

Essentials of t

Teuvo Kohonen
Aalto University, School of S

ARTICLE I

Keywords:
Self-organizing map
SOM
Data analysis
Brain map
Similarity
Vector quantizat

from the other by racemization
D- and L-enantiomers for a
identical ¹⁴C ages.
Older, more poorly preserved
amenable to the determination
stable isotope method, as th
ment of indigenous amino
acids increases with time.
undergo racemization, the
isotopic compositions and
organisms. Such a circum
to recognize as more infor
the distribution and stabl
acid constituents of mod
Also, AMS dates on inc
in some cases, help to cl
when stratigraphic con
age range for the fossi
Finally, the develop
stable isotopic comp
enable us to establish
carbonaceous meteor
terrestrial contamin
M.H.E. thanks th
EAR-8352055) and
Young Investigato

Learning r by back-pr

David E. Rumelhart
& Ronald J. McClelland

* Institute for
San Diego, CA
† Department of
Pittsburgh, PA

We desc
networks

net an
adjust
or out

and t

of th
guis

Portable SIMD Performance with OpenMP* 4.x Compiler Directives

Florian Wende¹(cor), Matthias Noack¹, Thomas Steinke¹, Michael Klemm²,
Chris J. Newburn³, and Georg Zitzlsberger²
¹ Zuse Institute Berlin, Takustraße 7, 14195 Berlin, Germany
{wende,noack,steinke}@zib.de
² Intel Deutschland GmbH, 85579 Neubiberg, Germany
{michael.klemm,georg.zitzlsberger}@intel.com
³ Intel Corporation, Santa Clara, USA
chris.newburn@intel.com

Abstract. Effective vectorization is becoming increasingly important for high performance and energy efficiency on processors with wide SIMD units. Compilers often require programmers to identify opportunities for vectorization, using directives to disprove data dependences. **OpenMP 4.x SIMD directives strive to provide portability.** We investigate the ability of current compilers (GNU, Clang, and Intel) to generate SIMD code for microbenchmarks that cover common patterns in scientific codes and for two kernels from the VASP and the MOM5/ERA-Interim applications. We explore coding strategies for improving SIMD performance across different compilers and platforms (Intel® Xeon® Phi™ processor and Intel® Xeon Phi™ (co)processor). We compare OpenMP SIMD vectorization with and without vector data types against intrinsics and C++ SIMD types. **Our experiments show that in cases portable performance can be achieved.** All microbenchmarks available as open source are a reference for programmers and experts to evaluate SIMD code generation.

