

A Multigrid Tutorial

Irad Yavneh

Department of Computer Science

Technion - Israel Institute of Technology

irad@cs.technion.ac.il

Some Relevant Books:

1. W.L. Briggs, V.E. Henson, and S.F. McCormick: "A Multigrid Tutorial", 2nd ed., SIAM, 2000.
2. U. Trottenberg, C.W. Oosterlee, and A. Schueller: "Multigrid", Academic Press, 2001.
3. Brandt, A., "1984 Guide with Applications to Fluid Dynamics", GMD-Studie Nr. 85, 1984.
4. Hackbusch, W., "Multigrid Methods and Applications", Springer, Berlin, 1985.
5. W. Hackbusch and U Trottenberg eds.: "Multigrid Methods", Springer-Verlag, Berlin, 1982.
6. Wienands, R., and Joppich, W., "Practical Fourier Analysis for Multigrid Methods", Chapman & Hall/CRC, 2004.

What's it about?

A framework of efficient iterative methods for solving problems with many variables and many scales.

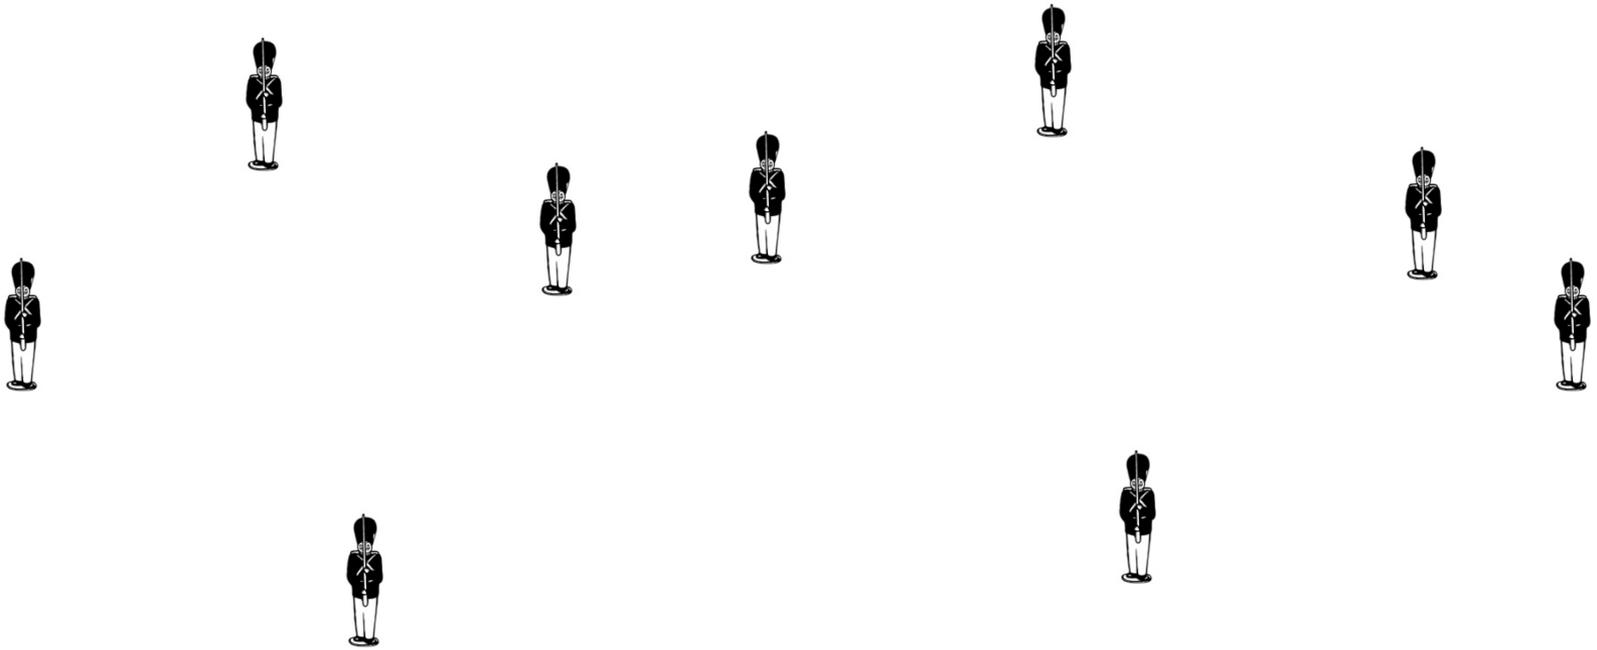
- **Framework**: common concept, different methods.
- **Efficient**: usually $O(N)$ or $O(N \log N)$ operations
The importance of efficient methods becomes greater as computers grow stronger!
- **Iterative**: most nontrivial problems in our field cannot be solved directly efficiently.
- **Solving**: approximately, subject to appropriate convergence criteria, constraints, etc.
- **Many variables**: the larger the number of variables, the greater the gain of efficient methods.
- **Many scales**: typical spatial and/or temporal sizes.

A framework of efficient iterative methods for solving problems with many variables and many scales.

Basic Concepts: Local vs. Global processing.

Imagine a large number of soldiers who need to be arranged in a **straight line and at equal distances** from each other.

The two soldiers at the ends of the line are fixed. Suppose we number the soldiers 0 to N , and that the length of the entire line is L .



Initial Position



Final Position

Global processing. Let soldier number j stand on the line connecting soldier 0 to soldier N at a distance jL/N from soldier number 0 .

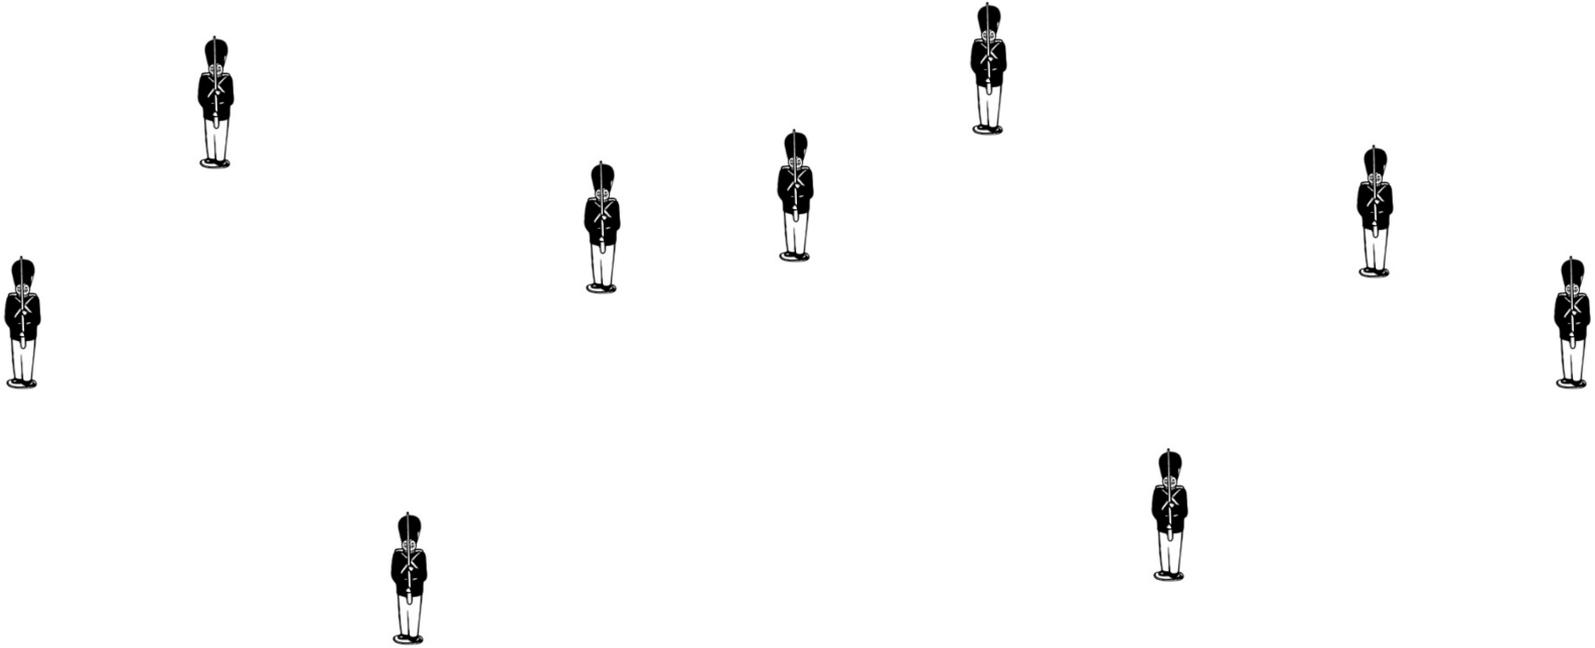


This method solves the problem **directly**, but it requires a high degree of **sophistication**: recognition of the extreme soldiers and some pretty fancy arithmetic.

Local processing (iterative method). Suppose that the inner soldiers' initial position is $\mathbf{x}^{(0)} = (x_1, x_2, \dots, x_{N-1})^{(0)}$. Then repeat for $i=1, 2, \dots$: Let each soldier $j, j=1, \dots, N-1$ at iteration i move to the point midway between the locations of soldier $j-1$ and soldier $j+1$ at iteration $i-1$:

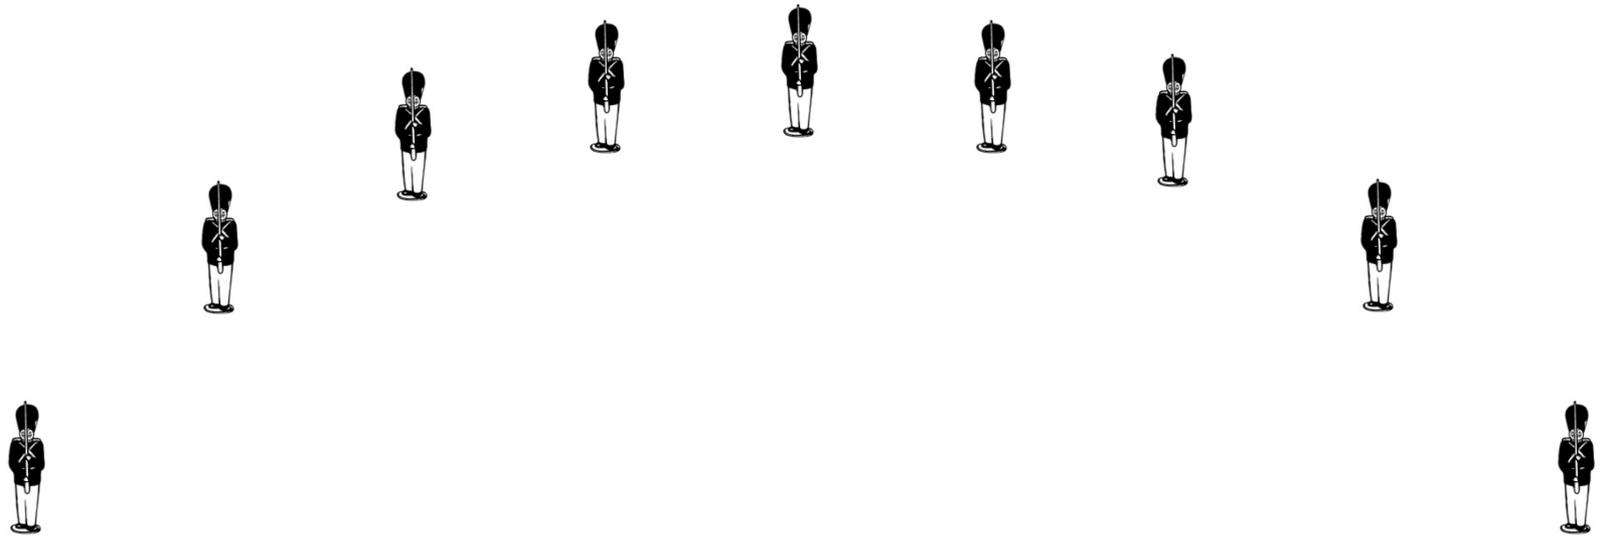
$$x_j^{(i)} = \frac{1}{2} \left(x_{j-1}^{(i-1)} + x_{j+1}^{(i-1)} \right)$$

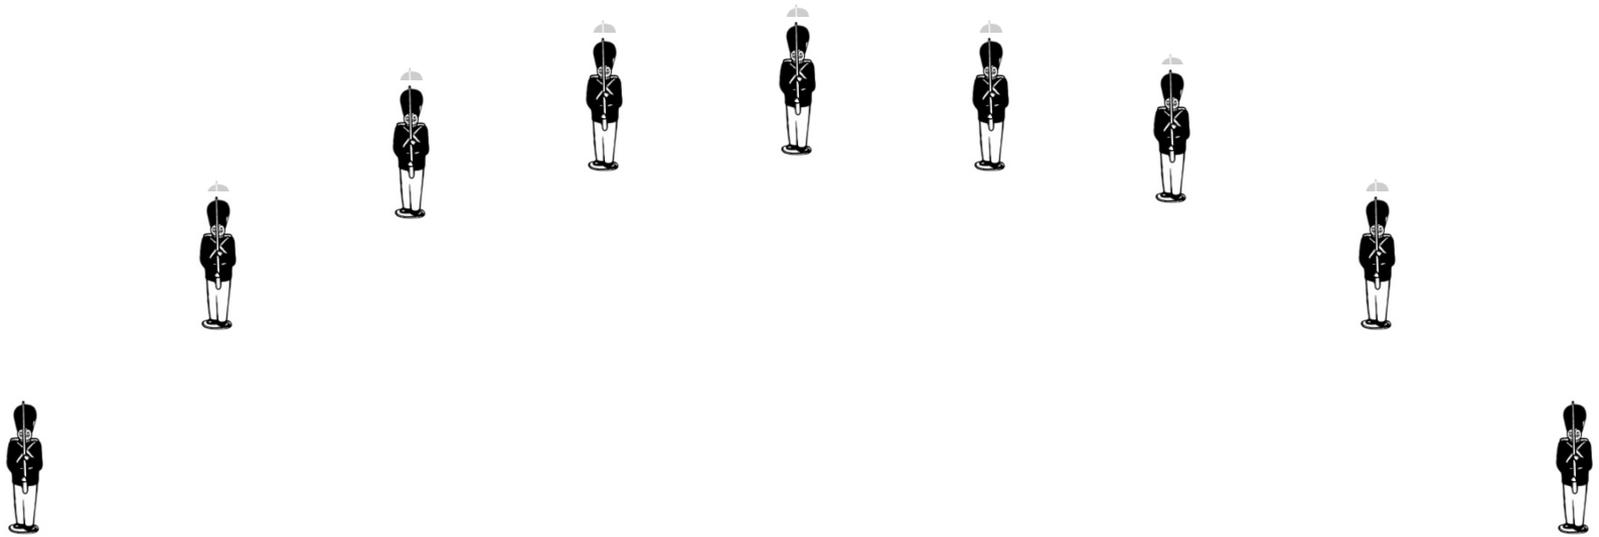
This is an **iterative** process. Each iteration brings us closer to the solution(?). The arithmetic is trivial.



