Monotone and second order consistent schemes for the Pucci and Monge-Ampere equations

Joseph Frédéric Bonnans, Guillaume Bonnet, Jean-Marie Mirebeau

Monotone and second order consistent schemes for the Pucci and Monge-Ampere equations

Joseph Frédéric Bonnans[*], Guillaume Bonnet[†], Jean-Marie Mirebeau[‡]

November 27, 2019

Abstract

We introduce a new strategy for the design of second-order accurate discretizations of non-linear second order operators of Bellman type, which preserves degenerate ellipticity. The approach relies on Selling's formula, a tool from lattice geometry, and is applied to the Pucci and Monge-Ampere equations, discretized on a two dimensional cartesian grid. In the case of the Monge-Ampere equation, our work is related to both the stable formulation [FJ17] and the second order accurate scheme [BCM16]. Numerical experiments illustrate the robustness and the accuracy of the method.

1 Introduction

Degenerate Ellipticity (DE) is a property of a class of partial differential operators, often non-linear and of order at most two. When satisfied, it implies a generalized comparison principle, from which can be deduced the existence, uniqueness and stability of a viscosity solution to the Partial Differential Equation (PDE), under mild additional assumptions [CIL92]. Discrete degenerate ellipticity is the corresponding property for numerical schemes, see Definition 2.3, which has similarly strong implications and often turn the convergence the analysis of solutions into a simple verification [Obe06]. It is therefore appealing to design PDE discretizations preserving the DE property, yet a strong limitation of the current approaches [BS91, Obe08, FJ17] is their low consistency order, usually below one. Filtered schemes [FO13] attempt to mitigate this issue by combining a DE scheme of low consistency order with a non-DE scheme of high consistency order, but their use requires careful parameter tuning, and theoretical results are lacking regarding their effective accuracy.

In this paper, we propose a new approach to develop second order accurate DE schemes, which is the highest achievable consistency order [Obe06], on two

[*]INRIA Saclay, Ecole Polytechnique, CMAP 91128 Palaiseau

[†]U. Paris-Sud, U. Paris-Saclay, 91405, Orsay, France

[‡]U. Paris-Sud, CNRS, U. Paris-Saclay, 91405, Orsay

This work was partly supported by ANR research grant MAGA, ANR-16-CE40-0014

dimensional Cartesian grids. The operator must be given in Bellman form as follows

$$\Lambda u(x) = \sup_{\alpha \in \mathcal{A}} a_\alpha + b_\alpha u(x) - \mathrm{Tr}(D_\alpha \nabla^2 u(x)), \tag{1}$$

where $\mathcal{A}$ is an abstract set of parameters, and the coefficients $a_\alpha \in \mathbb{R}$, $b_\alpha \geq 0$, and symmetric matrix $D_\alpha \succ 0$ may additionally depend on the position x. A specific feature of our approach, that is tied to the structure of the addressed problems, is that the parameter space $\mathcal{A}$ is not discretized. Special handling is required when the condition number of the matrices $(D_\alpha)_{\alpha \in \mathcal{A}}$ is unbounded, see §2.3. We apply this approach to the Pucci and Monge-Ampere equations in dimension two:

$$\lambda_{\min}(\nabla^2 u(x)) + \mu\, \lambda_{\max}(\nabla^2 u(x)) = f_{\mathrm{P}}(x), \qquad \det(\nabla^2 u(x)) = f^0_{\mathrm{MA}}(x), \tag{2}$$

both having Dirichlet boundary conditions, where $\lambda_{\min}$ and $\lambda_{\max}$ denote the smallest and largest eigenvalue of a symmetric matrix, and where $\mu > 0$ and $f^0_{\mathrm{MA}} \geq 0$. For simplicity, we assume that $\mu \leq 1$. These PDEs admit the following Bellman form, taken from [FJ17] in the Monge-Ampere case:

$$\max_{\theta \in [0,\pi]} -\mathrm{Tr}(D(\theta,\mu)\nabla^2 u(x)) = -f_{\mathrm{P}}(x) \qquad \text{where } D(\theta,\mu) := R_\theta \begin{pmatrix} 1 & 0 \\ 0 & \mu \end{pmatrix} R_\theta^T, \tag{3}$$

$$\sup_{\substack{\mathrm{Tr}(D)=1 \\ D \succ 0}} f_{\mathrm{MA}}(x)\sqrt{\det(D)} - \mathrm{Tr}(D\nabla^2 u(x)) = 0 \qquad \text{where } f_{\mathrm{MA}} := 2\sqrt{f^0_{\mathrm{MA}}}. \tag{4}$$

