

PhD Course on Information Geometry in Learning and Optimization

<http://image.diku.dk/MLLab/IG.php>

University of Copenhagen

September 22-26, 2014



Applications of Information Geometry to Combinatorial Optimization

Luigi Malagò

Shinshu University, Nagano, Japan – malago@shinshu-u.ac.jp
(09/14 → 09/15 @ INRIA TAO, Paris)



26 Sept 2014

The following slide contains spoilers!

In this lecture we discuss first order optimization over manifolds of distributions



We optimize functionals defined over statistical manifolds by means of (stochastic) natural gradient descent

We focus on discrete statistical models in the exponential family, even if most of the discussion applies also to the continuous case

Motivation

The main motivation is **stochastic optimization**, i.e., randomized search methods which introduce randomness into the search process

In particular we are interested in **model-based optimization**, i.e., a broad family of algorithms which employ statistical models to guide the search for the optimum of a function

Model-based algorithms often generate minimizing sequences of probability distributions in a statistical model

Some Notation

- $f(\mathbf{x}) : \Omega \rightarrow \mathbb{R}$ the objective function
- Ω a finite search space
- $p(\mathbf{x})$ a probability distribution over the sample space Ω
- p_0 the uniform distribution over Ω
- Δ the probability simplex
- $\mathcal{M} = \{p(\mathbf{x}; \boldsymbol{\xi}) : \boldsymbol{\xi} \in \Xi\} \subset \Delta$ a parametrized statistical model
- $\boldsymbol{\xi}$ a parameter vector for p

Stochastic Relaxation

Consider the minimization problem

$$(P) \quad \min_{x \in \Omega} f(x)$$

Let $F(p) = \mathbb{E}_p[f]$, we look for the minimum of f by optimizing the Stochastic Relaxation (SR) of f

$$(SR) \quad \inf_{p \in \mathcal{M}} F(p)$$

[Remark] We take \inf , since in general \mathcal{M} may not be closed

We get candidate solutions for P by sampling optimal solutions of the SR

We introduce a chart ξ over $\mathcal{M} = \{p(x; \xi) : \xi \in \Xi\}$, let $F(\xi) = \mathbb{E}_\xi[f]$, we have a parametric representation (in coordinates) of the SR

$$(SR) \quad \inf_{\xi \in \Xi} F(\xi)$$

A Few Remarks

We move the search to the space of probability distributions, from a discrete optimization problem over Ω to a continuous optimization problem over \mathcal{M}

In the parametric representation of F , the parameters ξ become the new variables of the SR

Since $\xi \in \Xi$, we may have a constrained formulation of the SR

[Disclaimer] The SR does not provide lower bounds for P, indeed

$$\min_{\mathbf{x} \in \Omega} f(\mathbf{x}) \leq F(p) \leq \max_{\mathbf{x} \in \Omega} f(\mathbf{x})$$

Let $\mathcal{M} = \Delta$, $\min_{\mathbf{x} \in \Omega} f(\mathbf{x}) = \min_{p \in \Delta} F(p)$

More in general, for $\mathcal{M} \subset \Delta$, $\min_{\mathbf{x} \in \Omega} f(\mathbf{x}) \leq \inf_{p \in \mathcal{M}} F(p)$

Closure of \mathcal{M}

We denote with $\overline{\mathcal{M}}$ the topological closure of \mathcal{M} , i.e., \mathcal{M} together with the limits of (weakly convergent) sequences of distributions in \mathcal{M}

Moreover, we suppose $\overline{\mathcal{M}}$ is compact so that by the extreme value theorem $F(p)$ attains its infimum over $\overline{\mathcal{M}}$

Equivalence of P and SR

Let us denote the optimal solutions

- $\mathbf{x}^* \in \Omega^* = \arg \min_{x \in \Omega} f(\mathbf{x})$
- $p^* \in P^* = \arg \min_{p \in \overline{\mathcal{M}}} F(p)$

The SR is equivalent to P if $p^*(\mathbf{x}^*) = 1$, i.e., the probability of sampling optimal solutions of P from optimal solutions of SR is equal to 1

In other words, there must exist a sequence $\{p_t\}$ in \mathcal{M} such that

$$\lim_{t \rightarrow \infty} p_t(\mathbf{x}^*) = 1$$

A sufficient condition for the equivalence of SR and P is that all Dirac distribution δ_x are included in $\overline{\mathcal{M}}$

“Why are we doing this?”

Let us clarify some points

- Let k be the cardinality of Ω , to parametrize Δ we need $k - 1$ parameters
- Minimizing $F(p)$ with $p \in \Delta$ is as computationally expensive as an exhaustive search

And so, why going from Ω to \mathcal{M} ?

- In practice we restrict the search to some lower-dimensional model
- The equivalence of P and SR can be easily guaranteed by low dimensional models
- We can develop efficient blackbox optimization algorithms
- It is possible to identify (e.g., learn) submodels for which “nice” properties hold (e.g., gradient descent converges to global optimum)

Which Model to Choose?

[Remark] The landscape of $F(p)$ depends on the choice of \mathcal{M} , i.e., the number of local minima of $F(p)$ depends on \mathcal{M}

Often it is convenient to work with graphical models

- There exists an interpretation of the model in terms of conditional independence assumptions among variables
- We can define hierarchical structures and control model dimension
- There is a strong literature about estimation, sampling, learning, ...

In the following we work (mainly) with exponential families

The Gibbs or Boltzmann Distribution

- The Gibbs distribution is the one-dimensional exponential family

$$p(\mathbf{x}; \beta) = \frac{e^{-\beta f(\mathbf{x})}}{\mathbb{E}_q[e^{-\beta f(\mathbf{x})}]} q(\mathbf{x}), \quad \beta > 0$$

- $f(\mathbf{x})$ the energy function
- β the inverse temperature
- $q(\mathbf{x})$ the reference measure
- The set of distributions is not weakly closed
- Suppose $f(\mathbf{x}) \geq 0$ and $f(\mathbf{x}) = 0$ for some $\mathbf{x} \in \Omega$, but not everywhere

$$\lim_{\beta \rightarrow 0} e^{-\beta f(\mathbf{x})} = q(\mathbf{x})$$

$$\lim_{\beta \rightarrow 0} \mathbb{E}_q[e^{-\beta f(\mathbf{x})}] = 1$$


$$\lim_{\beta \rightarrow \infty} e^{-\beta f(\mathbf{x})} = \begin{cases} q(\mathbf{x}) & \text{if } f(\mathbf{x}) = 0, \\ 0 & \text{otherwise} \end{cases}$$

$$\lim_{\beta \rightarrow \infty} \mathbb{E}_q[e^{-\beta f(\mathbf{x})}] = \sum_{\Omega^*} q(\mathbf{x})$$