A step in the right direction

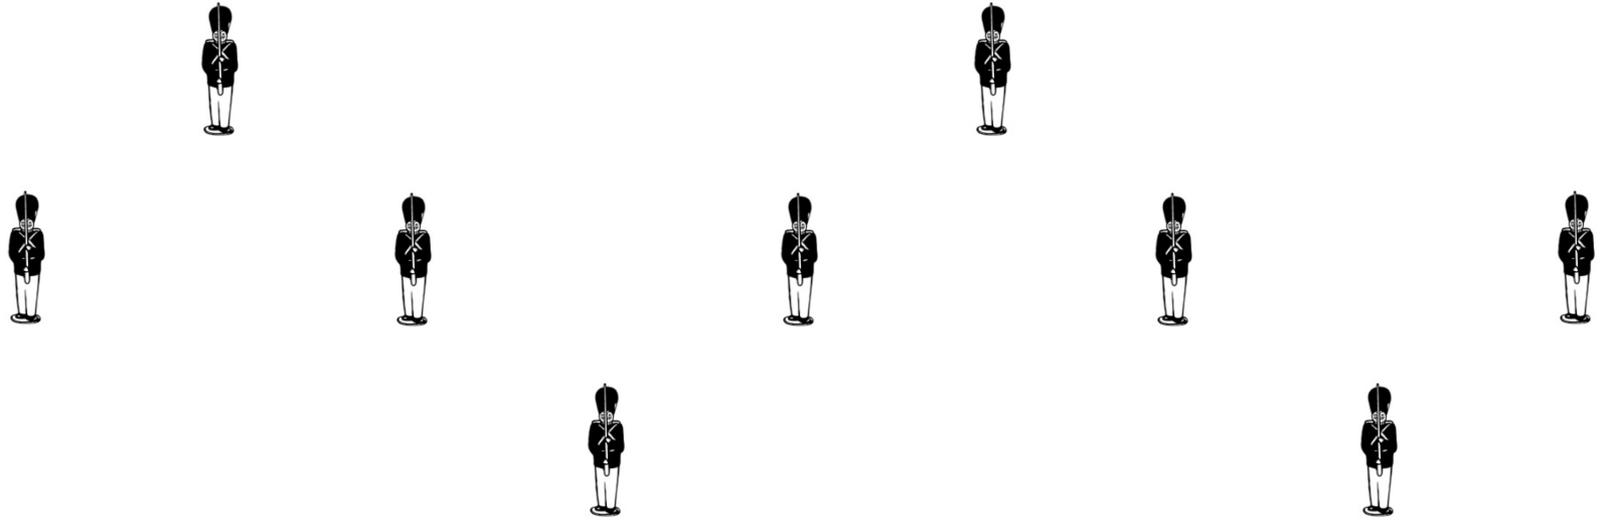


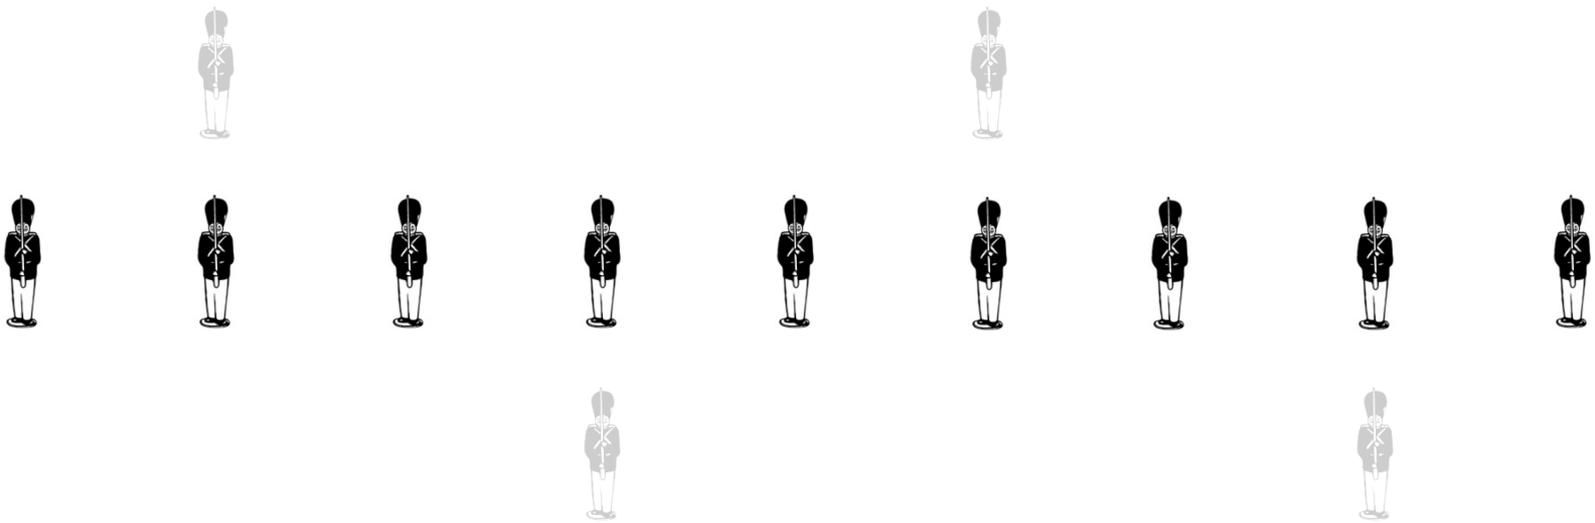




Slow convergence



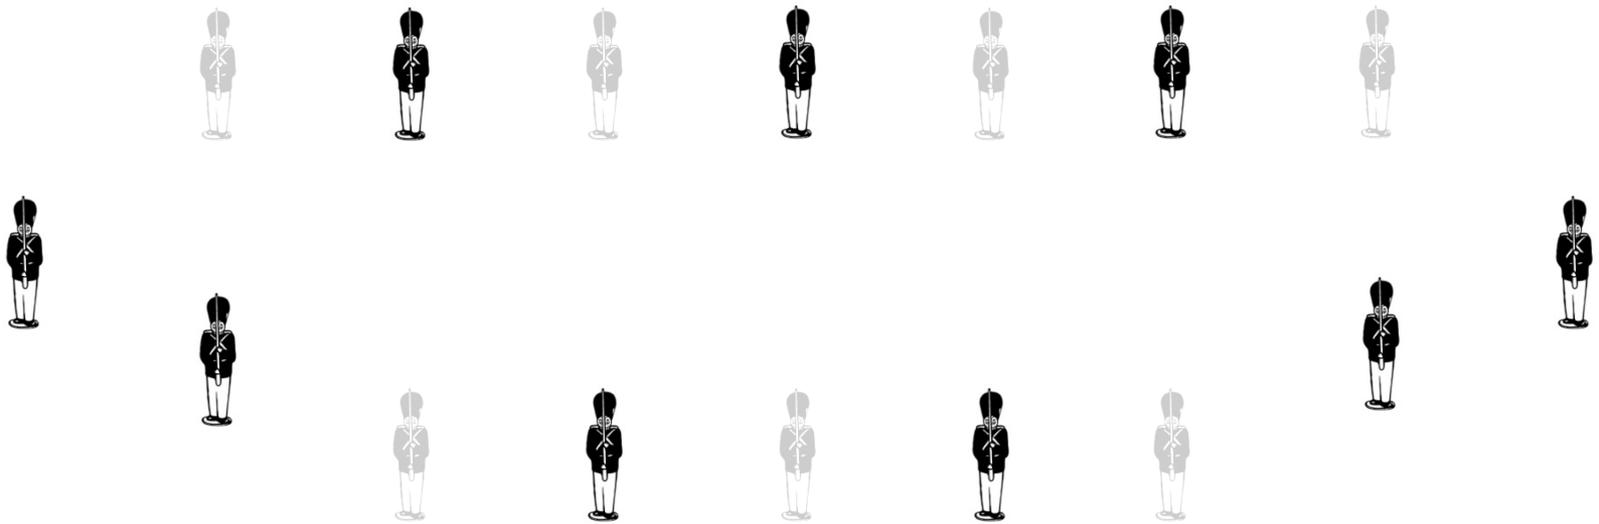




Fast convergence

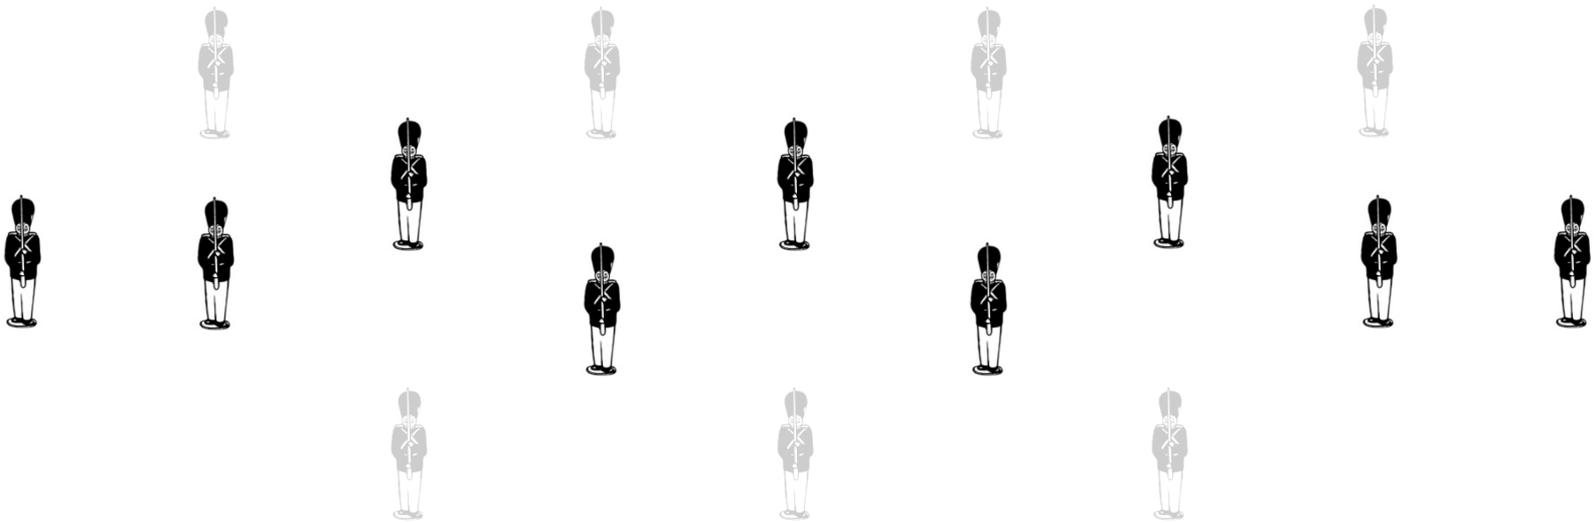






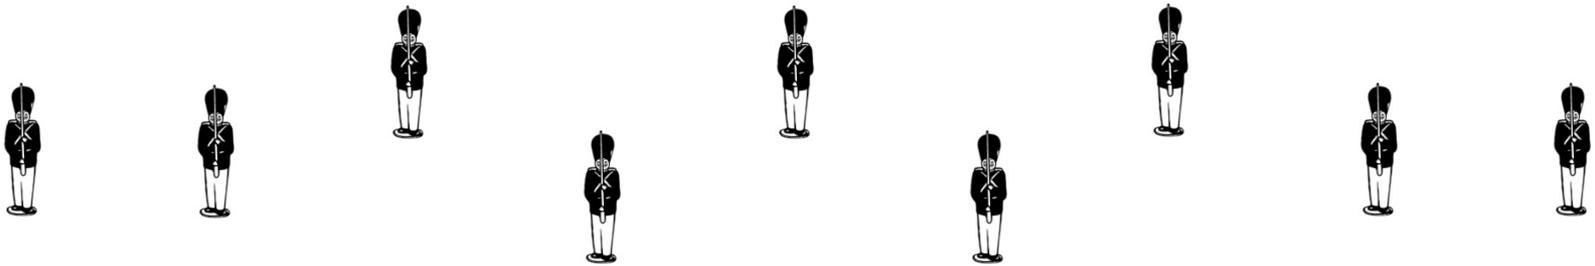
Slow convergence





Local solution: damping





Local solution: damping





Local solution: damping

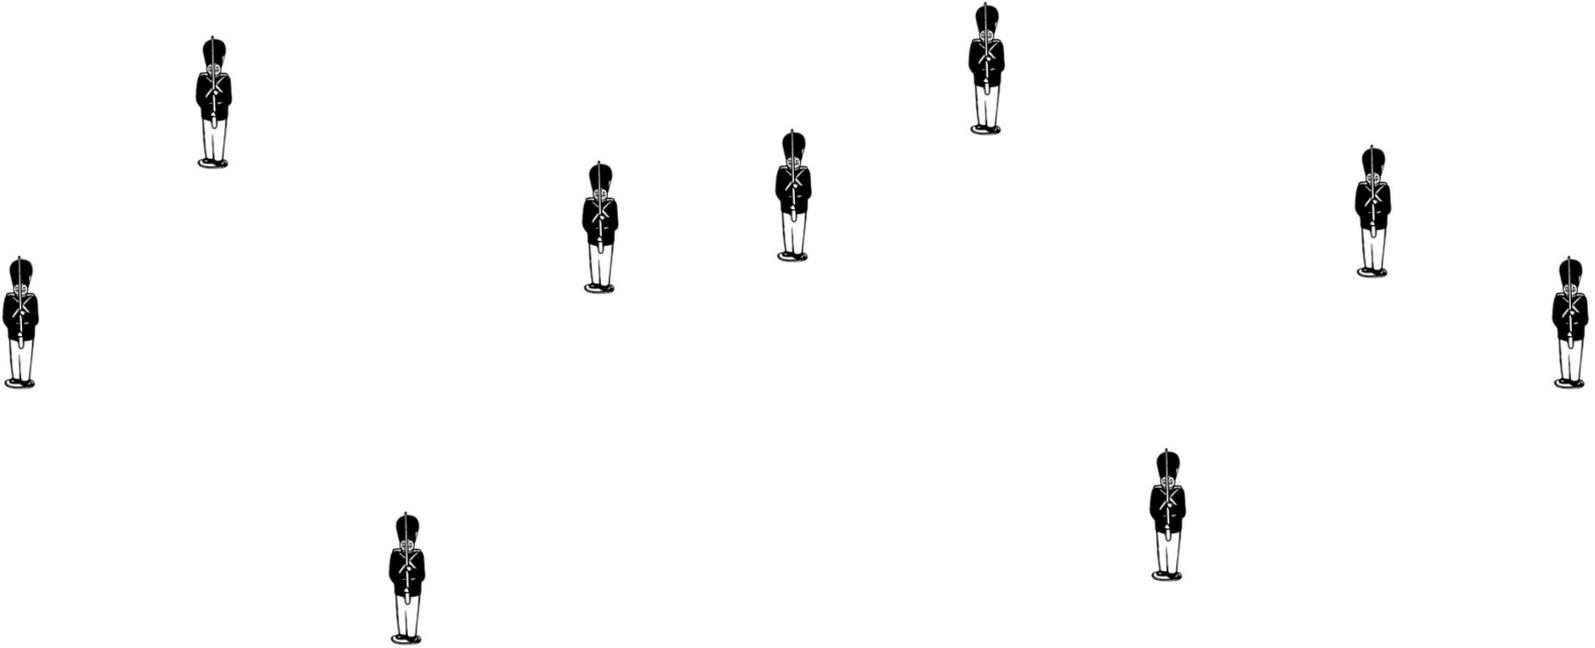


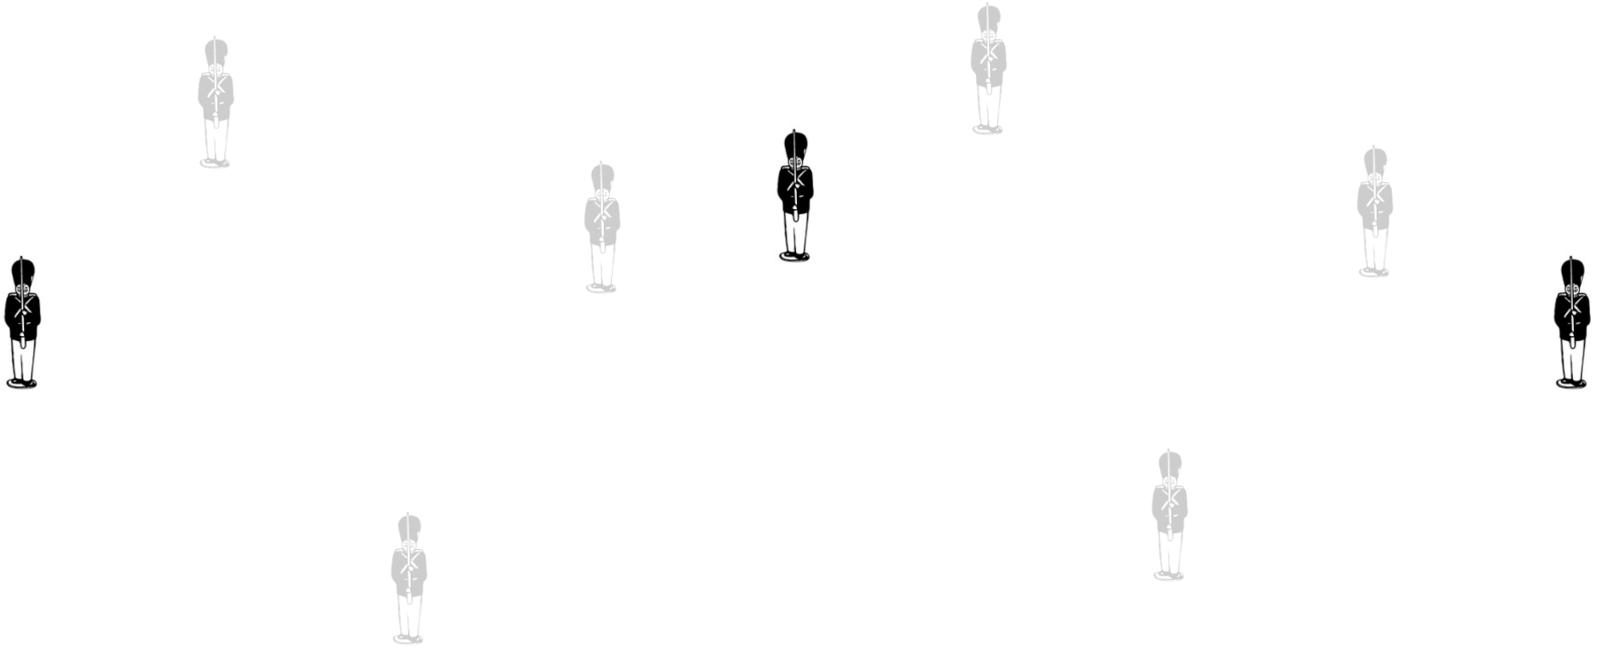


Local solution: damping



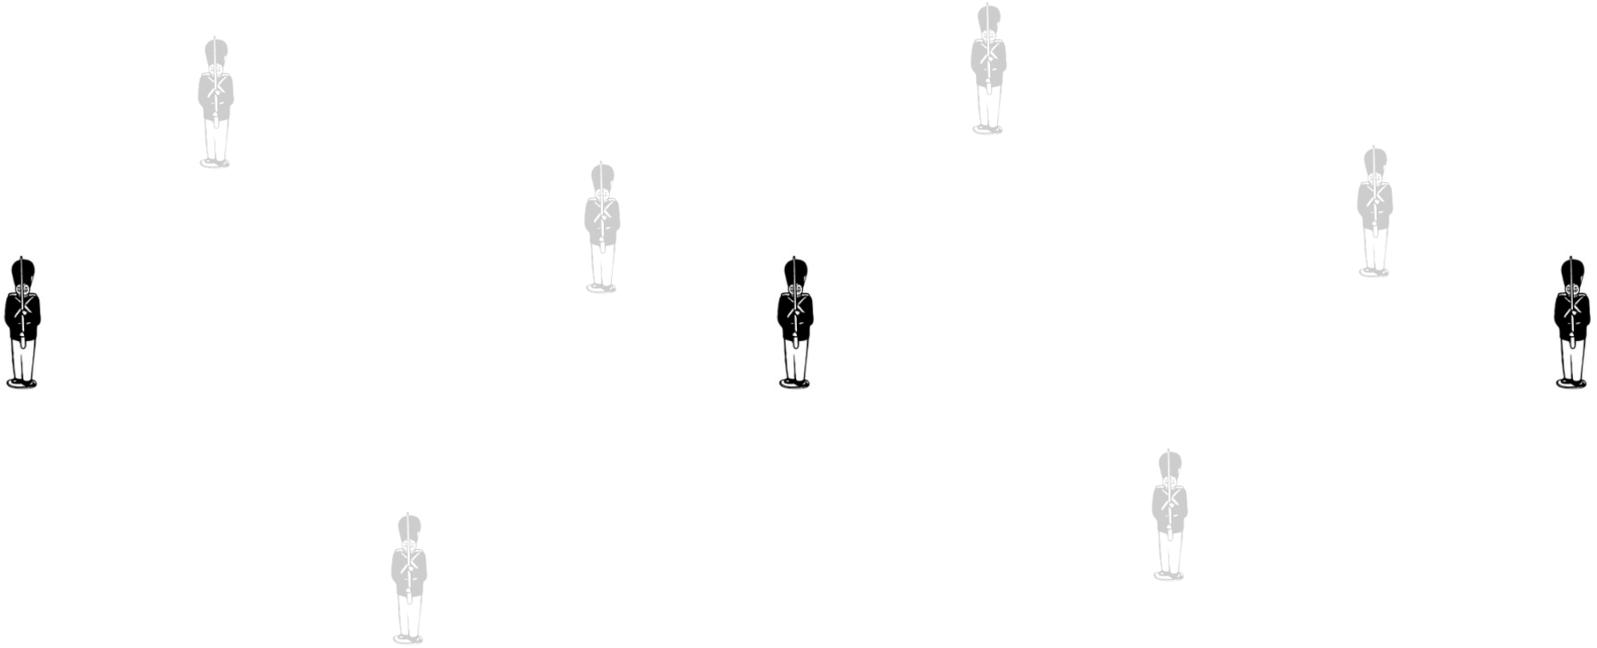
The multiscale idea: Employ the local processing with simple arithmetic. But do this on all the different scales.





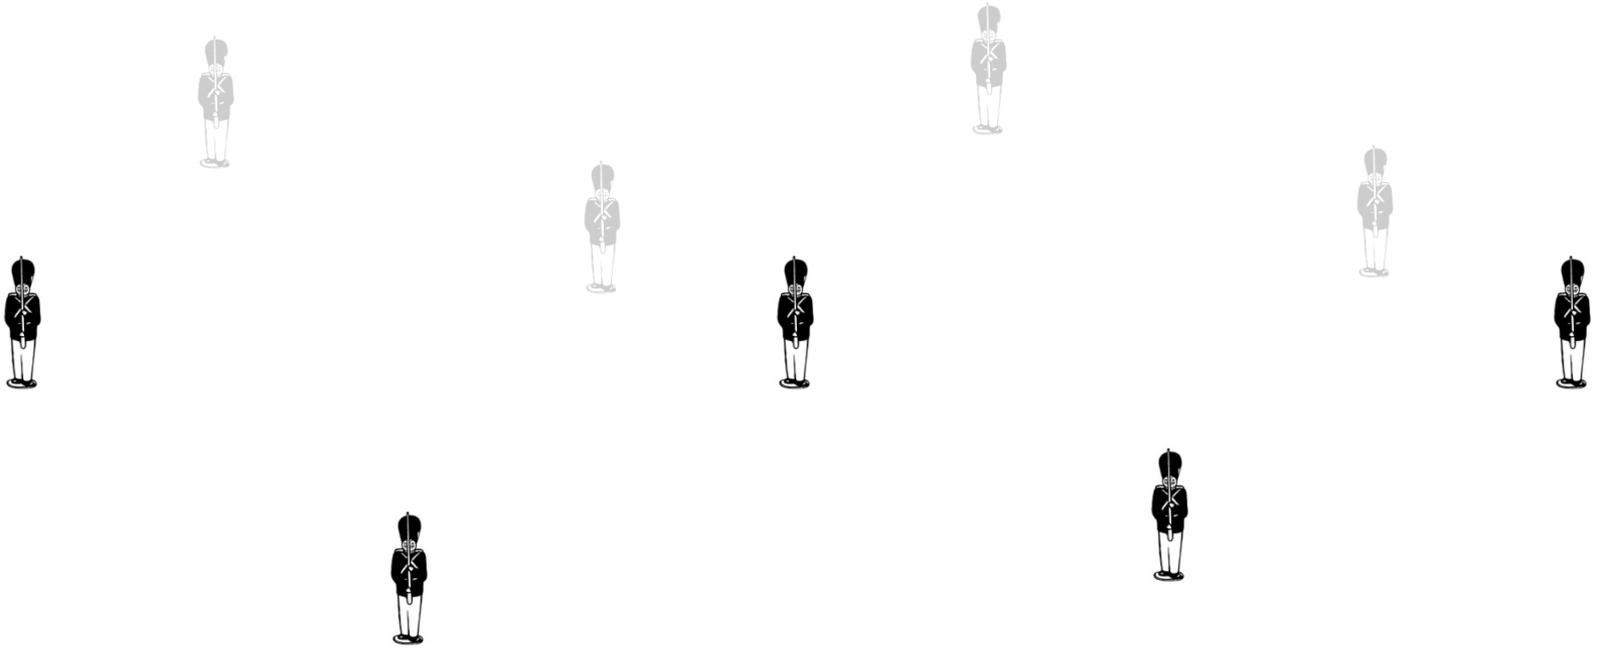
Large scale





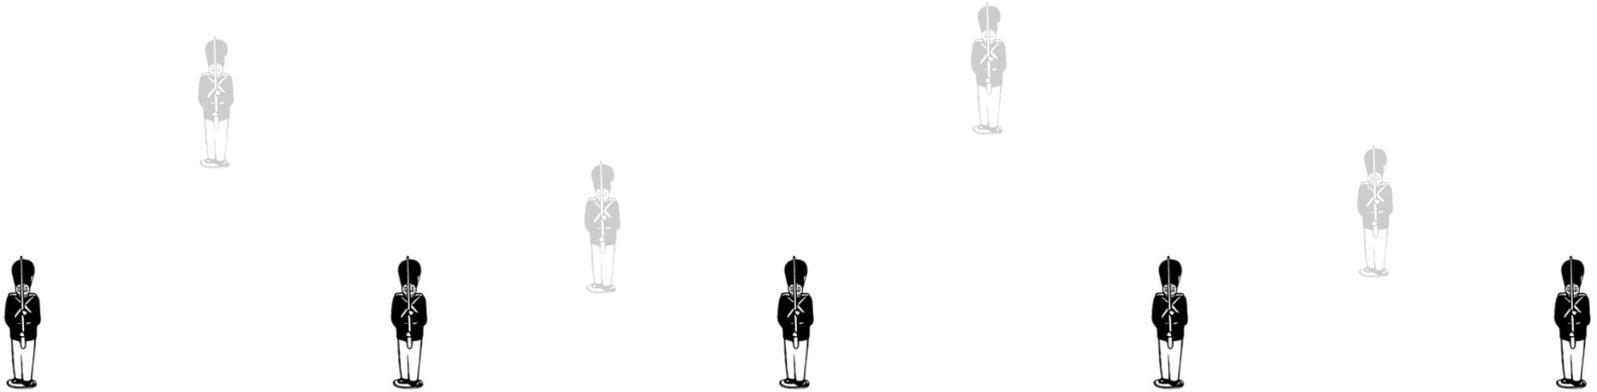
Large scale





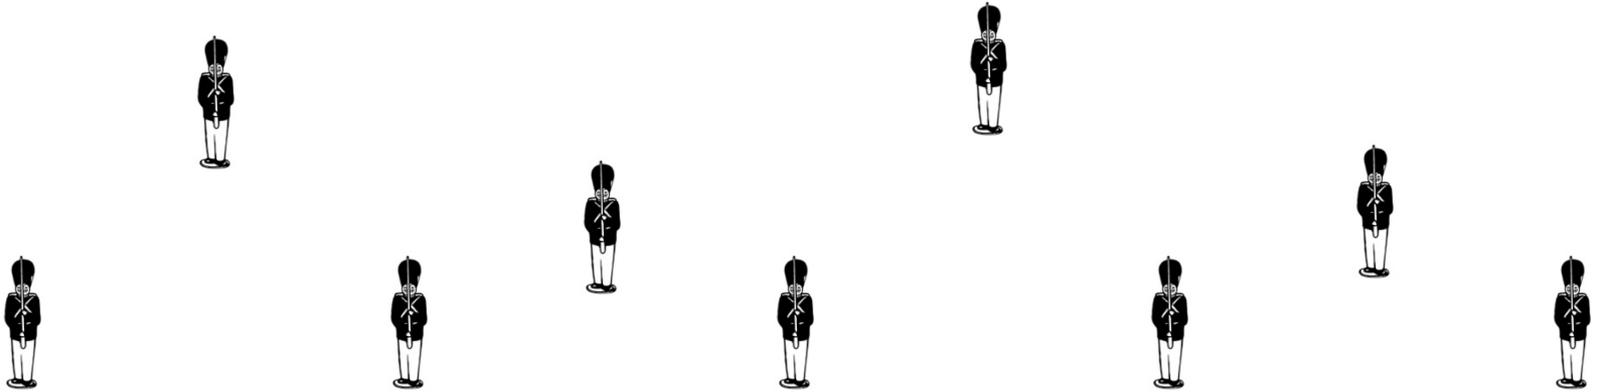
Intermediate scale





Intermediate scale





Small scale





This matrix \mathbf{S} has $N - 1$ linearly independent eigenvectors, \mathbf{v}^k , and corresponding real eigenvalues, λ_k

