AUGMENTED LAGRANGIAN METHODS pdf

1: Augmented Lagrangian method - Wikipedia

Augmented Lagrangian methods are a certain class of algorithms for solving constrained optimization problems. They have similarities to penalty methods in that they replace a constrained optimization problem by a series of unconstrained problems and add a penalty term to the objective; the difference is that the augmented Lagrangian method adds yet another term, designed to mimic a Lagrange.

These algorithms are listed below, including links to the original source code if any and citations to the relevant articles in the literature see Citing NLopt. I apologize in advance to the authors for any new bugs I may have inadvertantly introduced into their code. Nomenclature Each algorithm in NLopt is identified by a named constant, which is passed to the NLopt routines in the various languages in order to select a particular algorithm. Many of the algorithms have several variants, which are grouped together below. Comparing algorithms For any given optimization problem, it is a good idea to compare several of the available algorithms that are applicable to that problemâ€”in general, one often finds that the "best" algorithm strongly depends upon the problem at hand. That is, ask how long it takes for the two algorithms to reach the same function value. Better yet, run some algorithm for a really long time until the minimum fM is located to high precision. Then run the different algorithms you want to compare with the termination test: Global optimization All of the global-optimization algorithms currently require you to specify bound constraints on all the optimization parameters. However, any of them can be applied to nonlinearly constrained problems by combining them with the augmented Lagrangian method below. Something you should consider is that, after running the global optimization, it is often worthwhile to then use the global optimum as a starting point for a local optimization to "polish" the optimum to a greater accuracy. Many of the global optimization algorithms devote more effort to searching the global parameter space than in finding the precise position of the local optimum accurately. Stuckmann, "Lipschitzian optimization without the lipschitz constant," J. Optimization Theory and Applications, vol. These are deterministic-search algorithms based on systematic division of the search domain into smaller and smaller hyperrectangles. The Gablonsky version makes the algorithm "more biased towards local search" so that it is more efficient for functions without too many local minima. NLopt contains several implementations of both of these algorithms. If your dimensions do not have equal weight, e. However, the unscaled variations make the most sense if any with the original DIRECT algorithm, since the design of DIRECT-L to some extent relies on the search region being a hypercube which causes the subdivided hyperrectangles to have only a small set of side lengths. Finally, NLopt also includes separate implementations based on the original Fortran code by Gablonsky et al. These implementations have a number of hard-coded limitations on things like the number of function evaluations; I removed several of these limitations, but some remain. Most of the above algorithms only handle bound constraints, and in fact require finite bound constraints they are not applicable to unconstrained problems. They do not handle arbitrary nonlinear constraints. Ali, "Some variants of the controlled random search algorithm for global optimization," J. The original CRS2 algorithm was described by: Price, "A controlled random search procedure for global optimization," in Towards Global Optimization 2, p. Szego North-Holland Press, Amsterdam, Price, "Global optimization by controlled random search," J. The CRS algorithms are sometimes compared to genetic algorithms, in that they start with a random "population" of points, and randomly "evolve" these points by heuristic rules. In this case, the "evolution" somewhat resembles a randomized Nelder-Mead algorithm. The published results for CRS seem to be largely empirical; limited analytical results about its convergence were derived in: Only bound-constrained problems are supported by this algorithm. Rinnooy Kan and G. Timmer, "Stochastic global optimization methods," Mathematical Programming, vol. Actually 2 papers â€" part I: Sergei Kucherenko and Yury Sytsko, "Application of deterministic low-discrepancy sequences in global optimization," Computational Optimization and Applications, vol. In either case, MLSL is a "multistart" algorithm: MLSL is distinguished, however by a "clustering" heuristic that helps it to avoid repeated searches of the same local optima, and has some theoretical guarantees of finding all local optima in a finite number of local minimizations. Note that it is perfectly reasonable to set a relatively large tolerance for these local