Closure of the Gibbs Distribution

$$\lim_{\beta \rightarrow 0} p(\mathbf{x}; \beta) = q(\mathbf{x})$$

$$\lim_{\beta \rightarrow \infty} p(\mathbf{x}; \beta) = \frac{q(\mathbf{x})}{\sum_{\Omega^*} q(\mathbf{x})}$$



The Gibbs distribution is in principle an optimal choice for the SR, indeed

- The limit for $\beta \rightarrow \infty$ has support over the minima of f , and in particular for $q(\mathbf{x}) = p_0$, it is the uniform distribution δ_{Ω^*}
- Since

$$\nabla F(\beta) = -\text{Var}_{\beta}(f) < 0$$

$F(\beta)$ decreases monotonically as $\beta \rightarrow \infty$

- The Gibbs distribution satisfies the principle of maximum entropy

However, evaluating the partition function is computationally infeasible

Pseudo-Boolean Optimization

In the following we restrict to $\Omega = \{+1, -1\}^n$, and we use the harmonic encoding $\{+1, -1\}$ for a binary variable

$$-1^0 = +1 \qquad -1^1 = -1$$

A pseudo-Boolean function f is a real-valued mapping

$$f(\mathbf{x}) : \Omega = \{+1, -1\}^n \rightarrow \mathbb{R}$$

Any f can be expanded uniquely as a square free polynomial

$$f(\mathbf{x}) = \sum_{\alpha \in L} c_{\alpha} \mathbf{x}^{\alpha},$$

by employing a multi-index notation. Let $L = \{0, 1\}^n$, then $\alpha = (\alpha_1, \dots, \alpha_n) \in L$ uniquely identifies the monomial \mathbf{x}^{α} by

$$\alpha \mapsto \prod_{i=1}^n x_i^{\alpha_i}$$

The degree of the monomials represents the order of the interactions in f

Monomial Representation of PS Functions

Let $A^n = \underbrace{A^1 \otimes \dots \otimes A^1}_{n \text{ times}}$, where \otimes denotes the Kronecker product and

$$A^1 = \begin{matrix} & \begin{matrix} 0 & 1 \end{matrix} \\ \begin{matrix} + \\ - \end{matrix} & \begin{bmatrix} +1 & +1 \\ +1 & -1 \end{bmatrix} \end{matrix}$$

let $\mathbf{a} = (f(\mathbf{x}))_{\mathbf{x} \in \Omega}$, we have $A^n \mathbf{c} = \mathbf{a}$, with $\mathbf{c} = (c_\alpha)_{\alpha \in L}$ and $\mathbf{c} = 2^{-n} A^n \mathbf{a}$

[Example] In case of two variables $\mathbf{x} = (x_1, x_2)$, we have

$$f(\mathbf{x}) = c_0 + c_1 x_1 + c_2 x_2 + c_{12} x_1 x_2$$

$$\begin{array}{cc|c} x_1 & x_2 & f(\mathbf{x}) \\ \hline +1 & +1 & a_{++} \\ +1 & -1 & a_{+-} \\ -1 & +1 & a_{-+} \\ -1 & -1 & a_{--} \end{array} \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ c_{12} \end{bmatrix} = \frac{1}{4} \times \begin{matrix} & \begin{matrix} 00 & 10 & 01 & 11 \end{matrix} \\ \begin{matrix} ++ \\ +- \\ -+ \\ -- \end{matrix} & \begin{bmatrix} +1 & +1 & +1 & +1 \\ +1 & +1 & -1 & -1 \\ +1 & -1 & +1 & -1 \\ +1 & -1 & -1 & +1 \end{bmatrix} \end{matrix} \begin{bmatrix} a_{++} \\ a_{+-} \\ a_{-+} \\ a_{--} \end{bmatrix}$$

The Independence Model

Let \mathcal{M}_1 be the independence model for $\mathbf{X} = (X_1, \dots, X_n)$

$$\mathcal{M}_1 = \left\{ p : p(\mathbf{x}) = \prod_{i=1}^n p_i(x_i) \right\}$$

with marginal probabilities $p_i(x_i) = \mathbb{P}(X_i = x_i)$

We parametrize \mathcal{M}_1 using ± 1 Bernoulli distributions for X_i



$$\begin{aligned} p(\mathbf{x}; \boldsymbol{\mu}) &= \prod_{i=1}^n \mu_i^{(1+x)/2} (1 - \mu_i)^{(1-x)/2} \\ &= \prod_{i=1}^n (2\mu_i x_i - x_i + 1) / 2 \end{aligned}$$

with $\boldsymbol{\mu} = (\mu_1, \dots, \mu_n) \in [0, 1]^n$ and

$$\mu_i = \mathbb{P}(X_i = +1)$$

$$1 - \mu_i = \mathbb{P}(X_i = -1)$$

Marginal Parameters for the Independence Model



\mathcal{M}_1 is a n -dimensional manifold embedded in the $2^n - 1$ dimensional probability simplex Δ

Stochastic Relaxation: Who?

The SR can be solved in different ways

- Sampling, Selection and Estimation paradigm (EDAs, Larrañaga and Lozano, 2002, CE method, Rubinstein, 1997)
- Fitness Modeling (DEUM, Shakya et al., 2005)
- Covariance Matrix Adaptation (CMA-ES, Hansen et al., 2001)
- GAs (Holland, 1975), ACO (Dorigo, 1992), ESs (Rechenberg, 1960)
- Boltzmann distribution and Gibbs sampler (Geman and Geman, 1984)
- Simulated Annealing and Boltzmann Machines (Aarts and Korst, 1989)
- Method of the Moments (SDP, Meziat et al., 2001)
- LP relaxation in pseudo-Boolean optimization (Boros and Hammer, 2001)
- ...and many others

Gradient Descent

In the following we will focus on gradient descent techniques of the form

$$\boldsymbol{\xi}^{t+1} = \boldsymbol{\xi}^t - \lambda \nabla F(\boldsymbol{\xi}), \quad \lambda > 0$$

