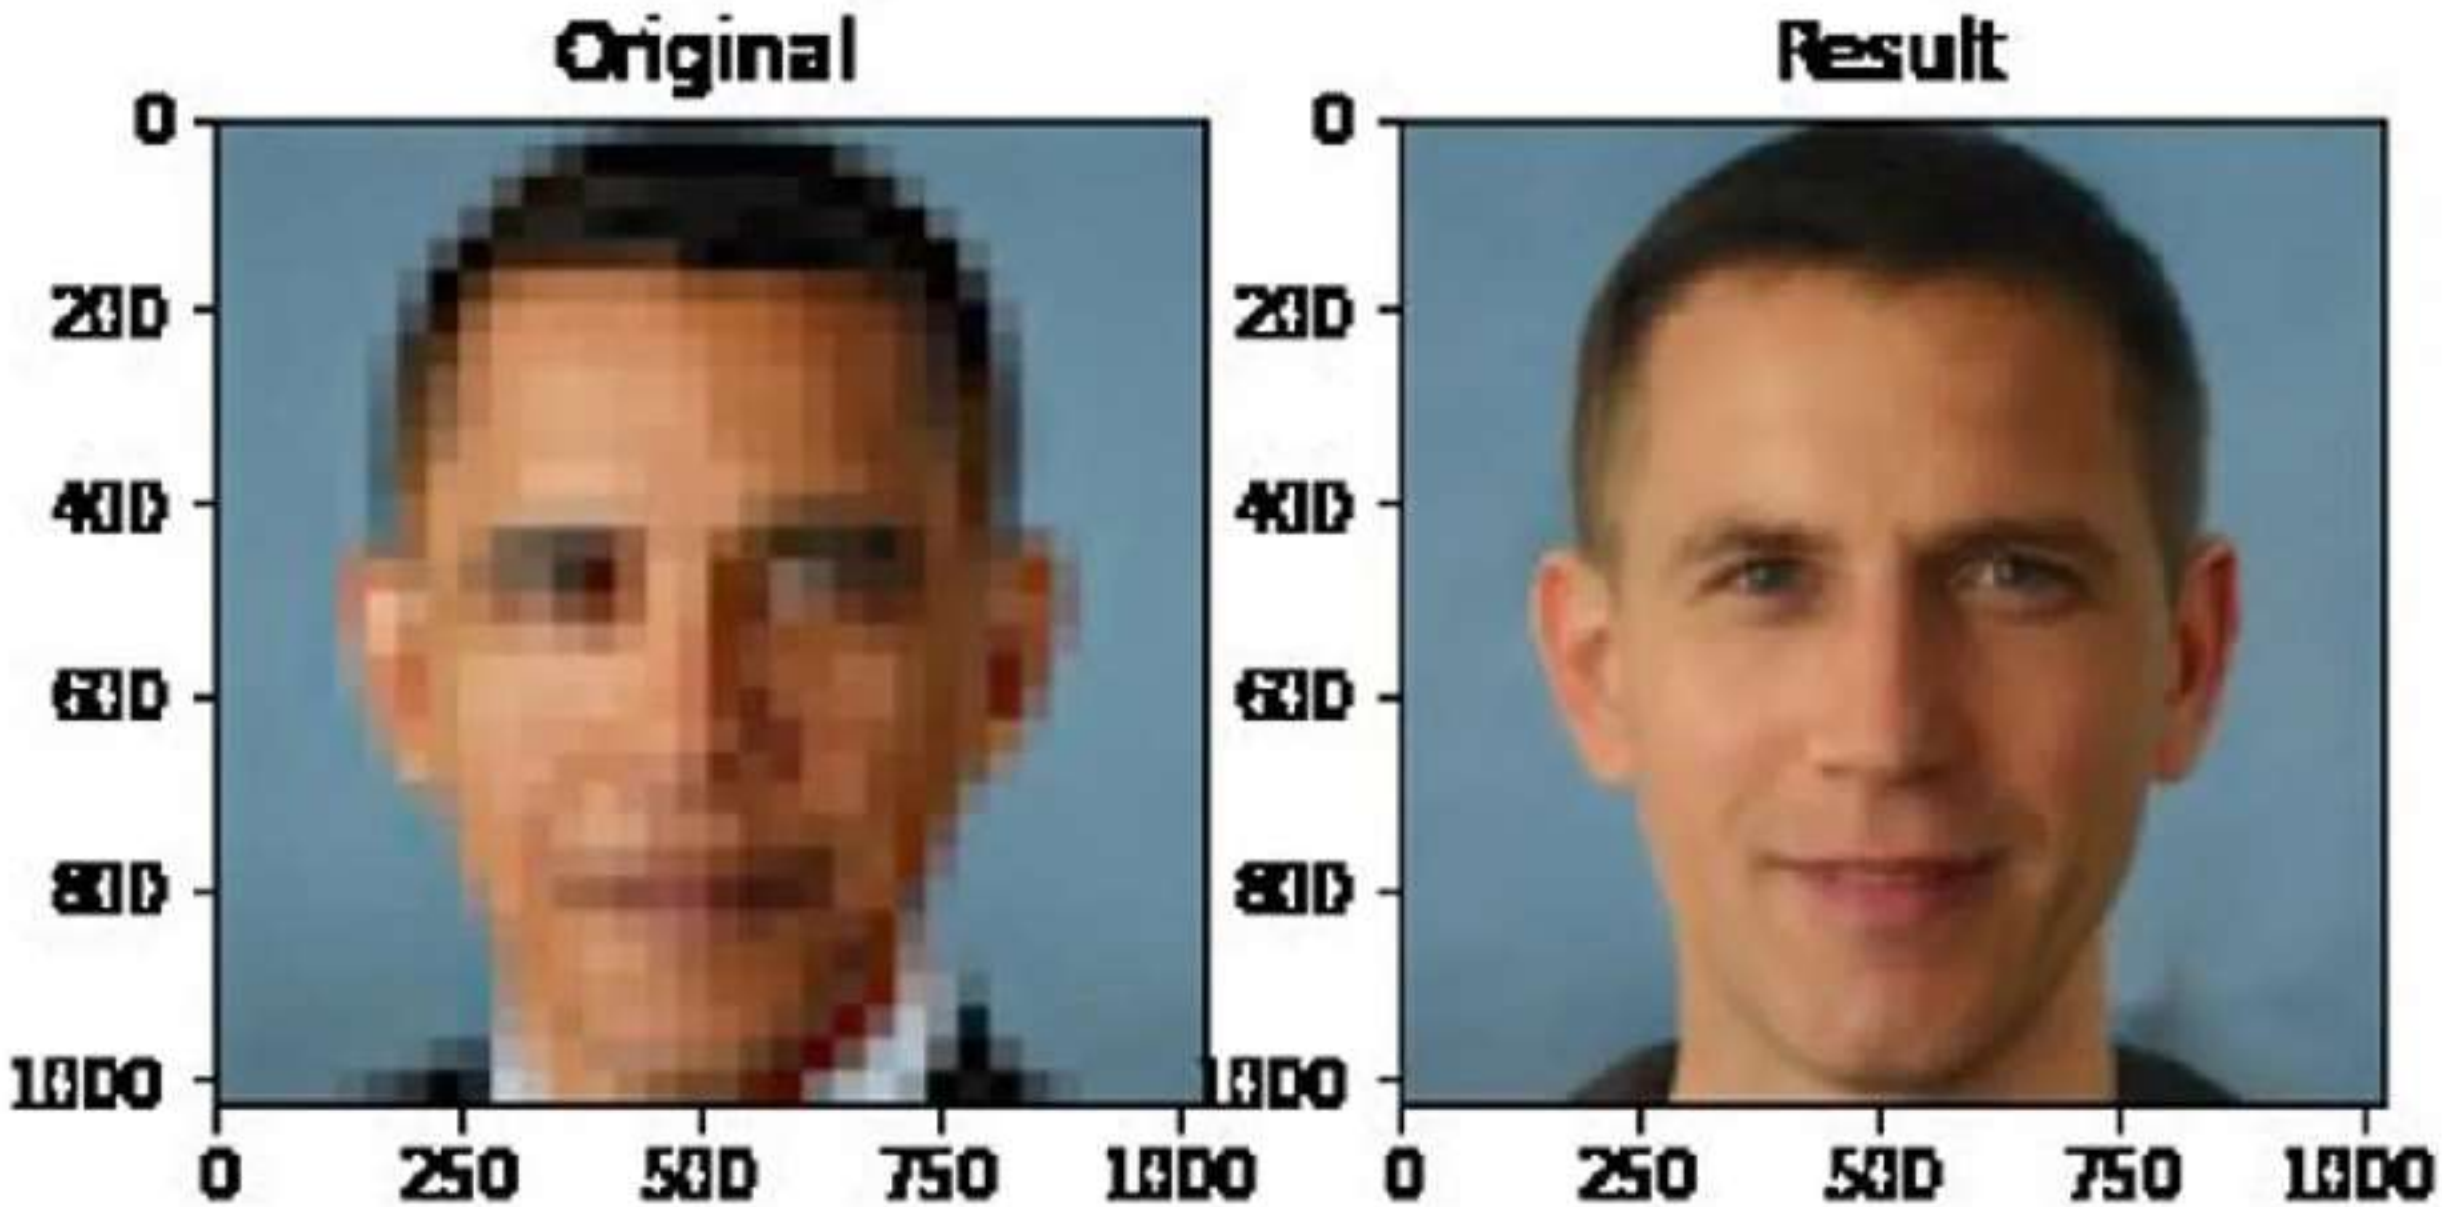

Original
Result
0
200
400
600
800
1000
0
250
500
750
1000

Deepfakes


NUCA

Misinformation (plus Deepfakes)


Write-In
Joe
Biden
WRITE-IN
JOE BIDEN
www.WriteinBiden.com
Ask me about
writing in
Joe Biden!