searches, run MLSL, and then at the end run another local optimization with a lower tolerance, using the MLSL result as a starting point, to "polish off" the optimum to high precision. StoGO is a global optimization algorithm that works by systematically dividing the search space which must be bound-constrained into smaller hyper-rectangles via a branch-and-bound technique, and searching them by a gradient-based local-search algorithm a BFGS variant , optionally including some randomness hence the "Sto", which stands for "stochastic" I believe. Some references on StoGO are: Gudmundsson, "Parallel Global Optimization," M. Zilinskas, "Global Optimization using Branch-and-Bound," unpublished A preprint of this paper is included in the stogo subdirectory of NLopt as paper. A copy of this report is included in the stogo subdirectory of NLopt as techreport. AGS This algorithm adapted from this repo. AGS can handle arbitrary objectives and nonlinear inequality constraints. Also bound constraints are required for this method. To guarantee convergence, objectives and constraints should satisfy the Lipschitz condition on the specified hyperrectangle. AGS is derivative-free and employs the Hilbert curve to reduce the source problem to the univariate one. The algorithm divides the univariate space into intervals, generating new points by using posterior probabilities. On each trial AGS tries to evaluate the constraints consequently one by one. If all constraints are preserved, i. Thus, some of constraints except the first one and objective can be partially undefined inside the search hyperrectangle. It supports up to 10 dimensions, but the method can stop early in case of 6 and more ones. Additional parameters of AGS which are not adjustable from the common NLOpt interface are declared and described in ags. Also an example of solving a constrained problem is given in the AGS source folder. An algorithm for solving global optimization problems with nonlinear constraints, Journal of Global Optimization, 7 4 , pp â€”, Strongin R. Global optimization with non-convex constraints. Sequential and parallel algorithms. Kluwer Academic Publishers, Dordrecht. Applications and Reviews, vol. It is a refinement of an earlier method described in: This is an independent implementation by S. Johnson based on the papers above. Runarsson also has his own Matlab implemention available from his web page here. The evolution strategy is based on a combination of a mutation rule with a log-normal step-size update and exponential smoothing and differential variation a Nelderâ€”Mead-like update rule. The fitness ranking is simply via the objective function for problems without nonlinear constraints, but when nonlinear constraints are included the stochastic ranking proposed by Runarsson and Yao is employed. The algorithm is adapted from ideas described in: Local derivative-free optimization Of these algorithms, only COBYLA currently supports arbitrary nonlinear inequality and equality constraints; the rest of them support bound-constrained or unconstrained problems only. Powell, "A direct search optimization method that models the objective and constraint functions by linear interpolation," in Advances in Optimization and Numerical Analysis, eds. Dordrecht, , p. Powell, "Direct search algorithms for optimization calculations," Acta Numerica 7, The original code itself was written in Fortran by Powell and was converted to C in by Jean-Sebastien Roy js jeannot. First, we incorporated all of the NLopt termination criteria. Second, we added explicit support for bound constraints although the original COBYLA could handle bound constraints as linear constraints, it would sometimes take a step that violated the bound constraints. Also, we support unequal initial-step sizes in the different parameters by the simple expedient of internally rescaling the parameters proportional to the initial steps , which is important when different parameters have very different scales. Equality constraints are automatically transformed into pairs of inequality constraints, which in the case of this algorithm seems not to cause problems. Powell, converted to C and modified for the NLopt stopping criteria. BOBYQA performs derivative-free bound-constrained optimization using an iteratively constructed quadratic approximation for the objective function. Because BOBYQA constructs a quadratic approximation of the objective, it may perform poorly for objective functions that are not twice-differentiable. The NLopt BOBYQA interface supports unequal initial-step sizes in the different parameters by the simple expedient of internally rescaling the parameters proportional to the initial steps , which is important when different parameters have very different scales. The original NEWUOA performs derivative-free unconstrained optimization using an iteratively constructed quadratic approximation for the objective function. Because NEWUOA constructs a quadratic approximation of the objective, it may perform poorly for objective functions that are not twice-differentiable. In my bound-constrained variant, we use the MMA algorithm for these subproblems to

solve them with both bound constraints and a spherical trust region. The appropriate reference seems to be: Reprinted by Dover, This seems to work, more-or-less, but appears to slow convergence significantly. Mead, "A simplex method for function minimization," The Computer Journal 7, p. This method is simple and has demonstrated enduring popularity, despite the later discovery that it fails to converge at all for some functions and examples may be constructed in which it converges to point that is not a local minimum. I would tend to recommend the Subplex method below instead, however. The main change compared to the paper is that I implemented explicit support for bound constraints, using essentially the method proposed in: Box, "A new method of constrained optimization and a comparison with other methods," Computer J. Kuester, "The complex method for constrained optimization," Commun. ACM 16 8 , In any case, this collapse of the simplex is somewhat ameliorated by restarting, such as when Nelder-Mead is used within the Subplex algorithm below. Since the algorithm is not too complicated, however, I just rewrote it. This seems to be a big improvement in the case where the optimum lies against one of the constraints. Krister Svanberg, " A class of globally convergent optimization methods based on conservative convex separable approximations ," SIAM J.

AUGMENTED LAGRANGIAN METHODS pdf

2: NLopt algorithms - NLopt Documentation

Augmented Lagrangian method is one of the algorithms in a class of methods for constrained optimization that seeks a solution by replacing the original constrained problem by a sequence of unconstrained subproblems.