In particular we refer to gradient-based algorithms such as


- Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES), Hansen et al., 2001
- Natural Evolutionary Strategies (NES), Wierstra et al., 2008
- Stochastic Natural Gradient Descent (SNGD), M. et al., 2011
- Information Geometry Optimization (IGO), Arnold et al., 2011

A Toy Example

Let $n = 2$, $\Omega = \{-1, +1\}^2$, we want to minimize

$$f(\mathbf{x}) = x_1 + 2x_2 + 3x_1x_2$$

x_1	x_2	$f(\mathbf{x})$
+1	+1	6
+1	-1	-4
-1	+1	-2
-1	-1	0



The gradient flow is the solution of the following differential equation

$$\dot{\xi} = \nabla F(\xi),$$

where we set the step size $\lambda \rightarrow \epsilon$


We are interested in studying gradient flows for different parameterization and different statistical models

Gradient Flows on the Independence Model


$$F(\boldsymbol{\mu}) = \sum_{\mathbf{x} \in \Omega} f(\mathbf{x}) p_1(x_1) p_2(x_2) = -4\mu_1 - 2\mu_2 + 12\mu_1\mu_2$$

$$\nabla F(\boldsymbol{\mu}) = (-4 + 12\mu_2, -2 + 12\mu_1)^T$$

Gradient flow in $\boldsymbol{\mu}$



Gradient vector in $\boldsymbol{\mu}$, $\lambda = 0.025$



$\nabla F(\boldsymbol{\eta})$ does not converge to (local) optima, a projection over the hyperplanes given by the constraints is required

Natural Parameters for the Independence Model

If we restrict to positive probabilities $p > 0$, we can represent the interior of the independence model as the exponential family

$$p(\mathbf{x}; \boldsymbol{\theta}) = \exp \left\{ \sum_{i=1}^n \theta_i x_i - \psi(\boldsymbol{\theta}) \right\}$$

where $\psi(\boldsymbol{\theta}) = \ln Z(\boldsymbol{\theta})$ is the log partition function

The natural parameters of the independence model \mathcal{M}_1 represented by an exponential family are $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n) \in \mathbb{R}^n$, with

$$p_i(x_i) = \frac{e^{\theta_i x_i}}{e^{\theta_i} + e^{-\theta_i}}$$

The mapping between marginal probabilities and natural parameters is one-to-one for $p > 0$


$$\theta_i = (\ln(\mu_i) - \ln(1 - \mu_i)) / 2 \qquad \mu_i = \frac{e^{\theta_i}}{e^{\theta_i} + e^{-\theta_i}}$$

Gradient Flows on the Independence Model


$$F(\boldsymbol{\theta}) = (-4e^{\theta_1 - \theta_2} - 2e^{-\theta_1 + \theta_2} + 6e^{\theta_1 + \theta_2}) / Z(\boldsymbol{\theta})$$

$$\nabla F(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}}[f(\mathbf{X} - \mathbb{E}_{\boldsymbol{\theta}}[\mathbf{X}])] = \text{Cov}_{\boldsymbol{\theta}}(f, \mathbf{X})$$

Gradient flow in $\boldsymbol{\theta}$




Gradient vectors in $\boldsymbol{\theta}$, $\lambda = 0.15$




In the $\boldsymbol{\theta}$ parameters, $\nabla F(\boldsymbol{\theta})$ vanishes over the plateaux

Gradient Flows on the Independence Model

Marginal probabilities μ



Natural parameters θ



Gradient flows $\nabla F(\xi)$ depend on the parameterization

In the η parameters, $\nabla F(\eta)$ does not convergence to the expected distribution δ_{x^*} over an optimum

Summary of the Intro

- Iterative algorithm generate sequences of distributions which can be compared to the gradient flow of $F(\xi)$
- The choice of the model and of the parameterization play a role in terms of appearance of local minima
- Euclidean geometry does not appear to be the proper geometry for statistical models

We need a more general mathematical framework able to deal with non-Euclidean geometries

The Exponential Family

In the following, we consider models in the exponential family \mathcal{E}

$$p(\mathbf{x}, \boldsymbol{\theta}) = \exp\left(\sum_{i=1}^m \theta_i T_i(\mathbf{x}) - \psi(\boldsymbol{\theta})\right)$$

- sufficient statistics $\mathbf{T} = (T_1(\mathbf{x}), \dots, T_m(\mathbf{x}))$
- natural parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m) \in \Theta \subset \mathbb{R}^m$
- log-partition function $\psi(\boldsymbol{\theta})$

Several statistical models belong to the exponential family, both in the continuous and discrete case, among them

- the independence model
- the Gibbs distribution
- Markov random fields
- multivariate Gaussians

Markov Random Fields

[Recall] The monomials $\{\mathbf{x}^\alpha\}$, $\alpha \in L$, define a basis for f

By choosing a subset of $\{\mathbf{x}^\alpha\}$ as sufficient statistics, we can identify a low-dimensional exponential family parametrized by θ

$$p(\mathbf{x}; \theta) = \exp\left(\sum_{\alpha \in M \subset L_0} \theta_\alpha \mathbf{x}^\alpha - \psi(\theta)\right), \quad L_0 = L \setminus \{0\}$$

Such models are known as


- log-linear models
- Markov random fields
- Boltzmann machines

There exists a statistical interpretation for the topology of the model

- The absence of edges in an undirected graphical model implies conditional independence among variables
- Joint probability distributions factorize over the cliques of the graph

In some cases it may be convenient to work with hierarchical models

A hierarchical model has the property that if x^α is a sufficient statistic of the exponential family then all $\{x^\beta : \beta \in L \setminus \{\mathbf{0}\} \wedge \beta \rightarrow \alpha\}$, must be sufficient statistics as well, where \rightarrow is the bitwise implication operator



Dual Parameterization for the Exponential Family

$$p(\mathbf{x}; \boldsymbol{\theta}) = \exp\left(\sum_{i=1}^m \boldsymbol{\theta}_i T_i(\mathbf{x}) - \psi(\boldsymbol{\theta})\right)$$