1 Introduction

Omar Awile (omar.awile@cern.ch)
27.1.2017

A worrying trend



Sensationalist or misleading news



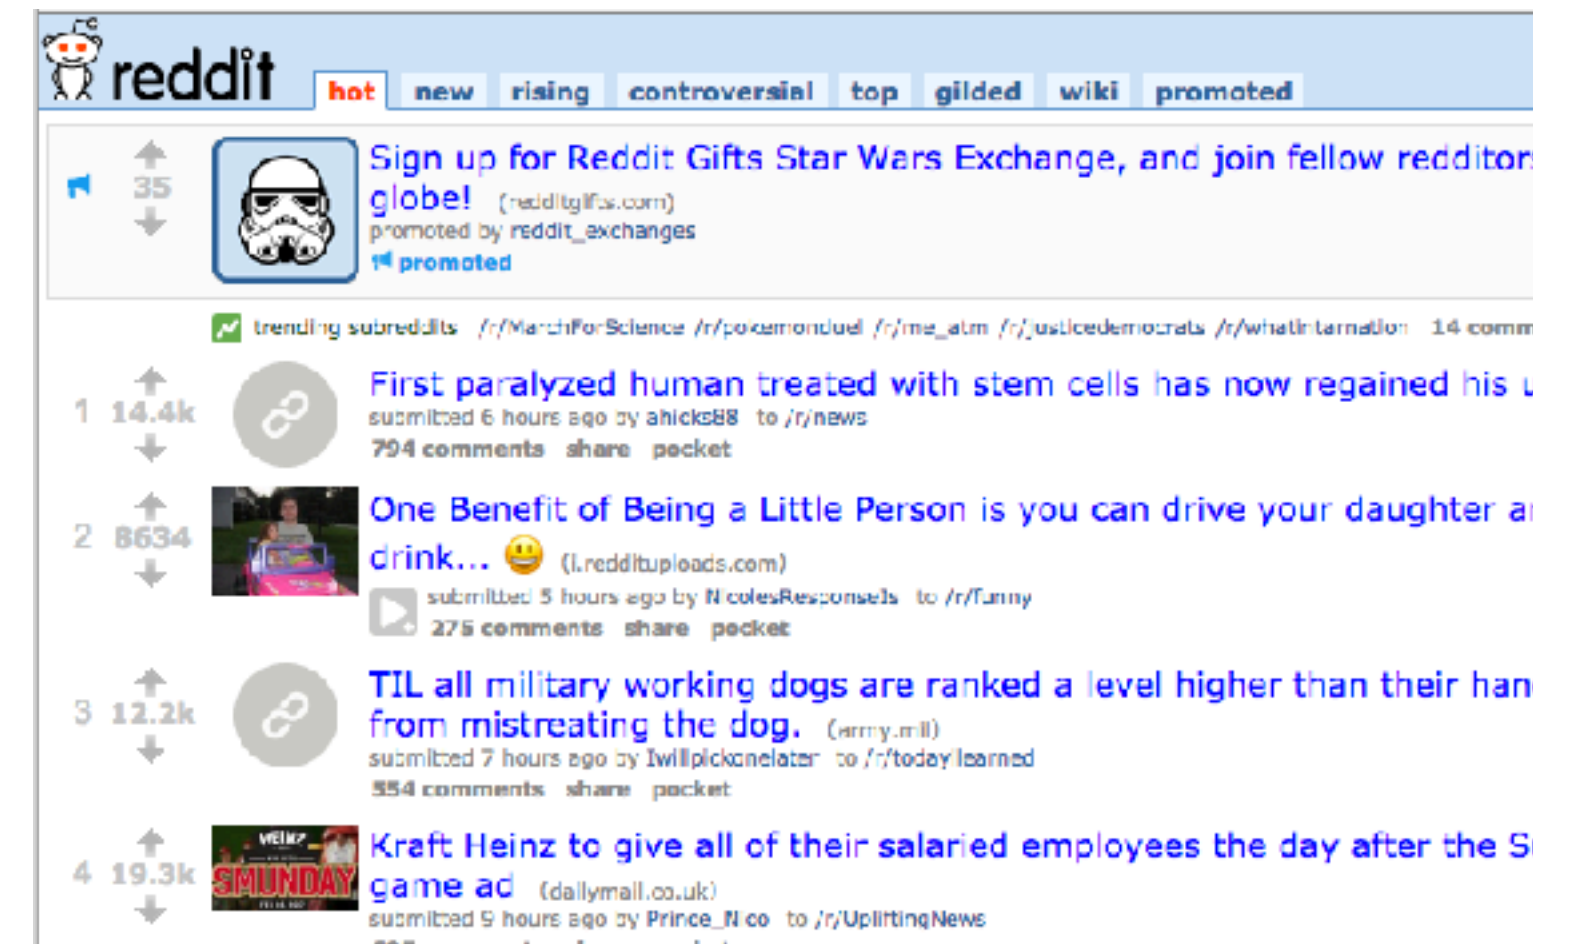
social media bubble and confirmation bias