2 Discretization

We rely on a tool from algorithmic lattice geometry, known as Selling's formula §2.1, which is particularly adequate for discretizing degenerate elliptic PDEs on cartesian grids of dimension two [BOZ04] or three [Mir18, Mir17, FM14]. This technique is applied to the Pucci operator in §2.2, and to the Monge-Ampere equation in §2.3. Throughout this section $\Omega \subset \mathbb{R}^2$ denotes a bounded domain, and $h > 0$ a grid scale. Define

$$\Omega_h := h\mathbb{Z}^2 \cap \Omega, \qquad \Delta^e_h u(x) := \frac{u(x+he) - 2u(x) + u(x-he)}{h^2}, \tag{5}$$

the discrete domain and the second order finite difference of a map $u : \Omega_h \cup \partial\Omega \to \mathbb{R}$ at $x \in \Omega_h$ in the direction $e \in \mathbb{Z}^2$. When x is close to $\partial\Omega$ the latter formula becomes

$$\Delta^e_h u(x) := \frac{2}{h_+ + h_-}\Big(\frac{u(x+h_+e) - u(x)}{h_+} + \frac{u(x-h_-e) - u(x)}{h_-}\Big), \tag{6}$$

where $h_\pm > 0$ is the least value such that $x \pm h_\pm e \in \Omega_h \cup \partial\Omega$. Note that (5, right) is a second order consistent approximation of $\langle e, \nabla^2 u(x) e\rangle$, whereas (6) is only first order consistent.

2.1 Selling's formula

Selling's decomposition of an element of the set S_2^{++} of symmetric positive definite 2×2 matrices, see Proposition 2.2, can be regarded as a variant of the eigenvector/eigenvalue decomposition, but whose vectors have *integer entries*. We rely on it to discretize non-divergence form linear (8) and non-linear (9) operators, in a manner that achieves discrete degenerate ellipticity, see Definition 2.3.

Definition 2.1. *A superbase of $\mathbb{Z}^2$ is a triplet $(e_0, e_1, e_2) \in (\mathbb{Z}^2)^3$ such that $e_0 + e_1 + e_2 = 0$ and $|\det(e_1, e_2)| = 1$. It is D-obtuse, where $D \in S_2^{++}$, iff $\langle e_i, De_j \rangle \leq 0$ for all $i \neq j$.*

Proposition 2.2 (Selling [Sel74])**.** *For each $D \in S_2^{++}$ there exists a D-obtuse superbase (e_0, e_1, e_2) of $\mathbb{Z}^2$, which can be obtained from Selling's algorithm. Furthermore one has Selling's formula*

$$D = \sum_{0 \leq i \leq 2} \rho_i v_i v_i^\top \qquad \text{with } \rho_i := -\langle e_{i-1}, De_{i+1} \rangle \geq 0, \quad v_i := e_i^\perp \in \mathbb{Z}^2, \tag{7}$$

where $e^\perp := (-b, a)^\top$ if $e = (a, b)^\top \in \mathbb{R}^2$. The set $\{(\rho_i, \pm v_i);\, 0 \leq i \leq 2,\, \rho_i > 0\}$ is uniquely determined.

Based on this formula, on can consider the following finite differences approximation, which is second-order accurate far enough from $\partial\Omega$:

$$\Delta_h^D u(x) := \sum_{0 \leq i \leq 2} \rho_i \Delta_h^{v_i} u(x), \qquad \Delta_h^D u(x) = \mathrm{Tr}(D\nabla^2 u(x)) + \mathcal{O}(h^2). \tag{8}$$

Only first order consistency is achieved when x is close to $\partial\Omega$, due to the use of (6).