$$\mathbf{S} \mathbf{v}^k = \lambda_k \mathbf{v}^k .$$

Since \mathbf{v}^k span the space \mathcal{R}^{N-1} , any initial configuration of the soldiers can be written as a linear combination:

$$\mathbf{x}^{(0)} = \sum_{k=1}^{N-1} c_k \mathbf{v}^k$$

with some coefficients, c_k .

Hence, we obtain after m iterations:

$$\begin{aligned}\mathbf{x}^{(m)} &= \mathbf{S}\mathbf{x}^{(m-1)} = \mathbf{S}^2\mathbf{x}^{(m-2)} = \\ \dots &= \mathbf{S}^m\mathbf{x}^{(0)} = \mathbf{S}^m \sum_k c_k \mathbf{v}^k = \sum_k c_k \lambda_k^m \mathbf{v}^k\end{aligned}$$

Conclusion:

$$\lim_{m \rightarrow \infty} \mathbf{x}^{(m)} \rightarrow 0 \quad \text{if} \quad |\lambda_k| < 1, \quad k = 1, \dots, N-1$$

The iteration converges if the spectral radius, ρ , of the iteration matrix, \mathbf{S} , is smaller than 1.

Observation: the eigenvectors and eigenvalues of the matrix \mathbf{S} are given by

$$\mathbf{v}^k = \left\{ \mathbf{v}_j^k \right\} = \sin \left(\frac{jk\pi}{N} \right), \quad j = 1, \dots, N-1,$$

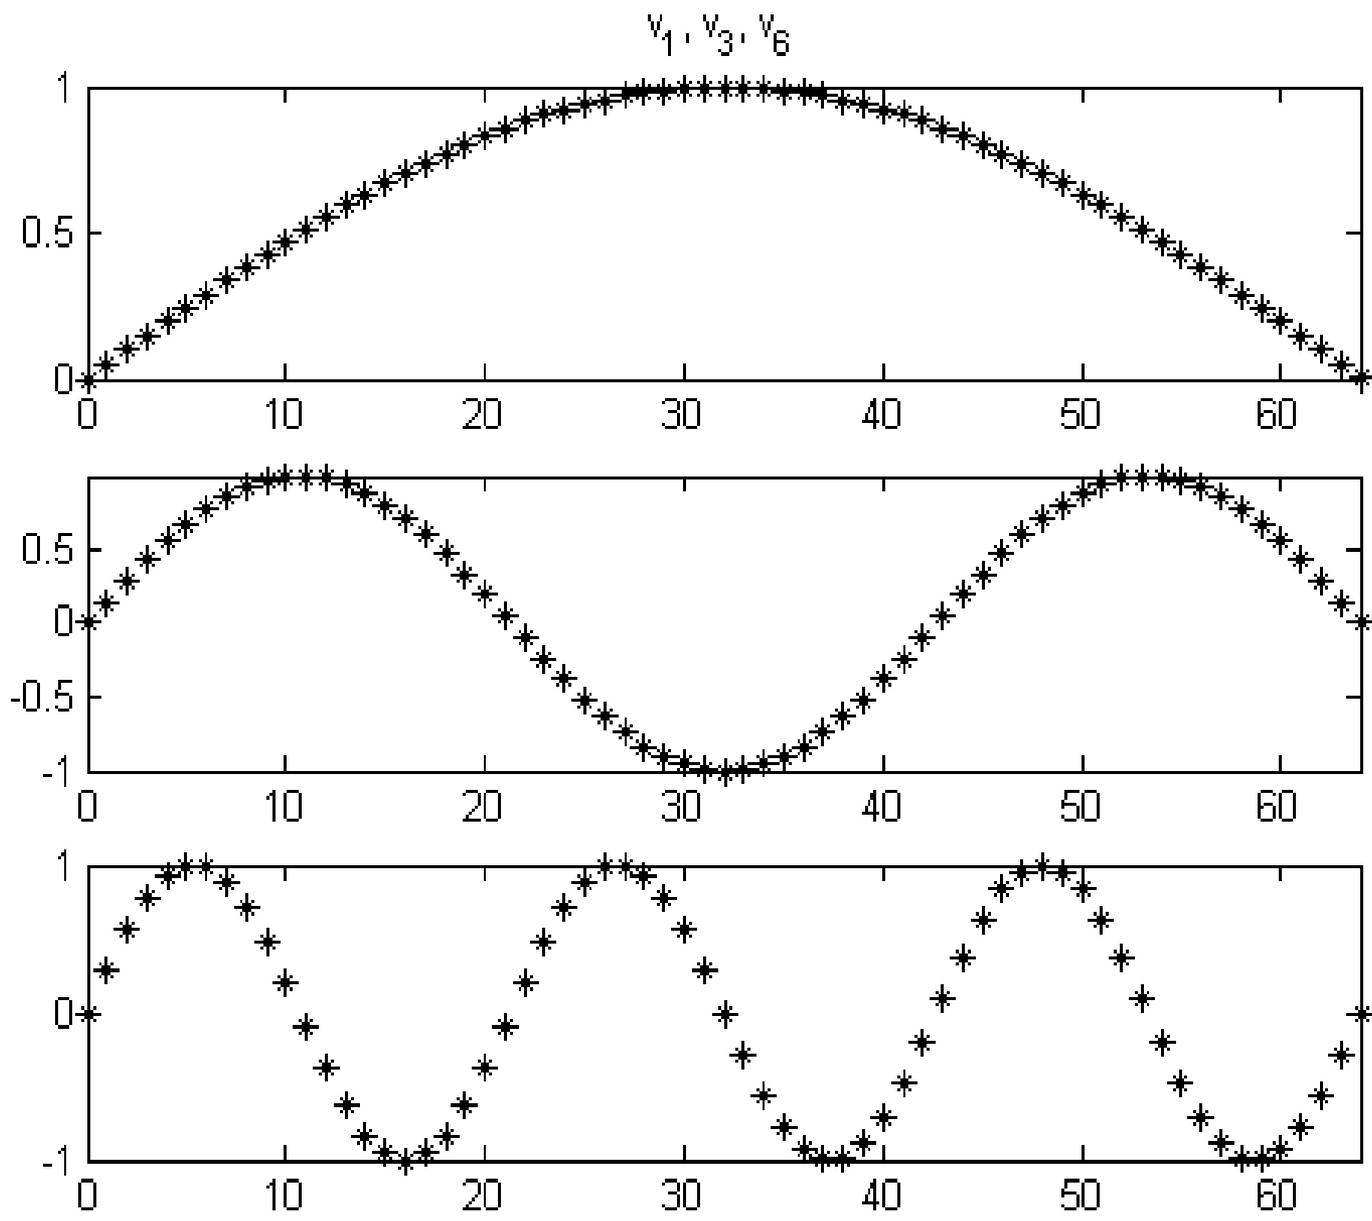
$$\lambda_k = \cos \left(\frac{k\pi}{N} \right),$$

with $k = 1, \dots, N-1$.

Proof: Using the trigonometric identity,

$$\frac{1}{2} \left[\sin \frac{(j-1)k\pi}{N} + \sin \frac{(j+1)k\pi}{N} \right] = \cos \frac{k\pi}{N} \sin \frac{jk\pi}{N},$$

and the fact that $\sin 0 = \sin \pi = 0$, we obtain by substitution, $\mathbf{S} \mathbf{v}^k = \lambda_k \mathbf{v}^k$.



Note: since $|\lambda_k| < 1$, the method converges. But, for some eigenvectors, $|\lambda_k|$ is close to 1, so convergence is slow. In particular, for $k\pi/N \ll 1$, we have,

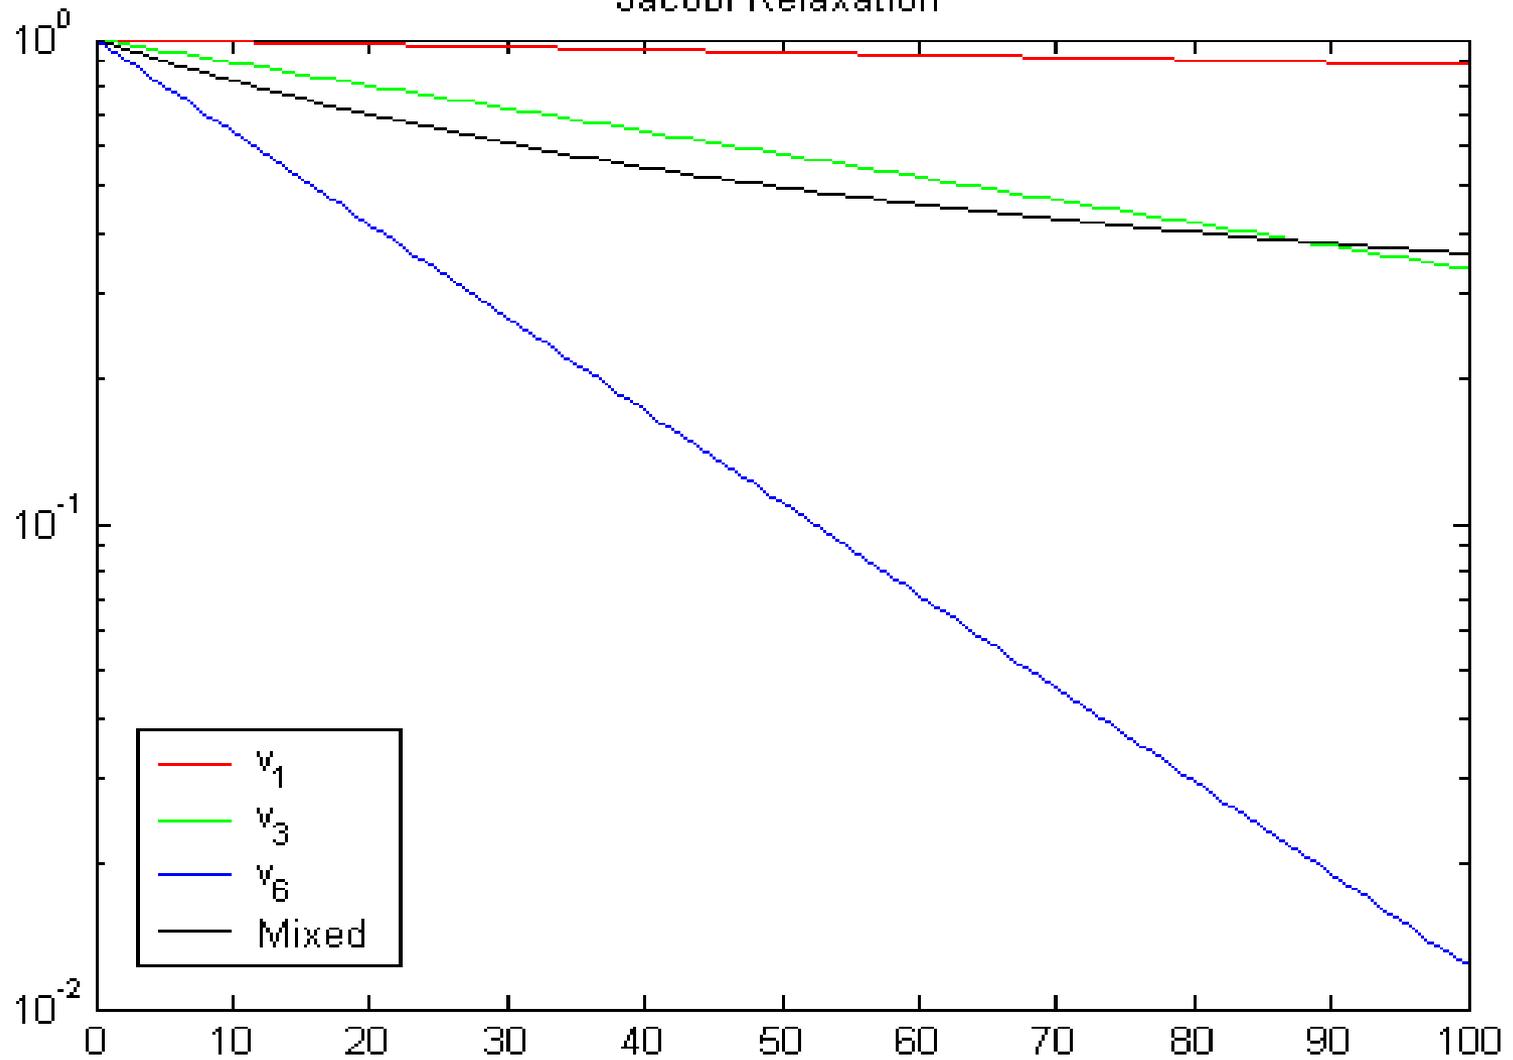
$$\lambda_k = \cos\left(\frac{k\pi}{N}\right) \approx 1 - \frac{1}{2} \left(\frac{k\pi}{N}\right)^2.$$

For $k=1$ we obtain

$$\lambda_1^m \approx \left[1 - \frac{1}{2} \left(\frac{\pi}{N}\right)^2\right]^m \approx e^{-\frac{1}{2}m \left(\frac{\pi}{N}\right)^2}.$$

Conclusion: $O(N^2)$ iterations are required to reduce such an error by an order of magnitude.

Jacobi Relaxation



How much work do we save?

Jacobi's method requires about N^2 iterations and $N^2 * N = N^3$ operations to improve the accuracy by an order of magnitude.

The multiscale approach solves the problem in about $\text{Log}_2(N)$ iterations (whistle blows) and only about N operations.

Example: for $N = 1000$ we require about:

10 iterations and 1000 operations

instead of about

1,000,000 iterations and 1,000,000,000 operations

How important is computational efficiency?

Suppose that we have **three different algorithms** for a given problem, with different computational complexities for input size N :

Algorithm 1: $10^6 N$ operations

Algorithm 2: $10^3 N^2$ operations

Algorithm 3: N^3 operations

Suppose that the problem size, N , is such that **Algorithm 1 requires one second.**

How long do the others require?

Computer Speed (ops/sec)	N	Algorithm 1 $O(N)$	Algorithm 2 $O(N^2)$	Algorithm 3 $O(N^3)$
1M ($\sim 10^6$) (1980's)	1	1 sec	0.001 sec	0.000001 sec
1G ($\sim 10^9$) (1990's)	1K	1 sec	1 sec	1 sec
1T ($\sim 10^{12}$) (2000's)	1M	1 sec	17 min	12 days
1P ($\sim 10^{15}$) (2010's)	1G	1 sec	12 days	31,710 years

Stronger Computers \Rightarrow

Greater Advantage of Efficient Algorithms!

The catch: in less trivial problems, we cannot construct appropriate equations on the large scales without first propagating information from the small scales.

Skill in developing efficient multilevel algorithms is required for:

1. Choosing a good local iteration.
2. Choosing appropriate coarse-scale variables.
3. Choosing inter-scale transfer operators.
4. Constructing coarse-scale approximations to the fine-scale problem.

Damping

Recall: the eigenvectors and eigenvalues of the iteration matrix \mathbf{S} are given by

$$\mathbf{v}^k = \{v_j^k\} = \sin\left(\frac{jk\pi}{N}\right), \quad j = 1, \dots, N-1,$$

$$\lambda_k = \cos\left(\frac{k\pi}{N}\right),$$

with $k = 1, \dots, N-1$.

Note that convergence is also slow for $k / N \approx 1$.

This slow convergence can be overcome by **damping**:

$$x_j^{(i)} = (1 - \omega)x_j^{(i-1)} + \omega \frac{1}{2} (x_{j-1}^{(i-1)} + x_{j+1}^{(i-1)}),$$

where ω is a parameter.

Then, $\mathbf{x}^{(i)} = \mathbf{S}_\omega \mathbf{x}^{(i-1)}$, where

$$\mathbf{S}_\omega = (1 - \omega)\mathbf{I} + \omega\mathbf{S}.$$

Note: \mathbf{v}^k are **eigenvectors** of \mathbf{S}_ω . The corresponding **eigenvalues** are now $\lambda_k^{(\omega)} = 1 - \omega + \omega\lambda_k = 1 - \omega(1 - \lambda_k)$.

For $0 < \omega \leq 1$, we have **convergence**, $|\lambda_k^{(\omega)}| < 1$.

Definition:

Eigenvectors \mathbf{v}^k with $1 \leq k < N/2$ are called **smooth** (low-frequency).

Those with $N/2 \leq k < N$ are called **rough** or oscillatory (high-frequency).

Recall that $\lambda_k = \cos\left(\frac{k\pi}{N}\right)$, so for **rough** eigenvectors,

$$\lambda_k \leq 0.$$

Exercise: Find $0 < \omega < 1$ which yields optimal convergence for the set of rough modes for arbitrary N :

$$\omega : \sup_N \max_{\frac{N}{2} \leq k < N} \left| \lambda_k^{(\omega)} \right| = \min!,$$

i.e.,

$$\omega : \sup_{\lambda \in (-1, 0]} |1 - \omega + \omega \lambda| = \min!,$$

What is then the bound on the convergence factor, $\left| \lambda_k^{(\omega)} \right|$, maximized over the rough modes? (Clues in my introductory paper.)

1D Model Problem

Find u which satisfies:

$$Lu = u''(x) = f(x) , \quad x \in (0, 1) , \quad (1)$$

$$u(0) = u_0 ,$$

$$u(1) = u_1 .$$

In the particular case where $f = 0$, the solution is a **straight line** that connects u_0 with u_1 .

Discrete approximation: Since closed-form solutions exist only for a small number of differential equations, we solve such equations approximately by a **discrete approximation**.

Define a grid: divide the domain $(0,1)$ into N intervals. Assume for simplicity a uniform grid of **mesh-size** $h=1/N$.

