

# Graph Drawing by Stress Majorization

Emden R. Gansner, Yehuda Koren, and Stephen North

AT&T Labs – Research,  
Florham Park, NJ 07932  
{erg,yehuda,north}@research.att.com

**Abstract.** One of the most popular graph drawing methods is based on achieving graph-theoretic target distances. This method was used by Kamada and Kawai [15], who formulated it as an energy optimization problem. Their energy is known in the multidimensional scaling (MDS) community as *the stress function*. In this work, we show how to draw graphs by stress majorization, adapting a technique known in the MDS community for more than two decades. It appears that majorization has advantages over the technique of Kamada and Kawai in running time and stability. We also found the majorization-based optimization being essential to a few extensions to the basic energy model. These extensions can improve layout quality and computation speed in practice.

## 1 Introduction

A graph is a structure  $G(V=\{1, \dots, n\}, E)$  representing a binary relation  $E$  over a set of nodes  $V$ . Visualizing graphs is a challenging problem, requiring algorithms that faithfully represent the graph's structure and the relative similarities of the nodes [4, 16]. Here we will focus on drawing undirected graphs with straight-line edges.

The most popular approach defines, sometimes implicitly, an energy, or cost function, based on some virtual physical model of the graph. Minimizing this function determines an optimal drawing. In the approach considered here, originally proposed by Kamada and Kawai [15], a nice drawing relates to good isometry. We have an ideal distance  $d_{ij}$  given for every pair of nodes  $i$  and  $j$ , modeled as a spring. Given a 2-D layout, where node  $i$  is placed at point  $X_i$ , the energy of the system is

$$\sum_{i < j} w_{ij} (\|X_i - X_j\| - d_{ij})^2. \quad (1)$$

We desire a layout that will minimize this function, thereby best approximating the target distances. Here, the distance  $d_{ij}$  is typically the graph-theoretical distance between nodes  $i$  and  $j$ . The normalization constant  $w_{ij}$  equals  $d_{ij}^{-\alpha}$ . Kamada and Kawai [15] chose  $\alpha = 2$ , whereas Cohen [6] also considered  $\alpha = 0$  and  $\alpha = 1$ . Moreover, Cohen suggested setting  $d_{ij}$  to the linear-network distance to convey the clustering structure of the graph.

The function (1), with  $\alpha = 0$ , appeared earlier as the stress function in multidimensional scaling (MDS) [5, 6, 18], where it was applied to graph drawing [17]. Whereas Kamada and Kawai proposed a localized 2-D Newton-Raphson process for minimizing the stress function, researchers in the MDS field have proposed a different, more

global approach called *majorization*. Majorization seems to offer some distinct advantages over localized processes like Newton-Raphson or gradient descent. These include guaranteed monotonic decrease of the energy value, improved robustness against local minima and shorter running times. The main contribution of this work is the introduction of this technique in the framework of graph layout.

Three useful extensions to stress optimization require the power and flexibility of majorization optimization. The first extension, described in Section 3, deals with weighting edge lengths in a way that better utilizes the drawing area, and is especially useful for drawing real-life graphs whose degree distribution follows a power law. We have found empirically that traditional stress optimization is unstable under such a weighting, while majorization works very well. The second extension deals with sparse stress functions, where only a small fraction of all pairwise distances are considered. This is essential for reducing the time and space complexity of stress optimization, and allows in-core layout of much larger graphs. We have found that sparse stress optimization is practically impossible when using the Kamada-Kawai technique (unless one has a very good initialization). Again, with majorization, it is easy to work with sparse models.

The last extension deals with obtaining an approximate drawing of the graph by constraining the layout axes to lie within a carefully selected small vector space. Such a technique was recently introduced by Koren [14] and can be integrated into layout algorithms based on matrix algebra. Fortunately, the algebraic nature of the majorization process allows us to perform rapid subspace-restricted stress minimization. The two latter extensions are described in the full version of this work.

