

Representing Piecewise Linear Functions by Functions with Small Arity

Christoph Koutschan · Bernhard Moser · Anton Ponomarchuk · Josef Schicho

Received: date / Accepted: date

Abstract A piecewise linear function can be described in different forms: as a nested expression of min- and max-functions, as a difference of two convex piecewise linear functions, or as a linear combination of maxima of affine-linear functions. In this paper, we provide two main results: first, we show that for every piecewise linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, there exists a linear combination of max-functions with at most $n + 1$ arguments, and give an algorithm for its computation. Moreover, these arguments are contained in the finite set of affine-linear functions that coincide with the given function in some open set. Second, we prove that the piecewise linear function $\max(0, x_1, \dots, x_n)$ cannot be represented as a linear combination of maxima of less than $n + 1$ affine-linear arguments. This was conjectured by Wang and Sun in 2005 in a paper on representations of piecewise linear functions as linear combination of maxima.

Keywords Piecewise linear function · Linear combination of max-functions · ReLU network · Convex polyhedra · Minkowski sum · Zero volume polytope

The research reported in this paper has been partly funded by BMK, BMDW, and the Province of Upper Austria in the frame of the COMET Programme managed by FFG in the COMET Module S3AI.

Christoph Koutschan
Johann Radon Institute for Computational and Applied Mathematics (RICAM)
E-mail: Christoph.Koutschan@ricam.oeaw.ac.at

Bernhard Moser
Software Competence Center Hagenberg (SCCH)
E-mail: Bernhard.Moser@scch.at

Anton Ponomarchuk
Johann Radon Institute for Computational and Applied Mathematics (RICAM)
E-mail: Anton.Ponomarchuk@ricam.oeaw.ac.at

Josef Schicho
Research Institute for Symbolic Computation, Johannes Kepler University (RISC JKU)
E-mail: Josef.Schicho@risc.jku.at

1 Introduction

The mathematical model of a neural network is a directed graph without cycles [3], where each vertex with in-degree 0 stands for an input parameter ranging over \mathbb{R} , and vertices with out-degree 0 stand for the output values. Each vertex with positive in-degree is called a neuron. Each neuron has finitely many input values corresponding to the incoming edges. The neuron applies an affine-linear function to the vector of these input values, followed by a non-linear activation function. We assume that the activation function is the function $x \mapsto \max(x, 0)$, also known as the ReLU function (Rectified Linear Unit, see [8]). The output of one neuron may be the input for another neuron, which is indicated by a directed edge between the two vertices in the graph. Figure 1.1 shows a ReLU network that computes a bivariate piecewise linear function.




Fig. 1.1 A neural network computing the piecewise linear function $(x_1, x_2) \mapsto 4 \max(-x_1 + 3x_2 + 2, 0) - 5 \max(2x_1 - 3, 0) + 6 \max(5x_2 + 1, 0) + 8$. The activation function for each neuron is the ReLU function $x \mapsto \max(0, x)$. The network has depth 1.

Algebraically, the depth of a ReLU network corresponds to the depth of nestings of ReLU functions in the expression determined by the network. The function $\max(x_1, \dots, x_n)$ can be written as a composition of binary max-function with nesting depth equal to $\lceil \log_2(n) \rceil$, and every binary max can be written in terms of a ReLU function by the identity $\max(x, y) = x + \max(y - x, 0)$. This shows that any piecewise linear function that can be written as a linear combination of maxima of at most n affine-linear functions can be realized by a ReLU network of depth at most $\lceil \log_2(n) \rceil$. This was already observed in [12].

It is well-known [6] that every piecewise linear function is the difference of two convex piecewise linear functions; see also [10] for a more efficient decomposition algorithm. It follows that every piecewise linear function is a linear combination of maxima of affine-linear functions, see [1, 2, 9, 11]. The paper [12] addresses the problem of minimizing the largest number of arguments of the maxima appearing in such a linear combination. The authors show that every piecewise linear function defined on \mathbb{R}^n can be written as a linear combination of maxima of at most $n + 1$ affine-linear functions. The authors also express their conviction that this bound is optimal. In their own words, “it seems impossible” that the function $\max(0, x_1, \dots, x_n)$ has an expression as a linear combination of maxima of at most n affine-linear functions.

Proposition 3.2 in [5], provides a proof that for any $n \geq 3$, the function $\max(0, x_1, x_2, \dots, x_{n-3}, \max(x_{n-2}, x_{n-1}) + \max(0, x_n))$ cannot be written as a linear combination of maxima of at most n affine-linear functions. To the best of our knowledge, it is the first proof for the conjecture in [12]. Using the same reasoning, one can show the same statement holds for the function $\max(0, x_1, \dots, x_n)$.

We provide an alternative proof for the conjecture in [12], which is valid for every $n \in \mathbb{N}$. Also, we prove that any continuous piecewise linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ has a representation as an integral linear combination of maxima of at most $n + 1$ affine-linear functions and from our proof, we derive an algorithm for computing the claimed representation. So, our proof is constructive in a way that it also proves the termination of the algorithm.

This paper is structured as follows. In Section 2 we show that for any piecewise linear function f there exists an integral linear combination of maxima of at most $n + 1$ affine-linear functions. Note that in the case $n = 1$ this corresponds to the fundamental theorem of tropical algebra. In Section 3 we give an algorithm for finding such linear integral combination of maxima with examples. In Section 4 we recall some notions about Minkowski-addition, duality between the set of non-empty convex polytopes in \mathbb{R}^n and the set of convex and positively homogeneous piecewise linear functions of degree 1. In Section 5 we give a proof for the conjecture in [12], thereby showing that their bound for the number of max-arguments is indeed optimal. The proof is based on properties of the Minkowski-addition of convex polytopes.

