

Index

- 5G, 433
- Accelerated Processing Unit (APU), 380
- Acceleration, 253
- Accelerator, 2, 263, 360
- Application-Specific Integrated Circuit (ASIC), 4
- Backpropagation, 139, 250, 256, 329, 330
- Bagging, 341
- Battery-powered device, 37, 46, 48, 52
- Bayes
 - naive Bayes classifier, 45, 316
 - optimization, 8, 95
- Bayesian network, 407
- Belief propagation, 407, 410, 411, 413
- Bernoulli experiment, 342
- Bidirectional Encoder Representations (BER), 163, 166, 167, 169, 174–176
- Bit
 - Error Rate (BER), 326, 330–334
 - flip, 10, 329–332, 334
- Boosting, 339, 340
- Cellular network, 433
- Central Processing Unit (CPU), 2, 26, 28, 31, 40, 41, 53, 259, 263, 264, 270, 271, 276, 277, 280, 281, 284, 287, 288, 291–304, 315, 340–343, 351, 352, 354–356, 406, 411, 419
 - architecture, 351
- Channel
 - distributed channel access, 425
- Classification, 5, 7, 81, 87, 174
 - multi-label, 272
- Classification And Regression Tree (CART), 351
- Cluster Feature Tree (CF-Tree), 219, 220, 222
- Clustering, 8, 10, 86, 87, 186, 187, 197, 228, 230, 290, 291
 - (k,l), 202
 - biclustering, 228
 - graph, 140
 - hierarchical, 215, 221
 - k-means, 200, 204, 212
 - k-medoids, 182, 185
- Code generation, 263, 350
- Communication
 - awareness, 74
 - network, 7, 425, 430
 - technologies, 423
- Compression, 11, 111, 112, 114, 145, 157–159
- Confidence interval, 79
- Coprocessor, 4, 271, 380
- Coresets, 10, 86–92, 95, 200, 212, 213
- Covariance, 94
- Cyber-physical system, 3, 45, 89, 285, 369, 370
- Data
 - acquisition, 16–20, 28, 49, 62, 66
 - packet, 52, 425
 - parallelism, 257
 - stream, 8, 10, 71, 74, 80, 85, 87, 88, 406
 - stream algorithms, 5, 10, 71, 74, 88, 89, 213
 - summary, 74, 86, 89
 - volume, 272, 427
- Deadlines, 291, 292, 367
- Decision tree, 7, 11, 45, 341, 343, 344, 350, 355
- Dependency Graph Approach (DGA), 360, 362
- Design space exploration, 426
- Directed-Acyclic Graph (DAG), 367
- Dynamic Power Management (DPM), 52
- Dynamic voltage scaling, 57
- Embedded system, 2–8, 15, 38, 47, 52, 60, 74, 89, 101, 285, 286, 363, 424
- Energy
 - awareness, 34, 35, 39, 41, 47, 48, 52, 54, 55, 59, 65, 66
 - consumption, 3–10, 32, 34, 37, 39–46, 48, 52–59, 61, 65, 66, 68, 74, 249, 281, 300, 302–304, 326, 327, 406, 407, 424–429
 - efficiency, 2, 6, 34, 36–38, 41, 45, 46, 48, 49, 56, 62, 67, 68
 - harvesting, 34, 37, 47–59, 61, 65–68, 424
 - measurement, 6, 40, 43, 50, 52