Finite-difference discretization; examples:

Forward differences:

$$u' = \frac{u(x+h) - u(x)}{h} + O(h).$$

Backward differences:

$$u' = \frac{u(x) - u(x-h)}{h} + O(h).$$

Central differences:

$$u' = \frac{u(x+h) - u(x-h)}{2h} + O(h^2).$$

Second derivative:

$$u''(x) = \frac{u(x-h) - 2u(x) + u(x+h)}{h^2} + O(h^2). \quad (2)$$

Derivation: by the **Taylor theorem**

We can thus approximate the differential equation by a set of algebraic difference equations:

$$L^h u^h = \frac{u_{i+1}^h - 2u_i^h + u_{i-1}^h}{h^2} = f_i^h,$$

$$i = 1, \dots, N - 1,$$

$$u_0^h = u_0,$$

$$u_N^h = u_1.$$

In matrix form:

$$\frac{1}{h^2} \begin{bmatrix} -2 & 1 & & & & \\ 1 & -2 & 1 & & & \\ & \dots & \dots & \dots & & \\ & & & 1 & -2 & 1 \\ & & & & 1 & -2 \end{bmatrix} \begin{bmatrix} u_1^h \\ u_2^h \\ \dots \\ u_{N-2}^h \\ u_{N-1}^h \end{bmatrix} =$$

$$\begin{bmatrix} f_1^h - u_0^h / h^2 \\ f_2^h \\ \dots \\ f_{N-2}^h \\ f_{N-1}^h - u_1^h / h^2 \end{bmatrix} \cdot$$

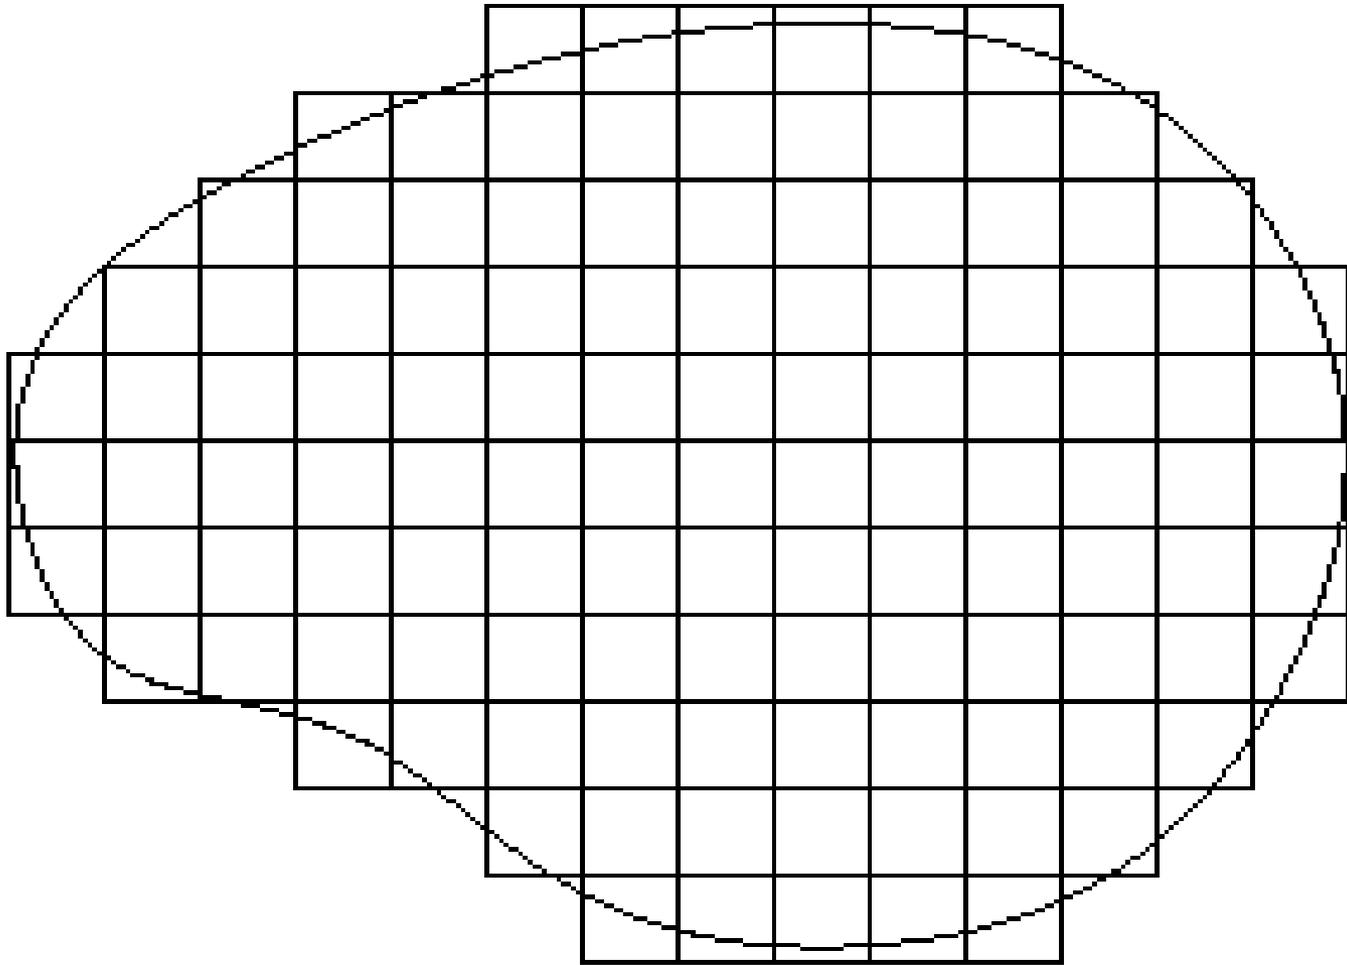
This is a **tridiagonal** system of equations which can be solved directly or iteratively.

2D Model Problem

Find u which satisfies:

$$\begin{aligned} Lu = u_{xx} + u_{yy} &= f(x, y), & (x, y) \in \Omega, \\ u &= g(x, y), & (x, y) \in \partial\Omega. \end{aligned} \quad (4)$$

This is the **2D Poisson equation**, with **Dirichlet boundary conditions**. It is an **elliptic** partial differential equation which appears in many models.



$$\Omega^h$$

Discrete approximation

Define a grid: $\Omega^h \subset \Omega$ (assumed to be uniform for simplicity, with mesh interval h).

Let u^h , g^h and f^h denote discrete approximations to u , g and f defined at the nodes of the grid.

Plug (2) for u_{xx} , and the analogous approximation for u_{yy} into (4), obtaining:

$$L^h u_{i,j}^h = \frac{u_{i-1,j}^h - 2u_{i,j}^h + u_{i+1,j}^h}{h^2} + \frac{u_{i,j-1}^h - 2u_{i,j}^h + u_{i,j+1}^h}{h^2} = f_{i,j}^h \quad \text{in } \Omega^h \quad (5)$$

$$u^h = g^h \quad \text{on } \partial^h \Omega^h$$

This yields a nonsingular linear system of equations for $u_{i,j}^h$ (the discrete operator satisfies a **maximum principle**.)

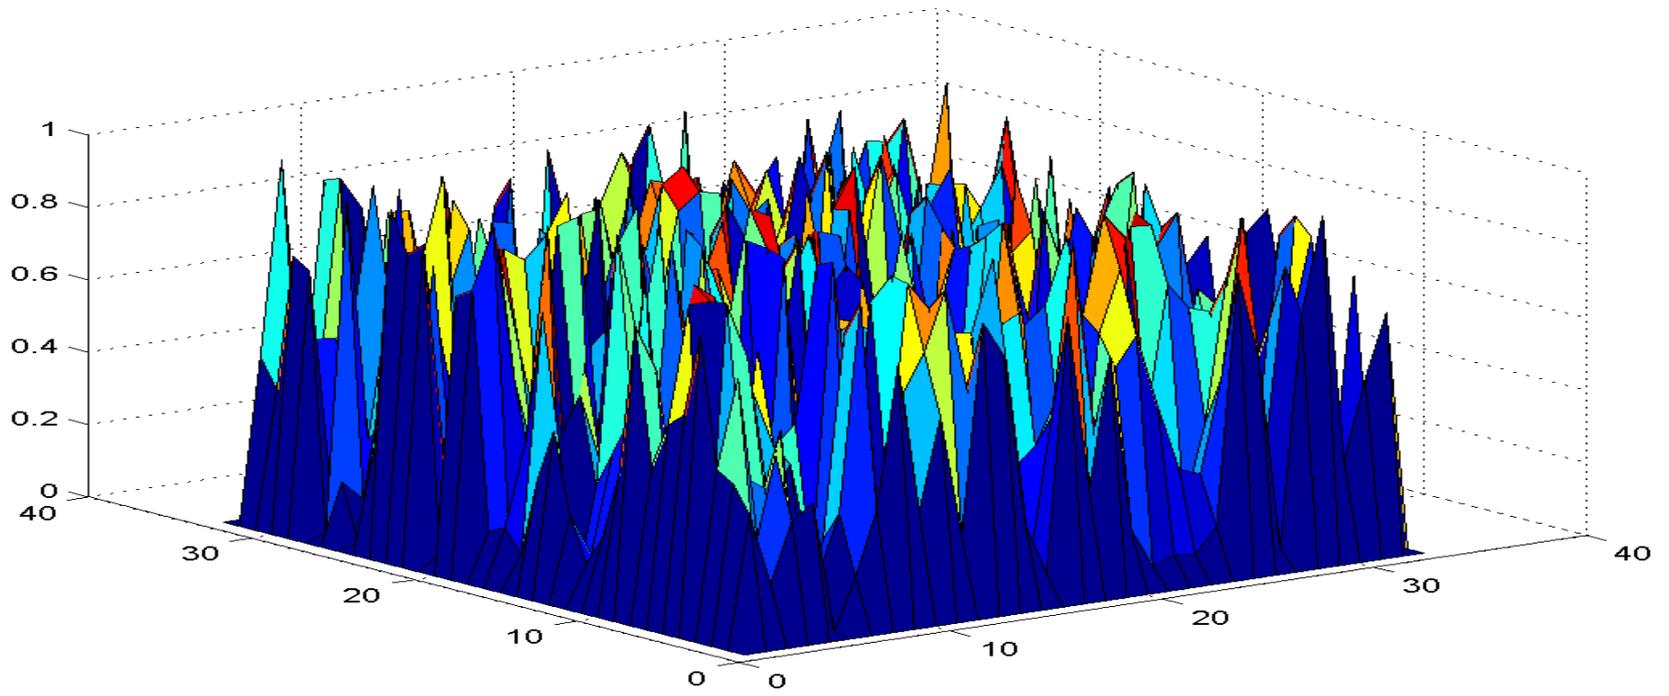
We consider solving this system by the classical approach of **Gauss-Seidel relaxation**.

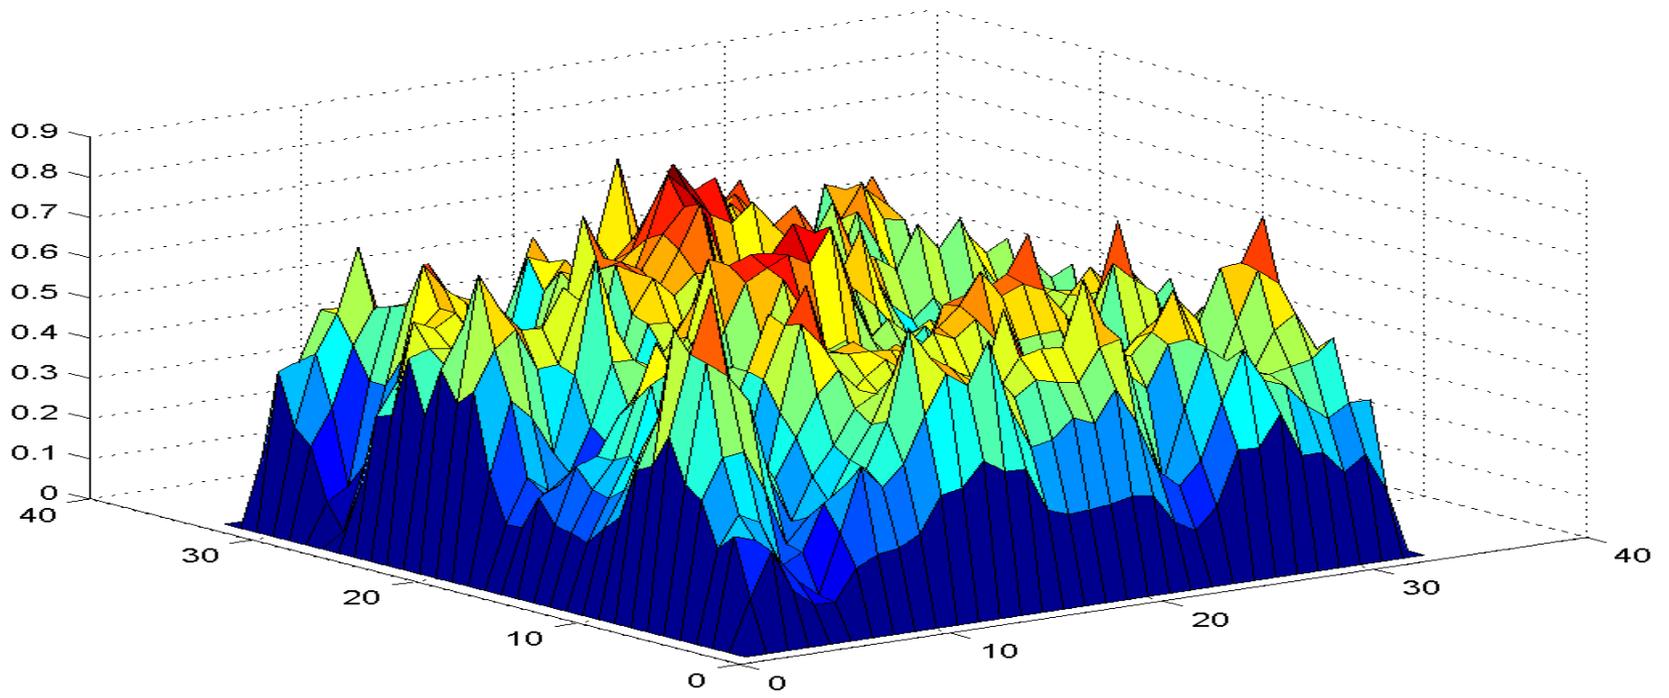
Gauss-Seidel (GS) Relaxation:

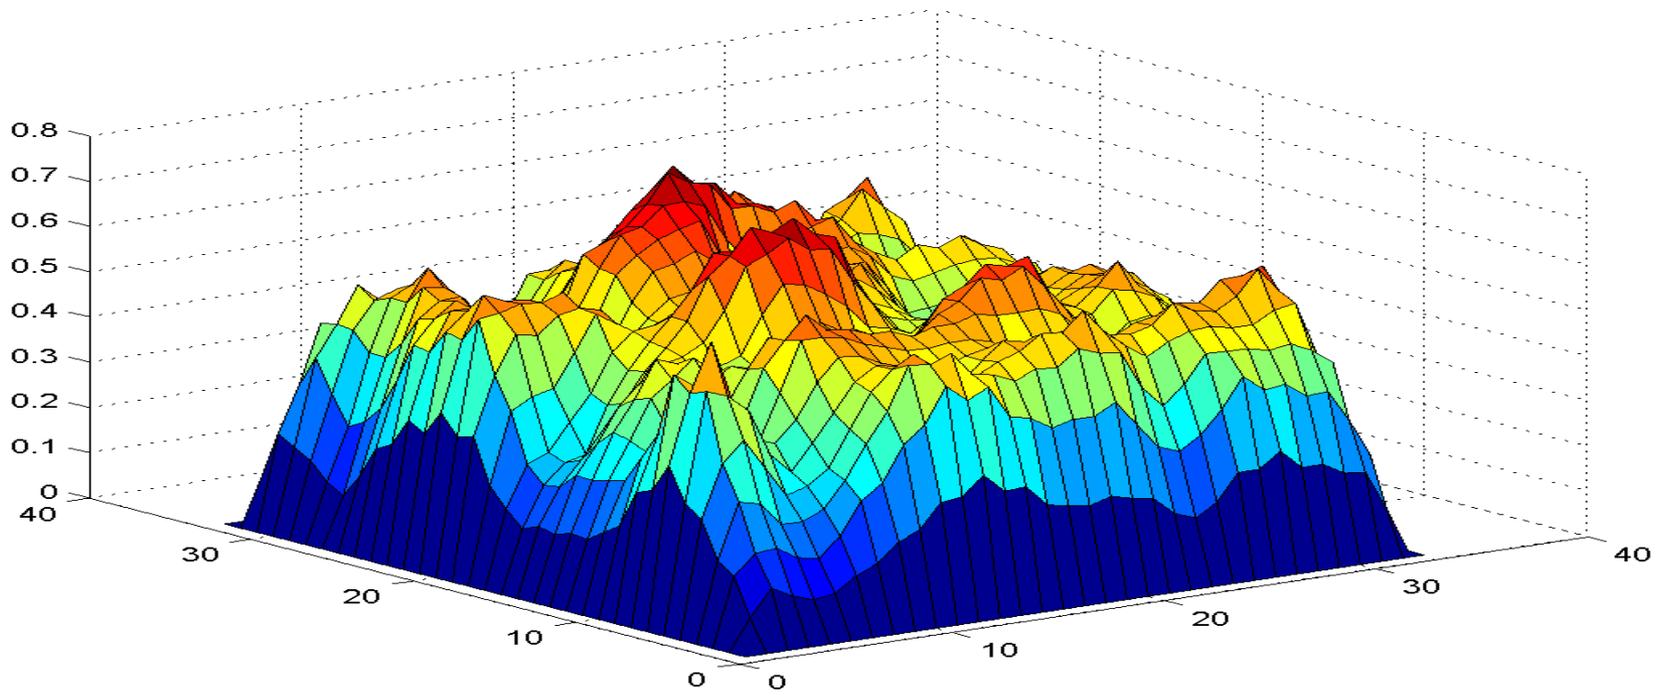
1. Choose initial guess, \tilde{u}^h .
2. Repeat until some convergence criterion is satisfied
{
Scan all variables in some prescribed order, and change each variable $\tilde{u}_{i,j}^h$ in turn so as to satisfy the (i,j) th equation.
}

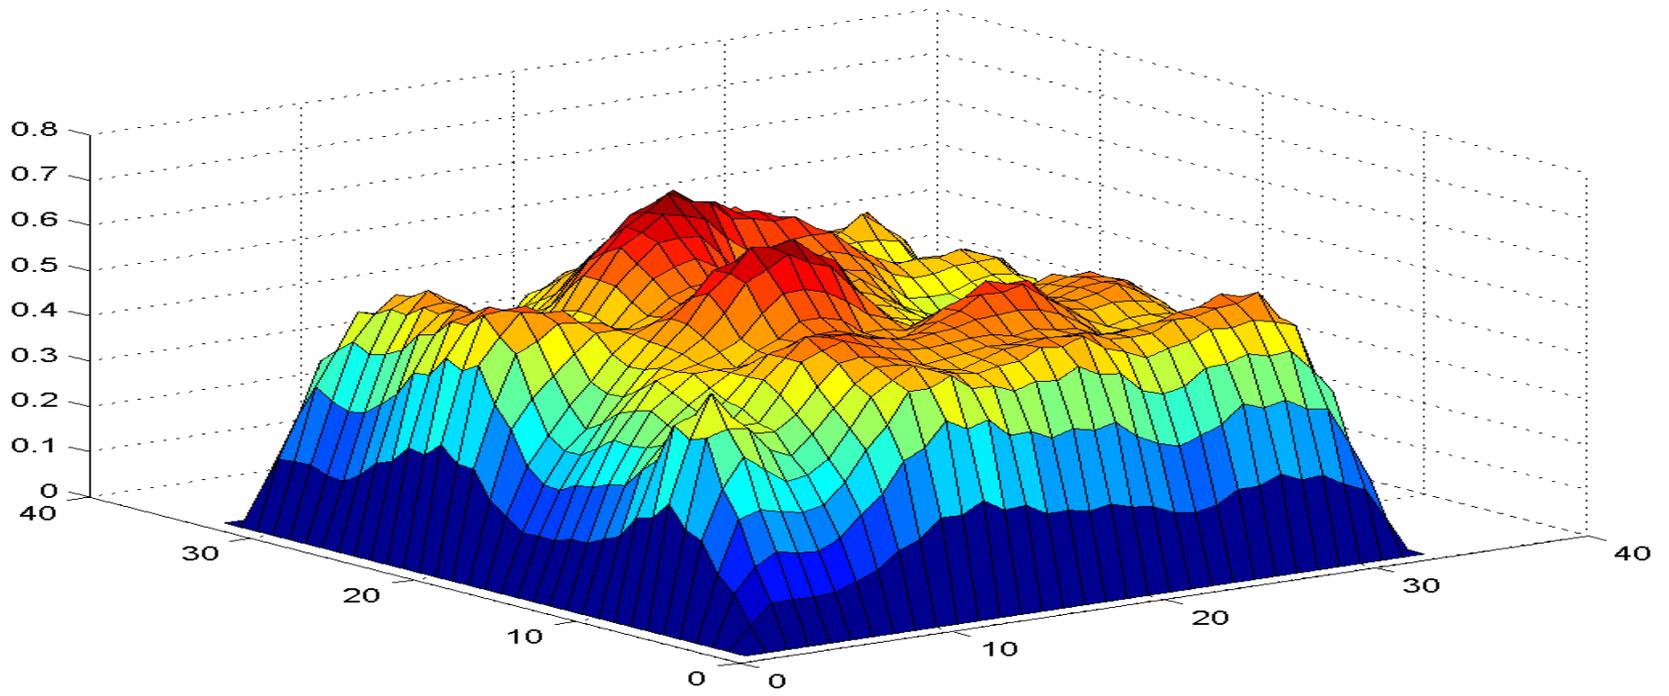
Observation: GS is a local process, because only near neighbors appear in each equation. Hence, it may be efficient for eliminating errors which can be detected locally. But large-scale ("smooth") errors are eliminated very slowly.

