

INPUT CONVEX KOLMOGOROV ARNOLD NETWORKS

THOMAS DESCHATRE AND XAVIER WARIN

ABSTRACT. This article presents an input convex neural network architecture using Kolmogorov-Arnold networks (ICKAN). Two specific networks are presented: the first is based on a low-order, linear-by-part, representation of functions, and a universal approximation theorem is provided. The second is based on cubic splines, for which only numerical results support convergence. We demonstrate on simple tests that these networks perform competitively with classical input convex neural networks (ICNNs). In a second part, we use the networks to solve some optimal transport problems needing a convex approximation of functions and demonstrate their effectiveness. Comparisons with ICNNs show that cubic ICKANs produce results similar to those of classical ICNNs.

1. INTRODUCTION

Recently, Kolmogorov-Arnold Networks (KANs) have been introduced as an alternative to multilayer perceptrons for high-dimensional function approximation, based on the Arnold-Kolmogorov representation theorem [18]. Arnold and Kolmogorov demonstrated [14] that a multivariate continuous smooth \mathbb{R} valued function f on a bounded domain can be expressed as a finite composition of sums of continuous single-variable functions. Specifically, if f is continuous on $[0, 1]^n$, then

$$(1) \quad f(x) = \sum_{i=1}^{2n+1} \psi_i \left(\sum_{j=1}^n \Phi_{i,j}(x_j) \right),$$

where $\Phi_{i,j} : [0, 1] \rightarrow \mathbb{R}$ and $\psi_i : \mathbb{R} \rightarrow \mathbb{R}$, $i = 1, \dots, 2n + 1$, $j = 1, \dots, n$. Since one-dimensional functions ψ or ϕ can be highly irregular or even fractal, it has been demonstrated that they may not be practically learnable [12, 23]. To address this issue, [18] propose extending this representation. Firstly, they suggest not limiting the outer sum in (1) to $2n + 1$ terms but to m terms and define a KAN l^{th} layer, $l = 1, \dots, L$, as an operator $\psi_{m,q}^l$ from $[0, 1]^m$ to \mathbb{R}^q :

$$(2) \quad (\psi_{m,q}^l(x))_k = \sum_{j=1}^m \Phi_{l,k,j}(x_j), \text{ for } k = 1, \dots, q.$$

Second, by stacking the layers, i.e., composing the operators ψ^l , $l = 0 \dots, L$, they define the KAN operator from $[0, 1]^m$ to \mathbb{R} :

$$(3) \quad K(x) = (\psi_{n_{L-1},d}^L \circ \psi_{n_{L-2},n_{L-1}}^{L-1} \circ \dots \circ \psi_{n_0,n_1}^1 \circ \psi_{m,n_0}^0)(x).$$

Since all ψ functions are one-dimensional, many classical methods are available to propose an easy-to-implement approximation. The original implementation proposed in [18] uses B-Splines but many other approximations based on wavelets [4], radial basis [17, 26], Chebyshev polynomials [25] can be used. All of these representations suffer from the same flaw: the output of a layer may not be in the grid initially chosen for the following layer. Sending back the output of a layer, a priori in \mathbb{R}^q to $[0, 1]^q$ using, for example, a

Date: March 2025.

sigmoid function is not a solution; it is clearly ineffective as the following layer discretizes $[0, 1]^q$ and only a small fraction of the compact is used. Aware of this problem, [18] proposes an adaptation technique, but this approach numerically fails.

Recently, using a P1 finite element flavour, [29] proposes a new architecture P1-KAN avoiding this flaw and proves its convergence : Numerical examples based on function approximation using a MSE criterion or real applications optimizing hydraulic valleys have shown that the P1-KAN is more effective than MLP and all tested KANs.

In this article, we extend the work of [29] and propose new architectures to approximate convex functions with a convex approximation. The approximation of a convex function while preserving its convexity property has been investigated in many articles. The approximation by cuts of a convex function using regression methods has been studied in [3, 10, 11] and [13], leading to max-affine approximations. A max-affine representation using group-max neural networks [28] has been proposed and proven to be convergent. The theoretical investigation of convex function approximation has also recently been explored, as demonstrated in [6], using a one-layer feedforward neural network with exponential activation functions in the inner layer and a logarithmic activation at the output. Numerically, this issue has been addressed in [2] through the development of the Input Convex Neural Network (ICNN) methodology, which is proven to be convergent in [7]. This effective approach has been widely applied in various fields requiring convex approximation, such as optimal transport problems [19, 15], optimal control problems [7, 1], inverse problems [22], and general optimization problems [8], among others.

In this article, we propose two Kolmogorov-Arnold Networks:

- The first version is based on a piecewise linear approximation of the 1D functions involved in the KAN approximation. It is based on the P1-KAN network recently developed in [29]. Different from other KAN networks, it inputs a grid G^1 defining the approximation domain and outputs a function value f approximating the data and a grid G^2 defining the image of the initial grid G^1 by f . This clear definition of a layer allow us to compose the layers easily. Imposing convexity in the one dimensional approximation, we develop a new network and we provide an Universal Approximation Theorem.
- The second one is based on the Hermite cubic-spline approximation where convexity is also imposed. This network enables us to provide a high-order approximation of a function, which can be a desired feature, for example, when the gradient of the convex approximation is the target. In this case a piecewise linear approximation provides a constant per mesh approximation which is no sufficient enough. However, no convergence property is given.

In Section 2, we describe the two networks belonging to the class of Input Convex Kolmogorov Networks (ICKAN), giving two variants, whether the grid is adapted or not, as in [29]. When the grid is adapted, we provide a universal approximation theorem. The different variations are tested using an MSE criterion, and we show that we can achieve results comparable to the best ones obtained by ICNNs.

In Section 3, we extend the previous networks to the case where the function to approximate is only partially convex with respect to the input, leading to Partial Input Convex Kolmogorov Networks (PICKAN). We again provide results comparing the efficiency of the PICKANs to the partial convex version of ICNNs (PICNN).

In Section 4, we use ICKANs to estimate the optimal transport map between two distributions μ and ν using the Wasserstein distance $W_2(\mu, \nu)$ on simulated data, showing the efficiency of the proposed networks.

2. ICKAN

We suppose here that the function to approximate is fully convex. In a first part we explain how to derive a one dimensional convex approximation of a function on $[0, 1]$ either using piecewise linear approximation

or a cubic-spline one. We show that adapting meshes can be an interesting feature on a numerical example. In a second part, we detail the different layers construction with the two networks and provide the Universal Approximation Theorem for the first network. In a third part, numerical examples are given.

2.1. Approximation of a 1D convex function f on $[0, 1]$.

We introduce the lattice $\{\hat{x}_0 := 0\} \cup \{\hat{x}_p\}_{1 \leq p \leq P-1} \cup \{\hat{x}_P := 1\}$ where the \hat{x}_p in $[0, 1]$ are increasing values with p .

2.1.1. *Piecewise linear approximation.* For $P > 0$, the degrees of freedom of a $P1$ approximation ϕ of f are the values of the function on the lattice leading to

$$(4) \quad \phi(x) = \sum_{p=0}^P a_p \Psi_p(x)$$

where $(a_p)_{p=0, P}$ are approximations of $(f(\hat{x}_p))_{p=0, \dots, P}$ and $(\Psi_p)_{p=0, \dots, P}$ is the basis of the shape functions: these functions have compact support in each interval $[\hat{x}_{p-1}, \hat{x}_{p+1}]$ for $p = 1, \dots, P-1$ and are defined as:

$$(5) \quad \Psi_p(x) = \begin{cases} \frac{x - \hat{x}_{p-1}}{\hat{x}_p - \hat{x}_{p-1}} & \text{for } x \in [\hat{x}_{p-1}, \hat{x}_p], \\ \frac{\hat{x}_{p+1} - x}{\hat{x}_{p+1} - \hat{x}_p} & \text{for } x \in [\hat{x}_p, \hat{x}_{p+1}], \end{cases}$$

for $p = 1, \dots, P-1$ and $\Psi_0(x) = \max(1 - \frac{x}{\hat{x}_1}, 0)$, $\Psi_P(x) = \max(\frac{x - \hat{x}_{P-1}}{1 - \hat{x}_{P-1}}, 0)$ (see figure 1).

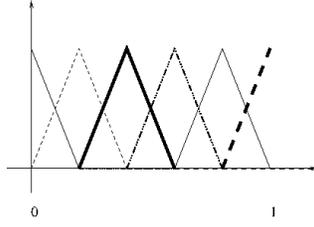


FIGURE 1. Uniform P1 basis functions on $[0, 1]$ with $P = 5$.

As we are interested in having a convex function, we want our piecewise linear approximation to have constant increasing derivatives on each mesh $[\hat{x}_p, \hat{x}_{p+1}]$ for $p = 0, \dots, P-1$. The trainable variables of the convex approximation are b , an approximation of $f(0)$, \hat{b} , an approximation of $f'(0)$, and $(d_i)_{i=1, \dots, P-1}$ an approximation of $f'(\hat{x}_i) - f'(\hat{x}_{i-1})$ for $i = 1, \dots, P$.