- Exponential families admit a dual parametrization to the natural parameters, given by the expectation parameters with $\boldsymbol{\eta} = \mathbb{E}_{\boldsymbol{\theta}}[\mathbf{T}]$
- Let $\varphi(\boldsymbol{\eta})$ be the negative entropy of p , then $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ are connected by the Legendre transform

$$\psi(\boldsymbol{\theta}) + \varphi(\boldsymbol{\eta}) - \langle \boldsymbol{\theta}, \boldsymbol{\eta} \rangle = 0$$

- Variable transformations are given by

$$\boldsymbol{\eta} = \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}) = (\nabla_{\boldsymbol{\eta}} \varphi)^{-1}(\boldsymbol{\theta})$$

$$\boldsymbol{\theta} = \nabla_{\boldsymbol{\eta}} \varphi(\boldsymbol{\eta}) = (\nabla_{\boldsymbol{\theta}} \psi)^{-1}(\boldsymbol{\eta})$$

Variable Transformations

[Recall] Let $A^n = \underbrace{A^1 \otimes \dots \otimes A^1}_{n \text{ times}}$, where \otimes denotes the Kronecker product

A probability distribution $p \in \Delta$ requires 2^n parameters $\rho = (p(\mathbf{x}))_{\mathbf{x} \in \Omega}$ to be uniquely identified, with constraints $0 \leq \rho_{\mathbf{x}} \leq 1$ and $\sum_{\mathbf{x} \in \Omega} \rho_{\mathbf{x}} = 1$

The expectation parameters $\eta = (\eta_{\alpha})$, $\alpha \in L$, provide an equivalent parameterization for p , and since $p(\mathbf{x})$ is a pseudo-Boolean function itself, we have

$$\rho = 2^{-n} A^n \eta \qquad \eta = A^n \rho$$

Positivity constraints and the fact that probabilities sum to one, give us $\eta_0 = 1$ and $A^n \eta \geq 0$.

The natural parameters $\theta = (\theta_{\alpha})$, $\alpha \in L$, can be obtained from raw probabilities, with the constraint $\theta_0 = -\log \mathbb{E}_{\theta}[\exp \sum_{\alpha \in L \setminus \{0\}} \theta_{\alpha} \mathbf{x}^{\alpha}]$

$$\ln \rho = 2^{-n} A^n \theta \qquad \theta = A^n \ln \rho$$

Mixed Parametrization for Markov Random Fields

An exponential family \mathcal{M} given by the sufficient statistics $\{\mathbf{x}^\alpha\}$, $\alpha \in M$, identifies a submanifold in Δ , parametrized by $\theta = ((\theta)_{\alpha \in M}; \mathbf{0})$

By the one-to-one correspondence between η and θ , \mathcal{M} can be parametrized by $\eta = (\eta_{\alpha \in M}; \eta_{\alpha \notin M}^*)$, where in general $\eta_{\alpha \notin M}^* \neq \mathbf{0}$

However, the $\eta_{\alpha \notin M}^*$ parameters are not free and it can be proved they are given by implicit polynomial algebraic equations in $\eta_{\alpha \in M}$

Due to the duality between θ and η , we can employ a mixed parametrization for \mathcal{M} and parametrize the model as $(\eta_{\alpha \in M}; \mathbf{0})$

[Remark] The study of statistical models using tools from computational algebra is called Algebraic Statistics (Pistone et al., 2000)

Algebraic Statistics: Invariants in ρ and η

[Example] Let $n = 2$, we consider the independence model parametrized by $(\theta_1, \theta_2; 0)$, with $\theta_{12} = 0$

The same model can be parametrized by $(\eta_1, \eta_2; 0)$, we show $\eta_{12} = \eta_1 \eta_2$

Since $\theta = A^n \ln \rho$, by imposing $\theta_{12} = 0$ we have

$$\ln \rho_{++} + \ln \rho_{--} = \ln \rho_{+-} + \ln \rho_{-+}$$

$$\rho_{++}\rho_{--} = \rho_{+-}\rho_{-+}$$

$$\begin{bmatrix} \rho_{++} \\ \rho_{+-} \\ \rho_{-+} \\ \rho_{--} \end{bmatrix} = \frac{1}{4} \times \begin{matrix} & \begin{matrix} 00 & 10 & 01 & 11 \end{matrix} \\ \begin{matrix} ++ \\ +- \\ -+ \\ -- \end{matrix} & \begin{bmatrix} +1 & +1 & +1 & +1 \\ +1 & +1 & -1 & -1 \\ +1 & -1 & +1 & -1 \\ +1 & -1 & -1 & +1 \end{bmatrix} \end{matrix} \begin{bmatrix} 1 \\ \eta_1 \\ \eta_2 \\ \eta_{12} \end{bmatrix}$$

$$(1 + \eta_1 + \eta_2 + \eta_{12})(1 - \eta_1 - \eta_2 + \eta_{12}) = (1 + \eta_1 - \eta_2 - \eta_{12})(1 - \eta_1 + \eta_2 - \eta_{12})$$

$$\eta_{12} = \eta_1 \eta_2$$

Marginal Polytope

The range of the expectation parameters $\eta = \mathbb{E}_{\theta}[\mathbf{T}]$ identifies a polytope M in \mathbb{R}^m called the **marginal polytope**

The marginal polytope can be obtained as the convex hull of $\mathbf{T}(\Omega)$, there \mathbf{T} is the vector of sufficient statistics of the model

[Example] Let $n = 2$, $\mathbf{T} = (x_1, x_1x_2)$

$$A = \begin{array}{cc} & \begin{array}{cc} x_1 & x_1x_2 \end{array} \\ \begin{array}{c} ++ \\ +- \\ -+ \\ -- \end{array} & \left[\begin{array}{cc} -1 & +1 \\ +1 & -1 \\ +1 & +1 \\ -1 & -1 \end{array} \right] \end{array}$$

Convex hull of

$$\begin{array}{c} (+1, +1) \\ (+1, -1) \\ (-1, -1) \\ (-1, +1) \end{array}$$

Marginal Polytope

The marginal polytope corresponds to the domain for the η parameters in the SR

- For the independence model $M = [-1, 1]^n$
- For the saturated model $M = \Delta$
- In the other cases, things can get very “nasty”, indeed the number of its faces can grow more than exponentially in n

[Example] Let $n = 3$, consider the exponential model with sufficient statistics given by

$$\{x_1, x_2, x_3, x_{12}, x_{23}, x_{13}\}$$


then the number of hyperplanes of M is 16

Information Geometry

The geometry of statistical models is not Euclidean

We need tools from differential geometry to define notions such as tangent vectors, shortest paths and distances between distributions