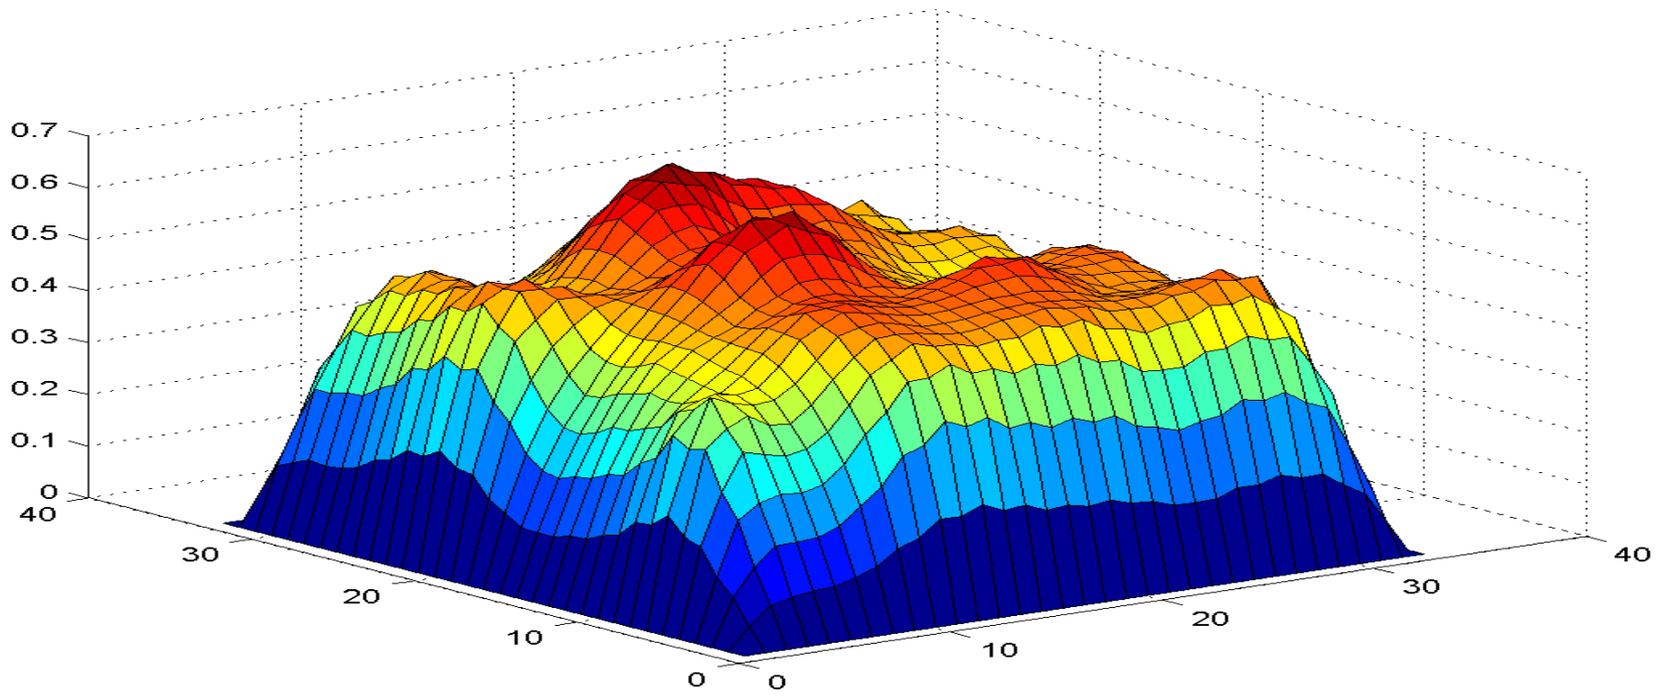
(The difference between GS and Jacobi is that old neighboring values are used in Jacobi, while the most updated values are used in GS.)