- Ensembles, 11, 339–342
 - tree, 11
- Error measure, 188
 - accuracy, 294, 295
 - mean average precision, 173
 - precision, 173
- Euclidean distance, 123, 216, 218, 220, 222, 290
- Event streams, 15, 19, 20, 26
- Evolutionary algorithms, 7, 286
- Exponential families, 9, 11, 90, 104, 409
- Ferroelectric Field-Effect Transistors (FeFET), 326–328
- Field-Programmable Gate Array (FPGA), 7, 8, 10, 249, 250, 340
- Finite State Machine (FSM), 41, 42, 65
- Floating point arithmetic, 328
- Gradient descent, 232, 415
 - stochastic, 329, 420
- Graph partitioning, 145, 152, 155
- Graphical models, 89, 92, 103, 106, 408, 411, 416
 - probabilistic, 408
- Graphics Processing Unit (GPU), 2, 8, 9, 140, 249, 253, 263–265, 270, 271, 300, 340, 341, 360, 361, 379, 407
- Heterogeneous processors, 300
- Idle state, 431
- Indicator function, 103
- Inference, 10, 11, 108, 137, 139, 166, 340, 341, 350
 - probabilistic, 105, 407, 410, 411, 413
 - quadrature-based, 105
 - variational, 105
- Instrumentation, 16, 17, 19, 26, 42
- Internet of Things (IoT), 3, 6, 38, 46–50, 74, 85, 423, 433
- k-nearest Neighbor (kNN), 45, 102, 316
- Kalman filter, 101
- Kernel functions, 5, 81, 96, 102, 118, 121–123, 294, 347, 355
 - graph, 116, 118, 120, 124, 125
 - hash graph, 124
 - RBF, 81
 - triangular, 168
 - Weisfeiler Leman graph, 120–125, 127
- Knowledge, 304
 - a priori, 93, 161
 - background, 239
- L1 norm, 95, 105, 111, 113
- L2 norm, 94, 105, 111, 113
- Latency, 30, 31, 277, 342, 343, 427
- Learning
 - contrastive, 174, 175
 - embeddings, 161, 162, 167
 - self-supervised, 161, 163, 167
 - supervised, 42, 272, 341
 - unsupervised, 10, 102
- Learning tasks, 5, 9, 81, 96, 100, 163
 - masking tasks, 163
- Leave-one-out, 171
- Light, 57, 59, 61, 68
- Linkage
 - centroid, 216–218
 - complete, 216, 290
 - median, 216–218
 - single, 216
 - Ward, 216
- Locking protocol, 361
- Log-likelihood, 105, 113
- Long Term Evolution (LTE), 423
- Markov Random Fields (MRFs), 9, 11, 102, 112–114, 408, 416, 420, 421
- Matrix factorization, 11
 - binary, 234
 - Boolean, 235
 - nonnegative, 229
 - objectives, 230
- Max-Cut, 144
- Maximum a posteriori, 105, 408, 409, 417, 420
- Maximum likelihood, 93, 105, 408, 409, 413, 419
- McQuitty's Weighted Pair-Group Method with Arithmetic mean (WPGMA), 216, 217, 222
- Memory, 9, 10, 74–78, 82, 83, 136, 157, 184, 222, 307, 308, 313, 327, 344, 347
 - allocation, 277, 284, 309–313
 - architecture, 276, 277, 342
 - bottleneck, 3, 277, 281, 284
 - cache, 3, 281, 284, 340–345, 347
 - capacity, 3
 - footprint, 108, 110, 111, 137, 138, 140, 307, 309, 313, 314, 329

- hierarchy, 3, 276, 281, 284, 343
- layout, 27, 313, 341, 342, 345
- locality, 276, 277, 342
- Magnetoresistive Random Access Memory (MRAM), 4
- Non-Uniform Memory Access (NUMA), 276
- non-volatile, 10, 326
- physical, 308, 309
- scratchpad, 3, 381
- shared, 145, 152, 158, 310, 381
- Static Random-Access Memory (SRAM), 4, 67, 326, 328
- virtual, 308–311
- Merge & Reduce, 87, 88
- Message passing, 8, 10, 105, 125, 130–132, 143, 410, 412, 413
- Model execution, 341
- Monte Carlo, 287, 288, 294, 297
- Multicore, 4, 8, 33, 249, 275, 285, 296, 304
- Multilayer perceptrons, 8, 137, 250