## 2 Stress Majorization

In this section, we review stress majorization as described in the MDS literature [3, 5]. We denote a  $d$ -dimensional layout by an  $n \times d$  matrix  $X$ . Thus, node  $i$  is located at  $X_i \in \mathbb{R}^d$  and the axes of the layout are  $X^{(1)}, \dots, X^{(d)} \in \mathbb{R}^n$ . The associated stress function is

$$\text{stress}(X) = \sum_{i < j} w_{ij} (\|X_i - X_j\| - d_{ij})^2. \quad (2)$$

We always take  $w_{ij} = d_{ij}^{-2}$ , which seems to produce the best drawings in most cases. Decompose (2) to obtain

$$\text{stress}(X) = \sum_{i < j} w_{ij} d_{ij}^2 + \sum_{i < j} w_{ij} \|X_i - X_j\|^2 - 2 \sum_{i < j} \delta_{ij} \|X_i - X_j\|, \quad (3)$$

where  $\delta_{ij} \stackrel{\text{def}}{=} w_{ij} d_{ij}$  for  $i, j = 1, \dots, n$ .

The first term of (3),  $\sum_{i < j} w_{ij} d_{ij}^2$ , is a constant independent of the current layout. The second term,  $\sum_{i < j} w_{ij} \|X_i - X_j\|^2$ , is a quadratic sum, and can be written using the quadratic form of the weighted Laplacian  $L^w$

$$\sum_{i < j} w_{ij} \|X_i - X_j\|^2 = \text{Tr}(X^T L^w X), \quad (4)$$

where the  $n \times n$  weighted Laplacian has its  $ij$  entry, for  $i, j = 1, \dots, n$ , defined as

$$L_{i,j}^w = \begin{cases} -w_{ij} & i \neq j \\ \sum_{k \neq i} w_{ik} & i = j \end{cases}.$$

The third term,  $\sum_{i < j} \delta_{ij} \|X_i - X_j\|$ , is more involved and we will bound it from below. We will make use of the Cauchy-Schwartz inequality

$$\|x\| \|y\| \geq x^T y$$

with equality when  $x = y$ . Consequently, given any  $n \times d$  matrix  $Z$ ,

$$\|X_i - X_j\| \|Z_i - Z_j\| \geq (X_i - X_j)^T (Z_i - Z_j)$$

with equality when  $X = Z$ . We can now bound the third term as follows

$$\sum_{i < j} \delta_{ij} \|X_i - X_j\| \geq \sum_{i < j} \delta_{ij} \text{inv}(\|Z_i - Z_j\|) (X_i - X_j)^T (Z_i - Z_j) \quad (5)$$

where  $\text{inv}(x) = 1/x$  when  $x \neq 0$  and 0 otherwise.

Inequality (5) can be written in a more convenient matrix form

$$\sum_{i < j} \delta_{ij} \|X_i - X_j\| \geq \text{Tr}(X^T L^Z Z),$$

where the  $n \times n$  matrix  $L^Z$  has its  $ij$  entry, for  $i, j = 1, \dots, n$ , defined as

$$L_{i,j}^Z = \begin{cases} -\delta_{ij} \text{inv}(\|Z_i - Z_j\|) & i \neq j \\ -\sum_{j \neq i} L_{i,j}^Z & i = j \end{cases}.$$

Combining all the above, we can bound the stress function using  $F^Z(X)$  defined as

$$F^Z(X) = \sum_{i < j} w_{ij} d_{ij}^2 + \text{Tr}(X^T L^w X) - 2\text{Tr}(X^T L^Z Z). \quad (6)$$

Thus, we have

$$\text{stress}(X) \leq F^Z(X) \quad (7)$$

with equality when  $Z = X$ .

Note that  $Z$  is a constant  $n \times d$  matrix. This way we have bounded the stress with a quadratic form  $F^Z(X)$ . We differentiate by  $X$  and find that the minima of  $F^Z(X)$  are given by solving

$$L^w X = L^Z Z.$$

Or, equivalently, for each axis we have to solve

$$L^w X^{(a)} = L^Z Z^{(a)}, \quad a = 1, \dots, d. \quad (8)$$

The characteristic of the minima is determined by the nature of the weighted Laplacian  $L^w$ , which is known to be positive semi-definite with a one-dimensional null space

spanned by  $1_n = (1, \dots, 1) \in \mathbb{R}^n$ . Hence,  $F^Z(X)$  has only global minima, which are invariant under translation (addition of  $\alpha \cdot 1_n$  is equivalent to translation). This makes sense, since the stress function is also invariant under translation.