Therefore $a_p \approx f(\hat{x}_p)$ in (4) is given by:

$$(6) \quad a_p = b + \sum_{j=1}^p (\hat{b} + \sum_{i=1}^{j-1} \max(d_i, 0)) (\hat{x}_j - \hat{x}_{j-1}), p=0, \dots, P$$

where the max is here to ensure that the approximation $f'(\hat{x}_p) \approx \hat{b} + \sum_{i=1}^p \max(d_i, 0)$ is increasing with p . As an example, it is possible to minimize the mean square error of a one dimensional convex function with its approximation (4), (6) training $b, \hat{b}, (d_i)_{i=1, \dots, P-1}$ supposing that the lattice is regular so $\hat{x}_j - \hat{x}_{j-1} = h$ for $j = 1, \dots, P$. It is also possible to adapt the grid as in [29], and in this case, $b, \hat{b}, (d_i)_{i=1, \dots, P-1}$ and $(\hat{x}_i)_{i=1, \dots, P-1}$ are trained by the algorithm method.

Noting θ the list of the trained variables, our approximation \tilde{f}^θ of a function f is parametrized by θ . We minimize:

$$(7) \quad \mathbb{E}[(f(X) - \tilde{f}^\theta(X))^2]$$

where X is a random variable sampled uniformly on $[-10, 10]$. We give the results approximating the functions f_i , $i = 1, \dots, 4$ on figure 2 where

- (1) $f_1(x) = x^2$,
- (2) $f_2(x) = x^2 + 10[(e^x - 1)1_{x < 0} + x1_{x \geq 0}]$,
- (3) $f_3(x) = (|x|^2 + 1)^2$,
- (4) $f_4(x) = |x|1_{|x| \leq 3} + \frac{x^2 - 3}{2}$.

The plots are obtained adapting the grid with $P = 10$ or $P = 20$. We also plot the adapted vertices.

2.1.2. *The cubic approximation.* On a mesh $[\hat{x}_p, \hat{x}_{p+1}]$, the function f is approximated for $x \in [\hat{x}_p, \hat{x}_{p+1}]$ using a cubic Hermite spline by $\phi_p \left(\frac{x - \hat{x}_p}{\hat{x}_{p+1} - \hat{x}_p} \right)$ with:

$$\phi_p(t) = f(\hat{x}_p)h_{00}(t) + f'(\hat{x}_p)h_{10}(t)(\hat{x}_{p+1} - \hat{x}_p) + f(\hat{x}_{p+1})h_{01}(t) + f'(\hat{x}_{p+1})h_{11}(t)(\hat{x}_{p+1} - \hat{x}_p)$$

and

$$\begin{aligned} h_{00}(t) &= 2t^3 - 3t^2 + 1, \\ h_{10}(t) &= t^3 - 2t^2 + t, \\ h_{01}(t) &= -2t^3 + 3t^2, \\ h_{11}(t) &= t^3 - t^2. \end{aligned}$$

As previously we require that the derivative are increasing. Under this condition, in order to get a convex approximation, we can check that the values $\{f(\hat{x}_p)\}_{p=0, \dots, P}$ have to satisfy for $p = 0, \dots, P - 1$:

$$f(\hat{x}_p) + \frac{\hat{x}_{p+1} - \hat{x}_p}{3} (2f'(\hat{x}_p) + f'(\hat{x}_{p+1})) \leq f(\hat{x}_{p+1}) \leq f(\hat{x}_p) + \frac{\hat{x}_{p+1} - \hat{x}_p}{3} (f'(\hat{x}_p) + 2f'(\hat{x}_{p+1})).$$

Therefore the convex approximation based on Cubic spline is defined on mesh $[\hat{x}_p, \hat{x}_{p+1}]$ for $p = 0, \dots, P - 1$ by $\phi_p \left(\frac{x - \hat{x}_p}{\hat{x}_{p+1} - \hat{x}_p} \right)$ with

$$(8) \quad \phi_p(t) = a_p^0 h_{00}(t) + a_p^1 h_{10}(t)(\hat{x}_{p+1} - \hat{x}_p) + a_{p+1}^0 h_{01}(t) + a_{p+1}^1 h_{11}(t)(\hat{x}_{p+1} - \hat{x}_p),$$

where :

$$(9) \quad \begin{aligned} a_p^1 &= \hat{b} + \sum_{i=1}^p \max(d_i, 0), \\ a_p^0 &= b + \sum_{i=1}^p \frac{\hat{x}_i - \hat{x}_{i-1}}{3} (2a_{i-1}^1 + a_i^1 + \sigma(e_i)(a_i^1 - a_{i-1}^1)), \end{aligned}$$

where σ is the sigmoid function and $\hat{b}, b, \{d_i\}_{i=1, \dots, p}, \{e_i\}_{i=1, \dots, p}$ are trainable. In the adapted version, $\{\hat{x}_p\}_{p=1, \dots, P-1}$ are also trained. As for the piecewise linear approximation we provide on figure 3 the vertices and an estimation obtained adapting the grid with $P = 5$ or $P = 10$.

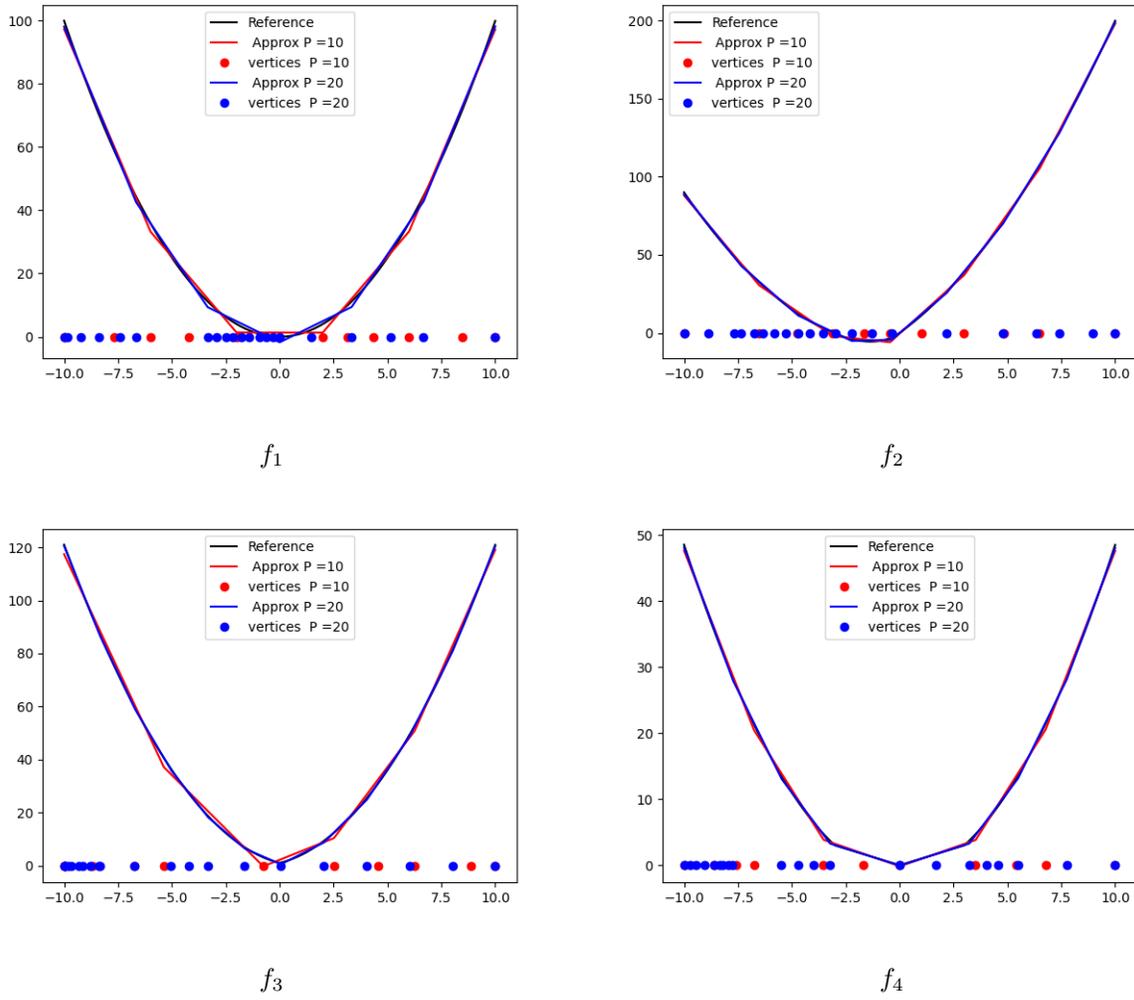


FIGURE 2. Piecewise linear approximation of a one dimensional function using (4), (6) with adaptation.

