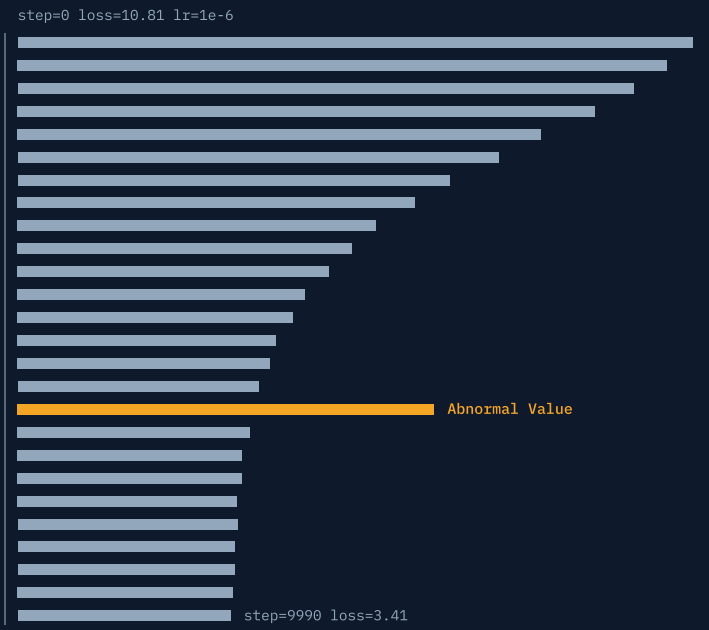


Machine Learning

Trace and Log Analysis Perspective

Dmitry Vostokov

OPENTASK

Machine Learning

Trace and Log Analysis Perspective

Dmitry Vostokov

OPENTASK

Published by OpenTask, Republic of Ireland

Copyright © 2026 by Dmitry Vostokov
Copyright © 2026 by OpenTask
Copyright © 2026 by Software Diagnostics Institute

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior written permission of the publisher.

Product and company names mentioned in this book may be trademarks of their owners.

OpenTask books are available through booksellers and distributors worldwide. For further information or comments, please send requests to press@opentask.com.

Developed with AI assistance (Claude Fable 5.2 and GPT-6 Astra). The concept, structure, and direction of the book are the author’s.

ISBN-13: 978-1-919135-10-6
Version 18 (November 2026)

Contents

Figures.....v
Preface.....vi

Part I. Traces of Machine Learning Systems.....1

1 The Inversion.....2
2 Trace Analysis Patterns in Brief.....8
3 Three Strata.....14
4 The Training Loop as a Trace.....20
5 Loss Curves and Counter Values.....30
6 Data Pipelines and Loaders.....40
7 Checkpoints, Resumption, and Truncated Traces.....48
8 Distributed Training as Threads of Activity.....56
9 Inference Serving Traces.....67
10 Agentic Runs.....75
11 Evaluation and Experiment Traces.....87
12 Anomalies in Machine Learning System Traces.....93

Part II. Machine Learning as Trace.....104

13 The Forward Pass as a Trace.....105
14 Tokenization as Text Trace.....114
15 Embeddings as Message Embedding.....120
16 Attention as Message Context.....125

17 The Backward Pass as an Adjoint Thread.....135
18 Parameter Updates as a Long Trace.....141
19 Dropout, Masks, and Silent Messages.....146
20 Batches as Strands of Activity.....151
21 Generation as Trace Production.....155
22 Generalization, Memorization, and Motifs.....162
23 What a Checkpoint Records.....167

Part III. Connections, Borrowings, and Limits….179

24 Interpretability as Trace Analysis.....180
25 Observability for Machine Learning.....187
26 Patterns the Language Borrowed from Machine Learning.....196
27 Four Directions, One Table.....202
28 Limits of Formalization.....206

Appendix A. Pattern Glosses.....215

Appendix B. Coverage of the Pattern Language.....239

Glossary.....258

Notes and Sources.....265

References.....273

Index.....281

Figures

2.1 One trace, five selections.....10
3.1 The three strata.....17
4.1 One reporting window, two reports.....25
5.1 Six shapes of a loss Signal.....33
6.1 A covered wait and two deliveries.....43
7.1 The join of a resumed run.....50
8.1 Four ranks at an all-reduce.....58
10.1 A late tool result, discarded and admitted.....82
11.1 The same score on different items.....89
12.1 Six histories of an absent message.....99
13.1 One computation at three granularities.....108
14.1 One text at three scales.....117
16.1 Evidence mass against its uniform baseline….133
17.1 Forward and backward traces, superimposed.....136
21.1 Decoding as path selection.....156
23.1 Two checkpoint saves.....176
28.1 The fields a question needs.....211

Preface

This book looks at machine learning from the point of view of trace analysis. The point of view is that of the trace and log analysis patterns of pattern-oriented software diagnostics: a language of more than two hundred named, recurrent structures in traces, each with a method of recognition and a diagnostic reading. The language was built for software traces in general. The book asks what it sees when it is turned on machine learning.

It sees two things, and the book is divided accordingly. The first half reads the traces that machine learning systems produce in the ordinary course of training, serving, and acting: training logs, loader logs, collective-communication logs, request logs, agent transcripts. Here the patterns apply literally, and the book's contribution is a catalog of readings. The second half reads machine learning itself, the forward pass, the backward pass, the optimization trajectory, the generation of output, and the trained model, as traces under stated encodings. Here the patterns apply structurally, and the book's contribution is the encodings and what the patterns say under them.

Every correspondence between a machine learning concept and a pattern is marked [L], [S], or [M] for literal, structural, or metaphorical, and Chapter 28

collects the marks. We have tried to make no claim that we cannot mark, and to mark no claim more strongly than it deserves.

The book has no exercises and proves nothing. It is a reading of one field in the vocabulary of another, written for the diagnostician who must debug a learning system and for the practitioner who suspects that the vocabulary of traces fits what they do. The book was written against the whole catalog of the pattern language; Appendix B records, for every pattern, whether and where it is used. The pattern glosses in Appendix A, the glossary, and the index are meant to be used while reading.

Trace excerpts are synthetic. They are modelled on common frameworks and are not copied from any real run.

The authorial plural throughout is the author's own.

The concept, structure, and direction of this book are the author's. The text was developed with AI assistance (Claude Fable 5.2 and GPT-6 Astra). This is Version 18.

Dublin, November 2026