Numerically, it is better to make the minimizer unique. Hence we recommend removing the translation degree-of-freedom by taking  $X_1 = 0$ . Therefore, we can remove the first row and column of  $L^w$ , as well as the first row of  $L^Z Z$ . The resulting  $(n-1) \times (n-1)$  matrix is strictly diagonal dominant and hence positive definite. This is very convenient, since methods like conjugate gradient, Gauss-Seidel, and Cholesky factorization are guaranteed to work [9].

### The Optimization Process

The above formulation leads to the following iterative optimization process. Given some layout  $X(t)$ , we want to compute a layout  $X(t+1)$  so that  $\text{stress}(X(t+1)) < \text{stress}(X(t))$ . We use the function  $F^{X(t)}(X)$  which satisfies  $F^{X(t)}(X(t)) = \text{stress}(X(t))$ .

We take  $X(t+1)$  as the minimizer of  $F^{X(t)}(X)$  by solving

$$L^w X(t+1)^{(a)} = L^{X(t)} X(t)^{(a)}, \quad a = 1, \dots, d. \quad (9)$$

At this point, if  $X(t+1) = X(t)$ , we terminate the process. Otherwise, we get

$$\text{stress}(X(t+1)) \leq F^{X(t)}(X(t+1)) < F^{X(t)}(X(t)) = \text{stress}(X(t)).$$

The first inequality is by (7) and the second inequality is by the uniqueness of the minimum.

In practice we terminate the process when

$$\frac{\text{stress}(X(t)) - \text{stress}(X(t+1))}{\text{stress}(X(t))} < \epsilon, \quad (10)$$

where  $\epsilon$  is the tolerance of the process. Typically,  $\epsilon \sim 10^{-4}$ .

To summarize, the majorization process involves iteratively solving (9). The matrix  $L^w$  is constant throughout the entire process, whereas the matrix  $L^{X(t)}$  would be re-computed at each iteration.

## 2.1 Equation Solvers

In practice we recommend using either Cholesky factorization or conjugate gradient (CG) [9] to solve (9) (by first fixing  $X_1 = 0$  as discussed above). Using Cholesky factorization implies that at a preprocessing stage we find the  $LL^T$  factorization of  $L^w$  using  $n^3/3$  flops (floating point operations). Then in each iteration we solve the linear system using back substitution in time  $O(n^2)$ . Hence, the significant cost in Cholesky factorization is independent of the number of iterations, making it suitable for graphs requiring many iterations of process (9).

On the other hand, CG optimization involves no preprocessing and its running time is evenly distributed among the iterations. Almost the entire solving time is devoted to performing matrix-vector multiplication. Each such multiplication takes  $n^2$  flops. Thus,

if the total number of matrix multiplications is less than about  $n/3$ , the CG process is expected to be faster than Cholesky factorization. Otherwise, Cholesky factorization is recommended. In practice, for most graphs we have experimented with, CG outperformed Cholesky since the total number of matrix-vector multiplications is typically less than  $n/3$ . Note that CG benefits by the fact that we have an initial approximate solution from the previous iteration. We observed that the overall number of iterations increases very moderately with the size of the graph. Therefore, for large graphs (over 10,000 nodes), we encountered cases where the total number of matrix-vector multiplications exceeded even  $n$ , so Cholesky factorization should do much better. In any case, all the results reported here employ CG.