Definition 2.3 (Discrete degenerate ellipticity [Obe06])**.** *A numerical scheme on a finite set X is a map $F : U \to U$, where $U := \mathbb{R}^X$ is the set of functions from X to $\mathbb{R}$, of the form:*

$$Fu(x) = F(x, u(x), (u(y) - u(x))_{y \neq x}),$$

for all $u \in U$, $x \in X$. It is Discrete Degenerate Elliptic (DDE) iff F is non-decreasing w.r.t. the second argument $u(x)$, and non-increasing w.r.t. each $u(y) - u(x)$, $y \neq x$.

The numerical scheme $-\Delta_h^D$ is DDE on Ω_h, thanks to the non-negativity of the weights $(\rho_i)_{0 \leq i \leq 2}$, and to the finite differences expression (5, right) and (6), where u is extended to $\partial\Omega$ with the provided Dirichlet boundary values. On this basis we obtain a DDE discretization of nonlinear second order operators in Bellman form (1), which is second-order accurate if x is far enough from $\partial\Omega$:

$$\Lambda_h u(x) := \sup_{\alpha \in \mathcal{A}} a_\alpha + b_\alpha u(x) - \Delta_h^{D_\alpha} u(x), \qquad \Lambda_h u(x) = \Lambda u(x) + \mathcal{O}(h^2). \tag{9}$$

Again, only first order consistency is achieved when x is close to $\partial\Omega$, depending on the width of the discretization stencil (7) of D_α, $\alpha \in \mathcal{A}$. As shown below, the supremum in (9, left) can be computed analytically in closed form, for the Pucci and Monge-Ampere PDEs, so that the numerical scheme Λ_h is explicit in terms of u.

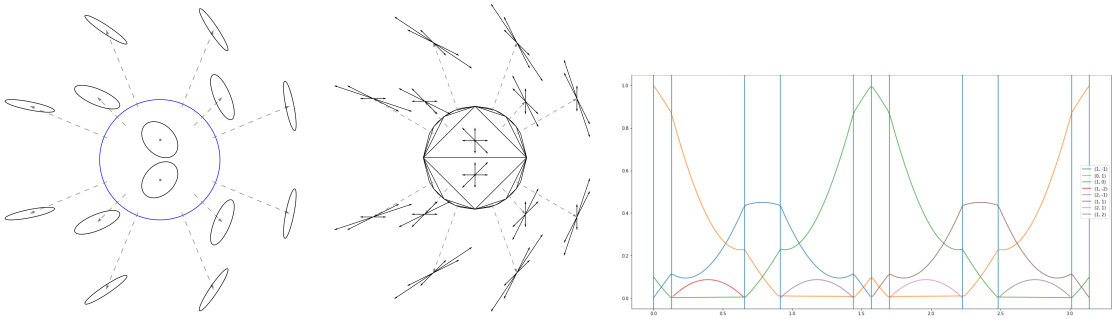

Figure 1: (Left) Ellipsoid $\{v \in \mathbb{R}^2;\, v^T \mathbf{D}(z) v = 1\}$ for some points z of the unit disc, see (10, right). Anisotropy degenerates as z moves toward the unit circle, shown blue. (Center) $\mathbf{D}(z)$-obtuse superbase, and opposites, for the same points z. This superbase is piecewise constant on an infinite triangulation of the unit disk [Sch09], shown black. (Right) Coefficients of the decomposition of $D(\theta, \mu)$ for $\theta \in [0, \pi]$ and $\mu = 0.1$.

2.2 The Pucci operator

The Bellman form of the Pucci operator (3) involves a family of matrices $D(\theta, \mu)$, parametrized by the inverse $0 < \mu \leq 1$ of their condition number and an angle $0 \leq \theta \leq \pi$. As a starter, we rewrite those in the following normal form

$$D(\theta, \mu) = \frac{1+\mu}{2} \mathbf{D}\Big(\frac{1-\mu}{1+\mu}\, \mathrm{n}(2\theta)\Big), \qquad \mathbf{D}(z) := \begin{pmatrix} 1+z_0 & z_1 \\ z_1 & 1-z_0 \end{pmatrix}, \tag{10}$$