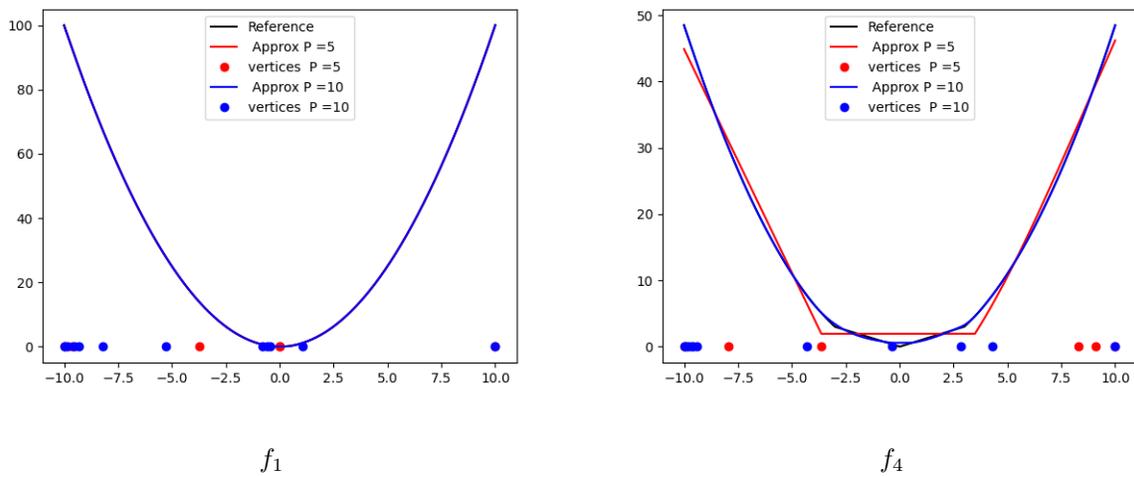


FIGURE 3. Cubic spline approximation of a one dimensional function using (8), (9) with adaptation.

2.2. The ICKAN layers. As for the P1-KAN layer, an ICKAN layer inputs an hypercube defining the domain and the x batched values, and it outputs an estimation of the function at x and an hypercube defined by the image of the initial hypercube. This allow us to concatenate the different layers in a coherent way.

2.2.1. The P1-ICKAN layers. As the first component of the first layer, we define $\hat{\kappa}_{n,m}^0$ for $x \in \mathbb{R}^n$ and the hypercube $G^0 := [0, 1]^n$ as:

(10)

$$\hat{\kappa}_{n,m}^0(x, G^0)_k = \sum_{j=1}^n \sum_{p=0}^P \left(b_{0,k,j} + \sum_{s=1}^p (\hat{b}_{0,k,j} + \sum_{i=1}^{s-1} \max(d_{0,k,j,i}, 0)) (\hat{x}_{0,j,s} - \hat{x}_{0,j,s-1}) \right) \Psi_p^{0,j}(x_j), \text{ for } k = 1, \dots, m,$$

where the $(\hat{x}_{0,j,s})_{s=0,\dots,P}$ correspond the one dimensional lattice of the dimension j of G^0 and the $\Psi_p^{0,j}$ the corresponding hat functions defined in Equation (5) and displayed in figure 1.

The image $G^1 = \prod_{k=1}^m [\underline{G}_k^1, \bar{G}_k^1]$ of G^0 by $\hat{\kappa}_{n,m}^0$ is exactly given by:

$$(11) \quad \begin{aligned} \underline{G}_k^1 &= \sum_{j=1}^n \min_{0 \leq p \leq P} \left(b_{0,k,j} + \sum_{s=1}^p (\hat{b}_{0,k,j} + \sum_{i=1}^{s-1} \max(d_{0,k,j,i}, 0)) (\hat{x}_{0,j,s} - \hat{x}_{0,j,s-1}) \right) \\ \bar{G}_k^1 &= \sum_{j=1}^{d_0} \max \left[b_{0,k,j}, b_{0,k,j} + \sum_{s=1}^P (\hat{b}_{0,k,j} + \sum_{i=1}^{s-1} \max(d_{0,k,j,i}, 0)) (\hat{x}_{0,j,s} - \hat{x}_{0,j,s-1}) \right]. \end{aligned}$$

for $1 \leq k \leq m$, and the layer is defined as

$$(12) \quad \kappa_{n,m}^0(x, G^0) = (\hat{\kappa}_{n,m}^0(x, G^0), G^1).$$

Remark 2.1. *As the approximation is convex, we have used the fact that the maximum \bar{G}_k^l is reached on the boundary of the domain.*

As a sum of convex 1D function in each dimension, the Hessian of $\hat{\kappa}_{n,m}^0(\cdot, G^0)$ is diagonal and definite positive.

As the first layer is convex in x , it remains to assure that it remains convex in x by composition of a layer $\hat{\kappa}_{m,q}^l(\cdot, G^l)$. As in [2], it is sufficient to have for $l \geq 1$, $\hat{\kappa}_{m,q}^l(x, G^l)_k$ convex and increasing in x and naturally we only have to impose that $\hat{b}_{l,k,j} \geq 0$ for all k and j . Therefore for $l \geq 1$, G^l being the hypercube defining the domain, we first define

$$\hat{\kappa}_{m,q}^l(x, G^l)_k = \sum_{j=1}^m \sum_{p=0}^P \left(b_{l,k,j} + \sum_{s=1}^p (\max(\hat{b}_{l,k,j}, 0) + \sum_{i=1}^{s-1} \max(d_{l,k,j,i}, 0)) (\hat{x}_{l,j,s} - \hat{x}_{l,j,s-1}) \right) \Psi_p^{l,j}(x_j),$$

for $k = 1, \dots, q$,

and the output hypercube G^{l+1} image of G^l is given by an expression similar to (11). Then the layer is defined by:

$$(13) \quad \kappa_{m,q}^l(x, G^l) = (\hat{\kappa}_{m,q}^l(x, G^l), G^{l+1}).$$

Concatenating the L layers, noting n_l the number of neurons of layer l (see (3)), supposing the number of mesh P constant, we define our global network as

$$(14) \quad K(x) = \hat{\kappa}_{n_{L-1},1}^L \circ \kappa_{n_{L-2},n_{L-1}}^{L-1} \circ \dots \circ \kappa_{n,n_0}^0(x, [0, 1]^n).$$

As in [29], we propose two versions :

- A first version, setting the lattice using an uniform meshing and the trainable variables are

$$\begin{aligned} \mathcal{A} := & (b_{0,k,j}, \hat{b}_{0,k,j}, (d_{0,k,j,i})_{i=1,\dots,P})_{k=1,\dots,n_0, j=0,\dots,n} \cup \\ & (b_{l,k,j}, \hat{b}_{l,k,j}, (d_{l,k,j,i})_{i=1,\dots,P})_{l=1,\dots,L-1, k=1,\dots,n_l, j=0,\dots,n_{l-1}} \cup \\ & (b_{L,1,j}, \hat{b}_{L,1,j}, (d_{L,1,j,i})_{i=1,\dots,P})_{j=0,\dots,n_{L-1}} \end{aligned}$$

- A second version also trains the lattice values inside the domain:

$$\mathcal{A} \cup (\hat{x}_{0,j,p})_{j=1,\dots,n, p=1,\dots,P-1} \cup (\hat{x}_{l,j,p})_{l=1,\dots,L, j=1,\dots,n_l, p=1,\dots,P-1}.$$

2.2.2. *The Cubic-ICKAN layers.* We define $\hat{\kappa}_{n,m}^0$ for $x \in \prod_{j=1}^n [\hat{x}_{0,j,p_j}, \hat{x}_{0,j,p_j+1}]$ where $p_j \in \{0, \dots, P-1\}$ for $j = 1, \dots, n$ and an hypercube $G^0 := [0, 1]^n$ by

$$(15) \quad \hat{\kappa}_{n,m}^0(x, G^0)_k = \sum_{j=1}^n a_{0,j,p_j}^0 h_{00}(t_{0,j}) + a_{0,j,p_j}^1 h_{10}(t_{0,j})(\hat{x}_{0,j,p_j+1} - \hat{x}_{0,j,p_j}) + \\ a_{0,j,p_j+1}^0 h_{01}(t_{0,j}) + a_{0,j,p_j+1}^1 h_{11}(t_{0,j})(\hat{x}_{0,j,p_j+1} - \hat{x}_{0,j,p_j}), \text{ for } k = 1, \dots, m,$$

where $t_{0,j} = \frac{x_j - \hat{x}_{0,j,p_j}}{\hat{x}_{0,j,p_j+1} - \hat{x}_{0,j,p_j}}$,

$$(16) \quad a_{0,j,p_j}^1 = \hat{b}_{0,j} + \sum_{i=1}^{p_j} \max(d_{0,j,i}, 0), \\ a_{0,j,p_j}^0 = b_{0,j} + \sum_{i=1}^{p_j} \frac{\hat{x}_{0,j,i} - \hat{x}_{0,j,i-1}}{3} (2a_{0,j,i}^1 + a_{0,j,i-1}^1 + \sigma(e_{0,j,i})(a_{0,j,i}^1 - a_{0,j,i-1}^1))$$

where the $(\hat{x}_{0,j,s})_{s=0,\dots,P}$ correspond the one dimensional lattice on G^0 of input dimension j .