## 2.2 Intuitive Interpretation

Let us concentrate on axis  $a$ , and denote the current coordinates by  $\hat{x} = X(t)^{(a)}$ . The majorization process determines the new coordinates  $x = X(t+1)^{(a)}$  by solving the system of equations (9). Eliminating  $x_i$  in equation  $i$ , we rewrite the system in an equivalent form

$$x_i = \frac{\sum_{j \neq i} w_{ij} (x_j + d_{ij}(\hat{x}_i - \hat{x}_j) \text{inv}(\|X(t)_i - X(t)_j\|))}{\sum_{j \neq i} w_{ij}}. \quad (11)$$

The intuitive interpretation of this process is simple. A node  $j$  located at  $x_j$  strives to place node  $i$  (on current axis  $a$ ) at  $x_j + d_{ij} \frac{\hat{x}_i - \hat{x}_j}{\|X(t)_i - X(t)_j\|}$ .

Based on the current placement, this is node  $j$ 's best strategy to assure that node  $i$  will be at distance  $d_{ij}$  from  $j$  in the full multidimensional layout. To see this, notice that the distance between the nodes depends on all the axes. Therefore, node  $j$ 's estimate of the contribution of axis  $a$  for the distance between  $i$  and  $j$  is the fraction  $\alpha = \left| \frac{\hat{x}_i - \hat{x}_j}{\|X(t)_i - X(t)_j\|} \right|$ . So the magnitude of displacement should be  $d_{ij}$  scaled down by  $\alpha$ . Now, after deciding the magnitude of the 1-D displacement, the direction must be decided: should we place  $x_i$  at  $x_j + \alpha d_{ij}$  or at  $x_j - \alpha d_{ij}$ ? Again, the decision is based on the current placement, whether currently  $\hat{x}_i < \hat{x}_j$  or vice versa.

This way, each node  $j$  votes for its desired placement of  $x_i$ . The final position is determined by taking the weighted average of the suggested positions. This intuition also suggests a localized optimization process, which we next describe.

## 2.3 Localized Optimization

Following the idea of Kamada and Kawai [15], we can fix the positions of all nodes, except some node  $i$ . Then, by the same argument given above for the full majorization process, it can be shown that the stress function is decreased by setting the position of  $i$  as follows

$$X_i^{(a)} \leftarrow \frac{\sum_{j \neq i} w_{ij} \left( X_j^{(a)} + d_{ij} (X_i^{(a)} - X_j^{(a)}) \text{inv}(\|X_i - X_j\|) \right)}{\sum_{j \neq i} w_{ij}}, \quad a = 1, \dots, d. \quad (12)$$

This way we can iterate through all nodes, and in each iteration relocate all the  $d$  coordinates of node  $i$  according to (12). Each iteration is guaranteed to strictly decrease the stress until convergence. Hence, oscillations and non-convergence are impossible.

In practice, we have only used the more involved global process (9) and have no experience yet with the local version. We provide this local version here mainly because it is simple and easy to implement, requiring no equation solver<sup>1</sup>.

## 2.4 Comparisons

A natural question is whether we should replace the traditional Kamada-Kawai based optimization with majorization. Based on several months of experimenting with both approaches, our definite answer is yes. We base this recommendation on several considerations.

We experimented with various example graphs. On each graph, we ran each of the two algorithms 25 times with different random initializations. At certain times during each execution, we measured the elapsed running time and the current value of the stress function, and averaged over all 25 executions. From this we obtained stress-vs.-time charts for the graphs. While it is impossible to present here all of the charts, we show a few representative ones in Figures 1-3. We can make some important observations.

**Layout Quality.** We observed that most of the time, the two methods eventually achieved about the same stress level. In certain cases, the Kamada-Kawai approach would yield a slightly better layout in terms of the stress value, but the difference was always small; see Figure 2. In other cases, however, the majorization approach yielded significantly better layouts as can be seen in Figure 3. Hence, probably due to its more global nature, majorization can be considered better in terms of layout quality.

**Monotonicity of Convergence.** A significant advantage of majorization is that iterations monotonically decrease the stress until convergence. This way, termination of the process is determined naturally by a condition like (10). However, our experience with the Kamada-Kawai approach, as implemented in Neato [7], shows that in some cases the latter process may cycle without converging, while the energy is oscillating. This requires an artificial or more convoluted termination condition.