Information overflow



Bad science journalism



Poor discussion culture

Practicing an essential skill in four simple steps

1

pick a text

Practicing an essential skill in four simple steps

1

pick a text



Practicing an essential skill in four simple steps

1

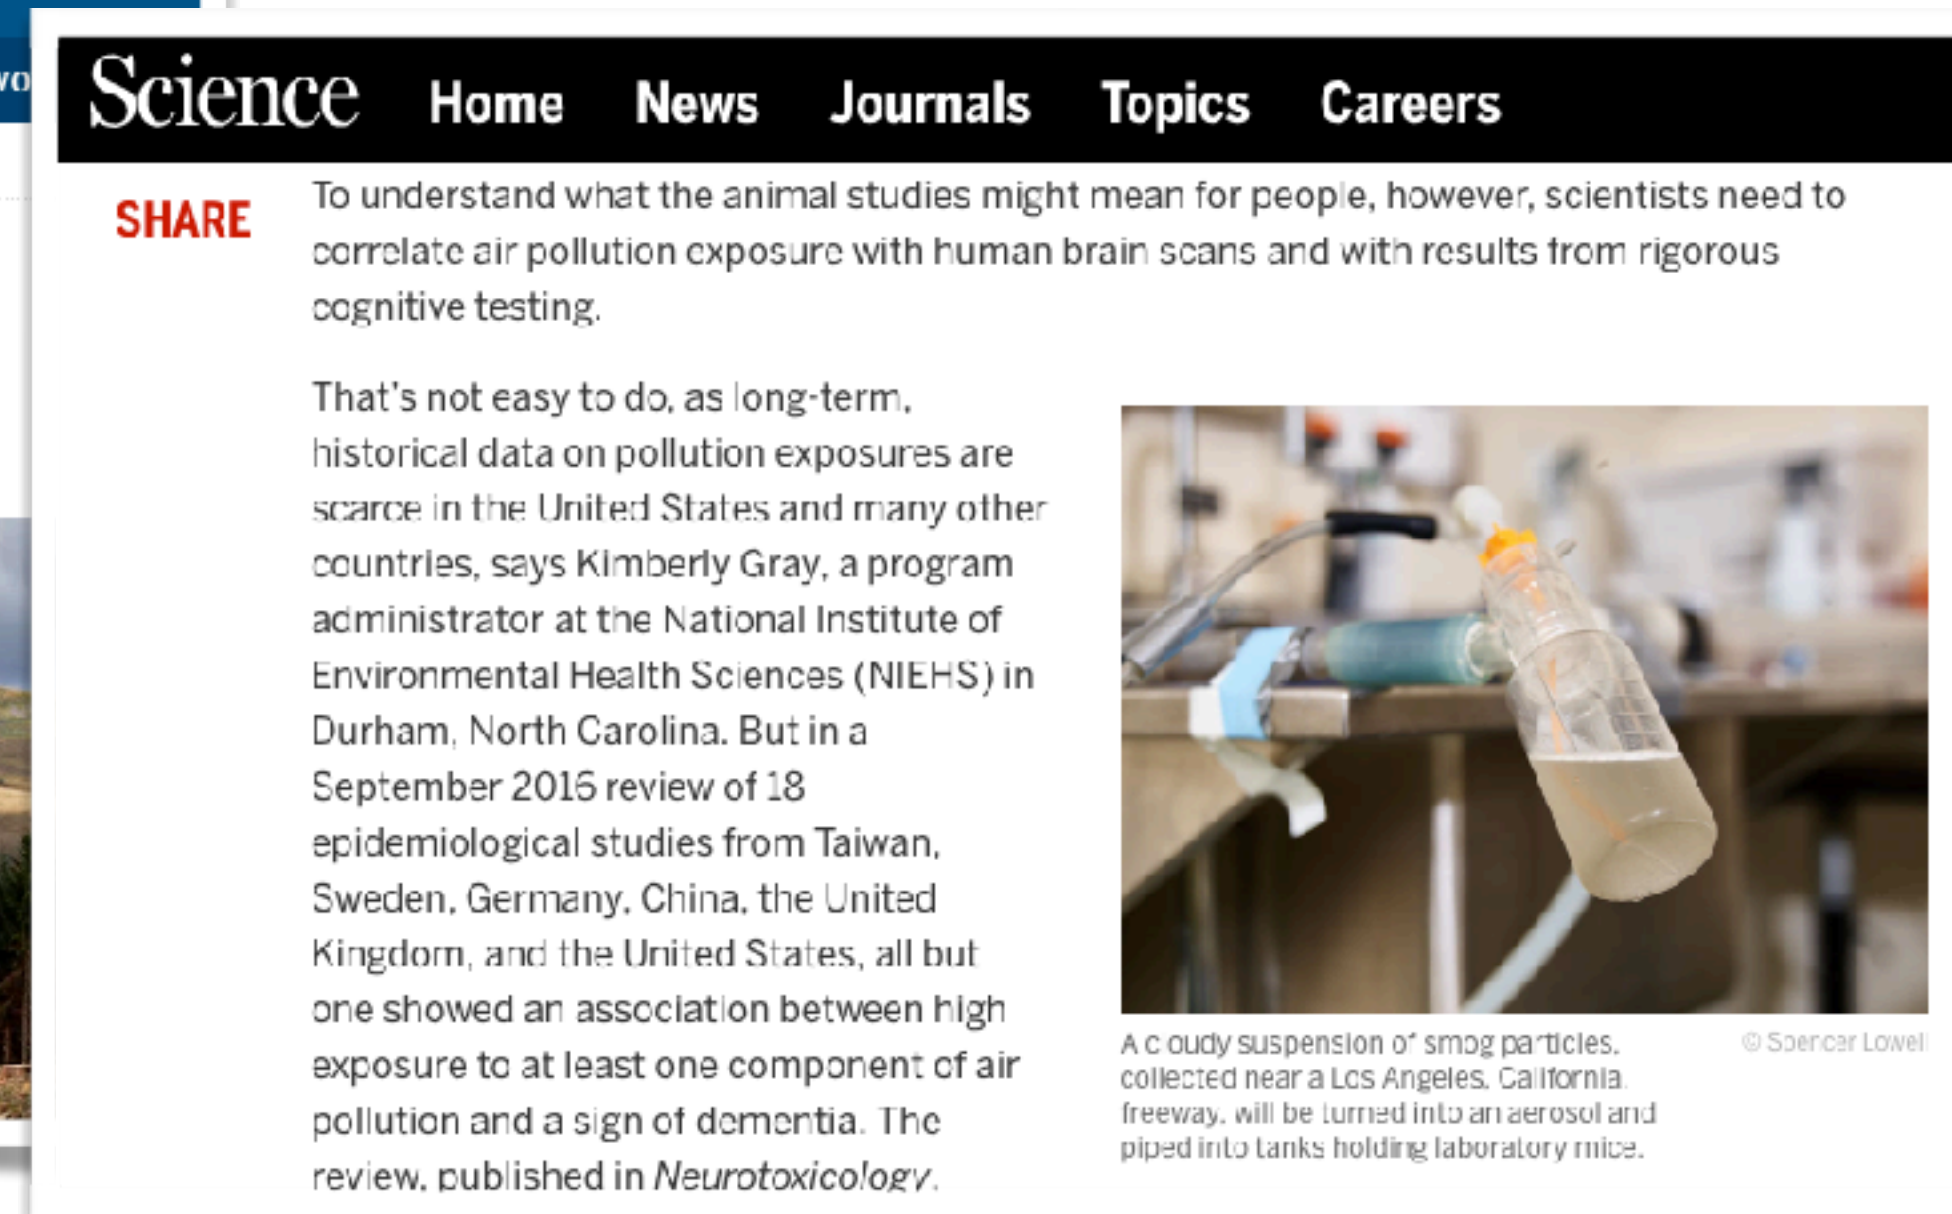

pick a text

The image shows two overlapping screenshots of news articles. The left screenshot is from the Guardian website, featuring a blue header with the logo and a navigation bar with links for 'home', 'environment', 'climate change', 'wildlife', 'energy', 'pollution', 'UK', and 'world'. The main article title is 'We may be closer than we thought to dangerous climate thresholds' with a sub-headline 'A new study identifies an extra 0.1°C of human-caused warming'. Below the text is a photograph of a landscape with rolling hills and a forest. The right screenshot is from a Science website, with a black header containing the logo and navigation links for 'Home', 'News', 'Journals', 'Topics', and 'Careers'. It features a 'SHARE' button and a paragraph of text discussing the correlation between air pollution and dementia. To the right of the text is a photograph of a laboratory setup with a pipette and a beaker. A caption below the photo reads: 'A cloudy suspension of smog particles, collected near a Los Angeles, California, freeway, will be turned into an aerosol and piped into tanks holding laboratory mice.' The photo credit is '© Spencer Lowell'.