where $\mathrm{n}(\varphi) = (\cos\varphi, \sin\varphi)$ left, and $z = (z_0, z_1)$ right. The anisotropy defined by $\mathbf{D}(z)$, $\|z\| < 1$, is illustrated on Figure 1 (left). A $\mathbf{D}(z)$-obtuse superbase is known explicitly for each $\|z\| < 1$, depending on a suitable triangulation $\mathcal{T}$ of the unit disc, see Figure 1 (center). Note that the argument of $\mathbf{D}$ in (10, left) describes a circle of fixed radius $\frac{1-\mu}{1+\mu}$, hence intersects finitely may elements of $\mathcal{T}$. Thus one can find $0 = \theta_0 < \cdots < \theta_N = \pi$ such that on each interval $[\theta_n, \theta_{n+1}]$ the superbase (e_0^n, e_1^n, e_2^n) is $D(\theta, \mu)$-obtuse and the coefficients in (7) take the form

$$\rho_i(\theta) = -\langle e_{i-1}^n, D(\theta, \mu) e_{i+1}^n \rangle = \alpha_i^n + \beta_i^n \cos(2\theta) + \gamma_i^n \sin(2\theta), \tag{11}$$

for suitable constants $\alpha_i^n, \beta_i^n, \gamma_i^n \in \mathbb{R}$, $0 \leq i \leq 2$, $0 \leq n < N$. See Figure 1 (right).

By linearity of (8) one also has $\Delta_h^{D(\theta,\mu)} u(x) = \alpha^n + \beta^n \cos(2\theta) + \gamma^n \sin(2\theta)$ for all $\theta \in [\theta_n, \theta_{n+1}]$, whose coefficients $\alpha^n, \beta^n, \gamma^n$ depend on ρ, h, u and x. Therefore, evaluating the discretized Bellman operator (9) associated with the Pucci equation (3) at a point $x \in \Omega_h$ amounts to solving a small number N of optimization problems, whose value is explicit:

$$\begin{aligned} &\max_{\varphi \in [\varphi_*, \varphi^*]} \alpha + \beta \cos\varphi + \gamma \sin\varphi \\ &= \begin{cases} \alpha + \sqrt{\beta^2 + \gamma^2} & \text{if } \arg(\beta + i\gamma) \in]\varphi_*, \varphi^*[, \\ \alpha + \max\{\beta\cos\varphi_* + \gamma \sin\varphi_*,\ \beta\cos\varphi^* + \gamma\sin\varphi^*\} & \text{else,} \end{cases} \end{aligned}$$

where $\arg(\omega)$ denotes the argument of $\omega \in \mathbb{C}$, taken in $[0, 2\pi[$. Choose $\varphi_* = 2\theta_n$, $\varphi^* = 2\theta_{n+1}$, $\alpha = \alpha^n$, $\beta = \beta^n$, and $\gamma = \gamma^n$, and take the largest value among $0 \leq n < N$. Numerical results are presented in Section 3.

2.3 The Monge-Ampere equation

For discretizing the Monge-Ampere equation, we use Selling's formula (7) to *generate* all $D \in S_2^{++}$, rather than to *decompose* them. Given a superbase $B = (e_0, e_1, e_2)$ and non-negative coefficients (ρ_0, ρ_1, ρ_2) consider the positive semi-definite tensor

$$D = \sum_{0\leq i\leq 2} \rho_i e_i e_i^\top, \text{ obeying } \operatorname{Tr}(D) = \sum_{0\leq i\leq 2} \rho_i \|e_i\|^2, \det(D) = \sum_{0\leq i\leq 2} \rho_i \rho_{i+1}. \quad (12)$$

By Selling's formula (7), any $D \in S_2^{++}$ is of this form. (We implicitly used that the vectors (v_0, v_1, v_2) in Proposition 2.2 are also a superbase.) The formula (12, right) is obtained by computing the determinant of D in the frame (e_1, e_2). If the superbase B is fixed, then the discretization (9) of Bellman's form of the Monge-Ampere equation (4) writes as:

$$\Lambda_h^B u(x) := \max_{\rho\in\mathbb{R}^3_+} f_{\mathrm{MA}}(x) \sqrt{\sum_{0\leq i\leq 2} \rho_i \rho_{i+1}} - \sum_{0\leq i\leq 2} \rho_i \Delta_h^{e_i} u(x), \qquad \text{s.t.} \sum_{0\leq i\leq 2} \rho_i \|e_i\|^2 = 1.$$