Our experiments show that, as expected, the majorization approach was always monotonic in decreasing the stress value. The non-monotonicity problem of the Kamada-Kawai method was extremely rare (remember that we averaged over 25 executions, lessening the impact of a single bad non-monotonic execution). We did observe this non-monotonic behavior when experimenting with the Qh882 graph [1]. The result is provided in Fig. 1, which compares the average behavior of both approaches on this graph. We should note that here we weighted edges as explained in Section 3. The reader can see that after 2 seconds of running, the stress value in the Kamada-Kawai approach increases for some period. Here, this did not prevent it from converging at about the same stress level as the majorization process.


---

<sup>1</sup> Process (12) should not be confused with the similar Gauss-Seidel process that can be used to solve (9).

**Running Time.** The running time of the majorization process is consistently less than that of the Kamada-Kawai process. In all runs, it can be observed that majorization reaches the low stress level much before Kamada-Kawai.

A partial explanation is that majorization's running time is dominated by matrix operations (matrix-vector multiplication or Cholesky factorization). These operations are implemented in libraries like BLAS and LAPACK which are highly optimized on the machine instruction level for common platforms. We are using the Intel Math Kernel Library [22]; another well-known implementation is Atlas [23].

For implementations not relying on special matrix software, we found the situation to be similar to that of the stress function. Sometimes the Kamada-Kawai approach would be marginally faster; on the other hand, when the majorization process was faster, it was significantly faster. And as the size of the graphs increased, the advantage swung completely to majorization.




**Fig. 1.** Stress function vs. running time for the graph Qh882 [1] ( $|V|=882$ ,  $|E|=1533$ ). Here both methods reached about the same stress. Interestingly, Kamada-Kawai is not monotonic.


Before leaving this topic, we must point out that our implementation of the Kamada and Kawai process on which we based our comparisons differs slightly from the implementation originally suggested [15]. We are using the more common implementation which replaces the two nested loops with a single loop; see [2, 11]. As noted in Brandenburg, Himsolt, and Rohrer [2], this leads to a significant speed-up over the original implementation. This more efficient implementation is also the one used in Neato [7] and GraphLet [21].

### 3 Weighting Edge Lengths

In many real life graphs, the degree distribution decays at a much lower rate than in random graphs. Usually this distribution follows a power law and is proportional to  $d^{-\lambda}$ . Setting desired edge lengths to a uniform length (typically 1) inevitably makes the neighborhood of high degree nodes too dense in the layout. Consequently, we suggest weighting edges by their neighborhood size.



**Fig. 2.** Stress function vs. running time for the graphs Bcspwr07 [1] ( $|V|=882$ ,  $|E|=1533$ ) and 516 [19] ( $|V|=516$ ,  $|E|=729$ ).



**Fig. 3.** Stress function vs. running time for the graphs Qh1484 [1] ( $|V|=1470$ ,  $|E|=6420$ ) and Plsk1919 [1] ( $|V|=1919$ ,  $|E|=4831$ ).

Specifically, we set the length of each edge  $\langle i, j \rangle \in E$  as


$$l_{ij} = |N_i \cup N_j| - |N_i \cap N_j|, \tag{13}$$

where  $N_i = \{j | \langle i, j \rangle \in E\}$ . Then, each target distance  $d_{ij}$  is the length of the shortest weighted path between  $i$  and  $j$ .

This simple change is surprisingly effective in many real life irregular graphs that have highly non-uniform degree distributions. We present here two examples. The first example is the 1138Bus graph ( $|V|=1138$ ,  $|E|=1458$ ) from the Matrix Market repository [1]. This graph models a network of high-voltage power distribution lines. Figure 4 shows two layouts of this graph. In one layout, edges were weighted according to (13). The other layout was made with unweighted edges. Nodes are much better dispersed in the weighted-edge-based layout. By weighting edges, more space is allocated to the dense areas, avoiding many of the edge crossings.

Another interesting example is a BGP connectivity graph representing communications between autonomous systems ( $|V|=3847$ ,  $|E|=11539$ ). This graph has a few nodes





**Fig. 4.** Two layouts of the 1138Bus graph [1].


of high degree (e.g., one node has degree 695 and a few others are around 100), as well as 3257 nodes of degree 1. We show two layouts of this graph in Figure 5. Again, it is clear that when weighting edges, the resulting layout is much more informative. For example, in both layouts the central node is the one of degree 695. In the weighted version, its neighborhood is placed far enough from it to make it fairly visible. In the unweighted version, however, all of its neighbors are positioned densely around it, hiding its structure completely.

