Tuning metrics for low-rank matrix completion

Bamdev Mishra and Rodolphe Sepulchre
Department of Electrical Engineering and Computer Science
University of Liège
Emails: {b.mishra, r.sepulchre}@ulg.ac.be

We discuss two new Riemannian geometries for fixed-rank matrices, and tailor them specifically to the low-rank matrix completion problem. Exploiting the degree of freedom of a quotient space, we tune the metric (notion of distance) on our search space to the particular least square cost function of the completion problem. At one level, the contribution illustrates in a novel way how to exploit the versatile framework of optimization on quotient manifold. At another level, our proposed algorithms can be considered as improved versions of LMaFit [WYZ10], a state-of-the-art Gauss-Seidel algorithm.

Low-rank matrix completion problem
The problem of low-rank matrix completion amounts to completing a matrix from a small number of entries by assuming a low-rank model for the matrix. A standard way of approaching the problem is by casting the low-rank matrix completion problem as a fixed-rank optimization problem with the assumption that the optimal rank r is known a priori as shown below

$$\begin{array}{ll} \min_{\mathbf{X}\in\mathbb{R}^{n\times m}} & \frac{1}{|\Omega|}\|\mathscr{P}_\Omega(\mathbf{X})-\mathscr{P}_\Omega(\widetilde{\mathbf{X}})\|_F^2 \\ \text{subjectto} & \operatorname{rank}(\mathbf{X}) = r, \end{array}$$

where $\widetilde{\mathbf{X}}\in\mathbb{R}^{n\times m}$ is a matrix whose entries known for indices if they belong to the subset $(i,j)\in\Omega$, where Ω is a subset of the complete set of indices $\{(i,j): i\in\{1,...,n\}\text{andj}\in\{1,...,\text{m}\}\}$. The operator $\mathscr{P}_\Omega(\mathbf{X}_{ij}) = \mathbf{X}_{ij}$ if $(i,j)\in\Omega$ and $\mathscr{P}_\Omega(\mathbf{X}_{ij}) = 0$ otherwise is called the *orthogonal sampling operator* and is a mathematically convenient way to represent the subset of entries. The objective function is, therefore, a mean least square objective function where $\|\cdot\|_F$ is *Frobenius* norm with $|\Omega|$ is the cardinality of the set Ω (equal to the number of known entries). The search space $\mathbb{R}_r^{n\times m}$ is the space of $r-$rank matrices of size $n\times m$, with $r\ll\min\{m,n\}$. The rank constraint correlates the known with the unknown entries.

Exploiting the quotient nature of matrix factorization and the problem structure
Efficient algorithms depend on properly exploiting both the structure of the constraints and the structure of the cost function. A natural way to exploit the rank constraint is through matrix factorization, which leads to a quotient structure of the search space.

In addition to this, we exploit the structure of the cost function is to select one particular metric (from multiple choices) on the quotient manifold. This is a novel step forward from earlier works which discussed only factorization models, i.e., the search space [MMBS12].

- We review two factorization models. A $r-$rank matrix $\mathbf{X}\in\mathbb{R}_r^{n\times m}$ is factorized as
$$\mathbf{X} = \mathbf{G}\mathbf{H}^T = \mathbf{U}\mathbf{R}\mathbf{V}^T$$
where $\mathbf{G}\in\mathbb{R}_*^{n\times r}$ and $\mathbf{H}\in\mathbb{R}_*^{m\times r}$ are full column-rank matrices, and $\mathbf{U}$, $\mathbf{V}$ are matrices of sizes $n\times r$ and $m\times r$ with orthonormal columns and $\mathbf{R}$ is *small* $r\times r$ matrix. Such factorization models are not unique and involve symmetries. We discuss the metrics in our search space to deal with symmetries.
- A further selection constraint for the metric is to look at the Hessian of the objective function. Scaling the search space with the Hessian information leads to efficient optimization algorithm like, for example, the Newton algorithm. In many cases, using the full Hessian information is computationally costly. A popular trade-off between convergence and numerical efficiency is to scale the gradient by the *diagonal elements* of the Hessian.

These two considerations lead to precisely identifying the metrics that best approximate the Hessian at low computational cost for the low-rank matrix completion problem. The simulation results support our notions. Part of this work has recently appeared in the technical report [MAAS12].

References

[MAAS12] B. Mishra, K. Adithya Apuroop, and R. Sepulchre, *A Riemannian geometry for low-rank matrix completion*, Tech. report, arXiv:1211.1550, 2012.

[MMBS12] B. Mishra, G. Meyer, S. Bonnabel, and R. Sepulchre, *Fixed-rank matrix factorizations and Riemannian low-rank optimization*, Tech. report, arXiv:1209.0430, 2012.

[WYZ10] Z. Wen, W. Yin, and Y. Zhang, *Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm*, Tech. report, Rice University, 2010.

Acknowledgments

This work was financially supported by the Belgian FRFC (Fonds de la Recherche Fondamentale Collective). BM is a research fellow of the Belgian National Fund for Scientific Research (FNRS). The scientific responsibility rests with its author(s).