Note that the linear constraints define a compact and convex set, actually a triangle $T_B \subset \mathbb{R}^3_+$ with vertices $(\|e_0\|^{-2}, 0, 0)$ $(0, \|e_1\|^{-2}, 0)$ $(0, 0, \|e_2\|^{-2})$. The optimized functional is concave, since $\sqrt{\det}$ is concave on S_2^{++}, hence it suffices to find a stationary point. Interestingly, a problem with the exact same algebraic form arises in discretizing the eikonal equation, see (4) in [Mir14]. The solution ρ_{opt} has a closed form, easily deduced from the optimality conditions. In more detail, if $\rho_{\mathrm{opt}} \in \mathbb{R}^3_+$ lies in the interior of T_B, then it is positively proportional to a vector defined in terms of the solution $t \in \mathbb{R}$ to a univariate quadratic equation (13, right):

$$\rho_{\mathrm{opt}} \propto Q^{-1}(t\omega + \delta) \qquad \text{where } (t\omega+\delta)^\top Q^{-1}(t\omega+\delta) = f_{\mathrm{MA}}(x)^2. \quad (13)$$

We denoted $\omega = (\|e_i\|^2)_{0\leq i\leq 2}$, $\delta = (\Delta_h^{e_i} u(x))_{0\leq i\leq 2}$, and $Q = \frac{1}{2}\begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}$ the matrix of the determinant (12, right) seen as a quadratic form. On the other hand, ρ_{opt} is easily determined if it lies on the boundary of T_B, which is the union of three segments.

Eventually, we fix a finite set $\mathcal{B}$ of superbases, and define the numerical scheme

$$\Lambda_h u(x) := \max_{B\in\mathcal{B}} \Lambda_h^B u(x).$$

This scheme is second order consistent (only first order near the boundary) provided the solution is smooth, uniformly convex, and the condition number of $\nabla^2 u(x)$ is bounded by a constant depending on $\mathcal{B}$. Extending the set $\mathcal{B}$ of superbases indeed allows to represent more matrices in the form (12), see Figure 1, and improves the consistency properties. On the opposite spectrum, the scheme degenerates if $f_{\mathrm{MA}}(x) = 0$ into a standard [Obe13, Obe08] discretization of $-\lambda_{\min}(\nabla^2 u(x))$:

$$\Lambda_h u(x) = \max_{B\in\mathcal{B}} \max_{0\leq i\leq 2} -\frac{\Delta_h^{e_i} u(x)}{\|e_i\|^2}.$$

3 Numerical experiments

We present numerical results for the discretization schemes described in this paper, which illustrate the qualitative behavior of the solutions in various examples, and

validate the scheme robustness and accuracy on synthetic problems with known solutions. Some of the domains considered are neither smooth nor convex, and the considered synthetic solutions range from smooth to singular.

The numerical schemes are implemented as described in the previous section, and a Newton method is used to solve the resulting coupled systems of non-linear equations. In practice, convergence to machine precision is achieved in a dozen of iterations, without damping, from an arbitrary guess. We take advantage of the envelope theorem to efficiently differentiate the numerical scheme (9) w.r.t. the unknown and assemble its Jacobian matrix. An open source Python® implementation is available on the third author's webpage: `Github.com/Mirebeau/AdaptiveGridDiscretizations`.

3.1 Pucci equation

We illustrate on Figure 2 the transition of the Pucci equation from a strongly elliptic Laplacian-like PDE to a combinatorial-type convex-envelope problem, as the parameter μ takes values $1/4$ and $1/400$ ($\alpha := \mu$ in image titles). The chosen domain is non-smooth and non-simply connected : $\Omega := U \backslash U'$ where $U := B(0,1) \cup (]0,1[\times]-1,1[)$ and $U' := kR_\theta(U)$ is its image under a scaling ($k = 0.4$) and a rotation ($\theta = \pi/3$). The boundary condition is 1 on ∂U, and 0 on $\partial U'$.

On figure 3, we reconstruct some known synthetic solutions from their image by the Pucci operator, with parameter $\mu = 0.2$, and their trace on the boundary. The examples are taken from the literature [FJ17, FO13], and the reconstruction errors are provided in the L^1 and L^∞ norm.