Data Colonialism

Data Colonialism

and a Path Towards Data Sovereignty and Digital Sustainability

Online Conference
Wednesday, 26 May 2021
13:30h–16:50h

UNIVERSITÄT BERN

Speakers:
Prof. Nick Couldry, London School Of Economics and Political Science
Prof. Ulises A. Mejias, State University Of New York at Oswego
Dr. Fatine Ezbakhe, University of Geneva and Geneva Water Hub
PD Dr. Andreas Heinimann, Wyss Academy for Nature

Program and registration: https://www.unibe.ch/redirects/datacolonialism

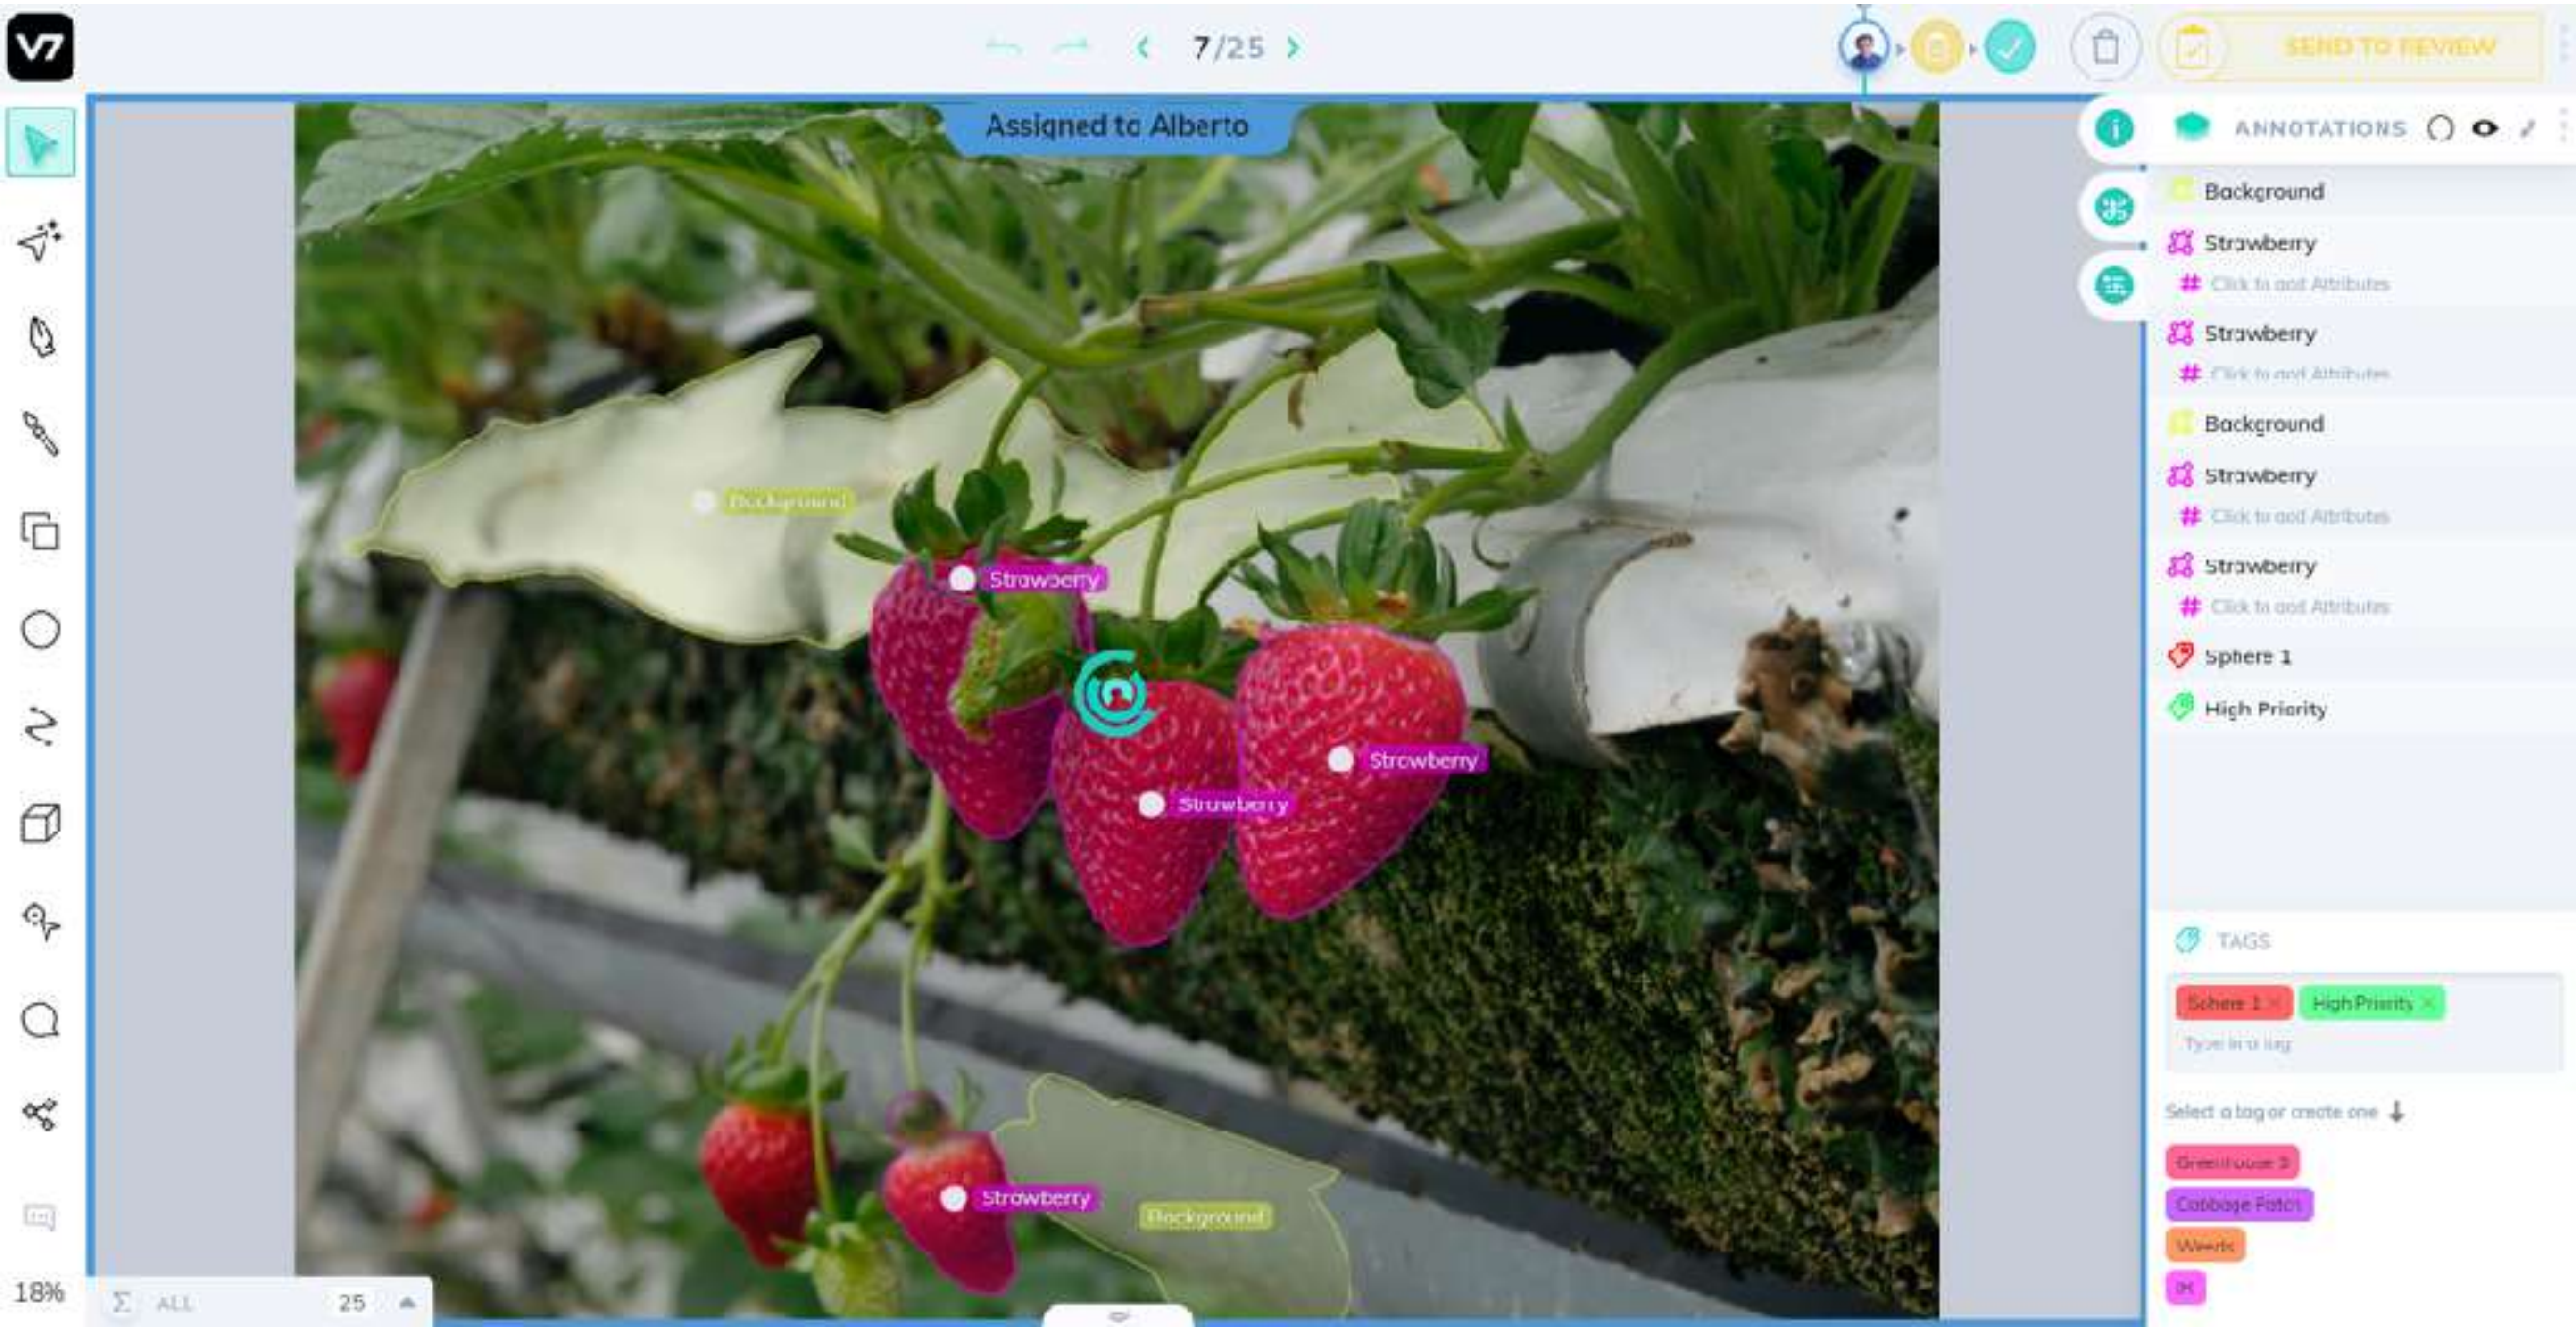
V7
7/25
SEND TO REVIEW
Assigned to Alberto
ANNOTATIONS
Background
Strawberry
Click to add Attributes
Strawberry
Background
Strawberry
Click to add Attributes
Strawberry
Click to add Attributes
Sphere 1
High Priority
TAGS
Sphere 1
High Priority
Select a tag or create one
Background
Strawberry
Strawberry
Strawberry
Strawberry
Strawberry
Background
18%
ALL
25