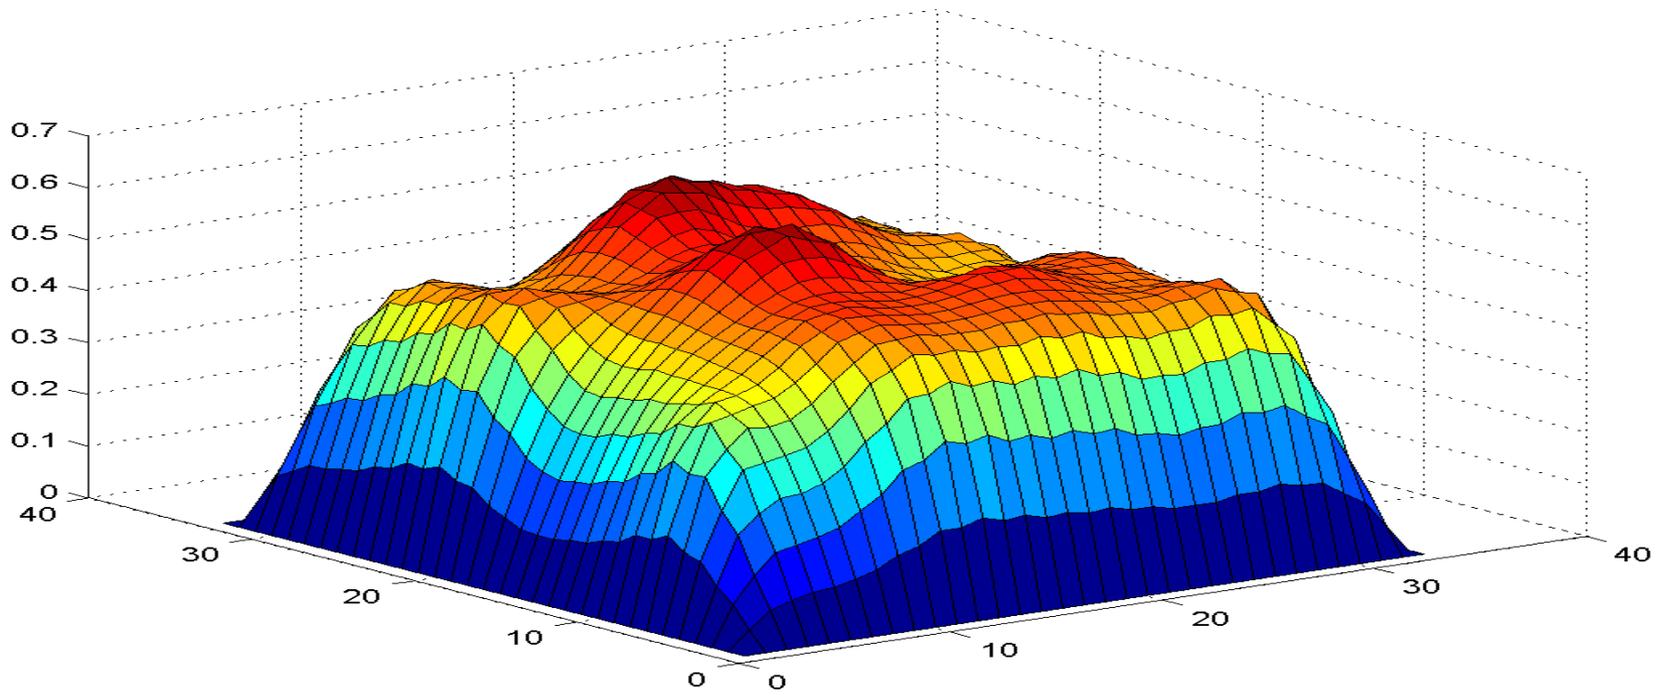


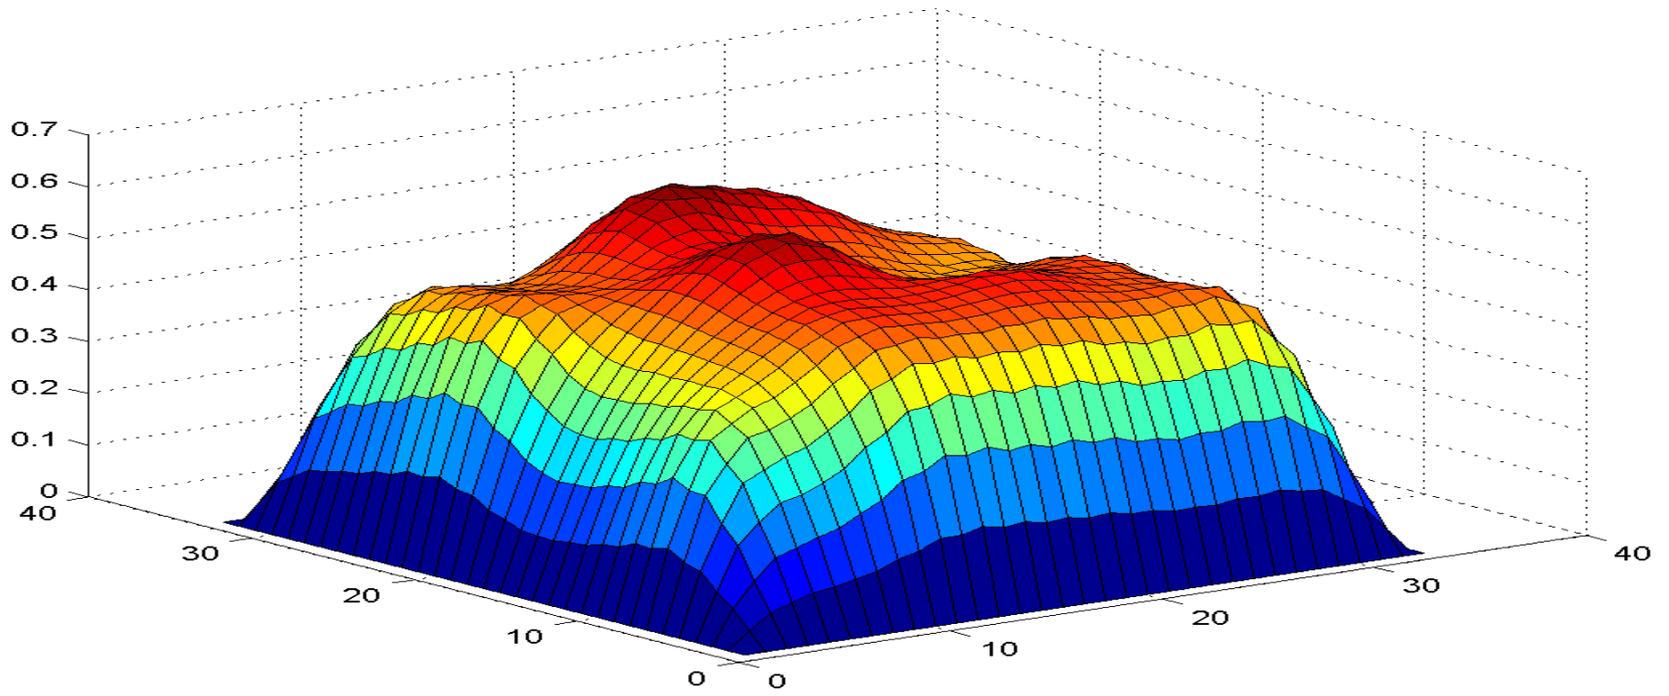


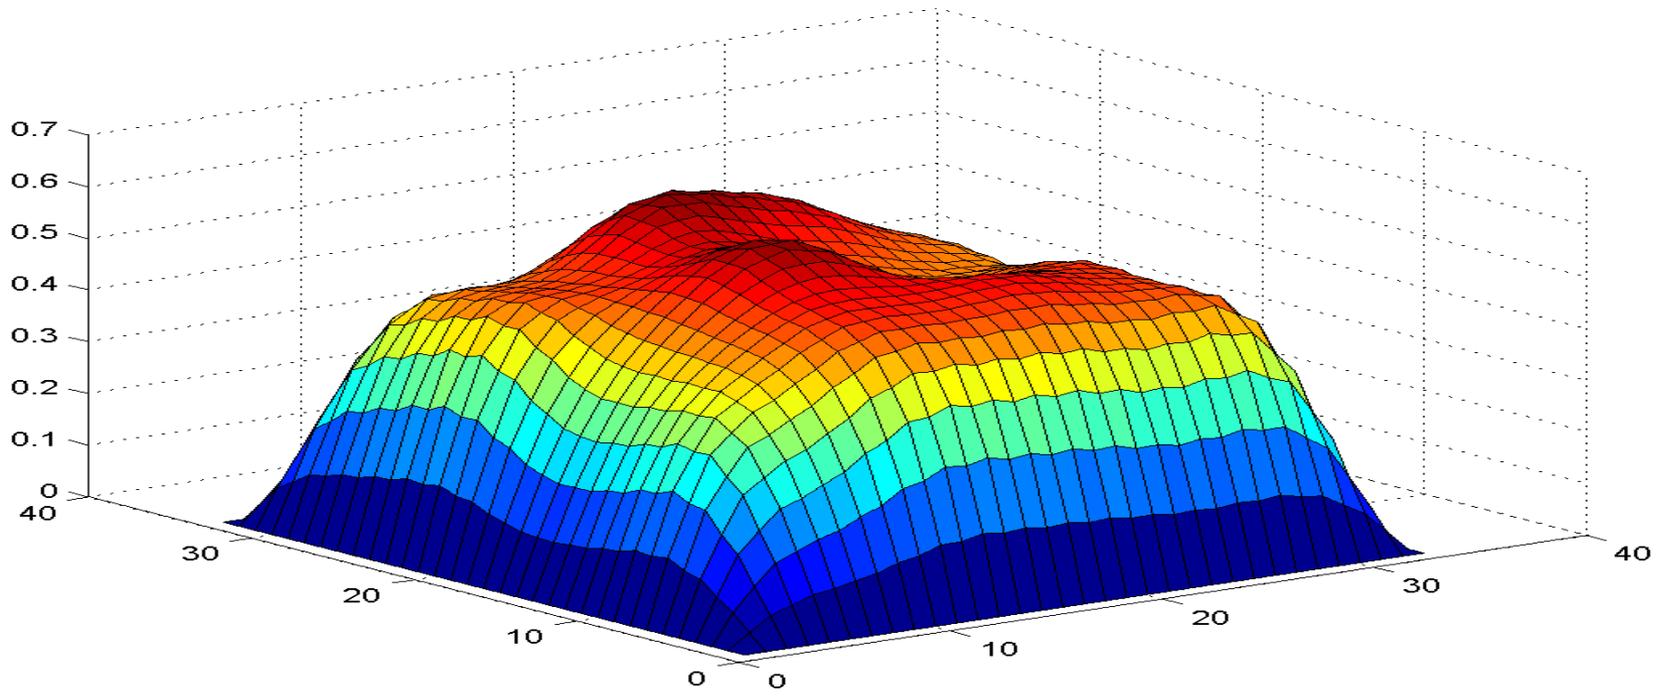


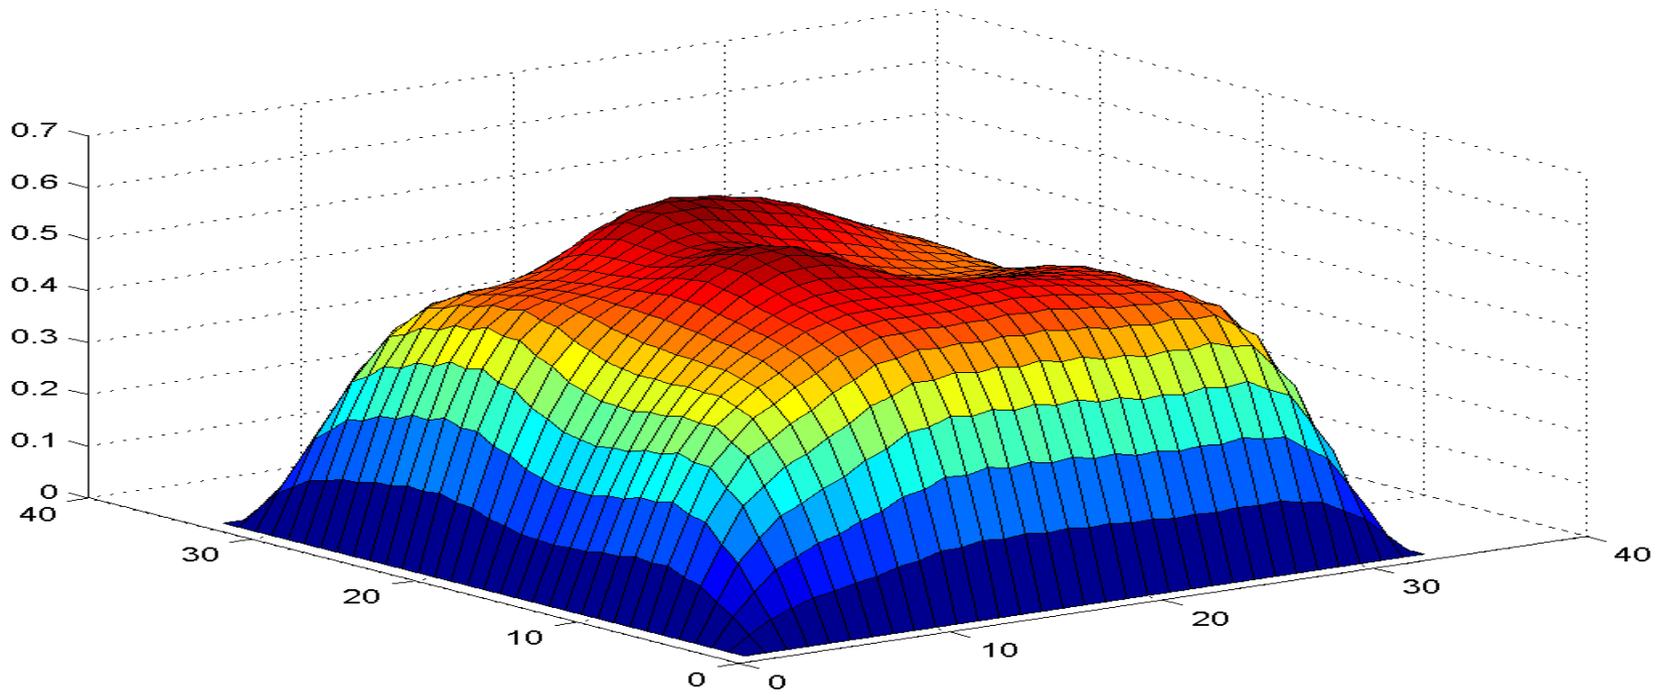


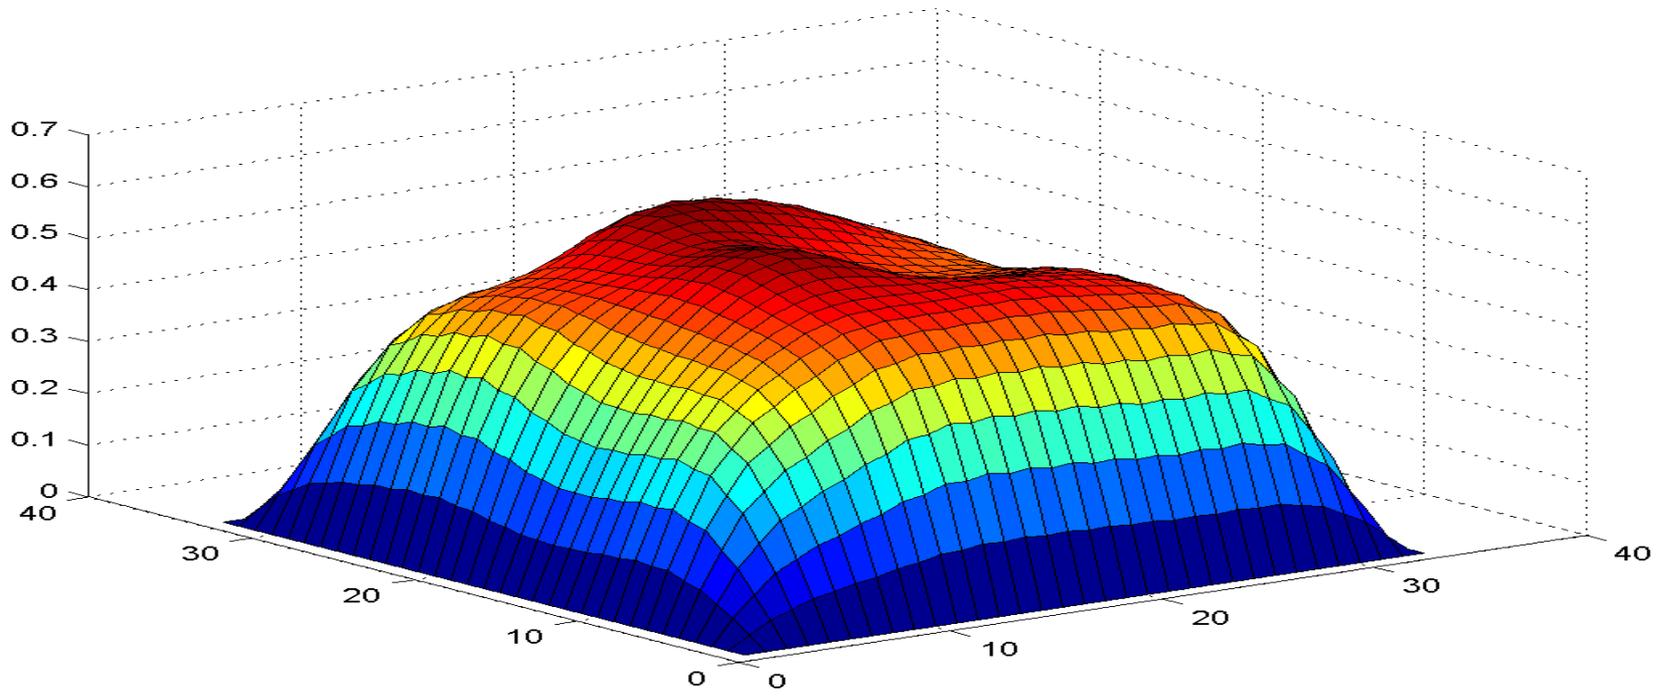


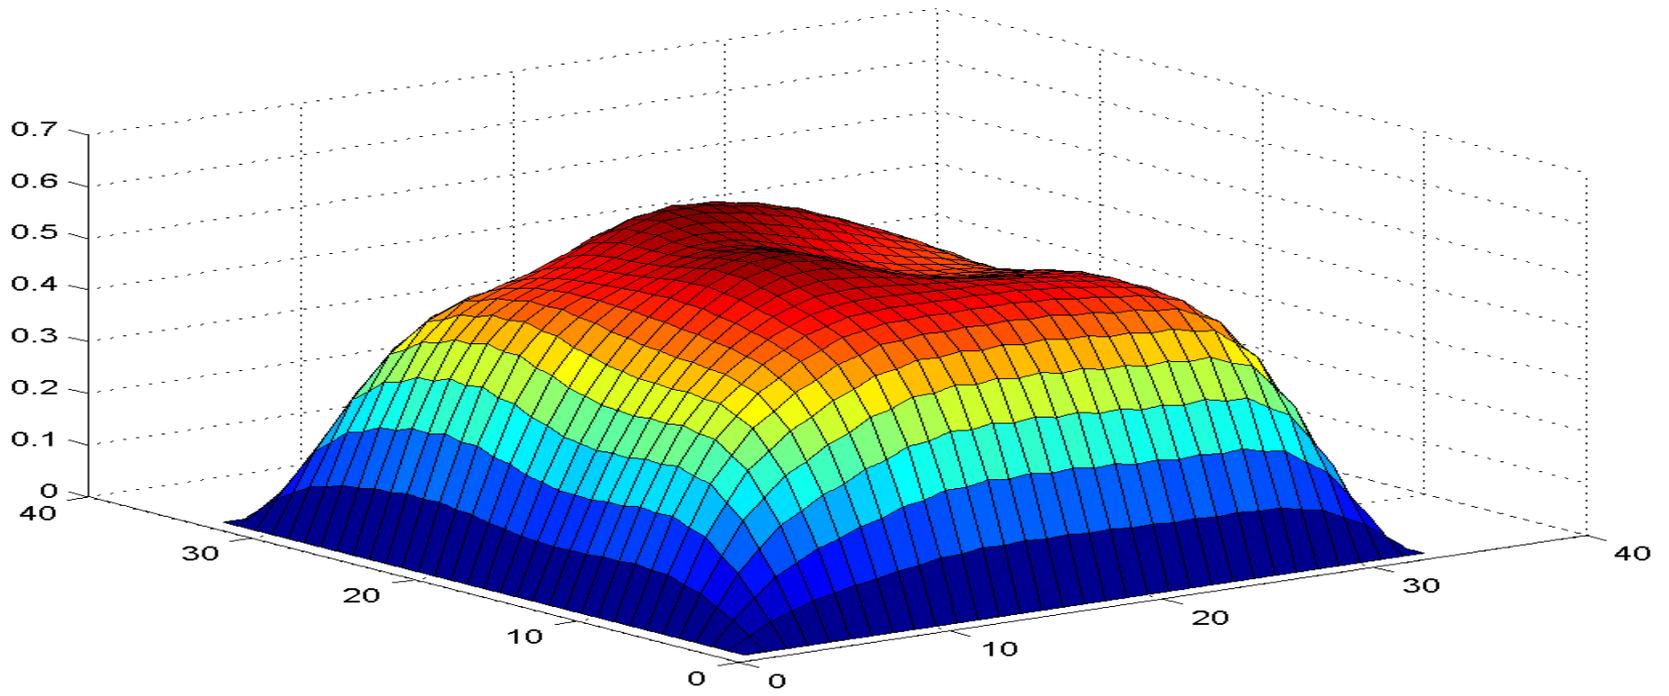


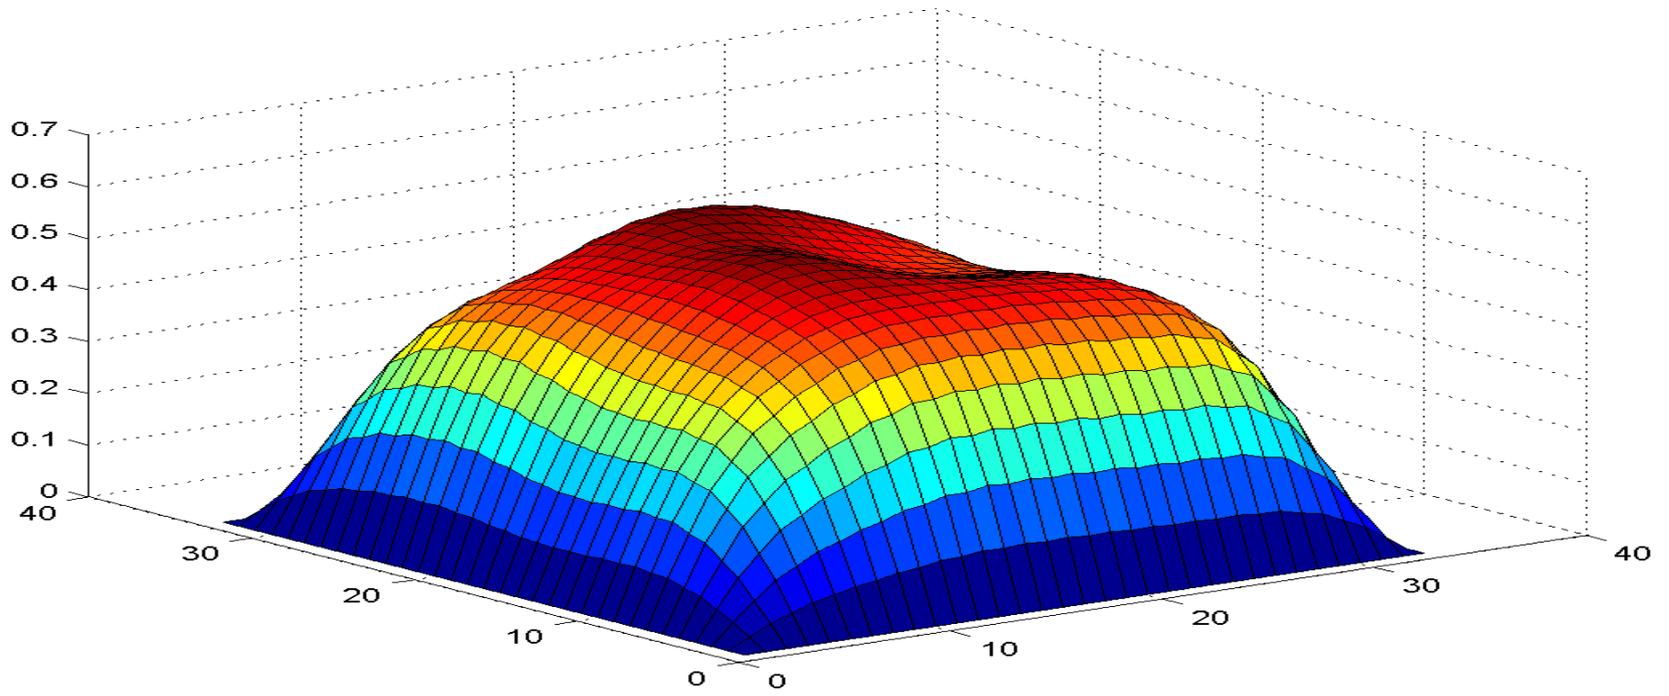


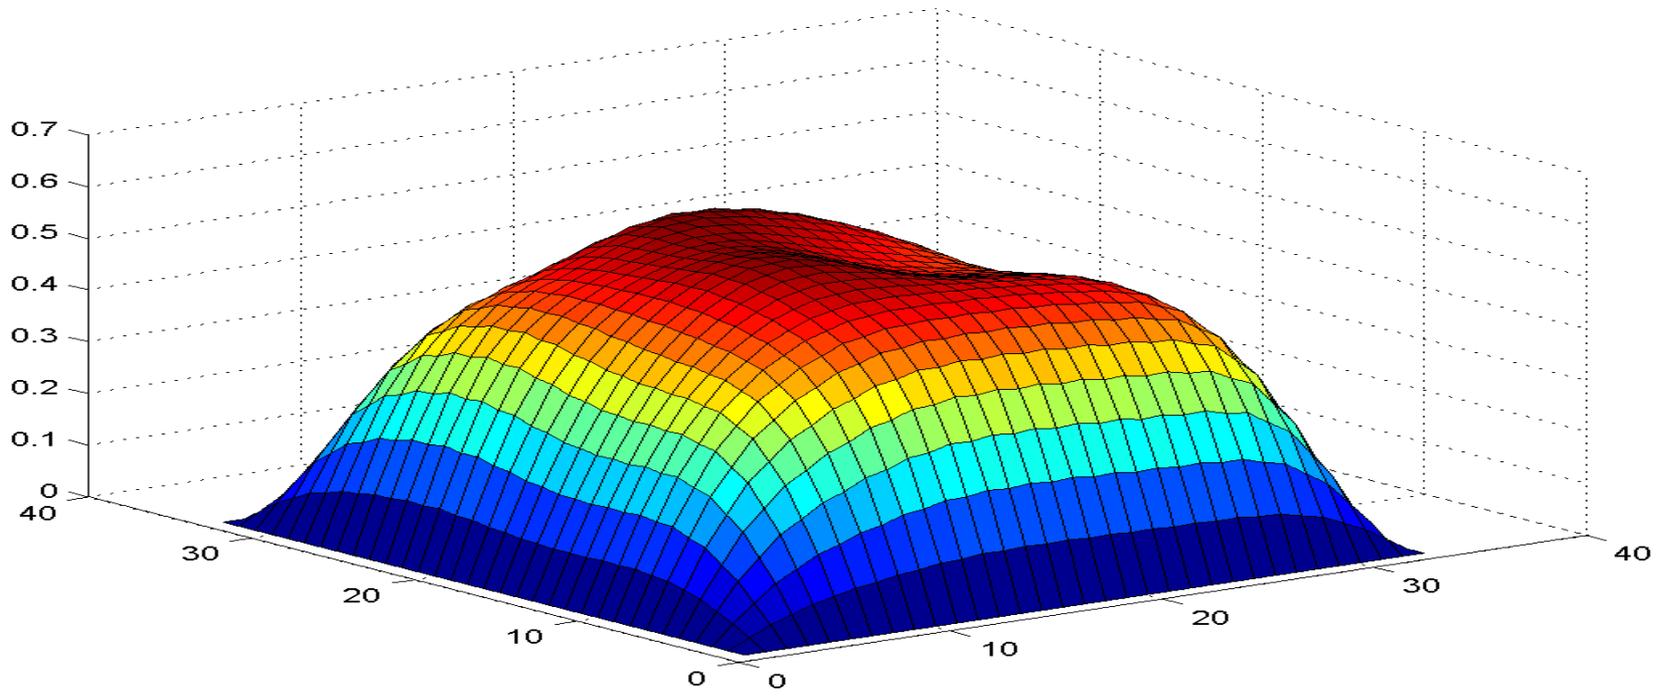


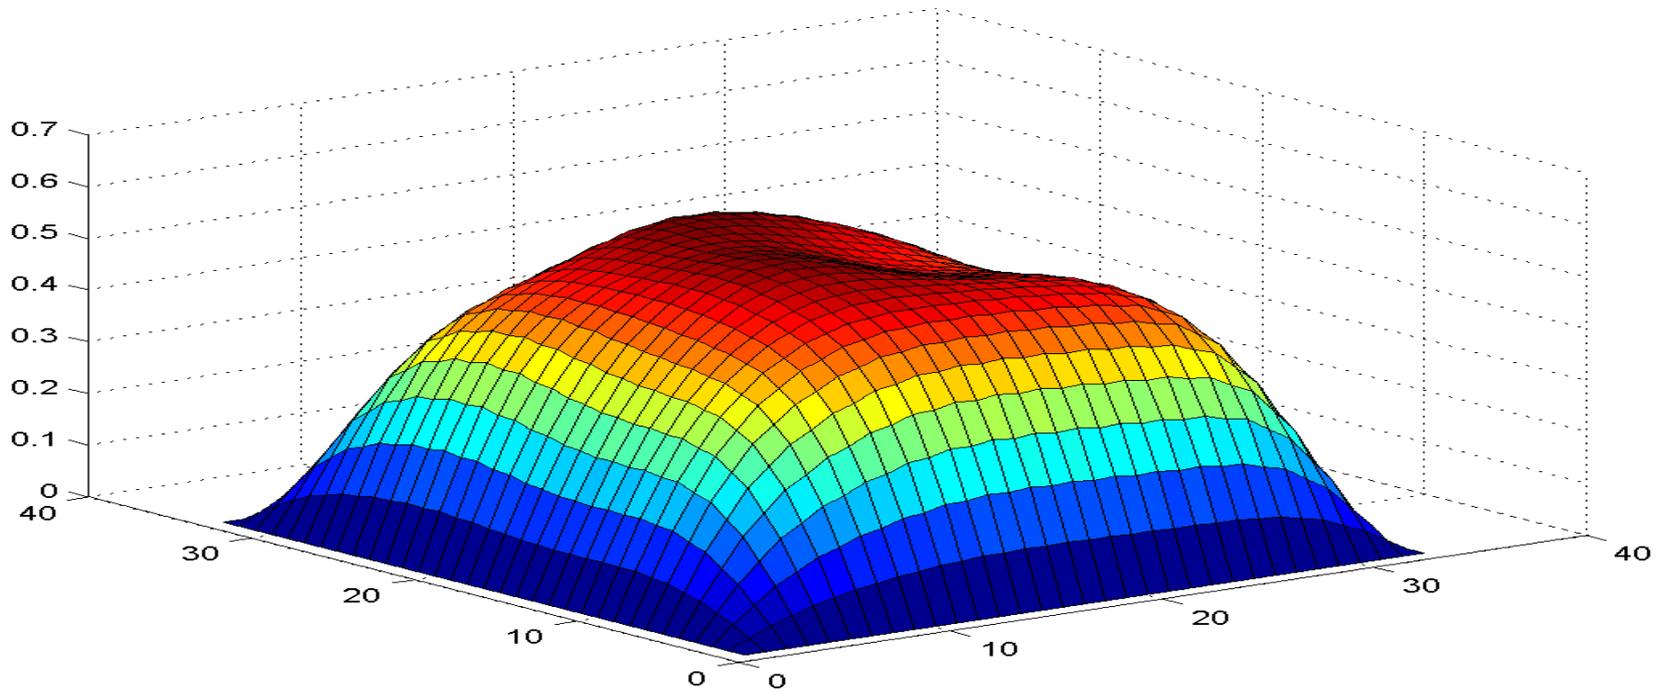


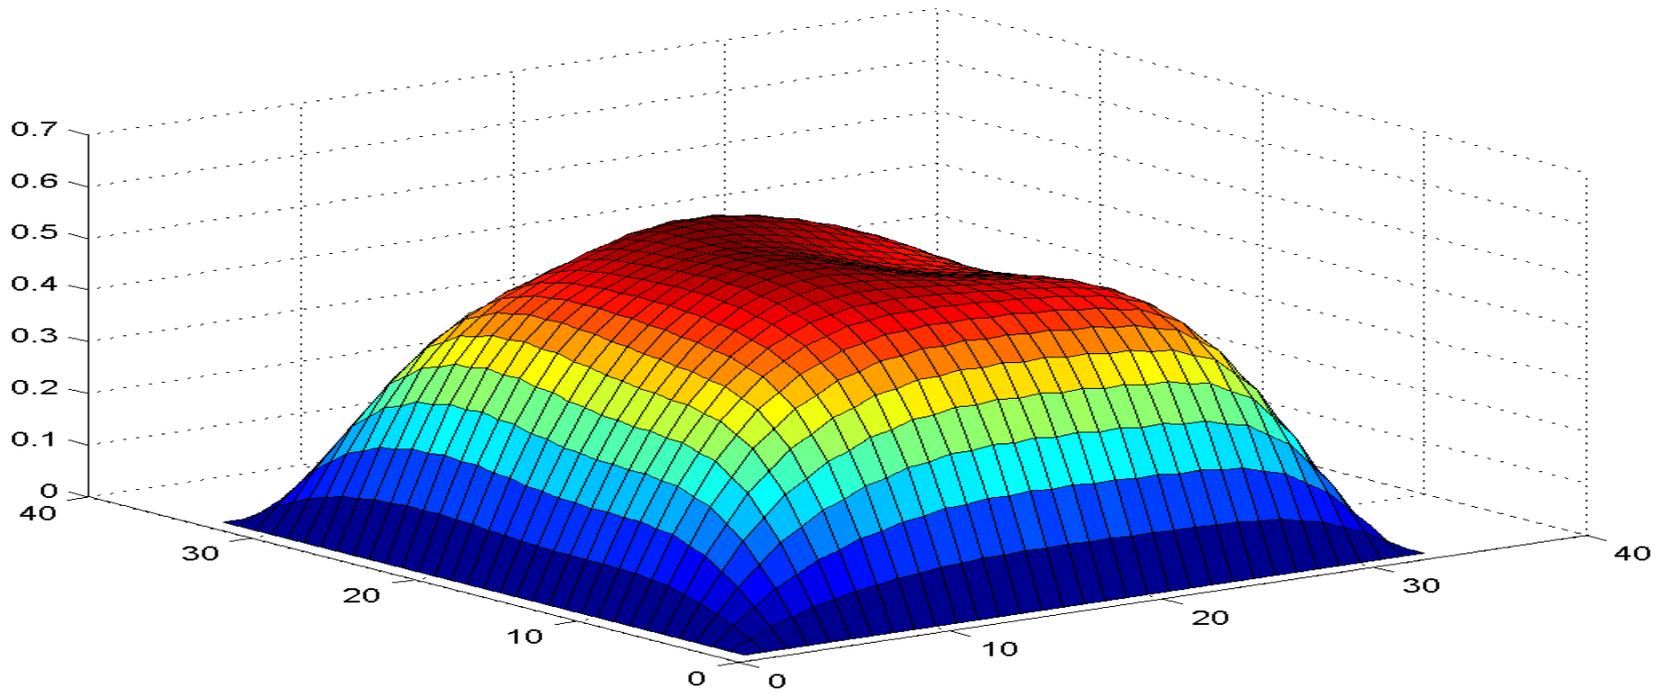


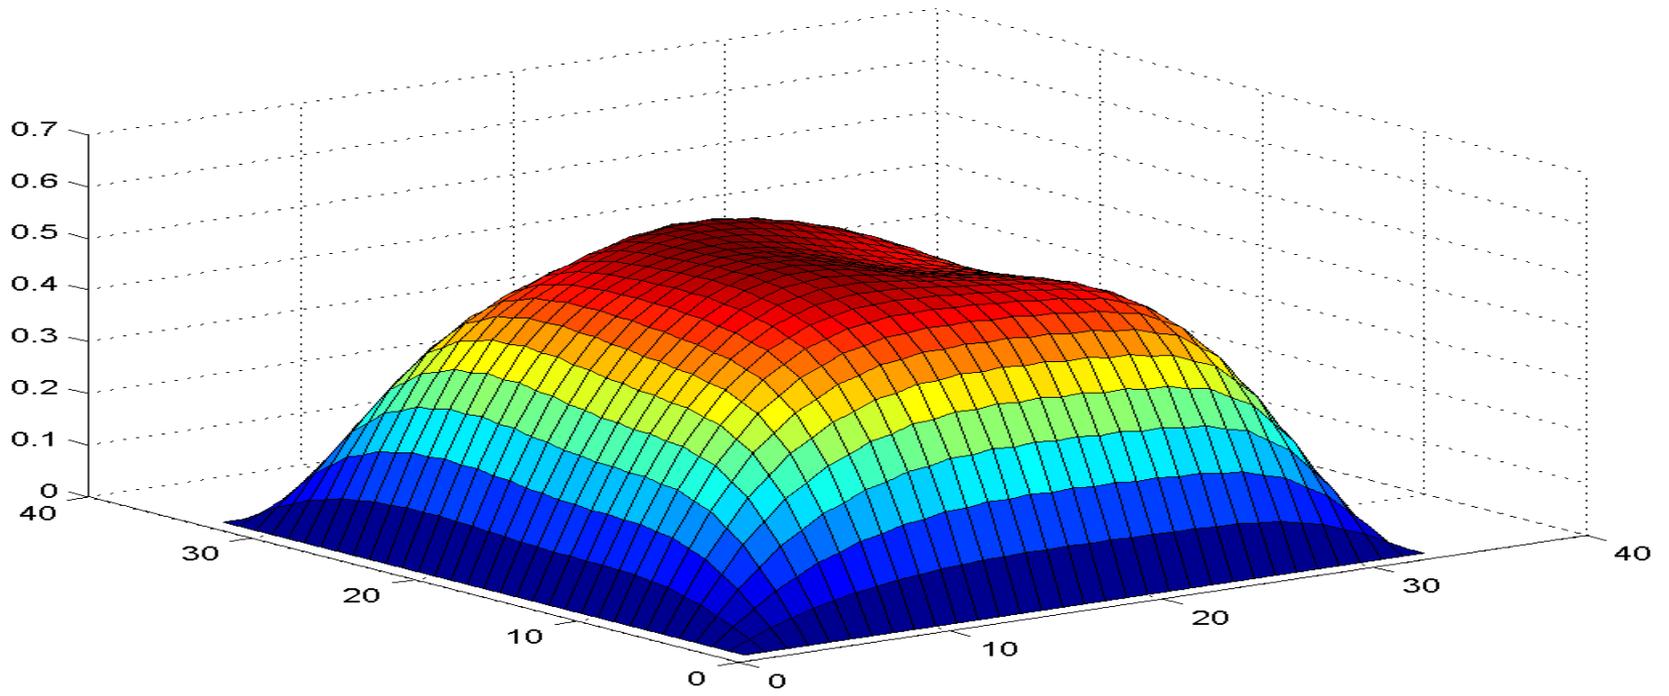


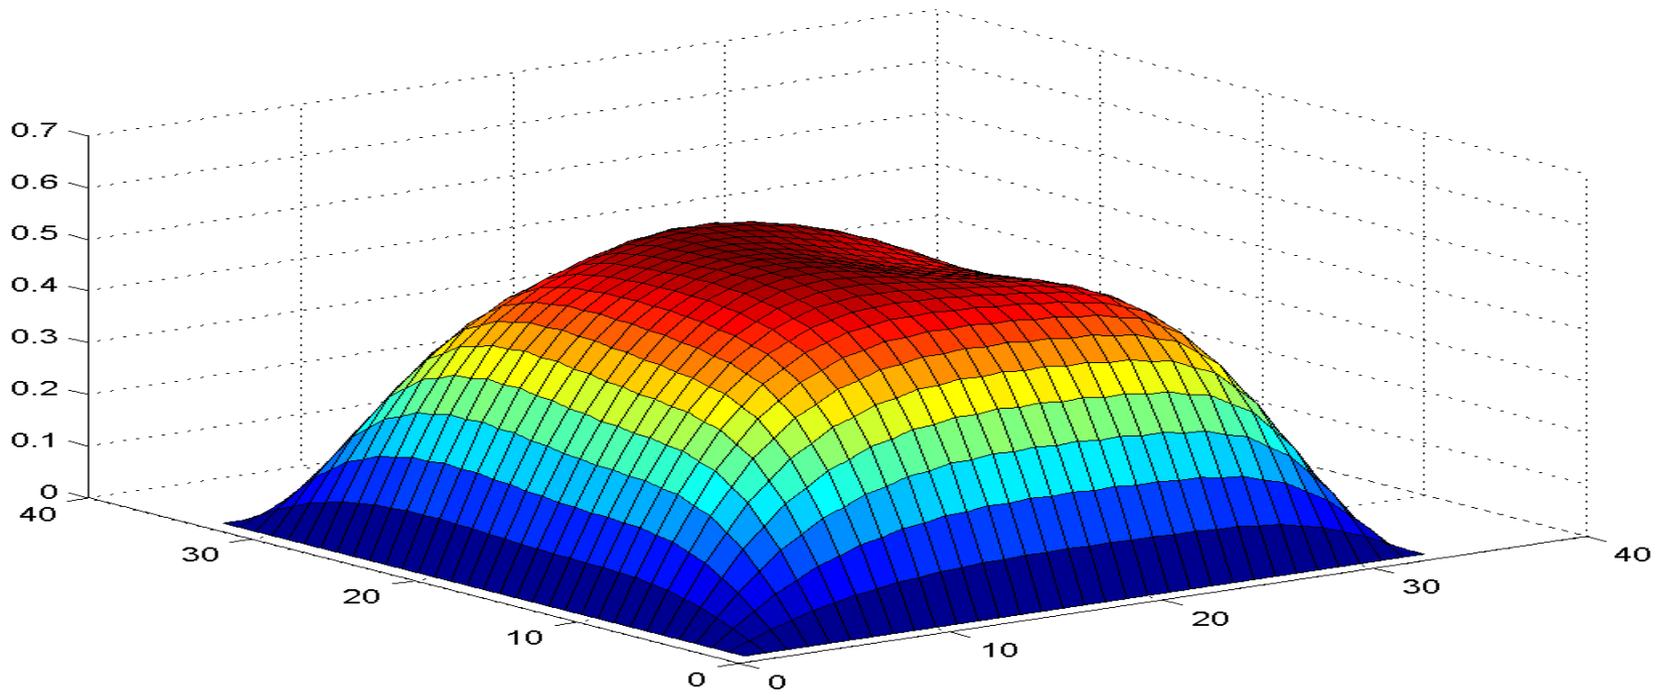


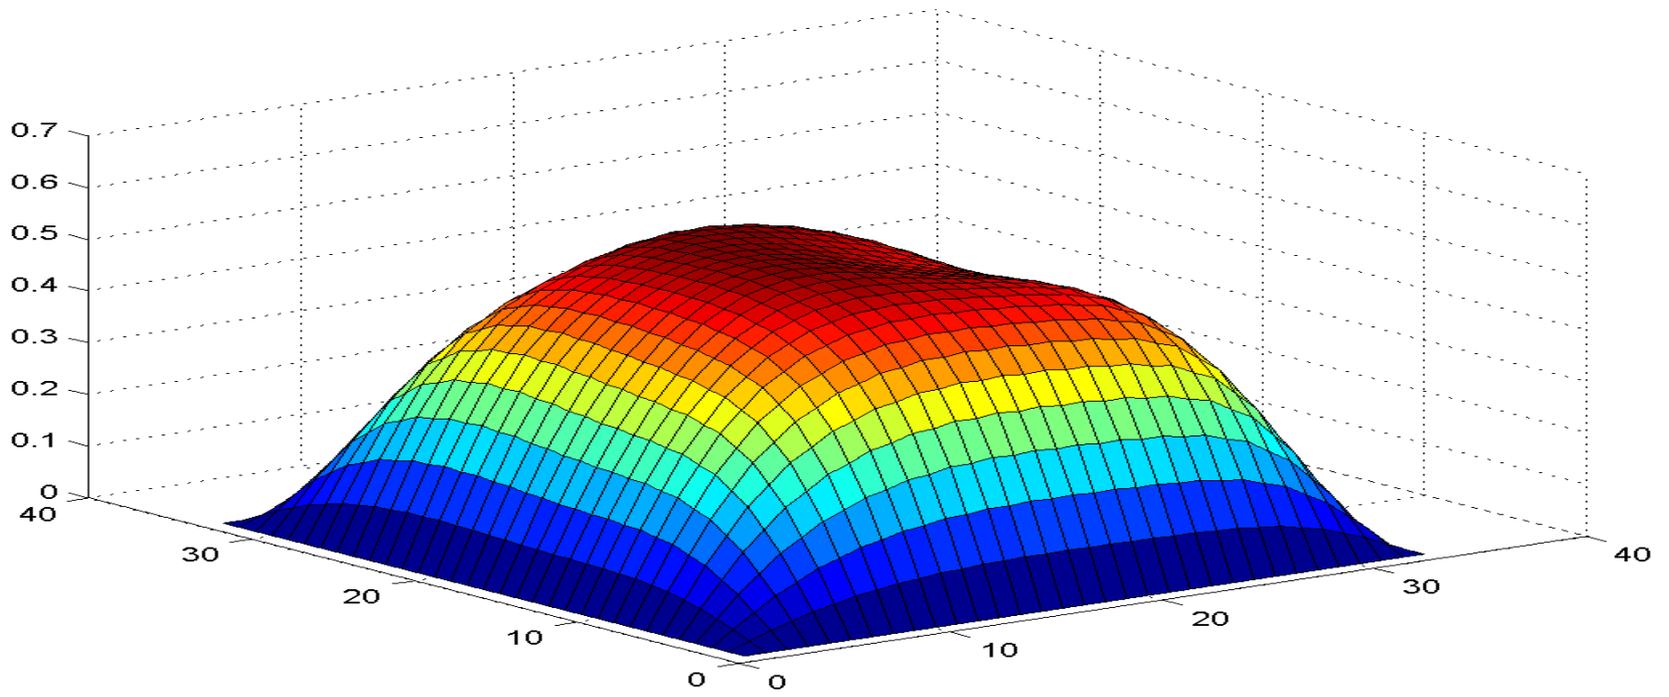


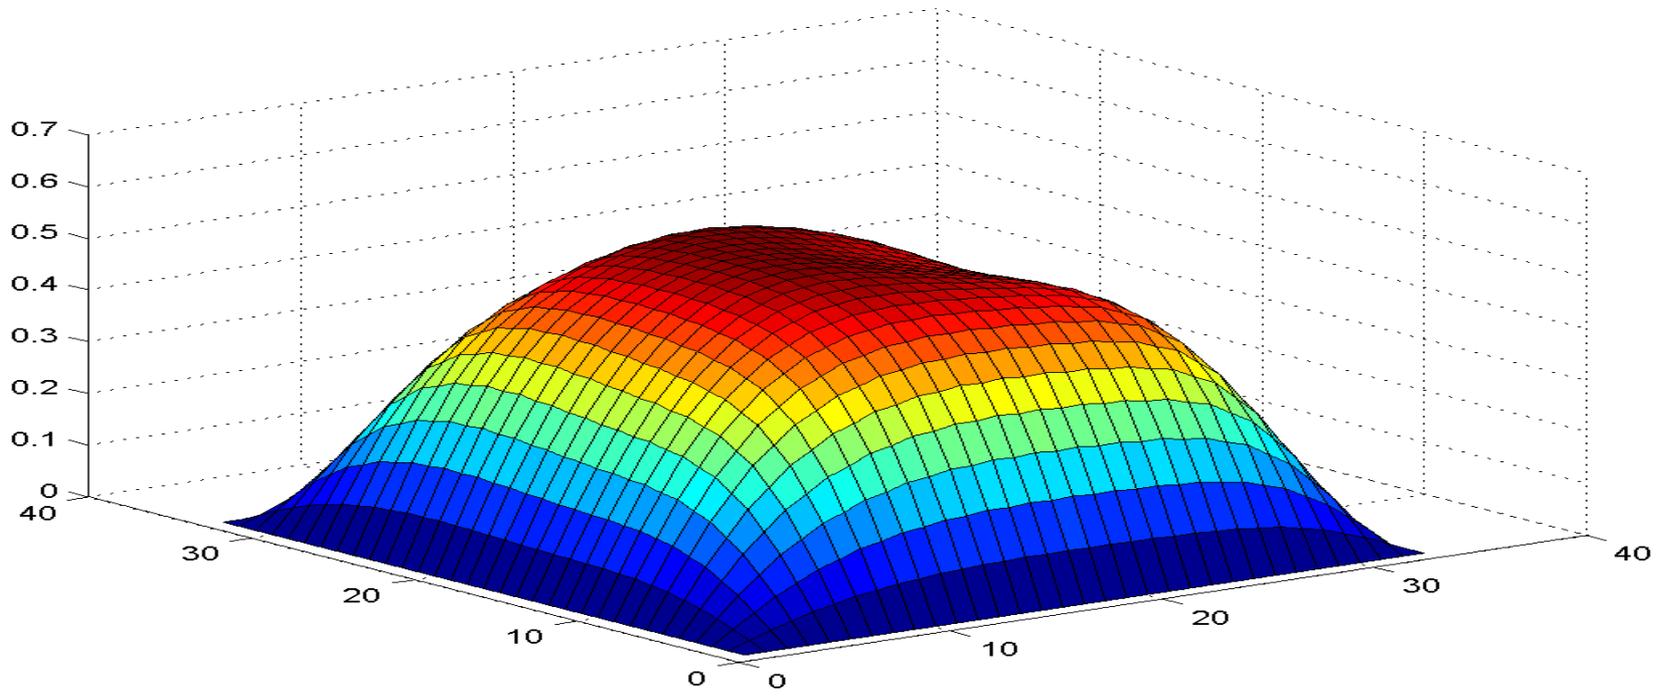


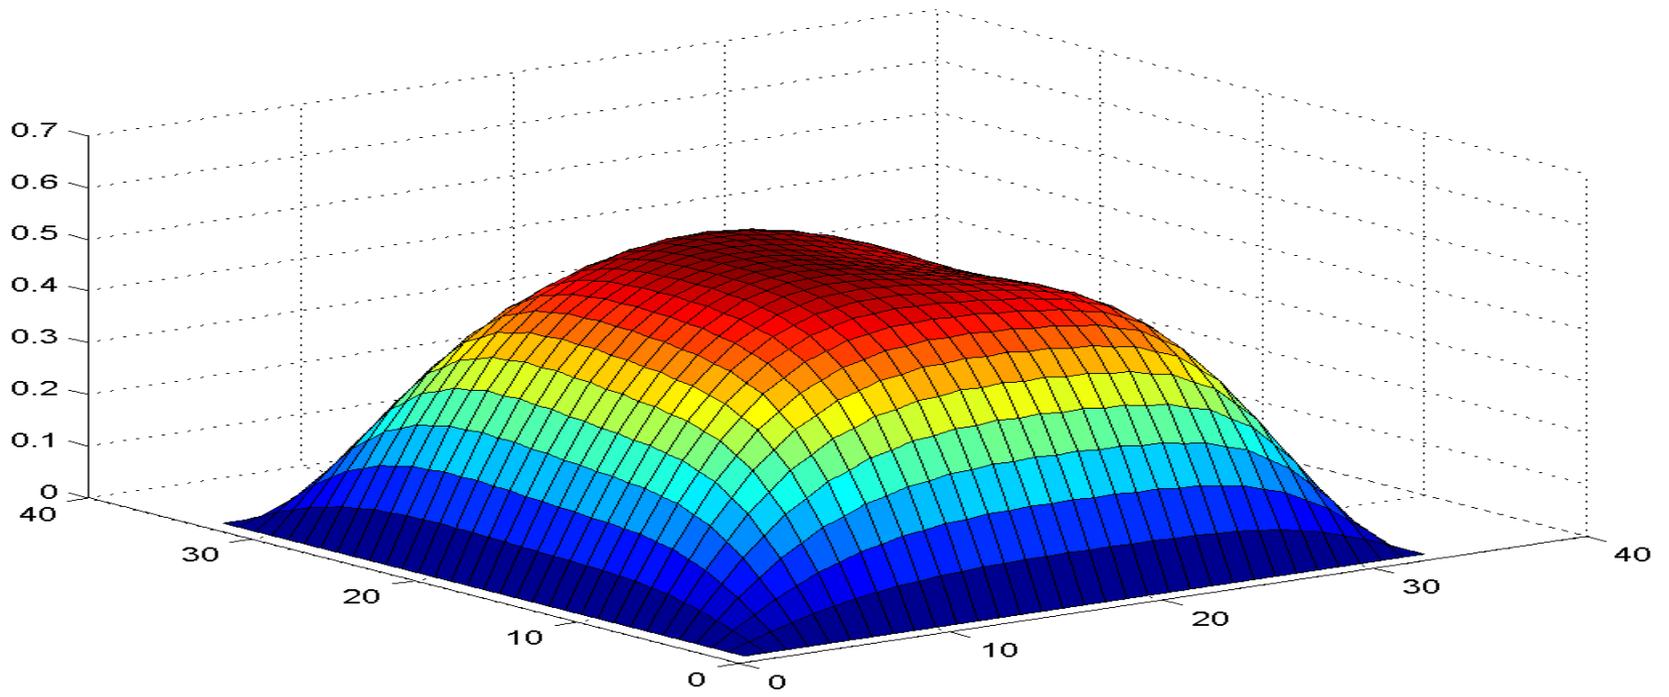


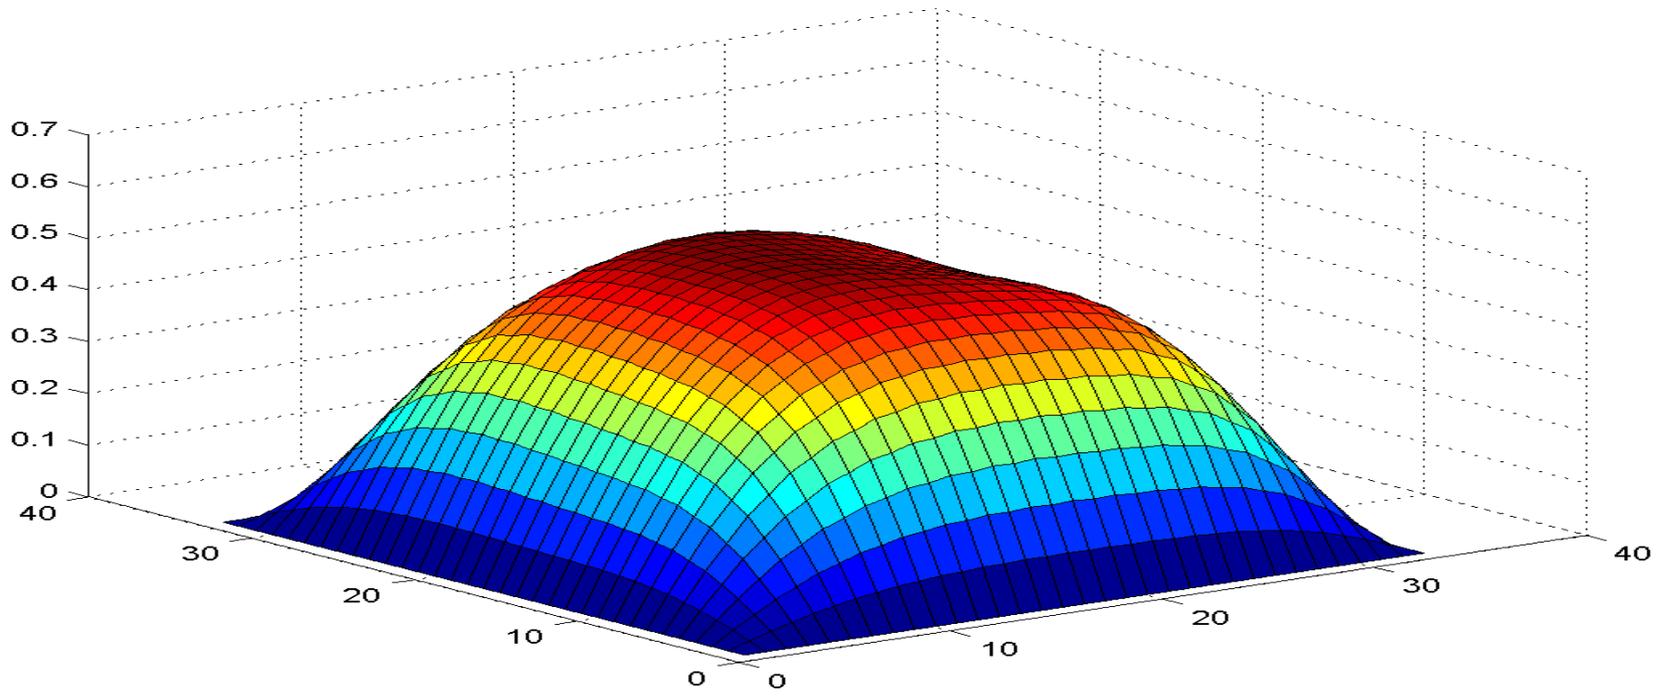












Key Observation re-worded: Relaxation cannot be generally efficient for reducing the error (i.e., the difference vector $\tilde{u}^h - u^h$). But relaxation may be extremely efficient for **smoothing the error relative to the grid.**

Practical conclusion:

1. A smooth error can be approximated well on a **coarser grid.**
2. A coarser grid implies less variables, hence **less computation.**
3. On the coarser grid the error is no longer as smooth relative to the grid, so **relaxation may once again be efficient.**

Grid-refinement algorithm

Define a sequence of progressively finer grids all covering the full domain. Then,

1. Define and solve the problem on the coarsest grid, say by relaxation.
2. Interpolate the solution to the next-finer grid. Apply several iterations of relaxation.
3. Interpolate the solution to the next-finer grid and continue in the same manner...