- (Smooth example [FJ17]) $u(x) = (x^2 + y^2)^2$ on $\Omega = B(0,1) \cup]0,1[^2$
- (C^1 example [FO13]) $u(x) = \max\{0, \|x - x_0\|^2 - 0.2)$ on $\Omega =]0,1[^2$.
- (Singular example [FO13]) $u(x) = \sqrt{2 - \|x\|^2}$ on $\Omega =]0,1[^2$.

Empirically, the L^1 numerical error behaves like $\mathcal{O}(h^2)$, where h is the grid scale (inverse of resolution in images). The L^∞ error behaves like $\mathcal{O}(h^2)$ in the smooth and C^1 examples, but decays more slowly for the singular solution. *Note: we rotated the Cartesian discretization grid by $\pi/3$ in these experiments, since otherwise the perfect alignment of the domain boundary with the coordinate axes gives an unfair advantage to grid based methods (like ours).*

3.2 Monge-Ampere equation

Our discretization of the Monge-Ampere equation requires to choose a finite set $\mathcal{B}$ of superbases. We consider three possibilities, featuring 2, 6 and 10 superbases, whose support defines the stencil of the numerical scheme, see Figure 4 .

The impact of the stencil choice is clearly visible when reconstructing the piecewise linear function $u(x) = |x + y/\sqrt{10}|$, by solving the degenerate Monge-Ampere equation $\det(\nabla^2 u) = 0$ (equivalent to the convex envelope problem) with Dirichlet boundary conditions, see Figure 4. In addition, the convergence curves on Figure 5, on the same test cases as for the Pucci operator, show that numerical error eventually stalls if an excessively small stencil is used. Empirical error is $\mathcal{O}(h^2)$ otherwise.

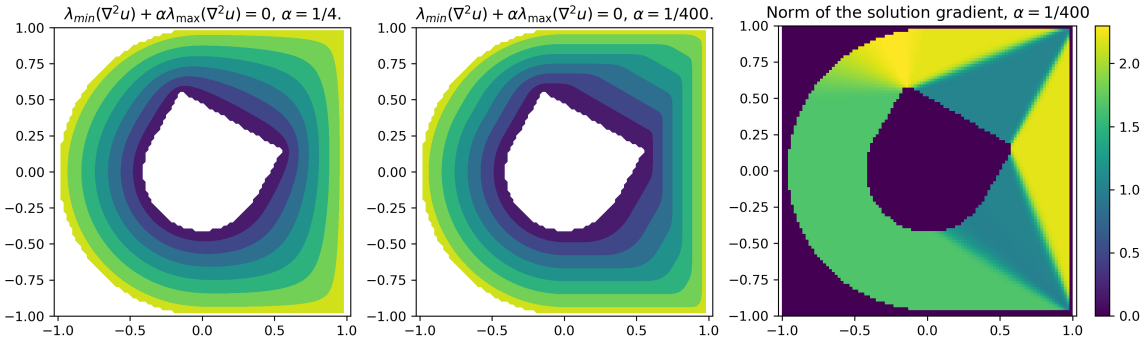

Figure 2: Solution of the Pucci PDE with $\alpha = 1/4$ (left), $\alpha = 1/400$ (center, right: gradient norm)

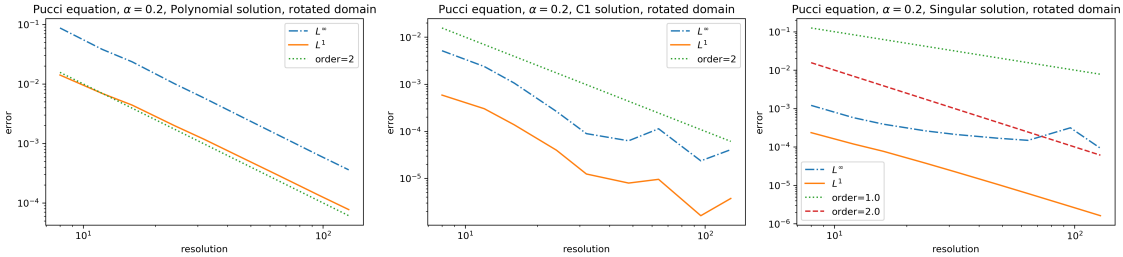

Figure 3: Numerical error as a function of grid size, for synthetic solutions to the Pucci equation.