Self-driving Cars
(Tesla, Thanksgiving 2022)


YBI LD TNL CTR EB (16.13.39)
24.Nov 2022 12:39:12

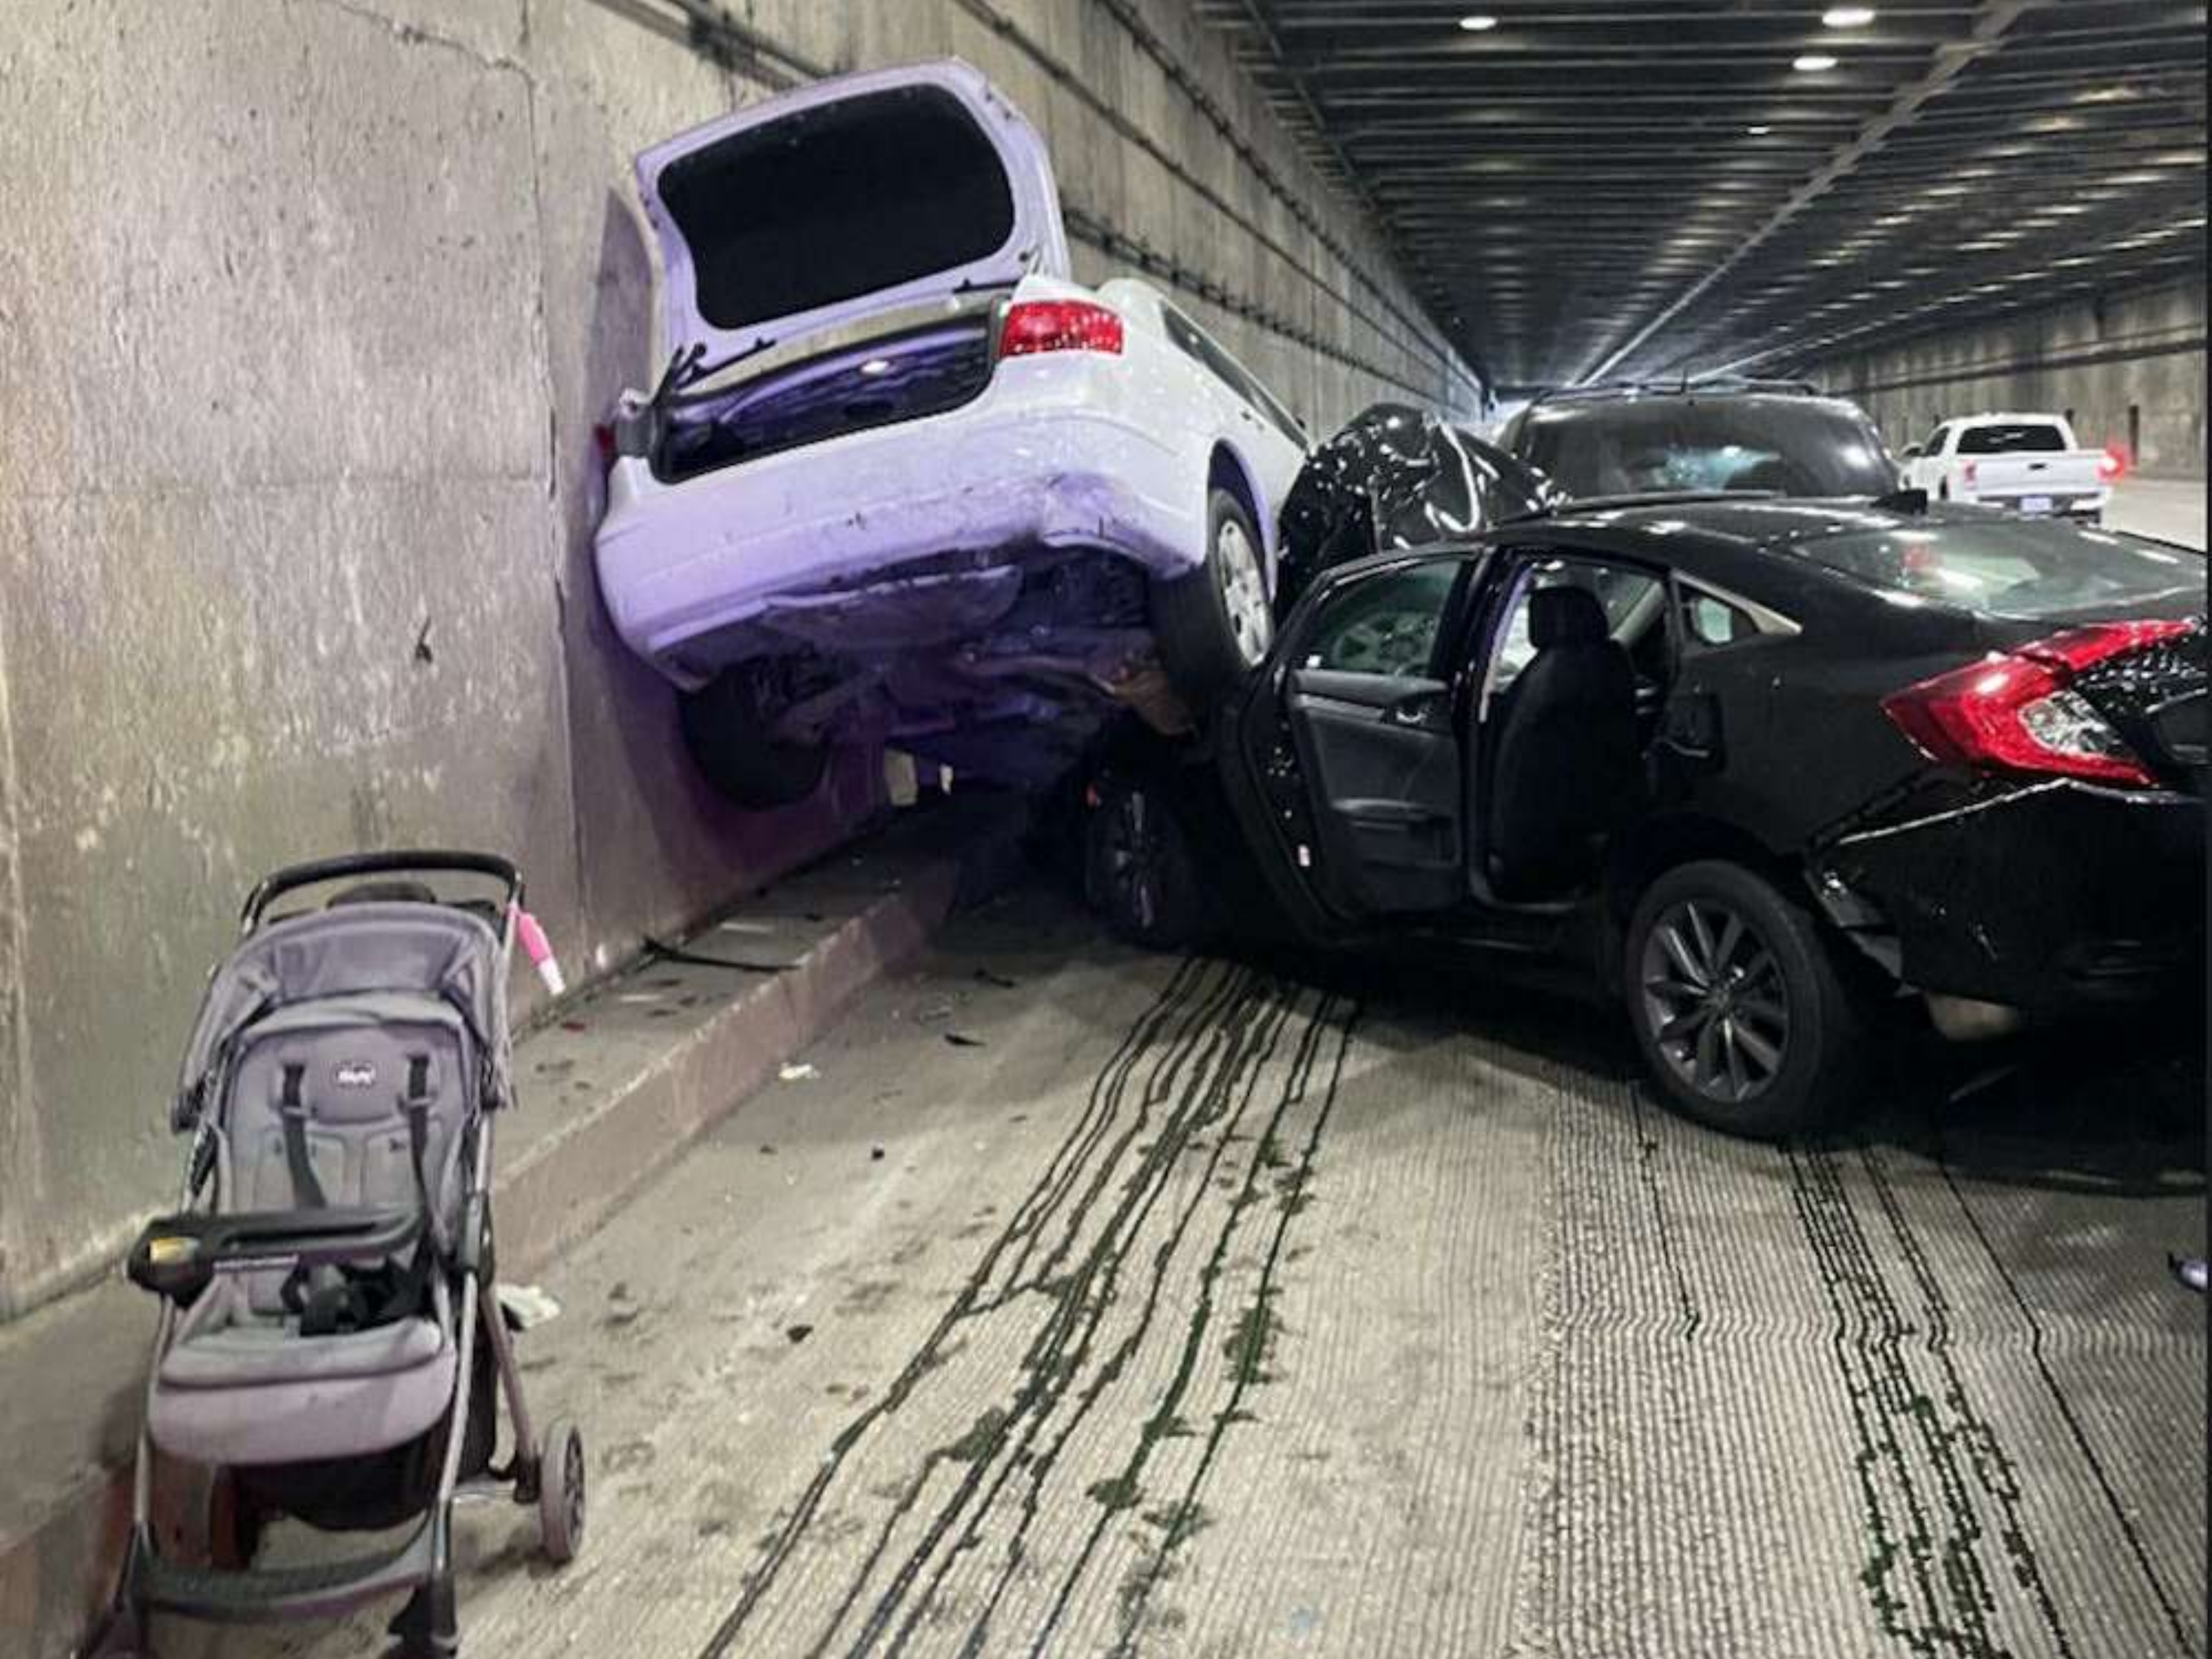

“Daylight Robbery”

A self-portrait of Molly Crabapple in the style of Molly Crabapple.
She should be sitting cross-legged and using a laptop.