We have frequently found that when there are large deviations in edge lengths, as in the BGP graph, classic Kamada-Kawai optimization fails to find a nice layout. The result of Kamada-Kawai optimization on the edge-weighted BGP graph is shown in Figure 6(a). It is clearly inferior to the majorization result shown in Figure 5. We also compare the average stress-vs.-time behavior of the two methods in Figure 6(b), where it is clear the Kamada-Kawai-type optimization is pretty helpless here. Although we do not fully understand this limitation of Kamada-Kawai optimization, it seems that its local nature somehow limits its ability to deal with significantly unbalanced edge lengths.

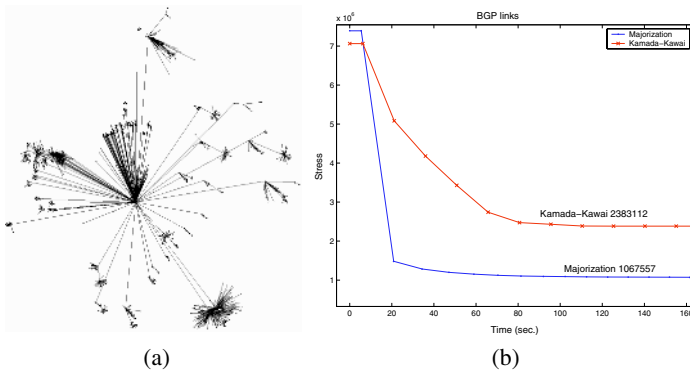
## 4 Related Work

Substantial work in statistical MDS deals with the properties of the majorization process, including proofs of its convergence rate [3]. The MDS literature suggests solving equation (9) by computing  $(L^w)^+$ , the Moore-Penrose inverse of the singular matrix  $L^w$ . Our suggestion to set  $X_1 = 0$  allows a much faster solution by Cholesky factorization.

Several studies in the graph drawing field suggest improving stress computation by multi-scale extensions [8, 10, 11], which approximate the graph by a smaller one, to quickly obtain an initial layout. We see these approaches as complementary to our



**Fig. 5.** Two majorization-based layouts of BGP connectivity, with a skewed degree distribution.



**Fig. 6. (a)** Layout of the edge weighted BGP connectivity graph using Kamada-Kawai optimization. **(b)** Stress-vs.-time behavior of majorization and Kamada-Kawai on weighted BGP connectivity example graph.

proposal, as one can apply majorization to optimizing the stress at each scale. In general, our recommendation is to get an initial placement either by multi-scale techniques or by subspace-restricted computation [14].

Recent work by Koren and Harel [13] describes an algorithm for monotonically decreasing the stress function in 1-D, and a heuristic extension to higher dimensions whose convergence properties are unknown. It is easy to prove that this 1-D algorithm is equivalent to 1-D majorization, although derived differently. Majorization, however, is more powerful as it can be generalized to higher dimensions. Interestingly, the optimization process of [13] is equivalent to the full,  $n$ -D Newton-Raphson process. Ac-

cordingly, we conclude that in 1-D, the majorization process is equivalent to the full,  $n$ -D Newton-Raphson process. This is unlike the Kamada-Kawai process which is based on a localized 2-D Newton-Raphson process.

## 5 Conclusions

Majorization, a technique developed in studies of statistical MDS, is relevant to practical graph drawing. The MDS community has studied it extensively from the standpoint of optimizing the stress function and escaping local minima. Further ideas along these lines may also prove useful in graph drawing.

The main algorithms discussed here are available in the Neato program in the Graphviz open source package [20].