$G^1 = \prod_{k=1}^m [\underline{G}_k^1, \bar{G}_k^1]$ defined with

$$(17) \quad \underline{G}_k^1 = \sum_{j=1}^n \min_{0 \leq p \leq P} a_{0,j,p}^0 \\ \bar{G}_k^1 = \sum_{j=1}^{d_0} \max [a_{0,j,0}^0, a_{0,j,P}^0],$$

is now included in the image of G^0 . In order that G^1 is exactly the image of G^0 , we can truncate the output of the layer $\hat{\kappa}$ and :

$$\hat{\kappa}_{n,m}^0(x, G^0) = ((\hat{\kappa}_{n,m}^0(x, G^0)_k \vee \underline{G}_k^1)_{k=1,\dots,m}, G^1).$$

Remark 2.2. *The truncation is not compulsory but it helps defining clearly a consistent layer. It potentially clips the approximation once per spline. By convexity only the min value has to be truncated.*

Similarly to the piecewise linear case, for $l \geq 1$, G^l being the hypercube defining the domain, the operator $\hat{\kappa}_{m,q}^l(x, G^l)_k$ for $l > 1$ is defined similarly to (15) using

$$a_{l,j,p_j}^1 = \max(\hat{b}_{l,j}, 0) + \sum_{i=1}^{p_j} \max(d_{l,j,i}, 0)$$

instead of (16).

The output hypercube G^{l+1} image of G^l is given by an expression similar to (17). Then the layer is defined by:

$$\kappa_{m,q}^l(x, G^l) = ((\hat{\kappa}_{m,q}^l(x, G^l)_k \vee \underline{G}_k^{l+1})_{k=1,\dots,q}, G^{l+1}).$$

The concatenation of the layers is still given by (14). As for the piecewise linear case, we can either use an uniform meshing or adapting one.

2.2.3. Convergence results. We provide two universal approximation theorem in the piecewise linear case for the two different cases.

Theorem 2.1. *The space spanned by the P1-ICKAN letting n_l for $l = 0, \dots, L - 1$ and L vary for $P > 1$ is dense in set of Lipschitz convex functions on $[0, 1]^n$ with the sup norm when adaptation is used.*

The proof of theorem 2.1 is given in appendix A. We give at last a theorem for the non adapted case.

Theorem 2.2. *The space spanned by the P1-ICKAN letting n_l for $l = 0, \dots, L - 1$, L and P vary is dense in set of Lipschitz convex functions on $[0, 1]^n$ with the sup norm when no adaptation is used.*

Proof. The idea of the demonstration is the same. The only difference is that as one uses an uniform grid, 0 is generally not in the lattice generated by the $(\hat{x}_{i,p}^l)_{p=0,\dots,P}$ and the ReLU function cannot be exactly generated but has to be approximated. In order to control the approximation error uniformly on a compact, one has to take P large. This explains why it is impossible to bound P while controlling the error. \square

2.2.4. Numerical Results. We estimate the function

$$(18) \quad f(x) = \sum_{i=1}^d (|x_i| + |1 - x_i|) + x^\top Ax$$

using an Input Convex Neural Network with a ReLU activation function and an Input Convex Kolmogorov Network approximation \tilde{f}^θ parametrized by θ by minimizing (7) where $X \sim U[-2, 2]^d$. We use ADAM optimizer with a learning rate equal to 10^{-3} , a batch size equal to 1000 and use 200000 iterations. Using 10 runs, the average MSE with a batch of 100000 and its standard deviation is calculated with different parametrization. For the ICNN, we use 2,3, 4 or 5 layers, a number of neurons in $\{10, 20, 40, 80, 160, 320\}$ and the ReLU activation function. Testing all configurations our ICNN reference is the one with the smallest averaged MSE. Results are given in tables 1 and 2 in dimensions 3 and 7. Results with the P1-ICKAN are stable with the number of layers, and P . Results with adaptation (P1-ICKAN adapt) are better than without adaptation (P1-ICKAN no adapt). The P1-ICKAN with adaptation gives results similar to the best ICNN, while using a smaller network.

method	nb Layers	nb neurons	P	Average	std
Best ICNN	2	320		9.37E-05	8.29E-05
P1-ICKAN no adapt	2	20	20	1.56E-04	7.39E-05
P1-ICKAN no adapt	2	20	40	1.14E-04	9.79E-05
P1-ICKAN no adapt	2	40	20	2.62E-04	3.68E-04
P1-ICKAN no adapt	2	40	40	2.69E-04	4.65E-04
P1-ICKAN no adapt	3	20	20	3.13E-04	3.10E-04
P1-ICKAN no adapt	3	20	40	6.83E-04	1.04E-03
P1-ICKAN no adapt	3	40	20	2.76E-04	1.90E-04
P1-ICKAN no adapt	3	40	40	1.67E-04	2.91E-04
P1-ICKAN adapt	2	20	20	1.23E-04	4.62E-05
P1-ICKAN adapt	2	20	40	1.93E-04	3.72E-04
P1-ICKAN adapt	2	40	20	2.62E-04	2.97E-04
P1-ICKAN adapt	2	40	40	1.13E-04	1.42E-04
P1-ICKAN adapt	3	20	20	9.82E-04	2.23E-03
P1-ICKAN adapt	3	20	40	1.93E-04	2.75E-04
P1-ICKAN adapt	3	40	20	3.63E-04	6.50E-04
P1-ICKAN adapt	3	40	40	5.35E-04	1.16E-03

TABLE 1. Results on 10 runs on the minimization of (7) in dimension 3 with function (18)
: ICCN versus P1-ICKAN.

method	nb Layers	nb neurons	P	Average	std
Best ICNN	2	320		2.39E-03	6.72E-04
P1-ICKAN no adapt	2	40	10	8.19E-03	9.23E-04
P1-ICKAN no adapt	2	40	20	3.71E-03	4.69E-03
P1-ICKAN no adapt	2	40	40	1.94E-03	1.20E-03
P1-ICKAN no adapt	3	40	10	9.05E-03	1.25E-03
P1-ICKAN no adapt	3	40	20	2.40E-03	8.88E-04
P1-ICKAN no adapt	3	40	40	3.45E-03	2.09E-03
P1-ICKAN adapt	2	40	10	2.31E-03	6.88E-04
P1-ICKAN adapt	2	40	20	2.72E-03	2.95E-03
P1-ICKAN adapt	2	40	40	1.33E-03	1.03E-03
P1-ICKAN adapt	3	40	10	2.69E-03	1.08E-03
P1-ICKAN adapt	3	40	20	2.87E-03	1.89E-03
P1-ICKAN adapt	3	40	40	2.24E-03	2.47E-03

TABLE 2. Results on 10 runs on the minimization of (7) in dimension 7 with function (18)
: ICCN versus P1-ICKAN.

As for computing time, on a GPU Tesla H100, in dimension 7, the best ICKAN (2 layers of 40 neurons, $P = 40$) take 4.3 times more than the best ICNN configuration (2 layers, 320 neurons). In table 3, results in dimension 7 using the Cubic-ICKAN show that the approximation is good:

method	nb Layers	nb neurons	P	Average	std
Best ICNN	2	320		2.39E-03	6.72E-04
Cubic-ICKAN no adapt	2	20	10	6.66E-03	2.06E-03
Cubic-ICKAN no adapt	2	20	20	3.59E-03	7.64E-04
Cubic-ICKAN no adapt	2	40	10	3.88E-03	1.06E-03
Cubic-ICKAN no adapt	3	20	10	6.19E-03	2.11E-03
Cubic-ICKAN no adapt	3	20	20	5.04E-03	2.81E-03
Cubic-ICKAN no adapt	3	40	10	3.82E-03	1.64E-03
Cubic-ICKAN adapt	2	20	10	1.43E-03	6.29E-04
Cubic-ICKAN adapt	2	20	20	1.19E-03	8.76E-04
Cubic-ICKAN adapt	2	40	10	1.38E-03	1.28E-03
Cubic-ICKAN adapt	3	20	10	3.06E-03	1.66E-03
Cubic-ICKAN adapt	3	20	20	1.95E-03	1.43E-03
Cubic-ICKAN adapt	3	40	10	1.49E-03	1.11E-03

TABLE 3. Results on 10 runs on the minimization of (7) in dimension 7 with function (18): ICNN versus Cubic-ICKAN.