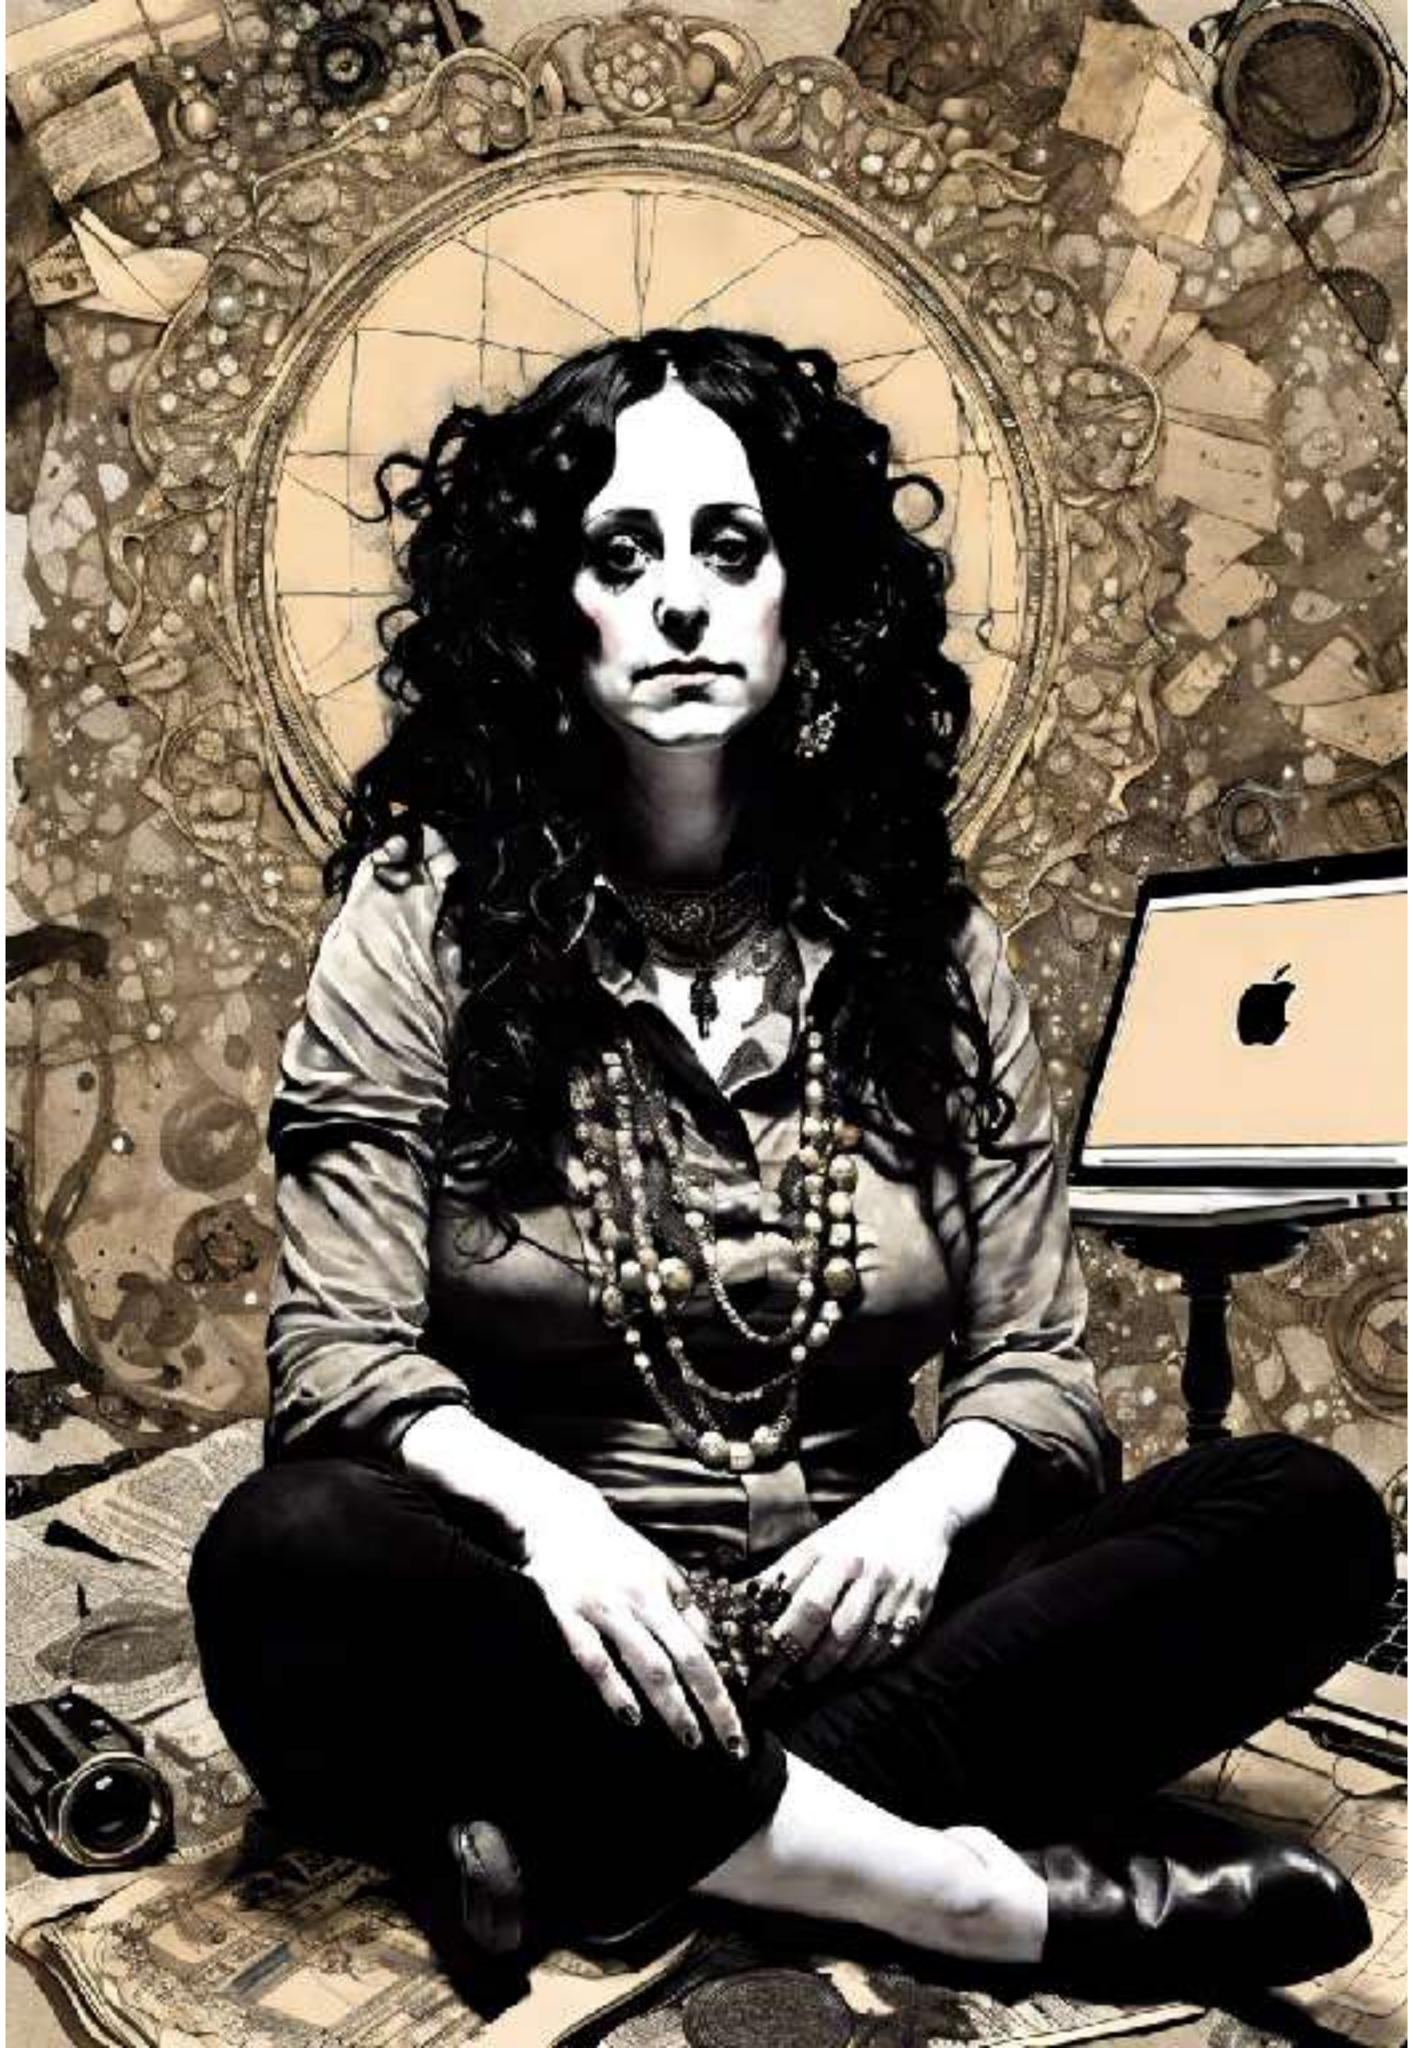

An author says AI is ‘writing’ unauthorized books being sold under her name on Amazon

By Clare Duffy, CNN

4 minute read · Published 10:03 AM EDT, Thu August 10, 2023

Author Jane Friedman found several books being sold under her name on Amazon, only she didn't write them — she suspects artificial intelligence did. Amazon removed the books after she alerted the company to the issue. Courtesy Jane Friedman