On figure 6 (left, center) we solve the Monge-Ampere equation $\det(\nabla^2 u) = 1$ with null boundary conditions, on both a convex and a non-convex domain. In the second case, the solution is only locally convex, and it is discontinuous at the boundary, as expected [FJ17]. Figure 6 (right) shows a solution to $\det(\nabla^2 u) = 1 - 1_U$ where U is an angular sector, thus the r.h.s. is discontinuous and is vanishing on part of the domain. Null boundary conditions are applied on the boundary of the unit disk.

References

[BCM16] Jean-David Benamou, Francis Collino, and Jean-Marie Mirebeau. Monotone and consistent discretization of the Monge-Ampere operator. *Mathematics of computation*, 85(302):2743–2775, 2016.

[BOZ04] J Frédéric Bonnans, Elisabeth Ottenwaelter, and Housnaa Zidani. A fast algorithm for the two dimensional HJB equation of stochastic control. *ESAIM: Mathematical Modelling and Numerical Analysis*, 38(4):723–735, 2004.

[BS91] Guy Barles and Panagiotis E Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic analysis*, 4(3):271–283, 1991.

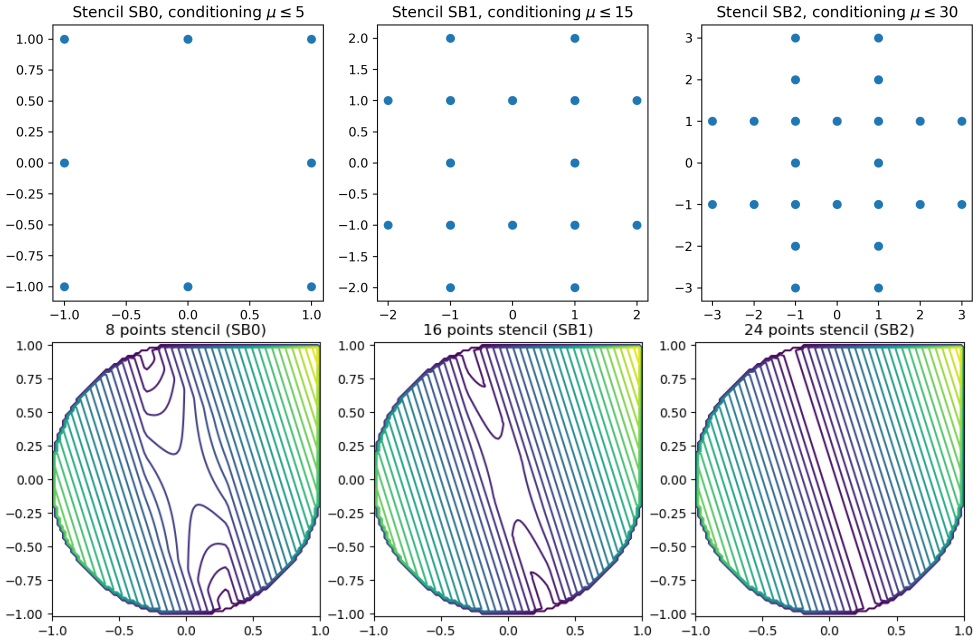

Figure 4: (top) Possible stencil choices for the Monge-Ampere equation. (bottom) Effect on the reconstruction of $u(x) = |x + y/\sqrt{10}|$ with the Monge-Ampere equation.

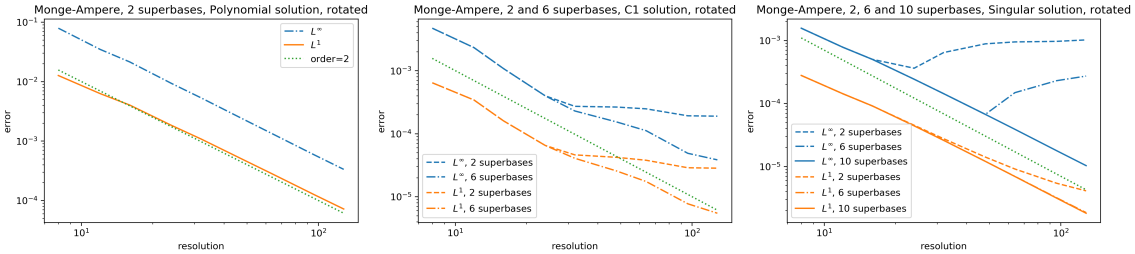

Figure 5: Numerical error as a function of grid size, for the Monge-Ampere equation, see text.