3. PICKAN

We suppose in this section that the function f to approximate is partially convex : $f(x, y)$ is convex in $y \in \mathbb{R}^{n_y}$ but not convex in $x \in \mathbb{R}^{n_x}$. We suppose that the function is defined on $G^0 = G_x^0 \times G_y^0$ where $G_x^0 = [0, 1]^{n_x}$, $G_y^0 = [0, 1]^{n_y}$. We only present the network using a piecewise linear approximation of a function but the network can be adapted to a Cubic approximation. We note P1-KANL the KAN layer defined in [29] and its corresponding operator $\rho_{p,q}^l$ for layer l , input dimension p , output dimension q . The first layer κ^0 defined by equation (12) is noted ICKANL0, while a layer κ^l with $l > 0$ defined by equation (13) is noted ICKANL1. The structure of the network is given on figure 4.

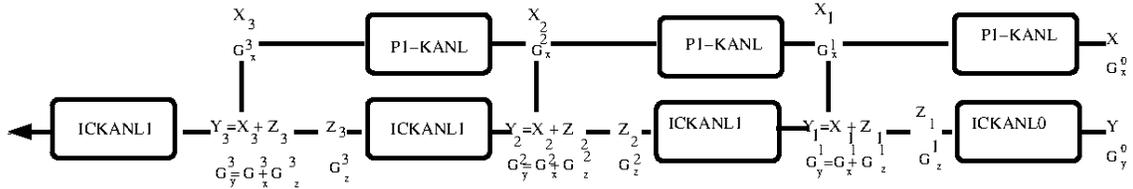


FIGURE 4. Partial Input Convex Kolmogorov Arnold Network using a piecewise linear approximation.

The recursion described in figure 4 is given for $M > 0$ by:

$$\begin{aligned}
(X_1, G_x^1) &= \rho_{n_x, M}^0(x, G_x^0), \\
(Y_1, G_y^1) &= (X_1, G_x^1) + \kappa_{n_y, M}^0(y, G_y^0), \\
(X_{i+1}, G_x^{i+1}) &= \rho_{M, M}^{i+1}(X_i, G_x^i), \\
(Y_{i+1}, G_y^{i+1}) &= (X_{i+1}, G_x^{i+1}) + \kappa_{M, M}^{i+1}(Y_i, G_y^i) \quad \text{for } i = 0, L-2, \\
(Y_L, G_y^L) &= \kappa_{M, 1}^L(Y_{L-1}, G_y^{L-1}).
\end{aligned}$$

Remark 3.1. Notice that when f is a non convex function imposing that $\kappa^l = 0$ for $l < L$ and $\kappa^L(x) = x_1$ gives back the P1-KAN network which has a clear convergence error when the Kolmogorov Arnold basis functions are Lipschitz.

When the function f is convex, nullifying the output of the P1-KANL layers gives back the ICKAN neural network.

Remark 3.2. In the PICKAN, we add the results from the convex part and the non convex part as done in the Kolmogorov Arnold layer (2).

Remark 3.3. Depending on the regularity of the function in y and x , it is possible the use a scheme in x with an approximation order different from the one used in y .

We estimate the function

$$(19) \quad f(x, y) = |y + 1||x + 2x^3| \text{ on } [-2, 2] \times [2, 2],$$

minimizing

$$\mathbb{E}[(f(X, Y) - \tilde{f}^\theta(X, Y))^2],$$

where \tilde{f}^θ is our neural network approximation parametrized by θ and $(X \times Y) \sim U[-2, 2]^2$. As before we use the ADAM descent method using 200000 iterations and a learning rate equal to 10^{-3} . For the ICNN with a ReLU activation function, we use 2,3, or 4 layers, a number of neurons in $\{10, 20, 40, 80, 160\}$. As before we compute for each configuration the averaged loss and the standard deviation obtained with 10 runs. Testing all configurations our PICNN (Partial Input Convex Neural Network in [2]) reference is the one with the smallest averaged MSE. Results are given in tables 4 and show that the accuracy is again similar to the best one obtained with an ICNN.

method	nb Layers	nb neurons	P	Average	std
Best PICNN	2	80		1.04E-03	7.37E-04
PP1-ICKAN no adapt	2	20	20	4.98E-03	1.10E-02
PP1-ICKAN no adapt	2	20	40	5.17E-03	4.87E-03
PP1-ICKAN no adapt	2	40	20	2.57E-03	3.29E-03
PP1-ICKAN no adapt	2	40	40	9.23E-04	5.66E-04
PP1-ICKAN no adapt	3	20	20	2.10E-03	3.55E-03
PP1-ICKAN no adapt	3	20	40	3.85E-03	3.69E-03
PP1-ICKAN no adapt	3	40	20	3.75E-03	5.97E-03
PP1-ICKAN no adapt	3	40	40	9.34E-03	1.81E-02
PP1-ICKAN adapt	2	20	20	2.17E-03	2.58E-03
PP1-ICKAN adapt	2	20	40	2.67E-03	4.25E-03
PP1-ICKAN adapt	2	40	20	1.02E-03	1.03E-03
PP1-ICKAN adapt	2	40	40	1.27E-03	1.23E-03
PP1-ICKAN adapt	3	20	20	2.71E-03	3.98E-03
PP1-ICKAN adapt	3	20	40	1.10E-02	1.98E-02
PP1-ICKAN adapt	3	40	20	8.11E-03	1.65E-02
PP1-ICKAN adapt	3	40	40	5.79E-03	5.07E-03

TABLE 4. Testing partial convexity for function (19) : average and standard deviation for 10 runs.

4. APPLICATION FOR OPTIMAL TRANSPORT

Neural networks preserving convexity can be used to solve optimal transport problems, see for example [19, 16] which use ICNN. In this section we use ICKAN models to solve optimal transport problems and compare their performance with ICNN.

4.1. Monge optimal transport and Brenier's theorem. Let us consider two probability measures with support $\Omega \subset \mathbb{R}^d$, μ and ν , such that $\int_{\Omega} \|x\|^2 d\mu(x)$, $\int_{\Omega} \|x\|^2 d\nu(x) < \infty$. Our goal is to estimate the optimal transport map between X and Y , $T : \Omega \rightarrow \Omega$, which is a solution to the Monge problem [21]

$$(20) \quad \inf_{T \# \mu = \nu} \int_{\Omega} \|x - T(x)\|^2 d\mu(x),$$

where $T \# P(\cdot) := P(T^{-1}(\cdot))$ for a probability measure P and $\|\cdot\|$ denotes the Euclidean norm on \mathbb{R}^d . The Brenier's theorem [5] guarantees the existence and uniqueness of the optimization problem (20) as soon as the probability measure μ is absolutely continuous with respect to the \mathbb{R}^d Lebesgue measure, that we assume in the following. Furthermore, the optimal transport map T solution of (20) writes $T = \nabla f$ with $f : \mathbb{R}^d \rightarrow \mathbb{R}$ a convex function differentiable almost everywhere. Furthermore, denoting $CVX(\mu)$ the set of all convex functions in $L^1(\mu)$, f is the solution of the optimization problem

$$\inf_{\varphi \in CVX(\mu)} \int_{\Omega} \varphi d\mu + \int_{\Omega} \varphi^* d\nu,$$

with $\varphi^*(y) = \sup_{x \in \Omega} \{\langle x, y \rangle - f(x)\}$ for $y \in \Omega$ the Legendre-Fenchel transform of φ , see [20, Theorem 1]. f is called the Brenier's potential. In [19], it is proven whenever ν admits a density, there exists an optimal pair (φ_0, ψ_0) solution of

$$(21) \quad \sup_{\substack{\varphi \in CVX(\mu) \\ \varphi^* \in L^1(\nu)}} \inf_{\psi \in CVX(\nu)} - \int \varphi(x) d\mu(x) - \int (\langle y, \nabla \psi(y) \rangle - \varphi(\nabla \psi(y))) d\nu(y),$$

where $\nabla \psi_0 = T^{-1}$.

4.2. Algorithm. Given samples $X_1, \dots, X_n \sim \mu$ and $Y_1, \dots, Y_n \sim \nu$ both identically and independently distributed (i.i.d.), the objective is to estimate the optimal transport map between a μ distribution and a ν distribution. We denote by $\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and $\hat{\nu}_n = \frac{1}{n} \sum_{i=1}^n \delta_{Y_i}$ the empirical distribution of X and Y respectively. We use the formulation (21) to solve the minimax problem associated parameterizing φ_{Θ} by $\Theta \in \mathbb{R}^l$ and ψ_{θ} by $\theta \in \mathbb{R}^l$ using an ICKAN or an ICNN where l is the number of parameters of the networks. The empirical version of (21) solves

$$\max_{\Theta} \min_{\theta} \frac{1}{n} \sum_i \varphi_{\Theta}(\nabla \psi_{\theta}(Y_i)) - \langle Y_i, \nabla \psi_{\theta}(Y_i) \rangle - \varphi_{\Theta}(X_i)$$

using the classical minimax algorithm 1.