Information Geometry (IG) consists of the study of statistical models as manifolds of distributions endowed with the Fisher information metric (Amari 1982, 2001)



Characterization of the Tangent Space of \mathcal{E}

Over the manifold of distributions we can introduce an affine chart in p such that any density q is locally represented w.r.t. to the reference measure p by $\frac{q}{p} - 1$

The tangent space at each point p is defined by

$$T_p = \{v : \mathbb{E}_p[v] = 0\}$$

Consider a curve $p(\theta)$ such that $p(0) = p$, then $\dot{p} \in T_p$

In a moving coordinate system, tangent (velocity) vectors in $T_{p(\theta)}$ to the curve are given by logarithmic derivative

$$\frac{\dot{p}(\theta)}{p(\theta)} = \frac{d}{d\theta} \log p(\theta)$$

Characterization of the Tangent Space of \mathcal{E}

The one dimensional model

$$p(\theta) = \exp\{\theta T - \psi(\theta)\}$$

is a curve in the manifold, with tangent (velocity) vector

$$\frac{\dot{p}(\theta)}{p(\theta)} = T - \frac{d}{d\theta}\psi(\theta)$$

On the other side, given a vector field, at each p we have a vector $U(p)$ tangent to some curve, then we obtain the differential equation

$$\frac{d}{d\theta} \log p(\theta) = U(p),$$

whose solutions are one-dimensional models in Δ

Natural Gradient

Let (\mathcal{M}, I) be a statistical manifold endowed with a metric $I = [g_{ij}]$, and let $F(p) : \mathcal{M} \mapsto \mathbb{R}$ be a smooth function

For each vector field U over \mathcal{M} , the natural gradient of $\tilde{\nabla}F(p)$, i.e., the direction of steepest descent of $F(p)$ is the unique vector that satisfies

$$g(\tilde{\nabla}F, U) = D_U F,$$

where $D_U F$ is the directional derivative of F in the direction U

Given a coordinate system ξ for \mathcal{M} we have

$$\tilde{\nabla}F(\xi) = \sum_{i=1}^n \sum_{j=1}^n g^{ij} \frac{\partial F}{\partial \xi_i} \frac{\partial}{\partial \xi_j} = I(\xi)^{-1} \nabla F(\xi)$$

[Remark] There is only one (natural) gradient given by the geometry of \mathcal{M}

We use $\tilde{\nabla}F(\xi)$ to distinguish the natural gradient from the vanilla gradient $\nabla F(\xi)$, i.e., the vector of partial derivatives of F w. r. t. ξ

Geometry of the Exponential Family

In case of a finite sample space Ω , we have

$$p(\mathbf{x}; \boldsymbol{\theta}) = \exp\left(\sum_{i=1}^m \theta_i T_i(\mathbf{x}) - \psi(\boldsymbol{\theta})\right) \quad \boldsymbol{\theta} \in \mathbb{R}^m$$

and

$$\mathbb{T}_{\boldsymbol{\theta}} = \left\{ v : v = \sum_{i=1}^k a_i (T_i(\mathbf{x}) - \mathbb{E}_{\boldsymbol{\theta}}[T_i]), a_i \in \mathbb{R} \right\}$$

Since $\nabla F(\boldsymbol{\theta}) = \text{Cov}_{\boldsymbol{\theta}}(f, T)$, if $f \in \mathbb{T}_p$, the steepest direction is given by $f - \mathbb{E}_{\boldsymbol{\theta}}[f]$, otherwise we take the projection \widehat{f} of f onto \mathbb{T}_p

$$\widehat{f} = \sum_{i=1}^m \widehat{a}_i (T_i(\mathbf{x}) - \mathbb{E}_{\boldsymbol{\theta}}[T_i]),$$

and obtain \widehat{f} by solving a system of linear equations

Geometry of Statistical Models

Since $f - \widehat{f}$ is orthogonal to T_p

$$\mathbb{E}_{\boldsymbol{\theta}}[(f - \widehat{f}_{\boldsymbol{\theta}})(T - \mathbb{E}_{\boldsymbol{\theta}}[T])] = \text{Cov}_{\boldsymbol{\theta}}(f - \widehat{f}_{\boldsymbol{\theta}}, T) = 0,$$

from which we obtain, for $i = 1, \dots, m$

$$\text{Cov}_{\boldsymbol{\theta}}(f, T_i) = \text{Cov}_{\boldsymbol{\theta}}(\widehat{f}_{\boldsymbol{\theta}}, T_i) = \sum_{j=1}^m \widehat{a}_j \text{Cov}_{\boldsymbol{\theta}}(T_j, T_i)$$

As the Hessian matrix of $\psi(\boldsymbol{\theta})$ is invertible, we have


$$\widehat{\mathbf{a}} = [\text{Cov}_{\boldsymbol{\theta}}(T_i, T_j)]^{-1} \text{Cov}_{\boldsymbol{\theta}}(f, T) = I(\boldsymbol{\theta})^{-1} \nabla F(\boldsymbol{\theta})$$



In case $f \in \text{Span}\{T_1, \dots, T_m\}$, then $\widehat{f}_{\boldsymbol{\theta}} = f$

By taking the projection of f onto T_p , we obtained the natural gradient $\widetilde{\nabla} F$, i.e., the gradient evaluated w.r.t. the Fisher information metric I



The Big Picture

If $f \notin T_p$, the projection \hat{f} may vanish, and local minima may appear



Vanilla vs Natural Gradient: $\eta, \lambda = 0.05$ Vanilla gradient $\nabla F(\boldsymbol{\eta})$ Natural gradient $\tilde{\nabla} F(\boldsymbol{\eta})$ 


In both cases there exist two basins of attraction, however $\tilde{\nabla} F(\boldsymbol{\eta})$ converges to $\delta_{\mathbf{x}}$ distributions, which are local optima for $F(\boldsymbol{\eta})$ and where $\tilde{\nabla} F(\delta_{\mathbf{x}}) = 0$

Vanilla vs Natural Gradient: $\theta, \lambda = 0.15$ Vanilla gradient $\nabla F(\theta)$ Natural gradient $\tilde{\nabla} F(\theta)$ 


In both cases there exist two basins of attraction, however in the natural parameters $\tilde{\nabla} F(\theta)$ “speeds up” over the plateaux