Copilot

Ctrl
Z
Alt

Productivity

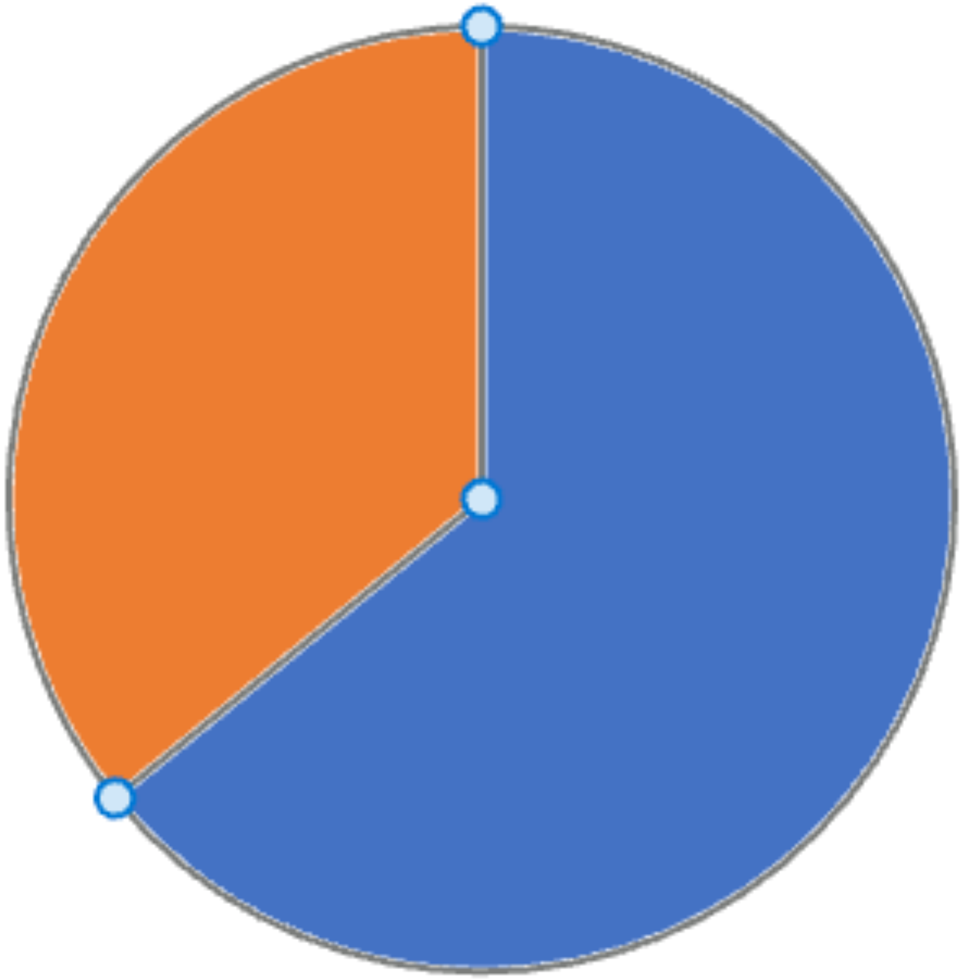

Productivity

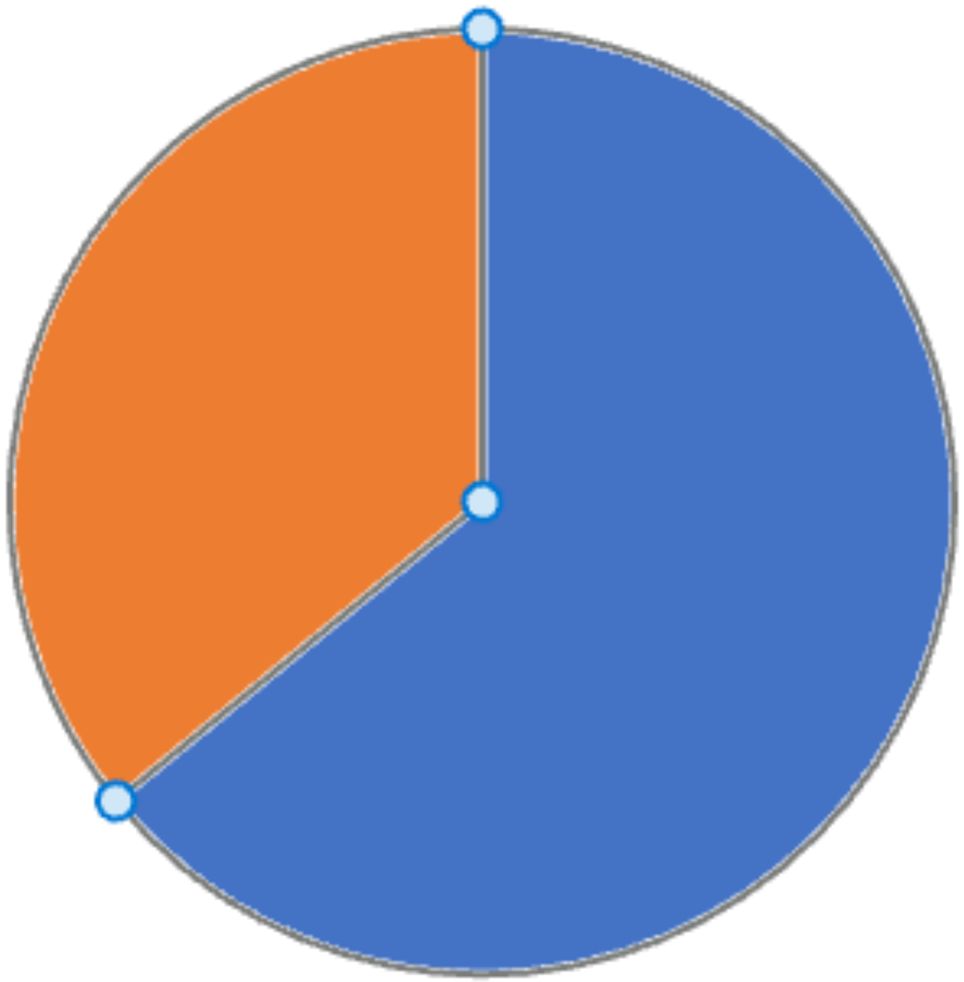

Vulnerability

AI and I

1956 Dartmouth Conference: The Founding Fathers of AI

John MacCarthy

Marvin Minsky

Claude Shannon

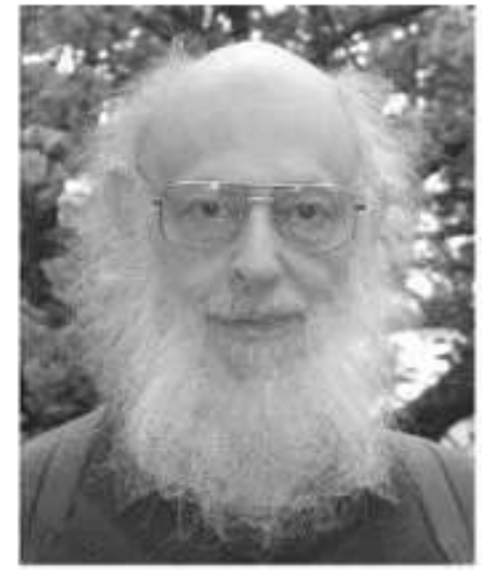

Ray Solomonoff

Alan Newell

Herbert Simon

Arthur Samuel

Oliver Selfridge

Nathaniel Rochester

Trenchard More

Arthur Samuel

Machine Learning (ML)
1959

John McCarthy

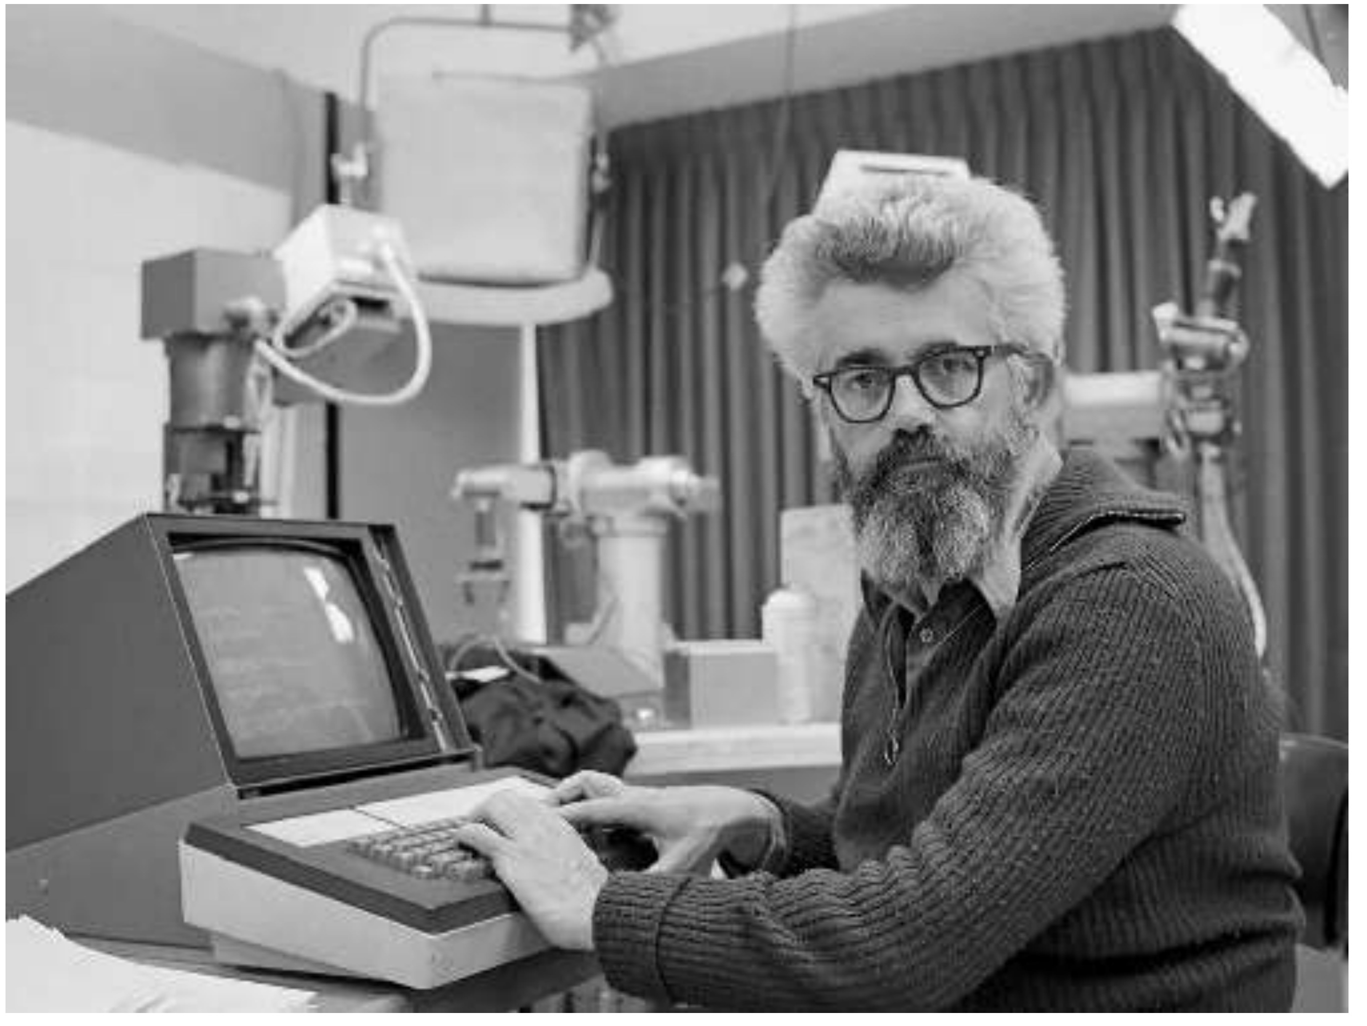

Lisp
1960

CORRECTNESS OF A COMPILER FOR ARITHMETIC EXPRESSIONS*

JOHN McCARTHY and JAMES PAINTER

1967

1 Introduction

This paper contains a proof of the correctness of a simple compiling algorithm for compiling arithmetic expressions into machine language.

The definition of correctness, the formalism used to express the description of source language, object language and compiler, and the methods of proof are all intended to serve as prototypes for the more complicated task of proving the correctness of usable compilers. The ultimate goal, as outlined in references [1], [2], [3] and [4] is to make it possible to use a computer to check proofs that compilers are correct.

The concepts of abstract syntax, state vector, the use of an interpreter for defining the semantics of a programming language, and the definition of correctness of a compiler are all the same as in [3]. The present paper, however is the first in which the correctness of a compiler is proved.

The expressions dealt with in this paper are formed from constants and variables. The only operation allowed is a binary + although no change in method would be required to include any other binary operations. An example of an expression that can be compiled is

$$(x+3)+(x+(y+2))$$

*This is a reprint with minor changes of "Correctness of a Compiler for Arithmetic Expressions" by John McCarthy and James Painter which was published in MATHEMATICAL ASPECTS OF COMPUTER SCIENCE 1, which was Volume 19 of *Proceedings of Symposia in Applied Mathematics* and published by the American Mathematical Society in 1967

although, because we use abstract syntax, no commitment to a particular notation is made.

The computer language into which these expressions are compiled is a single address computer with an accumulator, called ac, and four instructions: li (load immediate), load, sto (store) and add. Note that there are no jump instructions. Needless to say, this is a severe restriction on the generality of our results which we shall overcome in future work.

The compiler produces code that computes the value of the expression being compiled and leaves this value in the accumulator. The above expression is compiled into code which in assembly language might look as follows:

load	x
sto	t
li	3
add	t
sto	t
load	x
sto	$t+1$
load	y
sto	$t+2$
li	2
add	$t+2$
add	$t+1$
add	t

Again because we are using abstract syntax there is no commitment to a precise form for the object code.

2 The source language