Algorithm 1 Minimax algorithm for transport problem

Require: Batch size n , external iteration I_{ext} , internal iteration I_{int} .

for $i = 1, \dots, I_{ext}$ **do**
 Sample $X_1, \dots, X_n \sim \mu$
 for $j = 1, \dots, I_{int}$ **do**
 Sample $Y_1, \dots, Y_n \sim \nu$
 Update θ minimizing (21) using Adam method
 end for
 Update Θ maximizing (21) using Adam method
end for

Note that since the ICKAN networks are used on a compact set, and since we cannot control the values taken by $\nabla\psi(y)$ during the resolution, the ψ_θ network linearly extrapolates the function outside its domain of definition.

4.3. Numerical results on synthetic data. Since it is often difficult to get the transport map for a pair (ν, μ) , we define two examples using μ and T , and we consider $\nu = T\#\mu$. The metric we consider for the error between the optimal transport map T^* and the estimated one \hat{T} is the percentage of unexplained variance [16] $UVP (\%) = 100 \frac{\int_{\Omega} \|T^* - \hat{T}\|^2 d\hat{\mu}_n}{\int_{\Omega} \|x\|^2 d\hat{\nu}_n - \int_{\Omega} x d\hat{\nu}_n \|^2}$ on a validation set X_1, \dots, X_n and Y_1, \dots, Y_n (an error of 100% corresponds to the constant map $T^C(x) = \int_{\Omega} x d\hat{\nu}_n(x)$). As in [16], we choose as benchmark the linear transport map estimated on the validation set defined by $\hat{T}^L(x) = \hat{A}(x - \hat{m}_1) + \hat{m}_2$ where $\hat{m}_1 = \int_{\Omega} x d\hat{\mu}_n(x)$, $\hat{m}_2 = \int_{\Omega} x d\hat{\nu}_n(x)$, $\hat{A} = \hat{\Sigma}_1^{-1/2} \left(\hat{\Sigma}_1^{1/2} \hat{\Sigma}_2 \hat{\Sigma}_1^{1/2} \right)^{1/2} \hat{\Sigma}_1^{-1/2}$, $\hat{\Sigma}_1 = \int_{\Omega} x x^\top d\hat{\mu}_n(x) - \hat{m}_1 \hat{m}_1^\top$, $\hat{\Sigma}_2 = \int_{\Omega} x x^\top d\hat{\nu}_n(x) - \hat{m}_2 \hat{m}_2^\top$, see [9]. We consider for the optimization the same hyperparameters than [16], that is

- $I_{ext} = 50000$ iterations and $I_{int} = 15$ iterations ;
- at each iteration, we simulate a new batch of size 1024 ;
- the learning rate is equal to 0.001 ;
- every 100 iterations, we evaluate the error UVP on a test set consisting of 4096 data and we keep the network with the lowest error ;
- the validation dataset has size 2^{14} .

For the ICKAN networks, the domain is defined from the minimum and maximum values from a first dataset with 2^{14} data. This framework is idealised, as in practice we by no means have this amount of data and we do not have access to the optimal map, i.e. the one we need to estimate in the test step. We also initialise the networks with θ_0 minimising $\|\nabla f_\theta(x) - x\|^2$ so that the map is closed to identity at initialisation.

Transport map of [16]. We consider the transport map defined in the work of Korotin et al. [16]: the authors consider the optimal transport map T_1 between a mixture of 3 Gaussian random variables μ and a mixture of 10 Gaussian random variables and the optimal transport map T_2 between the same mixture of 3 Gaussian random variables and another mixture of 10 Gaussian random variables. They can then define the optimal transport map $\frac{1}{2}(T_1 + T_2)$. Since T_1 and T_2 are not explicitly known, they learn T_1 and T_2 with an ICNN, \hat{T}_1 and \hat{T}_2 respectively, that solves the optimal transport problem. The target distribution considered by Korotin et al. [16] and by us is then $\frac{1}{2}(\hat{T}_1 + \hat{T}_2)\#\mu$. We consider for the ICNN¹ 3 layers

¹We use the implementation at <https://github.com/iamaalexkorotin/Wasserstein2Benchmark/>, which uses a CELU activation function.

with 64, 64 and 32 neurons (same parametrisation as in [16]) and for the ICKANs a network with 2 layers with 64 and 32 neurons or with 10 and 5 neurons. The results are given in Table 5 for $d \in \{2, 4, 8, 16, 32\}$. Our optimization framework and network parametrization correspond to the [MMv2] case in [16]. For the Cubic ICKAN with adapted mesh and with 64 and 32 neurons, the errors are much smaller than for the linear map, but can be slightly larger than those obtained by the ICNN parametrization in some case, with similar orders of magnitude. As mentioned in [16], the optimal transport map itself is learned with an ICNN, which can be advantageous for the ICNN parametrization. Furthermore, we did not search for the optimal parametrization of the Cubic ICKAN (nor for the ICNN). The number of meshes $P \in \{10, 20, 40\}$ does not have much influence on the results, as the use of an unadapted grid does. However, using a linear mesh instead of a cubic one significantly degrades the results, but the network still outperforms the linear map. Using a smaller network of 10 and 5 neurons for the Cubic ICKAN results in larger errors, and this difference increases with dimension. The obtained distributions are displayed in Figure 5 for the case $d = 2$ and the Cubic ICKAN with $P = 10$.

Method	Neurons / Dim	2	4	8	16	32
Linear		13.93	14.96	27.29	42.05	55.48
Cubic ICKAN P=10	64 32	0.06	0.58	3.00	7.16	9.89
Cubic ICKAN P=20	64 32	0.05	0.51	2.44	5.97	7.66
Cubic ICKAN P=40	64 32	0.05	0.52	2.82	6.57	5.90
Cubic ICKAN P=40 No adapt	64 32	0.06	0.65	3.66	6.63	10.98
Linear ICKAN P=40	64 32	1.01	6.62	19.99	19.90	27.95
Cubic ICKAN P=40	10 5	0.16	1.84	11.56	47.25	101.83
ICNN	64 64 32	0.07	0.27	0.74	1.98	3.01

TABLE 5. Percentage of unexplained variance UVP (%) for the linear map, the map parametrized by the Cubic ICKAN with adapted mesh and $P \in \{10, 20, 40\}$, the Cubic ICKAN with non adapted mesh and $P = 40$, the linear ICKAN with adapted mesh and $P = 40$ and the ICNN map when the true map is the one in [16]. For the Cubic ICKAN, we consider networks with 2 layers and 64 and 32 neurons or 10 and 5 neurons while for the ICKAN we consider 3 layers with 64, 64, and 32 neurons.

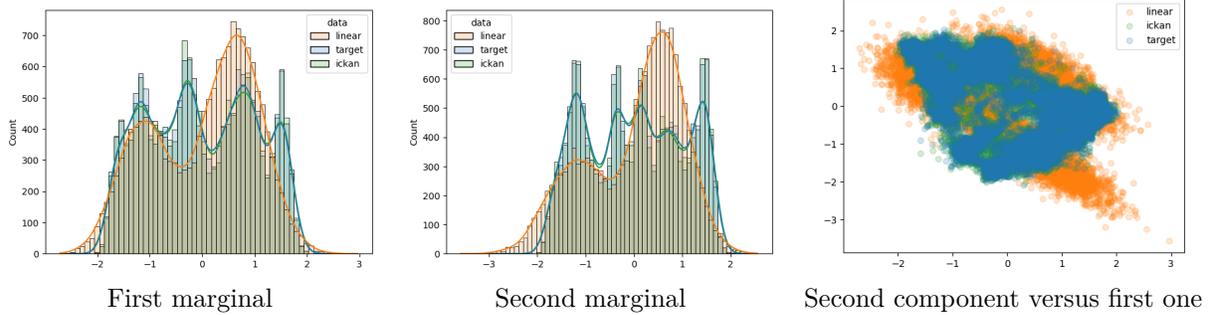


FIGURE 5. Distribution of the true target distribution as well as the one obtained by linear transport or Cubic ICKAN transport with adapted mesh, $P = 10$, and 64 and 32 neurons for the map in [16] and $d = 2$. The first two figures include the empirical histogram as well as a Gaussian kernel density estimator with bandwidth selected using Scott’s rule [24].