Vanilla vs Natural Gradient

Expectation parameters η



Natural parameters θ



Vanilla gradient ∇F vs Natural gradient $\tilde{\nabla} F$

The natural gradient flow is invariant to parameterization

Stochastic Natural Gradient Descent

In the exponential family, the natural gradient descent updating rule reads

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \lambda I(\boldsymbol{\theta})^{-1} \nabla F(\boldsymbol{\theta}), \quad \lambda > 0$$

Unfortunately, exact gradients cannot be computed efficiently

- in general the partition function must be evaluated
- or a change of parametrization from $\boldsymbol{\theta}$ to $\boldsymbol{\eta}$ is required

However, due to the properties of the exponential family, natural gradient can be evaluated by means of covariances

$$\nabla F(\boldsymbol{\theta}) = \text{Cov}_{\boldsymbol{\theta}}(f, \mathbf{T}) \quad I(\boldsymbol{\theta}) = \text{Cov}_{\boldsymbol{\theta}}(\mathbf{T}, \mathbf{T})$$

As a consequence, stochastic natural gradient can be estimated by replacing exact gradients with empirical estimates, so that

$$\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \lambda \widehat{\text{Cov}}_{\boldsymbol{\theta}^t}(\mathbf{T}, \mathbf{T})^{-1} \widehat{\text{Cov}}_{\boldsymbol{\theta}^t}(f, \mathbf{T}), \quad \lambda > 0$$

Stochastic Natural Gradient Descent

We can now move from gradient flows to iterative algorithms that discretize the flow

A naïve SNGD (M. et al., 2011) implements the following pseudo-code

- 0 $t = 0, \boldsymbol{\theta}_t = \mathbf{0}$
- 1 Start with an initial random sample \mathcal{P}_t
- 2 Evaluate f for each $\mathbf{x} \in \mathcal{P}_t$
- 3 (Subsample \mathcal{P}_t)
- 4 $\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t - \lambda \widehat{\text{Cov}}_{\boldsymbol{\theta}}(\mathbf{T}, \mathbf{T})^{-1} \widehat{\text{Cov}}_{\boldsymbol{\theta}}(f, \mathbf{T})$
- 5 Sample from $p(\mathbf{x}; \boldsymbol{\theta}^{t+1})$
- 6 Repeat 2-5 until convergence

Sampling from $p(\mathbf{x}; \boldsymbol{\theta}^{t+1})$ can be done using the Gibbs sampler or other sampling algorithms, for instance exploiting special conditional structure and properties of p

Gibbs Sampler with Cooling Scheme

Input: θ, N

Optional: \mathcal{P}_0, T_0

1: **function** GibbsSampler($\theta, N, \mathcal{P}_0, T_0$)

2: $\mathcal{P} \leftarrow \emptyset$

3: $t \leftarrow 1$

4: **repeat**

5: $\mathbf{x} \leftarrow \text{Random}(\mathcal{P}_0)$

6: $T \leftarrow T_0$

7: **repeat**

8: $i \leftarrow \text{Random}(\{1, \dots, n\})$

9: $\mathbf{x}_{\setminus i} \leftarrow (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$

10:
$$p_i(x_i | \mathbf{x}_{\setminus i}; \theta) = \frac{1}{1 + \exp\{2T^{-1}x_i \sum_{\alpha \in M_i} \theta_{\alpha \setminus i} \mathbf{x}^{\alpha \setminus i}\}}$$

11:
$$x_i \leftarrow \begin{cases} +1, & \text{with } \mathbb{P}_i(X_i = 1 | \mathbf{X}_{\setminus i} = \mathbf{x}_{\setminus i}; \theta) \\ -1, & \text{otherwise} \end{cases}$$

12: $T \leftarrow \text{CoolingScheme}(T)$

13: **until** StoppingCriteria()

14: $\mathcal{P} \leftarrow \mathcal{P} \cup \{\mathbf{x}\}$

15: $t \leftarrow t + 1$

16: **until** $t = N$

17: **return** \mathcal{P}

18: **end function**

- ▷ natural parameters, sample size
- ▷ pool of samples, initial temperature

- ▷ random point if $\mathcal{P}_0 = \emptyset$
- ▷ initial temp, $T = 1$ default value

- ▷ random variable

- ▷ decrease T

- ▷ add new point

Theorem 1 (M. et al., FOGA 2011)

If the sufficient statistics $\{\mathbf{x}^\alpha\}$ of the exponential family

$$p(\mathbf{x}; \boldsymbol{\theta}) = \exp\left(\sum_{\alpha \in M \subset L_0} \theta_\alpha \mathbf{x}^\alpha - \psi(\boldsymbol{\theta})\right)$$

form a basis for

$$f(\mathbf{x}) = \sum_{\alpha \in I} c_\alpha \mathbf{x}^\alpha$$

i.e., $I \setminus \{0\} \subset M$, then the natural gradient flows from every distribution p in the exponential family converge to the global optimum of the SR

Theorem 2 (M. et al., CEC 2013)

If the sufficient statistics $\{\mathbf{x}^\alpha\}$ of the exponential family

$$p(\mathbf{x}; \boldsymbol{\theta}) = \exp\left(\sum_{\alpha \in M \subset L_0} \theta_\alpha \mathbf{x}^\alpha - \psi(\boldsymbol{\theta})\right)$$

are centered in $\boldsymbol{\theta}$, i.e., $\mathbb{E}_\theta[\mathbf{x}^\alpha] = \boldsymbol{\theta}$ for every $\alpha \in M$, then the least squares estimator $\widehat{\mathbf{a}}_N$ of the regression model

$$\widehat{f}(\mathbf{x}) = \sum_{\alpha \in M \subset L_0} a_\alpha \mathbf{x}^\alpha$$

estimated from a random sample \mathcal{P} i. i. d. $\sim p(\mathbf{x}; \boldsymbol{\theta})$ converges to the natural gradient $\widetilde{\nabla} F(\boldsymbol{\theta})$, as $N \rightarrow \infty$

Natural Gradient and Fitness Modeling

Proof. Let A be the design matrix, e.g.,

$$A = \begin{array}{cccc} & x_1 & x_2 & x_1 x_2 & \dots \\ +---+- & +1 & -1 & -1 & \dots \\ --++++ & -1 & -1 & +1 & \dots \\ +-+--+ & +1 & +1 & +1 & \dots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{array}$$