Show Context Citation Context We assume the solution set of 1. As in [8], with the attempt to treat the original ADM in [5] and the split inexact Uzawa method in [9, 10] uniformly, we study the following ADM We describe an effective solver for the discrete Oseen problem based on an augmented Lagrangian formulation of the corresponding saddle point system. The proposed method is a block triangular preconditioner used with a Krylov subspace iteration like BiCGStab. The crucial ingredient is a novel mult Our analysis indicates that this approach results in fast convergence, independent of the mesh size and largely insensitive to the viscosity. We present experimental evidence for both isoP2-P0 and isoP2-P1 finite elements in support of our conclusions. We also show results of a comparison with two state-of-the-art preconditioners, showing the competitiveness of our approach. Sparse inverse covariance selection via alternating linearization methods by Katya Scheinberg, Shiqian Ma, Donald Goldfarb , " Gaussian graphical models are of great interest in statistical learning. Because the conditional independencies between different nodes correspond to zero entries in the inverse covariance matrix of the Gaussian distribution, one can learn the structure of the graph by estimating a sparse inverse co Numerical experiments on both synthetic and real data from gene association networks show that a practical version of this algorithm outperforms other competitive algorithms. Olshanskii, Zhen Wang , " We study different variants of the augmented Lagrangian-based block triangular preconditioner introduced by the first two authors in [SIAM J. The preconditioners are used to accelerate the convergence of the Generalized Minimal Residual GMRES method applied to various finite element and MAC discretizations of the Oseen problem in two and three space dimensions. Both steady and unsteady problems are considered. Numerical experiments show the effectiveness of the proposed preconditioners for a wide range of problem parameters. Implementation on parallel architectures is also considered. The augmented Lagrangian-based approach is further generalized to deal with linear systems from stabilized finite element discretizations. Other relevant work includes the A preconditioner for generalized saddle point problems by Michele Benzi, Gene, H. Appl , " In this paper we consider the solution of linear systems of saddle point type by preconditioned Krylov subspace methods. The potential of this approach is illustrated by numerical Show Context Citation Context We further assume that matrices A, B, and C are large and sparse. Systems of the form 1. Discretization of the Stokes equations produces a symmetric indefinite system of linear equations. For stable discretizations, a variety of numerical methods have been proposed that have rates of convergence independent of the mesh size used in the discretization. In this paper, we compare the perfo In this paper, we compare the performance of four such methods: The results indicate that where it is applicable, multigrid with smoothing based on incomplete factorizaton is more efficient than the other methods, but typically by no more than a factor of two. The conjugate residual method has the advantages of being both independent of iteration parameters and widely applicable. This system, the Stokes equations, is a fundamental problem arising in computationalsuid dynamics, see e. Discretization of 1 bysnite dierence orsnite element techniques leads to a linear system of equations of the form A B Combettes , " The principle underlying this paper is the basic observation that the problem of simultaneously solving a large class of composite monotone inclusions and their duals can be reduced to that of finding a zero of the sum of a maximally monotone operator and a linear skew-adjoint operator. An algorithmic framework is developed for solving this generic problem in a Hilbert space setting. New primal-dual splitting algorithms are derived from this framework for inclusions involving composite monotone operators, and convergence results are established. These algorithms draw their simplicity and efficacy from the fact that they operate in a fully decomposed fashion in the sense that the monotone operators and the linear transformations involved are activated separately at each iteration. Comparisons with existing methods are made and applications to composite variational problems are demonstrated. We now turn our attention to problems involving the sum of m composite functions.

AUGMENTED LAGRANGIAN METHODS pdf

Managerial accounting irwin accounting 16th edition Aesthetics of alienation The best of the Joy of painting with Bob Ross, Americas favorite art instructor A World of Art and Museums Its Time for Thanksgiving Pennsylvania and Middle Atlantic States genealogical manuscripts Professional Java JDK 6 Edition A Nation of Shamans (The Shamanic Library, Vol 1) Six leadership principles from the Gospels Immigrants in the city Fortress stealth metal detector manual 15. Berberidaceae A discourse on the agriculture of the state of Connecticut, and the means of making it more beneficial to What can play files The age of polymorphous perversity : can civilization survive? Construction Contracts Questions and Answers Major bus operators The seller of silk shirts. Children of France, 1914-1940 Open-Ended Questions, Grade 3 (Test Connection) Proceedings of Sigmod86 International Conference on Management of Data, Washington, D.C. May 28-30, 1986 Internet and Email Tips in Laymans Terms Matrimony, as taught by phrenology and physiology Values of venereal disease patients Private rights and public illusions The school of niklaus wirth the art of simplicity Enterprise restructuring and economic policy in Russia Microsoft project 2000 manual A manual for cleaning women selected stories Sediment records of biomass burning and global change Life span human development 7th edition boyd Spring, by Maria T. Earle: February, March, April, May, by Margaret Waterfield. Psychodynamic approach to therapeutic change I am part of every poem Some modern sculptors. V. 5. Hand book on pressurized irrigation Giant Dinosaurs/Big Book A Distant Music (The Mountain Song Legacy #1) Pt. 1. The tomb of Tehuti-Hatep Percy E. Newberry Zeila, the fair maid of Cabul.