Tensorized case [27]. We consider the example in Vacher et al. [27] where the transport map is defined by $T(x) = (T_i(x_i))_{i=1,\dots,d}$ with $T_i(x_i) = x_i + \frac{1}{6 - \cos(2\pi x_i)} - 0.2$, $i = 1, \dots, d$,² for $x = (x_i)_{i=1,\dots,d} \in [0, 1]^d$. μ is a uniform law on $[0, 1]^d$. We consider a 3-layers ICNN with 64, 64, 32 neurons and a 2-layers Cubic ICKAN with $\max(2d, 10)$ and $\max(d, 5)$ neurons, an adapted mesh and $P \in \{10, 20, 40\}$. The results are given in Table 6 for $d \in \{1, 2, 4, 8\}$. The ICNN performs very poorly for $d \in \{4, 8\}$, giving errors of the same order of magnitude as the linear transport map. The Cubic ICKAN performs very well for all values of P . The obtained distributions are displayed in Figure 6 for the case $d = 2$ case and the Cubic ICKAN with $P = 10$. We also display Figure 7 the first component of $\hat{T}(x, 0.5)$ (0.5 is chosen arbitrarily, the first component of $x_2 \rightarrow \hat{T}(x, x_2)$ being constant and equal to $\hat{T}_1(x)$) and the second component of $\hat{T}(0.5, x)$. We get similar figures for larger dimensions. The neural network reproduces the target distribution and the transport map very well. The better performance of ICKAN over ICNN is probably due to the structure of the Brenier map, which is of the form $f(x) = \sum_{i=1}^d f_i(x_i)$ with $f_i(x_i) = \int_0^{x_i} T_i(s) ds$, $i = 1, \dots, d$: the functions in the Arnold-Kolmogorov representation (3) then have a high degree of smoothness which can lead to a faster rate of convergence, see Proposition 2.2 and 2.3 in [29] in the non-convex case.

Method / Dim	1	2	4	8
Linear	0.49	0.54	0.52	0.53
Cubic ICKAN P=10	0.00	0.01	0.01	0.02
Cubic ICKAN P=20	0.00	0.02	0.02	0.02
Cubic ICKAN P=40	0.01	0.02	0.02	0.02
ICNN	0.04	0.05	0.34	0.51

TABLE 6. Percentage of unexplained variance UVP (%) for the linear map, the map parametrized by the Cubic ICKAN with adapted mesh and $P \in \{10, 20, 40\}$, and the ICNN map and different dimensions when the true map is $T(x) = (T_i(x_i))_{i=1,\dots,d}$ with $T_i(x_i) = x_i + \frac{1}{6 - \cos(2\pi x_i)} - 0.2$, $i = 1, \dots, d$.

²The example considered is the one from the code <https://github.com/litlboy/OT-Model-Selection/blob/main/Synth-XP/Tensorised/sinkhorn.py> which differs from the one of the paper [27].

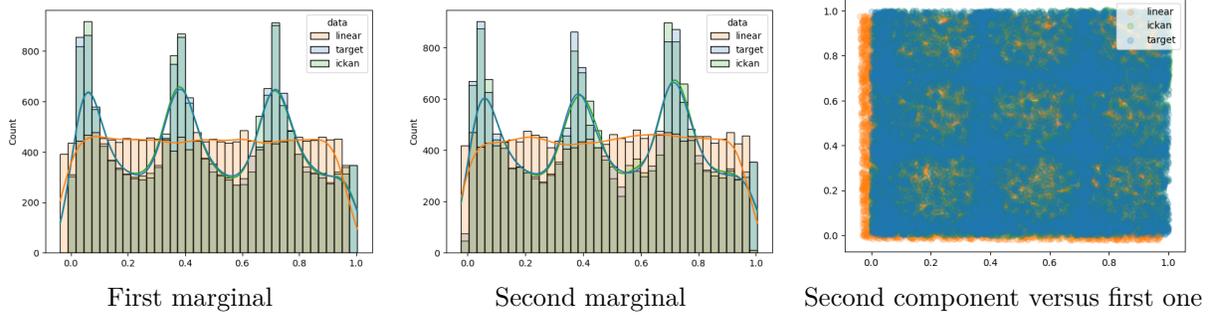


FIGURE 6. Distribution of the true target distribution as well as the one obtained by linear transport or Cubic ICKAN transport with adapted mesh and $P = 10$, for the map $T(x) = (T_i(x_i))_{i=1,2}$ with $T_i(x_i) = x_i + \frac{1}{6 - \cos(2\pi x_i)} - 0.2$, $i = 1, 2$. The first two figures include the empirical histogram as well as a Gaussian kernel density estimator with bandwidth selected using Scott's rule [24].

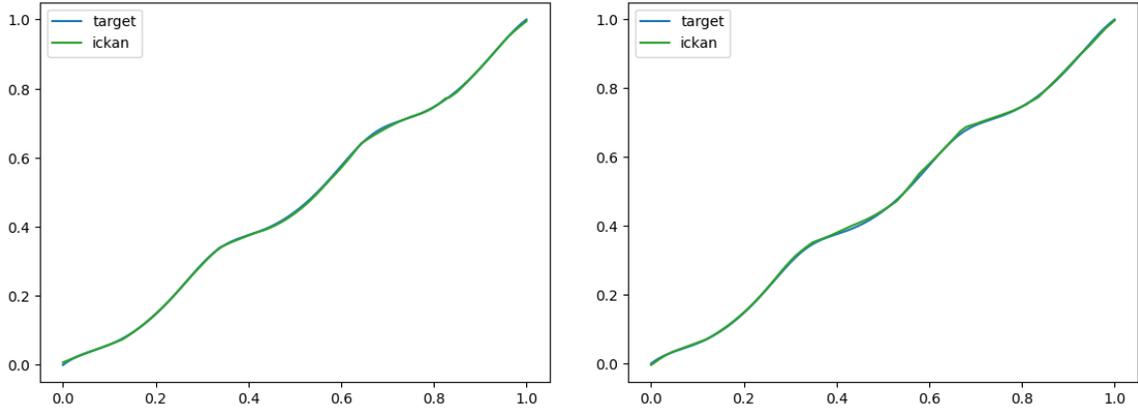


FIGURE 7. $x \mapsto T_1(x)$ (left) and $x \mapsto T_2(x)$ (right) for the map $T(x) = (T_i(x_i))_{i=1,2}$ with $T_i(x_i) = x_i + \frac{1}{6 - \cos(2\pi x_i)} - 0.2$, $i = 1, 2$ as well as the first component of $x \mapsto \hat{T}(x, 0.5)$ (left) and the second component of $x \mapsto \hat{T}(0.5, x)$ with \hat{T} the estimated map parametrized by the Cubic ICKAN with adapted mesh and $P = 10$.

Product case. In this last case, we consider the map $T = \nabla f$ with f the convex function defined on $[0, 1]^d$ by

$$(22) \quad f(x) = 3^{-d} \prod_{i=1}^d (x_i^2 + x_i + 1),$$

and μ is the uniform law on $[0, 1]^d$. The networks have the same architecture as in the second case (tensorized map). The results are given in Table 7: ICKAN and ICNN give similar results and outperform the linear map benchmark (except for the case $d = 1$ where the optimal transport map is linear).

Method / Dim	1	2	4	8
Linear	0.00	6.90	16.43	33.57
Cubic ICKAN P=40	0.01	3.15	2.33	2.85
ICNN	0.00	2.77	1.81	1.30

TABLE 7. Percentage of unexplained variance UVP (%) for the linear map, the map parametrized by the Cubic ICKAN with adapted mesh and $P = 40$, and the ICNN map and different dimensions when the true map is $T(x) = \nabla f(x)$ with f given in (22).