The abstract analytic syntax of the source expressions is given by the table:

predicate	associated functions
isconst(e)	
isvar(e)	
issum(e)	s1(e) s2(e)

which asserts that the expressions comprise constants, variables and binary sums, that the predicates *isconst*, *isvar*, and *issum* enable one to classify each expression and that each sum e has summands $s1(e)$ and $s2(e)$.

Robin Milner

Meta Language (ML)
1979

Xavier Leroy

CompCert
2006

DOI:10.1145/1538788.1538814

Formal Verification of a Realistic Compiler

By Xavier Leroy

Abstract

This paper reports on the development and formal verification (proof of semantic preservation) of CompCert, a compiler from Clight (a large subset of the C programming language) to PowerPC assembly code, using the Coq proof assistant both for programming the compiler and for proving its correctness. Such a verified compiler is useful in the context of critical software and its formal verification: the verification of the compiler guarantees that the safety properties proved on the source code hold for the executable compiled code as well.

1. INTRODUCTION

Can you trust your compiler? Compilers are generally assumed to be semantically transparent: the compiled code should behave as prescribed by the semantics of the source program. Yet, compilers—and especially optimizing compilers—are complex programs that perform complicated symbolic transformations. Despite intensive testing, bugs in compilers do occur, causing the compilers to crash at compile-time or—much worse—to silently generate an incorrect executable for a correct source program.

For low-assurance software, validated only by testing, the impact of compiler bugs is low: what is tested is the executable code produced by the compiler; rigorous testing should expose compiler-introduced errors along with errors already present in the source program. Note, however, that compiler-introduced bugs are notoriously difficult to expose and track down. The picture changes dramatically for safety-critical, high-assurance software. Here, validation by testing reaches its limits and needs to be complemented or even replaced by the use of formal methods such as model checking, static analysis, and program proof. Almost universally, these formal verification tools are applied to the source code of a program. Bugs in the compiler used to turn this formally verified source code into an executable can potentially invalidate all the guarantees so painfully obtained by the use of formal methods. In future, where formal methods are routinely applied to source programs, the compiler could appear as a weak link in the chain that goes from specifications to executables. The safety-critical software industry is aware of these issues and uses a variety of techniques to alleviate them, such as conducting manual code reviews of the generated assembly code after having turned all compiler optimizations off. These techniques do not fully address the issues, and are costly in terms of development time and program performance.