- Nearest-neighbor chain algorithm, 222
- Nearest-neighbor search, 174
- Neural Network
 - Binarized Neural Networks (BNN), 10, 326, 328–333
 - Convolutional Neural Networks (CNN), 62, 130, 162, 163, 333
 - Deep Neural Networks (DNN), 15, 61, 328, 434
 - Graph Neural Networks (GNN), 8–10, 116, 117, 125–128, 130–135, 137–141, 161–163, 165
- Neural network
 - Convolutional Neural Networks (CNN), 254
- Offloading, 4, 369, 378
- One-vs-Rest Classification, 274
- Operating system kernel, 17–20, 22, 23, 27–30, 32, 33, 310
- data, 17, 22
- Parallelism, 8, 249, 275, 287, 386
- Partition function, 104
- Peripheral Component Interconnect express (PCIe), 380
- Permission, 32, 61, 310
- Poisson dependency network, 92
- Potential function, 105, 411
- Power
 - consumption, 2, 3, 6, 7, 48, 55, 58, 64, 65, 256, 300, 301, 426, 429–433
 - measurement, 42, 45, 47, 49, 50, 58, 59, 300, 301
- Probability density function (pdf), 95, 104
- Processor utilization, 288
- Proximal Optimization, 233
- Pruning, 340
- Quantization, 10
- Query compilation, 386
- Radio, 49, 425, 433
 - mobile radio networks, 429
 - Software-Defined Radio (SDR), 433
- Random forest, 45, 341, 434
- Real-time system, 34, 360, 361, 370, 378
- Regression, 5, 6, 10, 43, 85–87, 90, 102, 285–292, 298, 316, 341, 434, 435
- Bayesian, 89, 93, 94
- generalized linear regression models, 89, 90, 92
- hierarchical, 94
- LASSO, 95
- linear, 6, 10, 86, 89, 90, 92, 93
- logistic, 90–92, 316
- ordinary least squares, 92
- Poisson, 92
- probit, 90, 95
- Regularization, 5, 11, 93, 103, 105, 110–114, 234, 236–238, 275, 331, 333
- Representation
 - Term-frequency inverse-document-frequency (Tf-idf), 272
- Representation learning, 129, 130, 167, 174, 176
- Resource
 - awareness, 288, 289, 292
 - efficiency, 227, 286
 - efficient transmission, 425
 - optimization, 429, 430
 - synchronization, 361
 - utilization, 16, 263, 268, 285, 287, 297, 298, 304, 430
- Resource Block (RB), 432
- Resource-constrained, 47, 61
- Resource-constrained learning, 101, 102,

- 108, 110, 111, 113, 136
- Rule of Three, 79, 80
- Runtime
 - estimation, 289, 291
- Sampling, 5, 6, 9, 10, 48, 64, 67, 85, 90–92, 130, 136, 139, 167, 168, 201, 203, 240, 287, 288, 294, 301
 - importance, 86, 89, 91, 93, 136
 - layer-wise, 136
 - node-wise, 136
 - reservoir, 76
 - sub-sampling, 137, 139
 - subgraph, 136
 - uniform, 93
- Scalability, 10, 87, 108, 130, 137, 140, 143, 223, 224, 293, 294, 304, 426
- Scheduling, 8, 9, 26, 27, 39, 285, 286, 288–292, 294, 295, 297–300, 302, 304, 366, 411, 432
 - Earliest Deadline First (EDF), 366
 - federated, 367
 - List-Earliest Deadline First (List-EDF), 366
 - non-preemptive fixed-priority, 363
 - Partitioned Earliest-Deadline-First (P-EDF), 367
 - preemptive fixed-priority, 375
 - Worst-Fit Partitioned Earliest-Deadline-First (WF-P-EDF), 368
- Security, 31, 68
- Sensor network, 6, 15, 34, 74, 101, 106
- Shannon entropy, 104
- Signal
 - digital processing, 430
 - Received Signal Strength Indicator (RSSI), 45
- Similarity measure, 117, 121, 162, 197, 198, 200, 204
 - cosine, 162, 174, 175
 - dynamic time warping, 199
 - Fréchet distance, 198, 199
- Single Instruction, Multiple Data (SIMD), 386
- Sketches, 10, 85, 86, 88–91, 93, 284
- Spatio-temporal state prediction, 100, 102, 103, 106, 111, 113, 114
- Spectral efficiency, 425
- Speedup, 184, 223, 279, 281, 282, 322
- Star Schema Benchmark (SSB), 381
- Submodular functions, 10, 74–77, 80, 84
- Sufficient statistics, 103, 104, 409, 410
- Support vector data description, 212
- Support Vector Machine (SVM), 10, 45, 316, 332
- Synchronization, 4, 140, 284
- Task period, 368
- Tensor Processing Unit (TPU), 2
- Test bed, 34–38, 40, 45
- Thread parallelism, 275
- Tiling, 214
- Time series data, 62, 65, 101, 197, 198
- Transformer, 163, 165, 166, 176, 283
- Transmission power, 44, 425, 430, 433
- Ultra-low power device, 9
- Ultra-low power state, 44, 52, 59, 67
- Uplink power, 433
- User Equipment (UE), 423, 425, 430, 431
- Variance, 96, 139, 291
- von Neumann architecture, 2, 3
- Wasserstein distance, 94, 122, 123
- Wireless Sensor Network (WSN), 35, 36
- Worst Case Execution Time (WCET), 363, 370
- Worst-Fit Heuristic, 367