Does this method converge fast?

1D Model Problem Revisited

Fine-grid (h) difference equation:

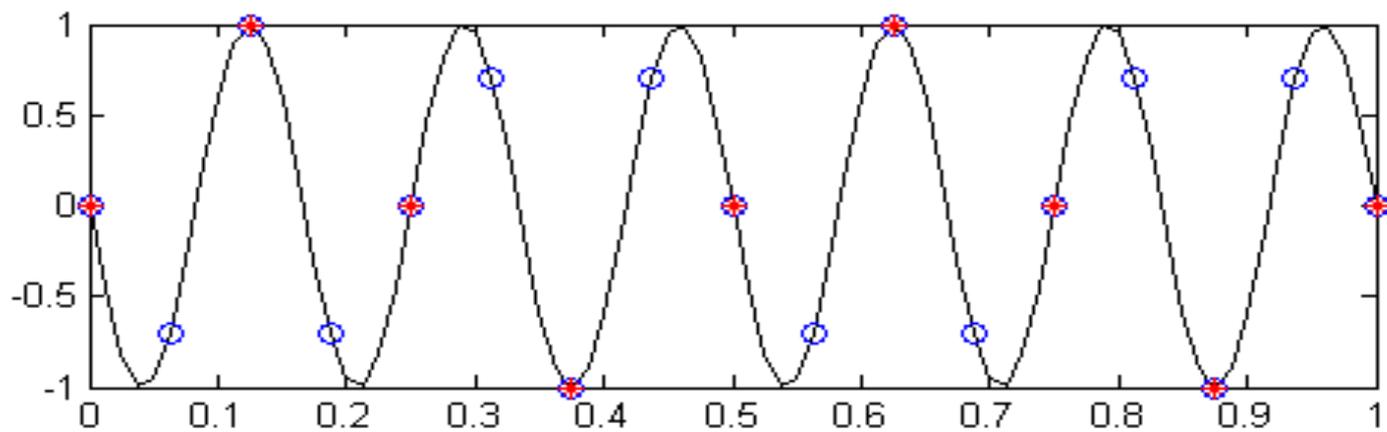
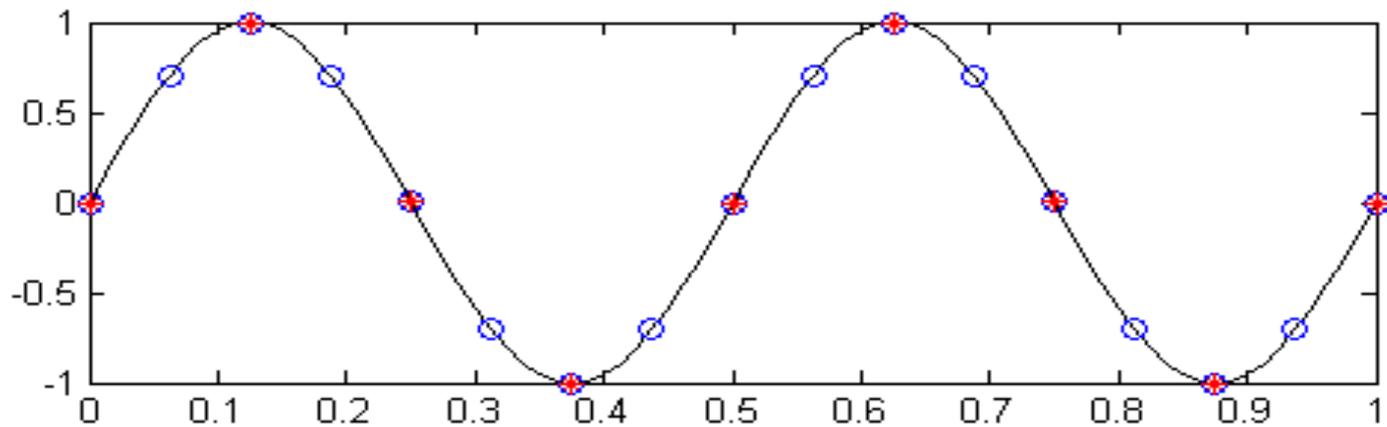
$$L^h u^h = \frac{u_{i+1}^h - 2u_i^h + u_{i-1}^h}{h^2} = f_i^h, \quad (6)$$
$$i = 1, \dots, N - 1,$$

$$u_0^h = u_0,$$

$$u_N^h = u_1.$$

The **eigenvectors** of L^h (like those of the Jacobi relaxation operation) are **Sine-function "waves"**:

$$\mathbf{v}^k = (\sin k\pi / N, \dots, \sin jk\pi / n, \dots, \sin(N-1)k\pi / N)^T \quad (7)$$



Aliasing

Smooth waves—with $k \ll N$ —have wavelengths large compared to h . Hence they can be approximated well on the coarse grids. But non-smooth eigenvectors alias with smooth components on the coarse grids.

Since the right-hand side, f^h , will generally have some non-smooth components, these will be “interpreted” as smooth components by the coarse grids, resulting in a smooth error.

Hence, when we interpolate a coarse-grid solution to the fine grid, we still have smooth errors in this solution. These cannot be corrected efficiently by relaxation.

Errors:

There is an important distinction here between the discretization error:

$$u - u^h,$$

and the algebraic error:

$$u^h - \tilde{u}^h,$$

Where \tilde{u}^h is our current approximation to u^h .

Note: Neither the solution, u^h , nor the discretization error are smoothed by relaxation, **only the algebraic error**. Hence, we formulate our problem in terms of this error.

Denote $v^h = u^h - \tilde{u}^h$.

Recall $L^h u^h = f^h$.

Subtract $L^h \tilde{u}^h$ from both sides, and use the linearity of L^h to obtain:

$$L^h v^h = f^h - L^h \tilde{u}^h \equiv r^h \quad (8)$$

As we have seen, we need to smooth the error v^h on the fine grid first, and only then solve the coarse-grid problem. Hence, we need two types of **integrid transfer operations**:

1. A **Restriction** (fine-to-coarse) operator: I_h^H .
2. A **Prolongation** (coarse-to-fine) operator: I_H^h .

For restriction we can often use simple injection, but full-weighted transfers are preferable.

For prolongation linear interpolation (bi-linear in **2D**) is simple and usually effective.

Two-grid Algorithm

- Relax several times on grid h , obtaining \tilde{u}^h with a smooth corresponding error.

- Calculate the residual: $r^h = f^h - L^h \tilde{u}^h$.

- Solve approximate error-equation on the coarse grid:

$$L^H v^H = f^H \equiv I_h^H r^h.$$

- Interpolate and add correction:

$$\tilde{u}^h \leftarrow \tilde{u}^h + I_H^h v^H.$$

- Relax again on grid h .

Multi-grid is obtained by recursion.

Multi-grid Cycle $V(v_1, v_2)$

Let u^{2h} approximate v^{2h} , u^{4h} approximate the error on grid $2h$, etc.

Relax on $L^h u^h = f^h$ v_1 times

Set $f^{2h} = I_h^{2h}(f^h - L^h u^h)$, $u^{2h} = 0$

Relax on $L^{2h} u^{2h} = f^{2h}$ v_1 times

Set $f^{4h} = I_{2h}^{4h}(f^{2h} - L^{2h} u^{2h})$, $u^{4h} = 0$

Relax on $L^{4h} u^{4h} = f^{4h}$ v_1 times

Set $f^{8h} = I_{4h}^{8h}(f^{4h} - L^{4h} u^{4h})$, $u^{8h} = 0$

...

Solve $L^{Mh} u^{Mh} = f^{Mh}$

...

Correct $u^{4h} \leftarrow u^{4h} + I_{8h}^{4h} u^{8h}$

Relax on $L^{4h} u^{4h} = f^{4h}$ v_2 times

Correct $u^{2h} \leftarrow u^{2h} + I_{4h}^{2h} u^{4h}$

Relax on $L^{2h} u^{2h} = f^{2h}$ v_2 times

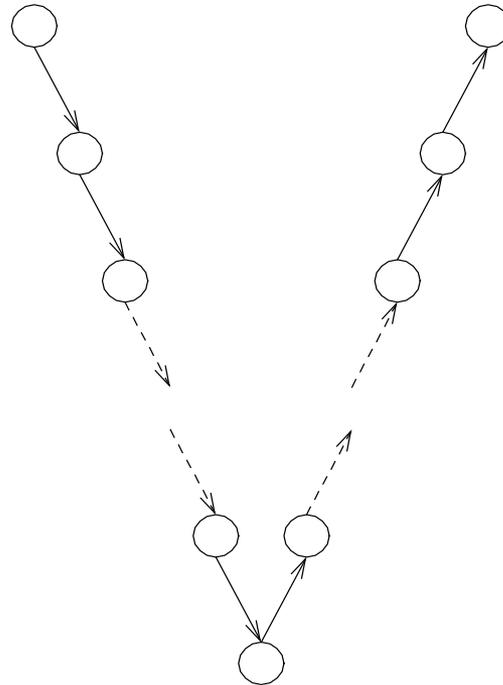
Correct $u^h \leftarrow u^h + I_{2h}^h u^{2h}$

Relax on $L^h u^h = f^h$ v_2 times

V cycle

Finest grid

Coarsest grid



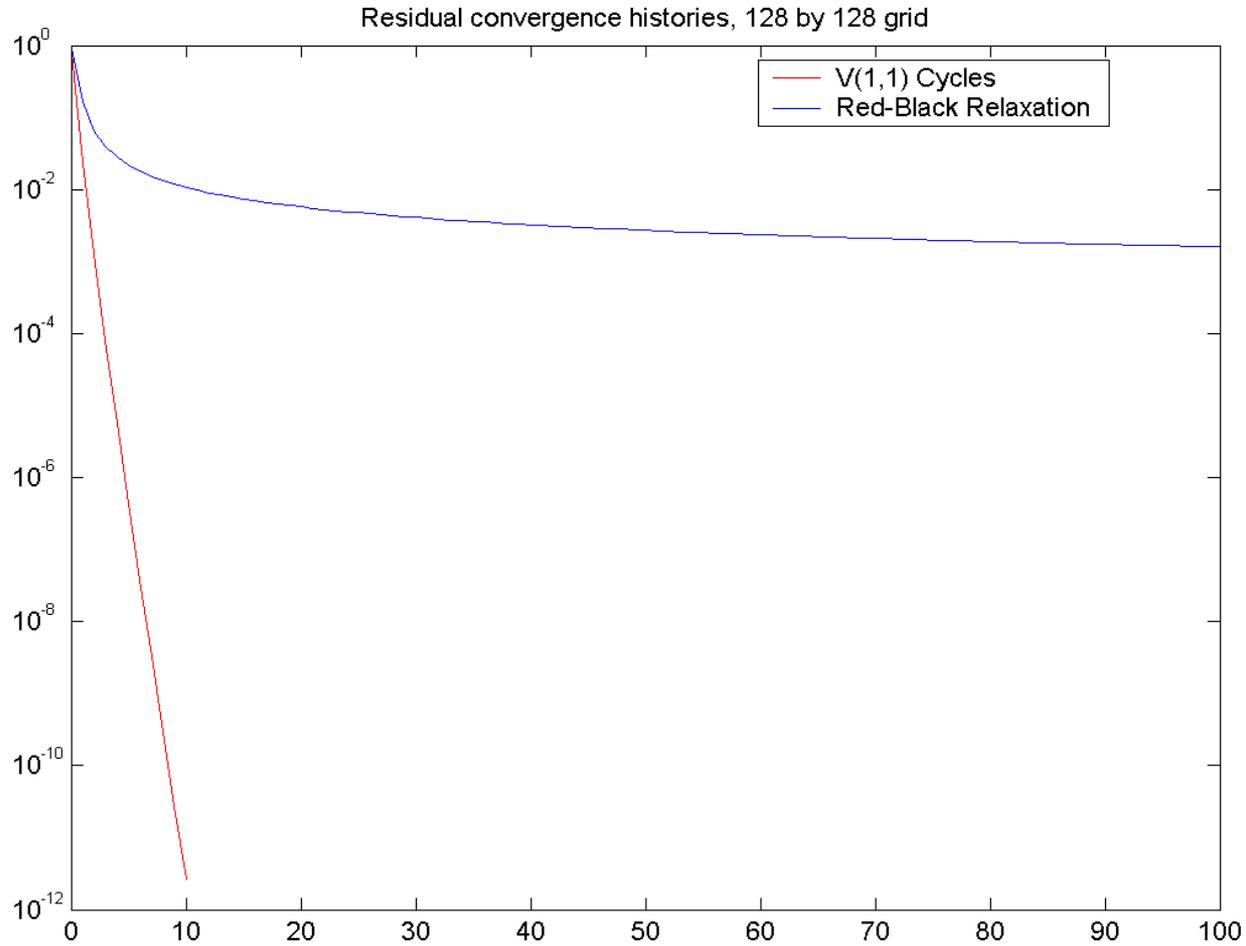
RELAXATION



RESTRICTION



PROLONGATION



Multigrid vs. Relaxation

Remarks:

1. Simple recursion yields a V cycle. More generally, we can choose a **cycle index γ** , and define a γ -cycle recursively as follows: Relax; transfer to next coarser grid; perform γ γ -cycles; interpolate and correct; Relax. (On the coarsest grid define the γ -cycle as an exact solution).
2. The best number of pre-relaxation + post-relaxation sweeps is normally **2** or **3**.
3. The boundary conditions for all coarse-grid problems is zero (because the coarse-grid variable is the error). **The initial guess for the coarse-grid solution must be zero.**

Algebrization of Multigrid

There are many problems for which multigrid is suitable in principle but cannot be applied in a straightforward way. For example,

1. Unstructured grids and complex geometries
2. Non-PDE applications

Such situations require **algebraic multigrid methods**.

The multigrid components can be expressed as **matrices**. Consider, for example, the **1D** model problem using linear interpolation and full-weighted residual transfers.

Given the fine-grid matrix, L^h , and the prolongation and restriction matrices, I_H^h , and I_h^H , how should we define the coarse-grid matrix, L^H ?

The coarse grid should be able to **correct smooth errors**. We use the following (algebraic) definition of smoothness: An error v_{before}^h is smooth if it is **in the range of interpolation**. That is, if there exists some coarse-grid function, w^H , such that

$$v_{before}^h = I_H^h w^H .$$

The error after the coarse-grid correction is given by

$$\mathbf{v}_{after}^h = \mathbf{C}^h \mathbf{v}_{before}^h$$

where

$$\mathbf{C}^h = \mathbf{I}^h - \mathbf{I}_H^h \left(\mathbf{L}^H \right)^{-1} \mathbf{I}_h^H \mathbf{L}^h.$$

Plugging in our smooth error we obtain:

$$\begin{aligned}
\mathbf{v}_{after}^h &= \mathbf{C}^h \mathbf{v}_{before}^h \\
&= \left[\mathbf{I}^h - \mathbf{I}_H^h \left(\mathbf{L}^H \right)^{-1} \mathbf{I}_h^H \mathbf{L}^h \right] \mathbf{I}_H^h \mathbf{w}^H \\
&= \left[\mathbf{I}_H^h - \mathbf{I}_H^h \left(\mathbf{L}^H \right)^{-1} \mathbf{I}_h^H \mathbf{L}^h \mathbf{I}_H^h \right] \mathbf{w}^H \\
&= \mathbf{I}_H^h \left[\mathbf{I}^H - \left(\mathbf{L}^H \right)^{-1} \mathbf{I}_h^H \mathbf{L}^h \mathbf{I}_H^h \right] \mathbf{w}^H .
\end{aligned}$$

In order to annihilate the error we must choose the Petrov-Galerkin coarse-grid operator:

$$\mathbf{L}^H = \mathbf{I}_h^H \mathbf{L}^h \mathbf{I}_H^h .$$

For symmetric problems especially, the preferred choice for the restriction is the transpose of the prolongation. Along with the Galerkin coarse-grid operator this yields so-called variational coarsening, which arises naturally in finite-element formulations.

It remains only to define the prolongation (and, implicitly, the set of variables which defines the coarse grid). The prolongation operator should produce good approximate fine-grid values from given coarse-grid values. Therefore, I_H^h needs to be determined using L^h . When used with appropriate coarse grids, such methods yield fast and robust algebraic solvers.

For tridiagonal matrices in **1D** the different algebraic methods become the same: an exact multigrid solver that is **equivalent to cyclic reduction**

If the fine-grid equations are

$$a_i u_{i-1} + b_i u_i + c_i u_{i+1} = f_i,$$

$I = 1, \dots, n-1$, with $a_1 = c_{n-1} \equiv 0$, we choose the prolongation matrix to be

Furthermore, we let $I_h^H = (I_H^h)^T$ and employ Galerkin coarsening. For smoothing we use half-Red-Black relaxation. That is, before restricting residuals we relax only on odd-indexed gridpoints, and after the coarse-grid correction only on even-indexed points.

Theorem: the two-level cycle is an exact solver. Furthermore, the coarse-grid equations are tridiagonal. Hence, the multigrid cycle is an exact solver.

Algebraic Multigrid (AMG)

Introduced by Brandt et al. (1983) and developed by Ruge and Stueben.

AMG takes the algebraization of multigrid to the limit. Here, a relaxation method is chosen (usually, point Gauss-Seidel), and then the coarse-grid variables are chosen by a heuristic graph algorithm such that **each fine-grid variable depends strongly on one or more coarse-grid variable** (i.e., with relatively large coefficient).

AMG enables us to handle **unstructured and non-PDE problems**.

An Abstract View of Algebraic Multigrid Methods

Consider the linear system

$$Au = f.$$

Suppose we partition the variables, u_i , into F variables and C variables, and permute the equations and variables to produce the following partitioning of the system:

$$Au = \begin{pmatrix} A_{FF} & A_{FC} \\ A_{CF} & A_{CC} \end{pmatrix} \begin{pmatrix} u_F \\ u_C \end{pmatrix} = \begin{pmatrix} f_F \\ f_C \end{pmatrix}.$$

An Abstract View of AMG

Given an approximate solution, \tilde{u} , define the error as

$$v = u - \tilde{u}.$$

Then, the partitioned equation for the error is

$$Av = \begin{pmatrix} A_{FF} & A_{FC} \\ A_{CF} & A_{CC} \end{pmatrix} \begin{pmatrix} v_F \\ v_C \end{pmatrix} = \begin{pmatrix} r_F \\ r_C \end{pmatrix},$$

where

$$r_F = f_F - A_{FF}\tilde{u}_F - A_{FC}\tilde{u}_C,$$

$$r_C = f_C - A_{CF}\tilde{u}_F - A_{CC}\tilde{u}_C.$$

An Abstract View of AMG

The upper block yields

$$\begin{aligned}A_{FF} v_F &= r_F - A_{FC} v_C, \\ \Rightarrow v_F &= A_{FF}^{-1} (r_F - A_{FC} v_C).\end{aligned}$$

Plugging this into the lower block yields

$$\begin{aligned}A_{CF} A_{FF}^{-1} (r_F - A_{FC} v_C) + A_{CC} v_C &= r_C, \\ \Rightarrow (A_{CC} - A_{CF} A_{FF}^{-1} A_{FC}) v_C &= r_C - A_{CF} A_{FF}^{-1} r_F.\end{aligned}$$

An Abstract View of AMG

Conclusion: the “ideal” prolongation and restriction are

$$P = \begin{pmatrix} -A_{FF}^{-1} A_{FC} \\ I \end{pmatrix}, \quad R = \begin{pmatrix} -A_{CF} A_{FF}^{-1} & I \end{pmatrix},$$

with the coarse-grid operator given by

$$A_C = RAP = A_{CC} - A_{CF} A_{FF}^{-1} A_{FC}.$$

An Abstract View of AMG

In particular, it is straightforward to verify that the two-level solution is exact in this case, provided that either a pre-relaxation or a post-relaxation eliminates r_F .

(If this is done by post-relaxation, only u_F should be relaxed.)

The problem: A_{FF}^{-1} is not sparse, and therefore, neither are P and R . Therefore, we generally look for good sparse approximations.

One exception is tri-diagonal systems, where A_{FF} is diagonal. In this case the multigrid V -cycle with the appropriate prolongation and restriction, and with relaxation only on u_F is an exact solver, equivalent to total reduction.