The least squares estimator reads

$$\begin{aligned} \hat{a}_N &= (A^\top A)^{-1} A^\top \mathbf{y} \\ &= \left[\frac{1}{N} \sum_{\mathbf{x} \in \mathcal{P}} \mathbf{x}^\alpha \mathbf{x}^\beta \right]^{-1} \left(\frac{1}{N} \sum_{\mathbf{x} \in \mathcal{P}} f(\mathbf{x}) \mathbf{x}^\alpha \right) \\ &= [\widehat{\text{Cov}}(\mathbf{x}^\alpha, \mathbf{x}^\beta) + \widehat{\mathbb{E}}[\mathbf{x}^\alpha] \widehat{\mathbb{E}}[\mathbf{x}^\beta]]^{-1} (\widehat{\text{Cov}}(f, \mathbf{x}^\alpha) + \widehat{\mathbb{E}}[f] \widehat{\mathbb{E}}[\mathbf{x}^\alpha]) \\ &= [\widehat{\text{Cov}}(\mathbf{x}^\alpha, \mathbf{x}^\beta)]^{-1} (\widehat{\text{Cov}}(f, \mathbf{x}^\alpha)) \end{aligned}$$

Model selection and Model Selection

The previous theorem determines a relationship between linear regression in machine learning and stochastic gradient estimation in optimization

- Natural gradient estimation can be solved by linear regression
- Stochastic natural gradient performs fitness modeling
- Model selection can be performed by subset selection during linear regression

Forward subset selection-style algorithms can be employed to simultaneously perform model selection and gradient estimation

- 1 Start with no sufficient statistics in the model
- 2 Choose among a set of candidates variables the one that improves the model the most
- 3 Evaluate the new residual vector $\mathbf{f} - \hat{\mathbf{f}}$
- 4 Repeat 2-3 until RSS is less than a threshold or the maximum number of variables is reached

Choice of \mathcal{M} in the SR

The choice of the statistical model \mathcal{M} determines the landscape of F

To remove the local minima of F over \mathcal{M} , we can add higher-order interactions x^α to the model

[Example] In our toy example, with

$$f(\mathbf{x}) = x_1 + 2x_2 + 3x_1x_2$$

we can remove the critical point of $\tilde{\nabla}F$ in \mathcal{M}_1 by adding the sufficient statistic X_{12} to the independence model and obtain

$$\{X_1, X_2\} \Rightarrow \{X_1, X_2, X_{12}\}$$

The independence model is a submanifold of the new (saturated) model: a distribution p in the independence model has the parameters associated to higher-order interactions set to 0

$$(\theta_1, \theta_2, 0) \Rightarrow (\theta_1, \theta_2, \theta_{12})$$

Variable transformation (Cucci et al., 2012)

An equivalent approach to change the landscape of F is based on variables transformation: we fix the model over a set of new variables $\mathbf{y} = (y_1, \dots, y_n) \in \Omega$ obtained from \mathbf{x}

For instance, consider the mapping $\tau_{ij} : \Omega \rightarrow \Omega : \mathbf{x} \mapsto \mathbf{y}$ with

$$y_k = \begin{cases} x_k x_j & \text{if } k = i \\ x_k & \text{otherwise} \end{cases}$$

A large class of transformations can be obtained by concatenating simple τ_{ij} transformations

Mapping \mathbf{y} back to \mathbf{x} is easy, since $\tau_{ij} = \tau_{ij}^{-1}$

Learning a transformation of \mathbf{x} which simplifies the landscape of F for a fixed \mathcal{M} over \mathbf{y} is equivalent to finding a good model for \mathbf{x}

Variables transformation for the independence model

[Example] Let us apply the transformation τ_{12} to $\mathbf{x} = (x_1, x_2)$

$$y_1 = x_1 x_2$$

$$y_2 = x_2$$

and consider the independence model over \mathbf{y}

$$p(\mathbf{y}; \boldsymbol{\zeta}) = \exp\{\zeta_1 y_1 + \zeta_2 y_2 - \psi(\boldsymbol{\zeta})\}$$

By representing the independence model for the \mathbf{y} variables as a model for \mathbf{x} in the probability simple Δ , we obtain the exponential model

$$p(\mathbf{x}; \boldsymbol{\zeta}) = \exp\{\zeta_2 x_2 + \zeta_1 x_1 x_2 - \psi(\boldsymbol{\zeta})\}$$

which can be expressed in $\boldsymbol{\theta}$ by renaming the variables

$$p(\mathbf{x}; \boldsymbol{\theta}) = \exp\{\theta_2 x_2 + \theta_{12} x_1 x_2 - \psi(\boldsymbol{\theta})\}$$

The independence models for \mathbf{x} and \mathbf{y} correspond to different models in Δ



Back to the Toy Example: Submanifolds in Δ

Independence model: $\theta = (\theta_1, \theta_2, 0)$

$$p(\mathbf{x}; \theta) = \exp\{\theta_1 x_1 + \theta_2 x_2 - \psi(\theta)\}$$

Exponential family: $\theta = (0, \theta_2, \theta_{12})$

$$p(\mathbf{x}; \theta) = \exp\{\theta_2 x_2 + \theta_{12} x_1 x_2 - \psi(\theta)\}$$



The lines represented in blue, together with all vertices of Δ , correspond to the distributions in the intersection of the closures of the models

Back to the Toy Example: Gradient Fields



Let $f(\mathbf{x}) = x_1 + 2x_2 + 3x_1x_2$

Independence model: $\theta = (\theta_1, \theta_2, 0)$

$$p(\mathbf{x}; \theta) = \exp\{\theta_1 x_1 + \theta_2 x_2 - \psi(\theta)\}$$

Exponential family: $\theta = (0, \theta_2, \theta_{12})$

$$p(\mathbf{x}; \theta) = \exp\{\theta_2 x_2 + \theta_{12} x_1 x_2 - \psi(\theta)\}$$



The dashed lines represent the intersection of the two models in Δ

In the new model there are no critical points for $\tilde{\nabla} F$

Back to the Toy Example: Gradient Flows



Let $f(\mathbf{x}) = x_1 + 2x_2 + 3x_1x_2$

Independence model: $\theta = (\theta_1, \theta_2, 0)$

$$p(\mathbf{x}; \theta) = \exp\{\theta_1 x_1 + \theta_2 x_2 - \psi(\theta)\}$$

Exponential family: $\theta = (0, \theta_2, \theta_{12})$

$$p(\mathbf{x}; \theta) = \exp\{\theta_2 x_2 + \theta_{12} x_1 x_2 - \psi(\theta)\}$$