## References

1. R. F. Boisvert et al., “The Matrix Market: A web resource for test matrix collections”, in *Quality of Numerical Software, Assessment and Enhancement*, R. F. Boisvert, ed., Chapman Hall, 1997, pp. 125–137. [math.nist.gov/MatrixMarket](http://math.nist.gov/MatrixMarket)
2. F.J. Brandenburg, M. Himsolt and C. Rohrer, “An Experimental Comparison of Force-Directed and Randomized Graph Drawing Algorithms”, *Proceedings of Graph Drawing '95*, LNCS 1027, pp. 76–87, Springer Verlag, 1995.
3. J. De Leeuw, “Convergence of the Majorization Method for Multidimensional Scaling”, *Journal of Classification* **5** (1988), pp. 163–180.
4. G. Di Battista, P. Eades, R. Tamassia and I. G. Tollis, *Graph Drawing: Algorithms for the Visualization of Graphs*, Prentice-Hall, 1999.
5. I. Borg and P. Groenen, *Modern Multidimensional Scaling: Theory and Applications*, Springer-Verlag, 1997.
6. J. D. Cohen, “Drawing Graphs to Convey Proximity: an Incremental Arrangement Method”, *ACM Transactions on Computer-Human Interaction* **4** (1997), pp. 197–229 .
7. E. R. Gansner and S. C. North, “Improved force-directed layouts”, *Proceedings of Graph Drawing '98*, LNCS 1547, pp. 364–373, Springer-Verlag, 1998.
8. P. Gajer, M. T. Goodrich and S. G. Kobourov, “A Multi-dimensional Approach to Force-Directed Layouts of Large Graphs”, *Proceedings of Graph Drawing 2000*, LNCS 1984, pp. 211–221, Springer-Verlag, 2000.
9. G. H. Golub and C. F. Van Loan, *Matrix Computations*, Johns Hopkins University Press, 1996.
10. R. Hadany and D. Harel, “A Multi-Scale Method for Drawing Graphs Nicely”, *Discrete Applied Mathematics* **113** (2001), pp. 3–21.
11. D. Harel and Y. Koren, “A Fast Multi-Scale Method for Drawing Large Graphs”, *Journal of Graph Algorithms and Applications* **6** (2002), pp. 179–202.
12. D. Harel and Y. Koren, “Graph Drawing by High-Dimensional Embedding”, *Proceedings of Graph Drawing 2002*, LNCS 2528, pp. 207–219, Springer-Verlag, 2002.
13. Y. Koren and D. Harel, “Axis-by-Axis Stress Minimization”, *Proceedings of Graph Drawing 2003*, Springer-Verlag, pp. 450–459, 2003.
14. Y. Koren, “Graph Drawing by Subspace Optimization”, *Proceedings 6th Joint Eurographics – IEEE TCVG Symposium Visualization (VisSym '04)*, pp. 65–74, Eurographics, 2004.
15. T. Kamada and S. Kawai, “An Algorithm for Drawing General Undirected Graphs”, *Information Processing Letters* **31** (1989), pp. 7–15.

16. M. Kaufmann and D. Wagner (Eds.), *Drawing Graphs: Methods and Models*, LNCS 2025, Springer-Verlag, 2001.
17. J. Kruskal and J. Seery, “Designing network diagrams”, *Proceedings First General Conference on Social Graphics* (1980), pp. 22–50.
18. J. W. Sammon, “A Nonlinear Mapping for Data Structure Analysis”, *IEEE Trans. on Computers* **18** (1969), pp. 401–409.
19. C. Walshaw, “A Multilevel Algorithm for Force-Directed Graph Drawing”, *Proceedings 8th Graph Drawing (GD’00)*, LNCS 1984, pp. 171–182, Springer-Verlag, 2000.
20. Graphviz. [www.research.att.com/sw/tools/graphviz/](http://www.research.att.com/sw/tools/graphviz/)
21. Graphlet. [www.infosun.fmi.uni-passau.de/Graphlet/](http://www.infosun.fmi.uni-passau.de/Graphlet/)
22. Intel Math Kernel Library. [www.intel.com/software/products/mkl/](http://www.intel.com/software/products/mkl/)
23. Automatically Tuned Linear Algebra Software (ATLAS).  
[math-atlas.sourceforge.net/](http://math-atlas.sourceforge.net/)