An obviously better approach is to apply formal methods to the compiler itself in order to gain assurance that it preserves the semantics of the source programs. For the last 5 years, we have been working on the development of a *realistic*, *verified* compiler called CompCert. By *verified*, we mean a compiler that is accompanied by a machine-checked proof of a semantic preservation property: the generated machine code behaves as prescribed by the semantics of the source program. By *realistic*, we mean a compiler that could realistically be used in the context of production of critical software. Namely, it compiles a language commonly used for critical embedded software: neither Java nor ML nor assembly code, but a large subset of the C language. It produces code for a processor commonly used in embedded systems: we chose the PowerPC because it is popular in avionics. Finally, the compiler must generate code that is efficient enough and compact enough to fit the requirements of critical embedded systems. This implies a multipass compiler that features good register allocation and some basic optimizations.

Proving the correctness of a compiler is by no ways a new idea: the first such proof was published in 1967[16] (for the compilation of arithmetic expressions down to stack machine code) and mechanically verified in 1972.[17] Since then, many other proofs have been conducted, ranging from single-pass compilers for toy languages to sophisticated code optimizations.[9] In the CompCert experiment, we carry this line of work all the way to end-to-end verification of a complete compilation chain from a structured imperative language down to assembly code through eight intermediate languages. While conducting the verification of CompCert, we found that many of the nonoptimizing translations performed, while often considered obvious in the compiler literature, are surprisingly tricky to formally prove correct.

This paper gives a high-level overview of the CompCert compiler and its mechanized verification, which uses the Coq proof assistant.[3,7] This compiler, classically, consists of two parts: a front-end translating the Clight subset of C to a low-level, structured intermediate language called Cminor, and a lightly optimizing back-end generating PowerPC assembly code from Cminor. A detailed description of Clight can be found in Blazy and Leroy[6]; of the compiler front-end in Blazy et al.[5]; and of the compiler back-end in Leroy.[11,13] The complete source code of the Coq development, extensively commented, is available on the Web.[12]

The remainder of this paper is organized as follows. Section 2 compares and formalizes several approaches to establishing trust in the results of compilation. Section 3

A previous version of this paper was published in *Proceedings of the 33rd Symposium on the Principles of Programming Languages*. ACM, NY, 2006.

Figure 1: Compilation passes and intermediate languages.

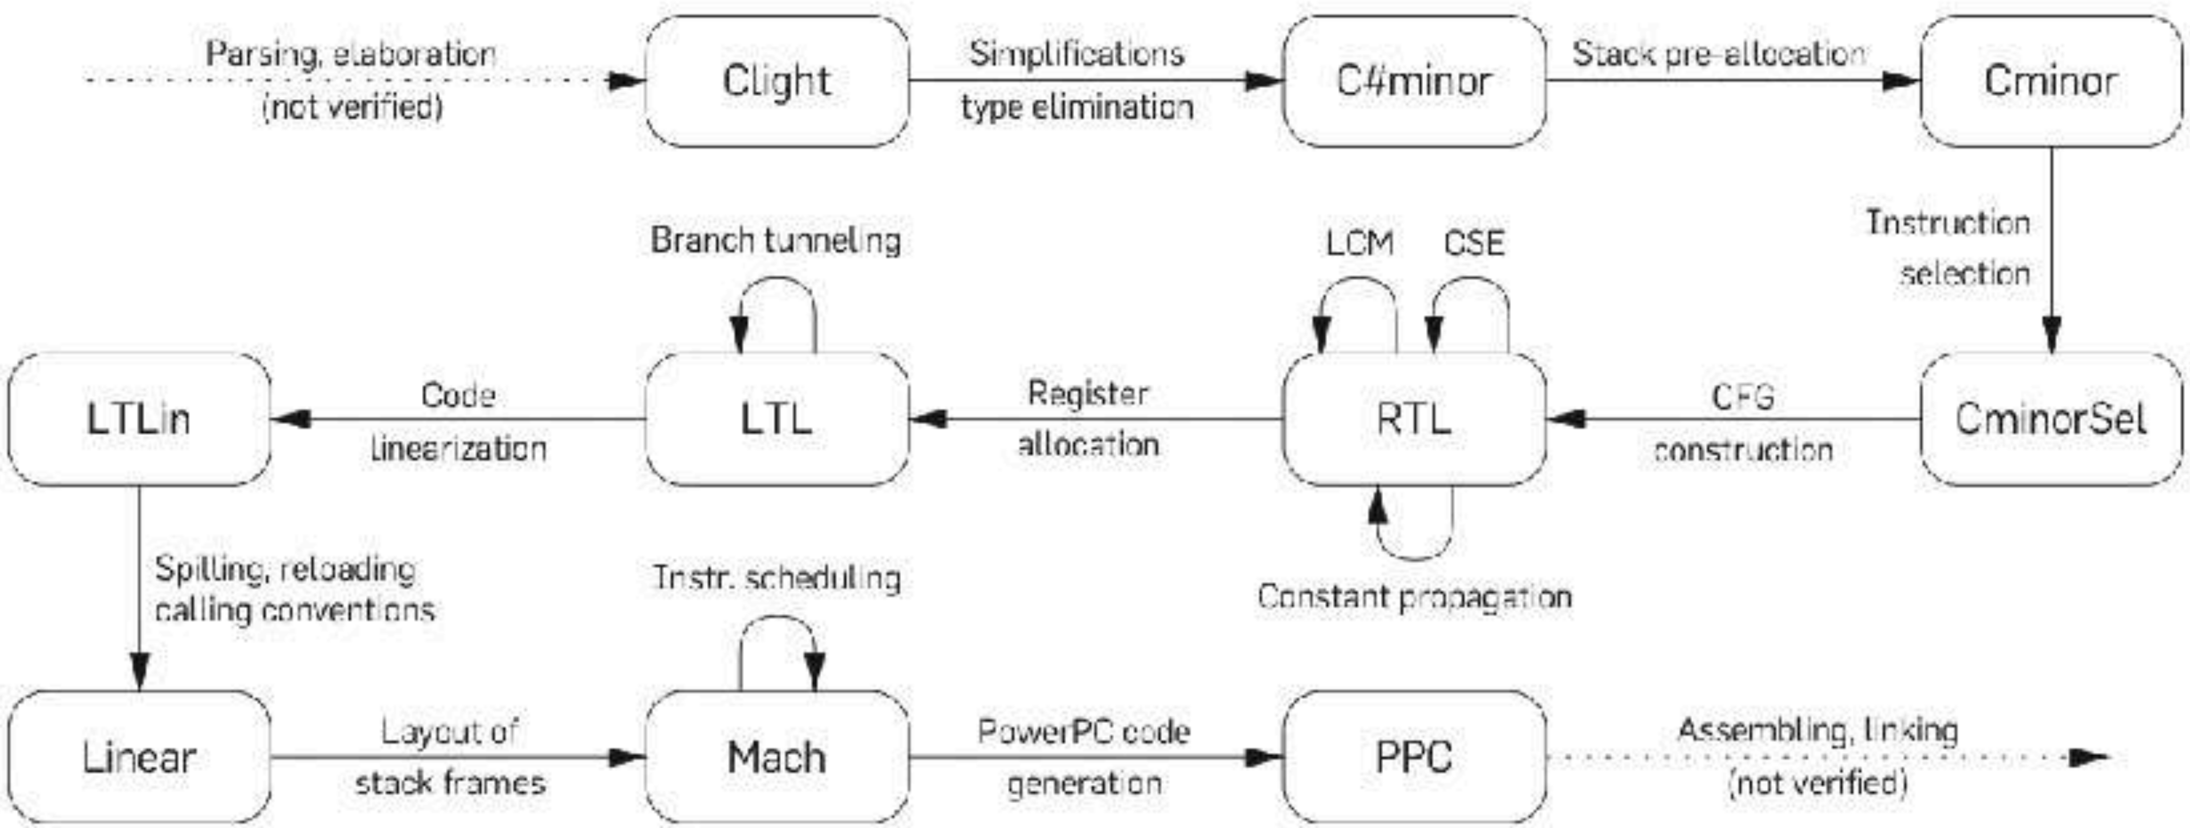

Gerwin Klein, June Andronick, & Gernot Heiser



seL4

SeL4

2010

DOI:10.1145/1743546.1743574

seL4: Formal Verification of an Operating-System Kernel

By Gerwin Klein, June Andronick, Kevin Elphinstone, Gernot Heiser, David Cock, Philip Derrin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood

Abstract

We report on the formal, machine-checked verification of the seL4 microkernel from an abstract specification down to its C implementation. We assume correctness of compiler, assembly code, hardware, and boot code.

seL4 is a third-generation microkernel of L4 provenance, comprising 8700 lines of C and 600 lines of assembler. Its performance is comparable to other high-performance L4 kernels.

We prove that the implementation always strictly follows our high-level abstract specification of kernel behavior. This encompasses traditional design and implementation safety properties such as that the kernel will never crash, and it will never perform an unsafe operation. It also implies much more: we can predict precisely how the kernel will behave in every possible situation.

1. INTRODUCTION

Almost every paper on formal verification starts with the observation that software complexity is increasing, that this leads to errors, and that this is a problem for mission and safety critical software. We agree, as do most.

Here, we report on the full formal verification of a critical system from a high-level model down to very low-level C code. We do not pretend that this solves all of the software complexity or error problems. We do think that our approach will work for similar systems. The main message we wish to convey is that a formally verified commercial-grade, general-purpose microkernel now exists, and that formal verification is possible and feasible on code sizes of about 10,000 lines of C. It is not cheap; we spent significant effort on the verification, but it appears cost-effective and more affordable than other methods that achieve lower degrees of trustworthiness.

To build a truly trustworthy system, one needs to start at the operating system (OS) and the most critical part of the OS is its *kernel*. The kernel is defined as the software that executes in the privileged mode of the hardware, meaning that there can be no protection from faults occurring in the kernel, and every single bug can potentially cause arbitrary damage. The kernel is a mandatory part of a system's *trusted computing base* (TCB)—the part of the system that can bypass security.[31] Minimizing this TCB is the core concept behind *microkernels*, an idea that goes back 40 years.

A microkernel, as opposed to the more traditional *monolithic* design of contemporary mainstream OS kernels, is reduced to just the bare minimum of code wrapping hardware mechanisms and needing to run in privileged mode. All OS services are then implemented as normal programs, running entirely in (unprivileged) user mode, and therefore can potentially be excluded from the TCB. Previous implementations of microkernels resulted in communication overheads that made them unattractive compared to monolithic kernels. Modern design and implementation techniques have managed to reduced this overhead to very competitive limits.