Practicing an essential skill in four simple steps

1


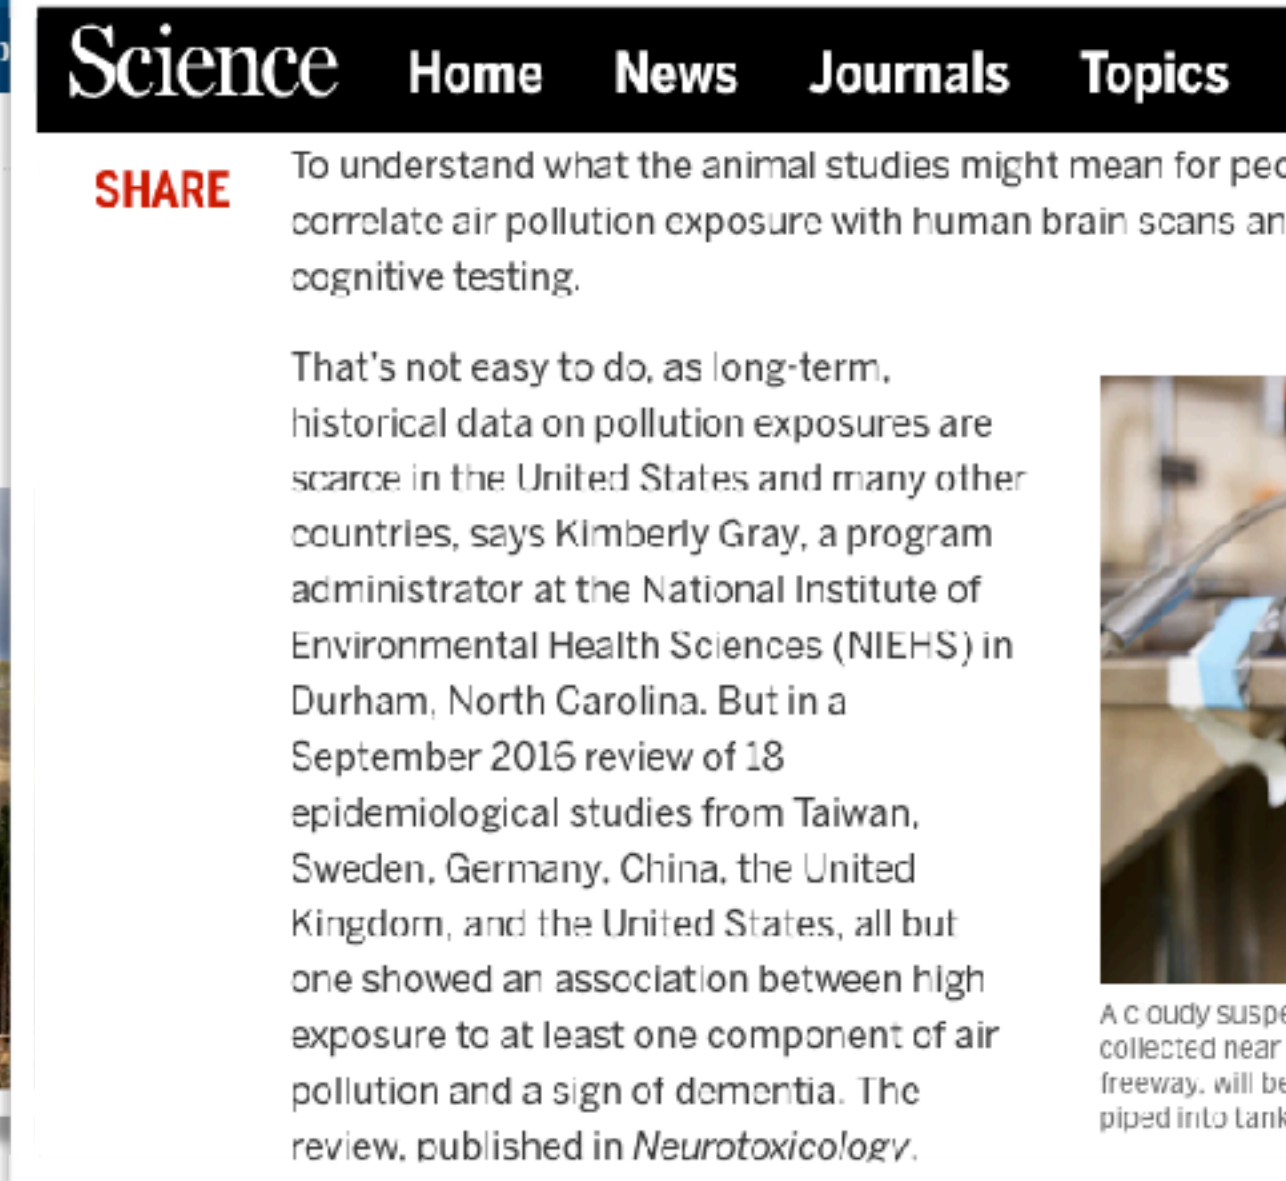

pick a text



Practicing an essential skill in four simple steps

1



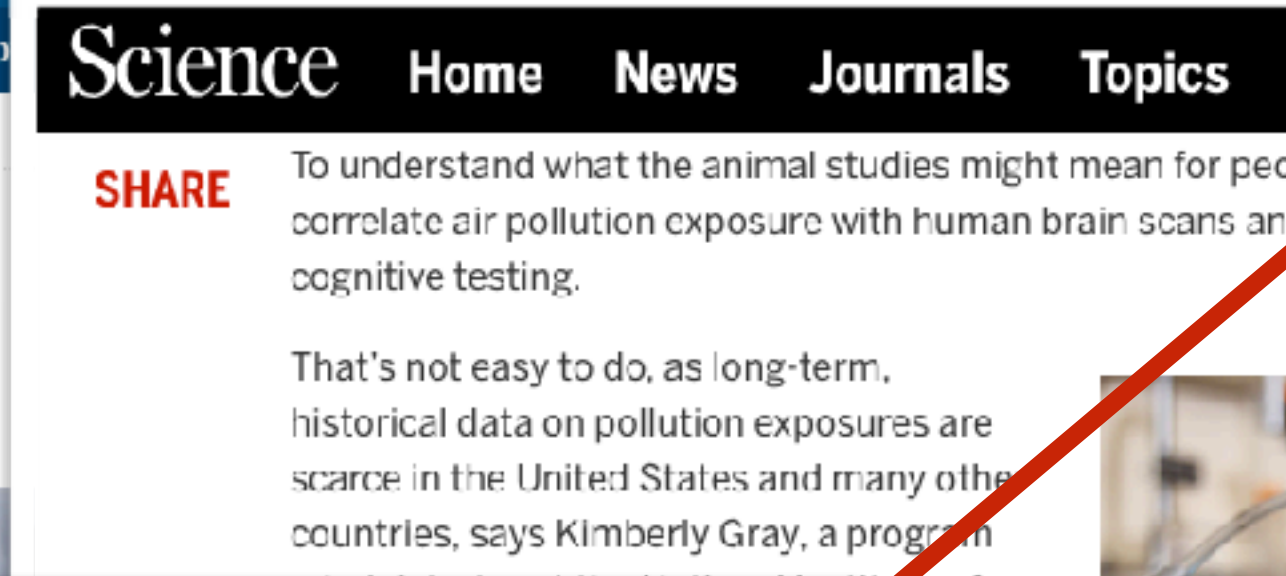

pick a text



Practicing an essential skill in four simple steps

1

pick a text



Practicing an essential skill in four simple steps


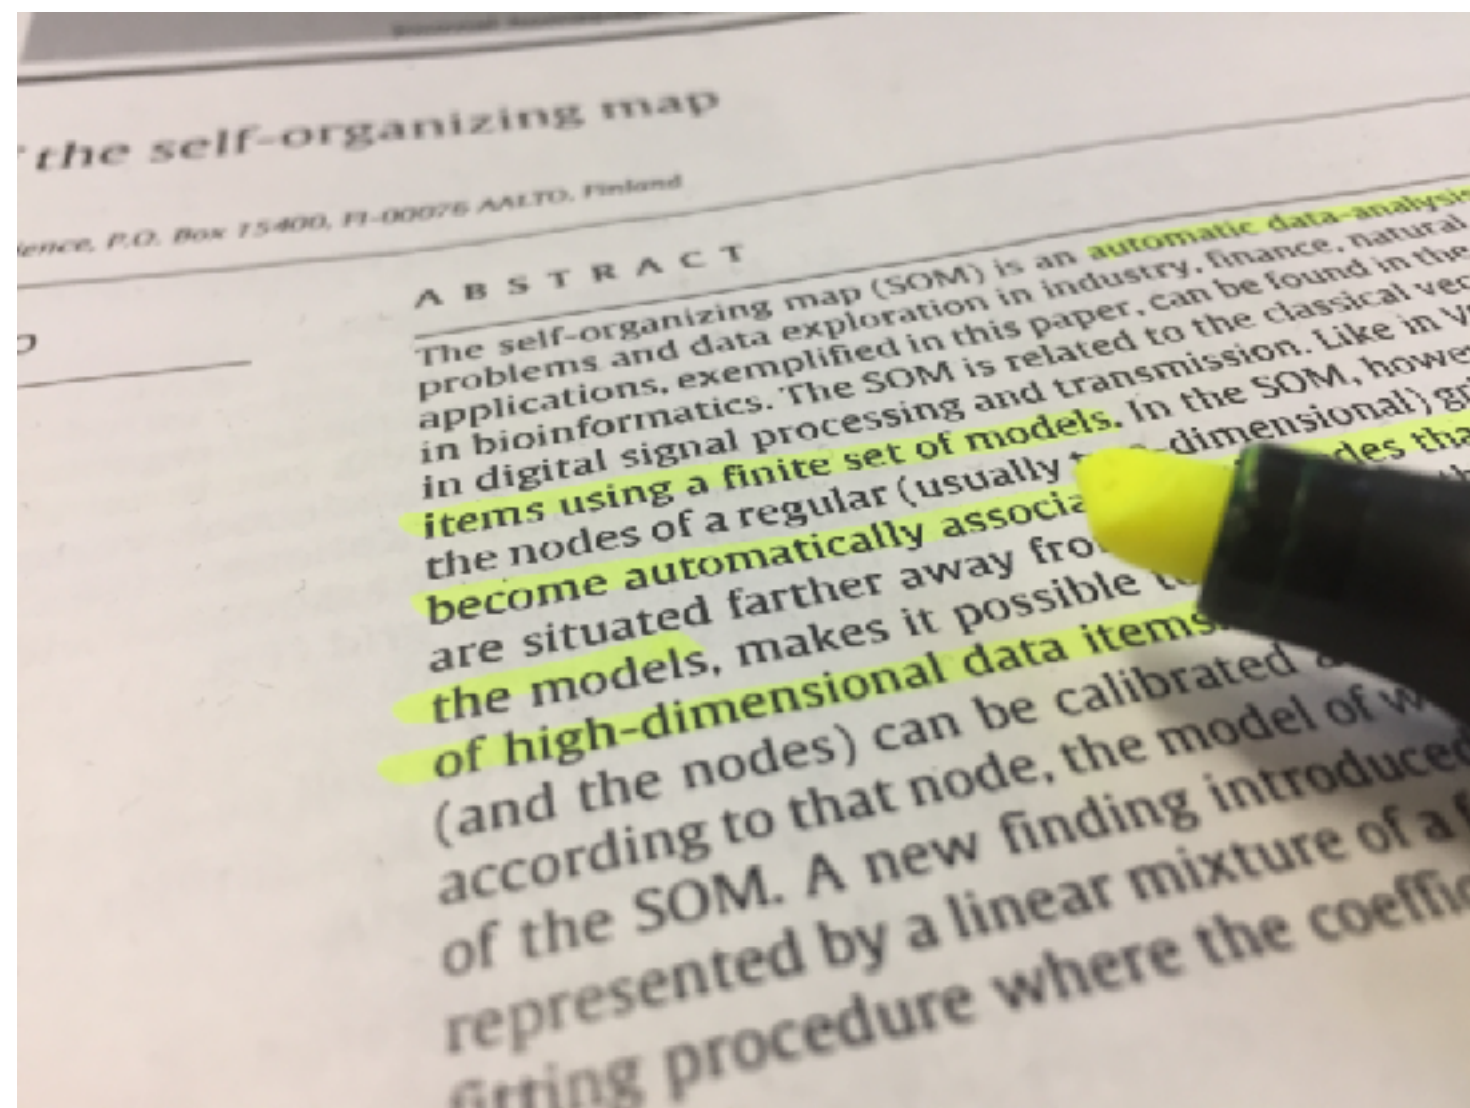
2

read consciously

Practicing an essential skill in four simple steps


2

read consciously



- highlight
- take notes
- make drawings

- check sources
- compare with current knowledge
- be critical
- ask questions (and write them down)



Practicing an essential skill in four simple steps

3

discuss in your group


Practicing an essential skill in four simple steps

3

discuss in your group

Every opinion counts!
start with a roundtable

Ask questions!



summarize the highlights
what did you like?
what was the key point?

Be critical but factual
why did you like or dislike
something?

Practicing an essential skill in four simple steps

4

open doors

Practicing an essential skill in four simple steps

4

open doors



Connect what you know with what
you've learned


As a group!

Spark an idea!



~~openlab~~HTCC reading club

- Started in October 2016
- Every 2 weeks, 1 hour discussion session
- 9 articles read
- Topics: Performance modeling, machine learning, distributed computing, parallel programming
- Average 5 participants



Join us! openlab reading club

- Started in October 2016
- Every 2 weeks, 1 hour discussion session
- 9 articles read
- Topics: Performance modeling, machine learning, distributed computing, parallel programming
- Average 5 participants

<https://e-groups.cern.ch/e-groups/Egroup.do?egroupId=10238392>

<https://mattermost.web.cern.ch/it-dep/channels/openlab-readingclub>

To receive announcements

To continue the discussion



WWW.PHDCOMICS.COM