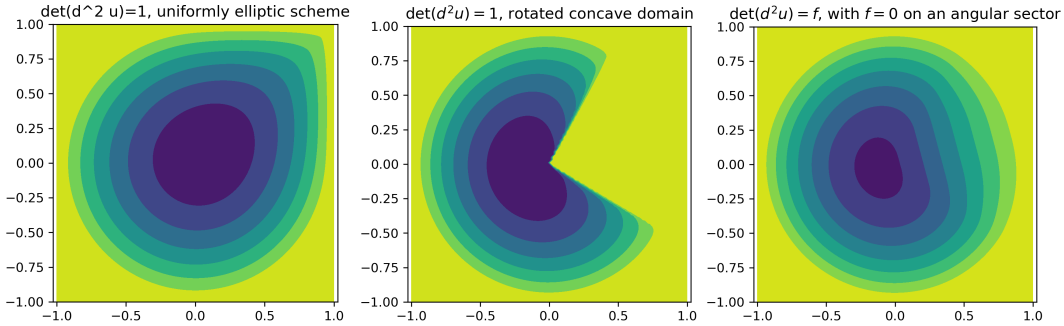

Figure 6: Solutions of the Monge-Ampere PDE, on various domains, see text.

[CIL92] Michael G Crandall, Hitoshi Ishii, and Pierre-Louis Lions. User's guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society*, 27(1):1–67, 1992.

[FJ17] Xiaobing Feng and Max Jensen. Convergent semi-Lagrangian methods for the Monge–Ampère equation on unstructured grids. *SIAM Journal on Numerical Analysis*, 55(2):691–712, 2017.

[FM14] Jérôme Fehrenbach and Jean-Marie Mirebeau. Sparse non-negative stencils for anisotropic diffusion. *Journal of Mathematical Imaging and Vision*, 49(1):123–147, 2014.

[FO13] Brittany D Froese and Adam M Oberman. Convergent Filtered Schemes for the Monge-Ampère Partial Differential Equation. *SIAM Journal on Numerical Analysis*, 51(1):423–444, January 2013.

[Mir14] Jean-Marie Mirebeau. Anisotropic Fast-Marching on cartesian grids using Lattice Basis Reduction. *SIAM Journal on Numerical Analysis*, 52(4):1573–1599, January 2014.

[Mir17] Jean-Marie Mirebeau. Riemannian fast-marching on cartesian grids using Voronoi's first reduction of quadratic forms. *HAL (Preprint)*, April 2017.

[Mir18] Jean-Marie Mirebeau. Fast-marching methods for curvature penalized shortest paths. *Journal of Mathematical Imaging and Vision*, 60(6):784–815, 2018.

[Obe06] Adam M Oberman. Convergent Difference Schemes for Degenerate Elliptic and Parabolic Equations: Hamilton-Jacobi Equations and Free Boundary Problems. *SIAM Journal on Numerical Analysis*, 44(2):879–895, January 2006.

[Obe08] Adam M Oberman. Wide stencil finite difference schemes for the elliptic Monge-Ampere equation and functions of the eigenvalues of the Hessian. *Discrete Contin Dyn Syst Ser B*, 2008.

[Obe13] Adam M Oberman. A Numerical Method for Variational Problems with Convexity Constraints. *SIAM Journal on Scientific Computing*, 35(1):A378–A396, January 2013.

[Sch09] Achill Schürmann. Computational geometry of positive definite quadratic forms. *University Lecture Series*, 49, 2009.

[Sel74] Eduard Selling. Ueber die binären und ternären quadratischen Formen. *Journal fur die Reine und Angewandte Mathematik*, 77:143–229, 1874.