Vanilla gradient ∇F vs Natural gradient $\tilde{\nabla} F$

Linear Programming Relaxation

- A standard method to solve pseudo-Boolean optimization is **linearization** (see Boros and Hammer, 2001 for a survey)
- Every monomial x^α in f is replaced by a new binary variable z_α
- The problem is translated into a continuous linear problem by relaxing integer constraints for z , i.e.,

$$\begin{array}{ll}
 \text{(P)} \quad \arg \min \sum_{\alpha \in I} c_\alpha x^\alpha & \Rightarrow \quad \text{(LP)} \quad \arg \min \sum_{\alpha \in I} c_\alpha z_\alpha \\
 \text{s.t. } x_i \in \{+1, -1\} & \text{s.t. } z_\alpha \in [+1, -1]
 \end{array}$$

- The problem becomes linear, but additional constraints are required to link the $z = (z_\alpha), \alpha \in L$ variables to the $x = (x_1, \dots, x_n)$ variables

$$z_\alpha = \prod_{i=1}^n z_i^{\alpha_i} = \prod_{\{i:\alpha_i=1\}} z_i$$

Linear Programming Relaxation

- The new constraints are nonlinear, but can be modeled by linear inequalities (see Buchheim and Rinaldi, 2007, for a review)
- [Remark] For $x_i \in \{0, 1\}$ a standard linearization is given by

$$z_{\alpha} \leq z_i \quad \forall i : \alpha_i = 1 \quad z_{\alpha} \geq \sum_{i:\alpha_i=1} z_i - \|\alpha\|_1 + 1$$

- As n increases, the LP relaxation requires a good (tight) half-space representation of the marginal polytope M

Theorem 3 (M. et al., DISCML 2009)

The LP relaxation in the new variables z corresponds to the SR in η

$$F(\theta) = \mathbb{E}_{\theta} \left[\sum_{\alpha \in L} c_{\alpha} x^{\alpha} \right] = \sum_{\alpha \in L} c_{\alpha} \mathbb{E}_{\theta} [x^{\alpha}] \rightarrow F(\eta) = \sum_{\alpha \in L} c_{\alpha} \eta_{\alpha} = \sum_{\alpha \in L} c_{\alpha} z_{\alpha}$$

Solving the LP relaxation we obtain a lower bound for the SR since it is usually defined over an approximation of the marginal polytope M

Back to the Toy Example

The landscape of $F(\boldsymbol{\eta})$ changes according to f and the choice of \mathcal{M}

[Example] Natural gradient flows in the $\boldsymbol{\eta}$ are given by

$$\dot{\eta}_1 = (1 - \eta_1^2)(a_1 + a_{12}\eta_2)$$

$$\dot{\eta}_2 = (1 - \eta_2^2)(a_2 + a_{12}\eta_1)$$

We fix \mathcal{M} as the independence model and study the flows for different a_{12}

The natural gradient vanishes over

- the vertices of the marginal polytope M
- $\mathbf{c} = (-a_2/a_{12}, -a_1/a_{12})^T$

The nature of the critical points can be determined by studying the eigenvalues of the Hessian

$$M = \begin{bmatrix} -2\eta_1(a_1 + a_{12}\eta_2) & a_{12}(1 - \eta_1^2) \\ a_{12}(1 - \eta_2^2) & -2\eta_2(a_2 + a_{12}\eta_1) \end{bmatrix}$$

Back to the Toy Example: Critical Points

The solutions of the differential equations associated to the flows can be studied for every value of η , even outside of M , where points do not represent distributions, since positivity constraints are violated

Let $v \in \{-1, +1\}^2$ be a vertex of M , the eigenvalues of H are

$$\lambda_1 = -2v_1(a_{12}v_2 + a_1)$$

$$\lambda_2 = -2v_2(a_{12}v_1 + a_2)$$

According to the signs of λ_1 and λ_2 , each vertex can be either a stable node (SN), an unstable node (UN) or a saddle point (SP)

For $c = (-a_2/a_{12}, -a_1/a_{12})^T$

$$\lambda_{1,2} = \pm \sqrt{(a_{12}^2 - a_2^2)(a_{12}^2 - a_1^2)}/a_{12}^2$$

Follows that c is saddle point for

$(|a_{12}| \geq |a_1| \wedge |a_{12}| \geq |a_2|) \vee (|a_{12}| \leq |a_1| \wedge |a_{12}| \leq |a_2|)$, in the other cases, it is center (C)

Back to the Toy Example: Bifurcation Diagram


We can interpret $|a_{12}|$ as the strength of the interaction among x_1 and x_2

For $|a_{12}| \neq 0$, c is a saddle point in the shaded regions, where there exist

- strong interactions, $|a_{12}| > |a_1| \wedge |a_{12}| > |a_2|$, i.e. $c \in M$
- weak interactions, $|a_{12}| < |a_1| \wedge |a_{12}| < |a_2|$, i.e., $c \notin M$


In the remaining cases c is a center

Projection of the bifurcation diagram (η_1, η_2, a_{12}) over (η_1, η_2) for arbitrary a_1, a_2 and $0 \leq a_{12} < \infty$




The coordinates of c depends on a_{12} , c is a SP on the dashed lines and a C on the dotted line; for $a_{12} \rightarrow \infty$, c converges to the center of M

Natural Gradient Flows over (η_1, η_2) for fixed a_{12}




($a_{12} = 0$) 1 **SN**, 1 **UN** and 2 SPs

No critical points besides the vertices of M , all trajectories in M converge to the global optimum




($a_{12} = 0.85$) 1 **SN**, 1 **UN** and 3 SPs

The interaction is weak, c is a SP and is outside of M so that all flows converge to the global optimum

Natural Gradient Flows over (η_1, η_2) for fixed a_{12} 

($a_{12} = 1.25$) 1 SN, 1 UN, SPs and 1 C

The interaction is not strong enough to have $c \in M$ and to generate local minima, we have period solutions



($a_{12} = 5$) 2 SNs, 2 UNs and 1 SP

The interaction is strong, c is a SP and belongs to M , flows converge to either local or global optimum