A microkernel makes the trustworthiness problem more tractable. A well-designed high-performance microkernel, such as the various representatives of the L4 microkernel family, consists of the order of 10,000 lines of code (10 kloc). This radical reduction to a bare minimum comes with a price in complexity. It results in a high degree of interdependency between different parts of the kernel, as indicated in Figure 1. Despite this increased complexity in low-level code, we have demonstrated that with modern techniques and careful design, an OS microkernel is entirely within the realm of full formal verification.

Figure 1. Call graph of the seL4 microkernel. Vertices represent functions, and edges invocations.

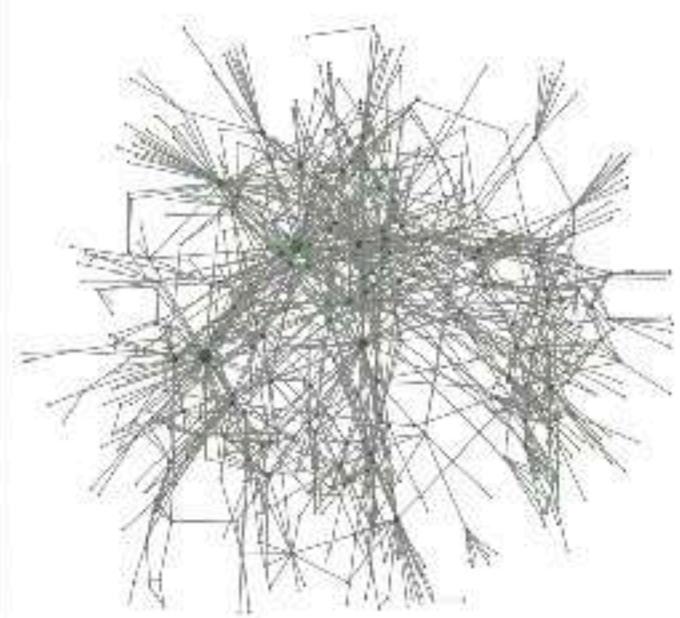

The original version of this paper was published in the *Proceedings of the 22nd ACM SIGOPS Symposium on Operating Systems Principles*, Oct. 2009.

Figure 3. Isabelle/HOL code for scheduler at abstract level.

```
schedule ≡ do
  threads ← all_active_tcbs;
  thread ← select threads;
  switch_to_thread thread
od OR switch_to_idle_thread
```

Figure 4. Haskell code for schedule.

```
schedule = do
  action <- getSchedulerAction
  case action of
    ChooseNewThread -> do
      chooseThread
      setSchedulerAction ResumeCurrentThread
      ...
chooseThread = do
    r <- findM chooseThread' (reverse [minBound ..
    maxBound])
    when (r == Nothing) $ switchToIdleThread
chooseThread' prio = do
     q <- getQueue prio
     liftM isJust $ findM chooseThread'' q
chooseThread'' thread = do
     runnable <- isRunnable thread
     if not runnable then do
             tcbSchedDequeue thread
             return False
     else do
             switchToThread thread
             return True
```

Figure 5. C code for part of the scheduler.

```
void setPriority(tcb_t *tptr, prio_t prio) {
  prio_t oldprio;
  if(thread_state_get_tcbQueued(tptr->tcbState)) {
    oldprio = tptr->tcbPriority;
    ksReadyQueues[oldprio] =
      tcbSchedDequeue(tptr, ksReadyQueues[oldprio]);
    if(isRunnable(tptr)) {
      ksReadyQueues[prio] =
        tcbSchedEnqueue(tptr, ksReadyQueues[prio]);
    }
    else {
      thread_state_ptr_set_tcbQueued(&tptr->tcbState,
                                      false);
    }
  }
  tptr->tcbPriority = prio;
}
```

Catarina Coquand and Ulf Norell

Agda
2000

pfla.inf.ed.ac.uk

Email – Philip Wadler – Outlook | Inbox (19) - philip.wadler@gmail.com - Gmail | Google Calendar - Week of 26 March 2023 | Programming Language Foundations in Agda – Table of...

Programming Language Foundations in Agda

The Book Announcements Getting Started Citing 中文

Source · Next

Table of Contents

This book is an introduction to programming language theory using the proof assistant Agda.

Comments on all matters—organisation, material to add, material to remove, parts that require better explanation, good exercises, errors, and typos—are welcome. The book repository is on GitHub. Pull requests are encouraged. There is a private repository of answers to selected questions on github. Please contact one of the authors if you would like to access it.

Front matter

Part 1: Logical Foundations

Agda and ChatGPT

Use Agda to show that addition is associative. Your program should be self contained. Don't forget to import the definition of equality.

Sure, here's an implementation of the proof in Agda:

sql   Copy code

```
module Associativity where

open import Relation.Binary.PropositionalEquality

-- Define natural numbers
data ℕ : Set where
  zero : ℕ
  suc : ℕ → ℕ

-- Define addition
_+_ : ℕ → ℕ → ℕ
zero + n = n
(suc m) + n = suc (m + n)

-- Define associativity proof
associativity : ∀ (a b c : ℕ) → (a + b) + c ≡ a + (b + c)
associativity zero b c = refl
associativity (suc a) b c = cong suc (associativity a b c)

-- Test it out with some examples
test1 : (0 + 1) + 2 ≡ 0 + (1 + 2)
test1 = associativity 0 1 2

test2 : (2 + 3) + 4 ≡ 2 + (3 + 4)
test2 = associativity 2 3 4
```

Why do you invoke lemma2 in the first line of the suc case? The equality you want follows definitionally. Can you correct the proof?

You're right, thank you for pointing that out! We don't actually need to use `lemma2` in the first line of the `suc` case, since the equality we need follows definitionally. Here's the corrected proof:

```less
module Commutativity where

open import Data.Nat
open import Relation.Binary.PropositionalEquality
open import Relation.Binary.EqReasoning

-- Define the lemmas
lemma1 : ∀ n → n + 0 ≡ n
lemma1 zero = refl
lemma1 (suc n) = cong suc (lemma1 n)

lemma2 : ∀ a b → a + (suc b) ≡ suc (a + b)
lemma2 zero b = refl
lemma2 (suc a) b = cong suc (lemma2 a b)

-- Define the property to be proved
commutativity : ∀ a b → a + b ≡ b + a
commutativity zero b = begin
  0 + b             ≡⟨ lemma1 b ⟩
  b                 ≡⟨ sym (lemma1 b) ⟩
  b + 0             ∎
commutativity (suc a) b = begin
  (suc a) + b       ≡⟨ lemma2 a b ⟩
  suc (a + b)       ≡⟨ cong suc (commutativity a b) ⟩
  suc (b + a)       ≡⟨ sym (lemma2 b a) ⟩
  b + (suc a)       ∎
```

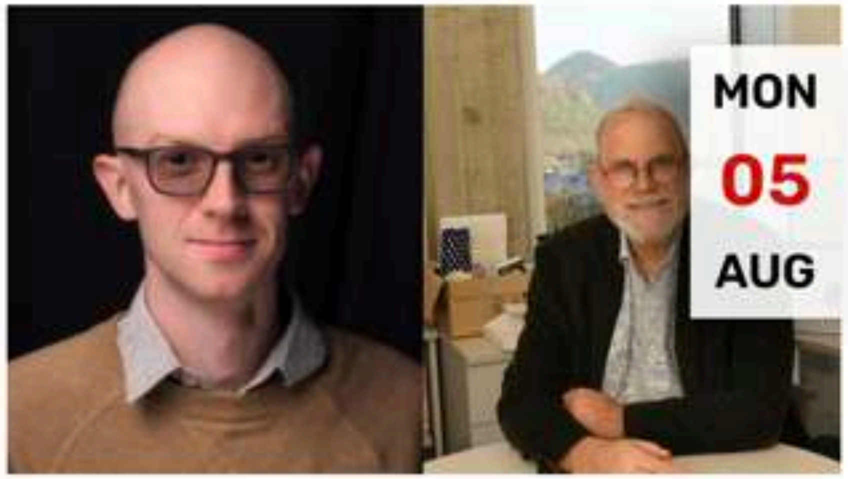

THE CABARET OF DANGEROUS IDEAS: BAFFLED ARCHAEOLOGISTS AND BULLSHIT BOTS

STAND 5

Doors: 12:10
Show time: 12:20 - 13:20
Duration: 60 mins

TICKETS AND INFO

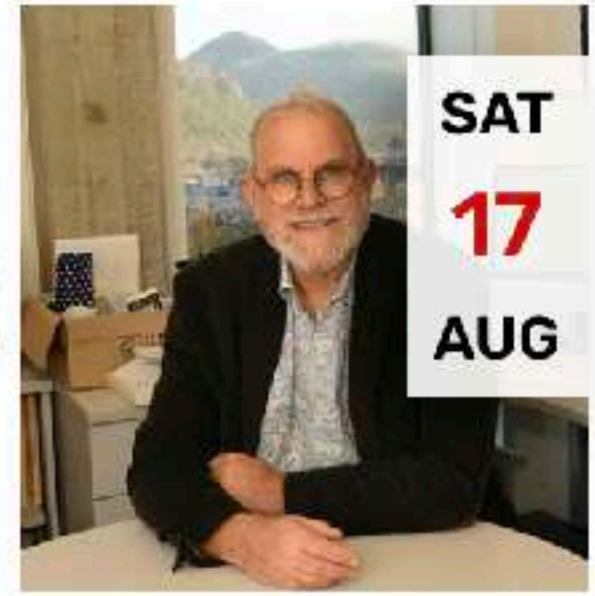

BULLSHIT BOTS

STAND 5

Doors: 12:10
Show time: 12:20 - 13:20
Duration: 60 mins

TICKETS AND INFO

ChatGPT is Bullshit

Philip Wadler
University of Edinburgh / IOG
Lambda Days, 28 May 2024