REFERENCES

- [1] Akshay Agrawal, Shane Barratt, Stephen Boyd, and Bartolomeo Stellato. Learning convex optimization control policies. In *Learning for Dynamics and Control*, pages 361–373. PMLR, 2020.
- [2] Brandon Amos, Lei Xu, and J Zico Kolter. Input convex neural networks. In *International conference on machine learning*, pages 146–155. PMLR, 2017.
- [3] Gábor Balázs, András György, and Csaba Szepesvári. Near-optimal max-affine estimators for convex regression. In *Artificial Intelligence and Statistics*, pages 56–64. PMLR, 2015.
- [4] Zavareh Bozorgasl and Hao Chen. Wav-kan: Wavelet kolmogorov-arnold networks. *arXiv preprint arXiv:2405.12832*, 2024.
- [5] Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.
- [6] Giuseppe C Calafiore, Stephane Gaubert, and Corrado Possieri. Log-sum-exp neural networks and posynomial models for convex and log-log-convex data. *IEEE transactions on neural networks and learning systems*, 31(3):827–838, 2019.
- [7] Yize Chen, Yuanyuan Shi, and Baosen Zhang. Optimal control via neural networks: A convex approach. *arXiv preprint arXiv:1805.11835*, 2018.
- [8] Yize Chen, Yuanyuan Shi, and Baosen Zhang. Input convex neural networks for optimal voltage regulation. *arXiv preprint arXiv:2002.08684*, 2020.
- [9] Rémi Flamary, Karim Lounici, and André Ferrari. Concentration bounds for linear monge mapping estimation and optimal transport domain adaptation. *arXiv preprint arXiv:1905.10155*, 2019.
- [10] Avishek Ghosh, Ashwin Pananjady, Adityanand Guntuboyina, and Kannan Ramchandran. Max-affine regression: Provable, tractable, and near-optimal statistical estimation. *arXiv preprint arXiv:1906.09255*, 2019.
- [11] Avishek Ghosh, Ashwin Pananjady, Adityanand Guntuboyina, and Kannan Ramchandran. Max-affine regression: Parameter estimation for gaussian designs. *IEEE Transactions on Information Theory*, 68(3):1851–1885, 2021.
- [12] Federico Girosi and Tomaso Poggio. Representation properties of networks: Kolmogorov’s theorem is irrelevant. *Neural Computation*, 1(4):465–469, 1989.
- [13] Seonho Kim and Kiryung Lee. Max-affine regression via first-order methods. *SIAM Journal on Mathematics of Data Science*, 6(2):534–552, 2024.
- [14] Andrei Nikolaevich Kolmogorov. *On the representation of continuous functions of several variables by superpositions of continuous functions of a smaller number of variables*. American Mathematical Society, 1961.
- [15] Alexander Korotin, Vage Egiazarian, Arip Asadulaev, Alexander Safin, and Evgeny Burnaev. Wasserstein-2 generative networks. *arXiv preprint arXiv:1909.13082*, 2019.
- [16] Alexander Korotin, Lingxiao Li, Aude Genevay, Justin M Solomon, Alexander Filippov, and Evgeny Burnaev. Do neural optimal transport solvers work? a continuous wasserstein-2 benchmark. *Advances in neural information processing systems*, 34:14593–14605, 2021.
- [17] Ziyao Li. Kolmogorov-arnold networks are radial basis function networks. 2024.
- [18] Ziming Liu, Yixuan Wang, Sachin Vaidya, Fabian Ruehle, James Halverson, Marin Soljačić, Thomas Y Hou, and Max Tegmark. Kan: Kolmogorov-arnold networks. *arXiv preprint arXiv:2404.19756*, 2024.
- [19] Ashok Makkuva, Amirhossein Taghvaei, Sewoong Oh, and Jason Lee. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning*, pages 6672–6681. PMLR, 2020.
- [20] Tudor Manole, Sivaraman Balakrishnan, Jonathan Niles-Weed, and Larry Wasserman. Plugin estimation of smooth optimal transport maps. *The Annals of Statistics*, 52(3):966–998, 2024.
- [21] Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pages 666–704, 1781.
- [22] Subhadip Mukherjee, Sören Dittmer, Zakhar Shumaylov, Sebastian Lunz, Ozan Öktem, and Carola-Bibiane Schönlieb. Learned convex regularizers for inverse problems. *arXiv preprint arXiv:2008.02839*, 2020.
- [23] Tomaso Poggio, Andrzej Banburski, and Qianli Liao. Theoretical issues in deep networks. *Proceedings of the National Academy of Sciences*, 117(48):30039–30045, 2020.
- [24] David W. Scott and Stephan R. Sain. 9 - multidimensional density estimation. In C.R. Rao, E.J. Wegman, and J.L. Solka, editors, *Data Mining and Data Visualization*, volume 24 of *Handbook of Statistics*, pages 229–261. Elsevier, 2005.
- [25] Sidharth SS. Chebyshev polynomial-based kolmogorov-arnold networks: An efficient architecture for nonlinear function approximation. *arXiv preprint arXiv:2405.07200*, 2024.
- [26] Hoang-Thang Ta. Bsrbf-kan: A combination of b-splines and radial basic functions in kolmogorov-arnold networks. *arXiv preprint arXiv:2406.11173*, 2024.
- [27] Adrien Vacher and François-Xavier Vialard. Parameter tuning and model selection in optimal transport with semi-dual brenier formulation. *Advances in Neural Information Processing Systems*, 35:23098–23108, 2022.

- [28] Xavier Warin. The groupmax neural network approximation of convex functions. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [29] Xavier Warin. P1-kan: an effective kolmogorov-arnold network with application to hydraulic valley optimization. *arXiv preprint arXiv:2410.03801*, 2024.

APPENDIX A. PROOF OF THEOREM 2.1

As usual, it is sufficient to show the result by fixing $P = 2$. The proof is constructive and is based on theorem 1 demonstration in [7]. First any Lipschitz convex function on $[0, 1]^n$ can be approximated within ϵ by the maximum of a finite set of affine functions (Lemma 1 in [7]). As a consequence, the proof of theorem 2.1 boils down to show that the P1-ICKAN network with adaptation with $P = 2$ can represent any maximum of a finite set of affine functions. Following [7], we begin to show that the linear KAN can represent the maximum of two affine functions f where for $x \in \mathbb{R}^n$,

$$\begin{aligned} f(x) &:= \max(\alpha_1^\top x + \beta_1, \alpha_2^\top x + \beta_2), \\ &= \max((\alpha_1^\top - \alpha_2^\top)x + \beta_1 - \beta_2, 0) + \alpha_2^\top x + \beta_2 \end{aligned}$$

with $\alpha_i \in \mathbb{R}^n$, $\beta_i \in \mathbb{R}$, $i = 1, 2$.

We consider a $L = 2$ -layers neural network. Take $n_1 = 2$ neurons for the first layer and $P = 1$ which is enough (and corresponds to a degenerate $P = 2$ case) and define in (10) the parameters

$$\begin{aligned} b_{0,1,1} &= \beta_1 - \beta_2, \\ b_{0,1,j} &= 0 \text{ for } 1 < j \leq n, \\ \hat{b}_{0,1,j} &= (\alpha_{1,j} - \alpha_{2,j}) \text{ for } 1 \leq j \leq n, \\ b_{0,2,1} &= \beta_2, \\ b_{0,2,j} &= 0 \text{ for } 1 < j \leq n, \\ \hat{b}_{0,2,j} &= \alpha_{2,j} \text{ for } 1 \leq j \leq n. \end{aligned} \tag{23}$$

One checks that

$$\begin{aligned} \hat{\kappa}_{n,2}^0(x, [0, 1]^n)_1 &= (\alpha_1^\top - \alpha_2^\top)x + \beta_1 - \beta_2, \\ \hat{\kappa}_{n,2}^0(x, [0, 1]^n)_2 &= \alpha_2^\top x + \beta_2. \end{aligned}$$

Let $I = \prod_{i=1}^2 [\hat{x}_{1,i,0}, \hat{x}_{1,i,2}]$ denote the image of $[0, 1]^n$ by the first layer. For the second layer, supposing that the non linearity is active for the first component of the first layer output, that is $0 \in]\hat{x}_{1,1,0}, \hat{x}_{1,1,2}[$, we take $\hat{x}_{1,1,1} = 0 \in]\hat{x}_{1,1,0}, \hat{x}_{1,1,2}[$, $\hat{x}_{1,2,1} \in]\hat{x}_{1,2,0}, \hat{x}_{1,2,2}[$, and define

$$\begin{aligned} b_{1,1,1} &= \hat{b}_{1,1,1} = 0, \\ d_{1,1,1,1} &= 1, \\ b_{1,1,2} &= \hat{x}_{1,2,0}, \\ \hat{b}_{1,1,2} &= 1, \\ d_{1,1,2,1} &= 0. \end{aligned}$$

The two first lines in the previous equation defines the ReLU with an active non linearity, while the three last defines the identity function and we get the desired result

$$\hat{\kappa}_{2,1}^1 \circ \kappa_{n,2}^0(x, [0, 1]^n)_1 = \max((\alpha_1^\top - \alpha_2^\top)x + \beta_1 - \beta_2, 0) + \alpha_2^\top x + \beta_2.$$

Of course, the case when the non linearity is not active can be treated using $P = 1$.

Similarly as in [7], on can extend iteratively the procedure for a given N for

$$f(x) := \max(\alpha_1^\top x + \beta_1, \dots, \alpha_N^\top x + \beta_N).$$

At last one cannot be sure that the term $\max(\hat{b}_{l,1,j}, 0)$ for $l > 1$ is really defined (it may be negative with our construction with $N > 2$ as it is for example defined as $\hat{b}_{0,1,j} = (\alpha_{1,j} - \alpha_{2,j})$ for $1 \leq j \leq n$ in the first layer in (23)). But, first by an affine transformation, on changes the initial problem on $[0, 1]^n$ in a problem on $[-1, 1]^n$ and by a expansion of the input, on defines a new input $\hat{x} = \begin{pmatrix} x \\ -x \end{pmatrix}$. With this new input it is possible to have positive \hat{b} values as shown in [7] (duplication trick).

THOMAS DESCHATRE, EDF LAB PARIS-SACLAY AND FiME, LABORATOIRE DE FINANCE DES MARCHÉS DE L'ENERGIE, 91120 PALAISEAU, FRANCE

Email address: `thomas-t.deschatre@edf.fr`

XAVIER WARIN, EDF LAB PARIS-SACLAY AND FiME, LABORATOIRE DE FINANCE DES MARCHÉS DE L'ENERGIE, 91120 PALAISEAU, FRANCE

Email address: `xavier.warin@edf.fr`