2 An Upper Bound for the Height

For a given piecewise linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}$, we define the height $H(f)$ as the smallest integer k such that f is a linear combination of maxima of at most $k + 1$ affine-linear functions. In this section we give a new proof for the bound $H(f) \leq n$ that was first shown in [12]. A constituent of the function f is an affine-linear function that coincides with f in some open subset of \mathbb{R}^n . We also show that the arguments for the maxima can be chosen among the constituents of f .

Lemma 2.1 *Let R be a finite set of affine-linear functions on \mathbb{R}^n such that $|R| \geq n + 2$. Then there exists a decomposition into two non-empty disjoint subsets, $R = S \uplus T$, such that for all points $x \in \mathbb{R}^n$, we have*

$$\max_{g \in T} g(x) \geq \min_{g \in S} g(x).$$

Proof Let $(\mathbb{R}^n)^*$ be the real vector space of all linear functions $\mathbb{R}^n \rightarrow \mathbb{R}$. The derivative of every function in R is a linear function in $(\mathbb{R}^n)^*$. The vector space $(\mathbb{R}^n)^*$ has dimension n . Therefore the set of derivatives $\{g' \mid g \in R\}$ is affinely dependent, i.e., there exist real numbers α_g , $g \in R$, not all equal to zero, such that

$$\sum_{g \in R} \alpha_g g' = 0, \quad \sum_{g \in R} \alpha_g = 0.$$

We set $T := \{g \mid \alpha_g > 0\}$ and $S := \{g \mid \alpha_g \leq 0\}$. Without loss of generality, we may assume $\sum_{g \in T} \alpha_g = 1$ and $\sum_{g \in S} \alpha_g = -1$ – if not, we multiply all α_g by a suitable positive constant. We also set $\beta_g := -\alpha_g$ for $g \in S$. Then the function $h := \sum_{g \in T} \alpha_g g - \sum_{g \in S} \beta_g g$ has derivative zero and therefore h equals to a constant c . Let us assume $c \geq 0$. Then we get

$$\max_{g \in T} g(x) \geq \sum_{g \in T} \alpha_g g(x) = \sum_{g \in S} \beta_g g(x) + c \geq \sum_{g \in S} \beta_g g(x) \geq \min_{g \in S} g(x),$$

for all $x \in \mathbb{R}^n$.

If $c < 0$, then we redefine $T := \{g \mid \alpha_g < 0\}$ and $S := \{g \mid \alpha_g \geq 0\}$ and get a similar chain of inequalities. \square

Lemma 2.2 *Let A be a finite non-empty set. Then*

$$\sum_{T \subseteq A} (-1)^{|T|} = 0.$$

Proof

$$\sum_{T \subseteq A} (-1)^{|T|} = \sum_{i=0}^{|A|} \binom{|A|}{i} (-1)^i = (1 - 1)^{|A|} = 0.$$

\square

Lemma 2.3 *With S, T as above, for all $x \in \mathbb{R}^n$, the following equality holds:*

$$\sum_{M \subseteq S} (-1)^{|M|} \max_{g \in M \cup T} g(x) = 0.$$

Proof For all $x \in \mathbb{R}^n$ except those in a finite union of hyperplanes, the values $g(x)$, $g \in R$ are pairwise distinct. We only need to prove the equality for x

under this assumption; then it follows for the remaining places by continuity. So, let us fix x such that the values $g(x)$, $g \in R$ are pairwise distinct. Then

$$\sum_{M \subseteq S} (-1)^{|M|} \max_{g \in M \cup T} g(x) = \sum_{u \in R} \left(\sum_{\substack{M \subseteq S \\ u \in M \cup T \\ \forall g \in M \cup T: g(x) \leq u(x)}} (-1)^{|M|} \right) u(x).$$

We claim the inside sum is 0, for each $u \in R$. We distinguish three cases.

Case 1: $u \in T$ and $\exists g \in T : g(x) > u(x)$. Then there is no M such that the conditions under the sum sign are fulfilled. Then the inner sum is the sum over the empty set which is 0 by definition.

Case 2: $u \in S$. We define S_u as the set of all $g \in S$ such that $g(x) \leq u(x)$. Note that S_u is non-empty since it contains u . Then by Lemma 2.2 the following equalities hold:

$$\sum_{\substack{M \subseteq S \\ u \in M \cup T \\ \forall g \in M \cup T: g(x) \leq u(x)}} (-1)^{|M|} = \sum_{M \subseteq S_u} (-1)^{|M|} = 0.$$

Case 3: $u \in T$ and $\forall g \in T : g(x) \leq u(x)$. We define S_u as the set of all $g \in S$ such that $g(x) \leq u(x)$. The set S_u is non-empty because of Lemma 2.1. As in the previous case, we get

$$\sum_{\substack{M \subseteq S \\ u \in M \cup T \\ \forall g \in M \cup T: g(x) \leq u(x)}} (-1)^{|M|} = \sum_{M \subseteq S_u} (-1)^{|M|} = 0.$$

□

Theorem 2.4 *Every piecewise linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ can be written as an integral linear combination of maxima of at most $n + 1$ affine-linear functions. Moreover, the affine-linear functions can be chosen among the constituents of f .*

Proof By [9,11], the function f can be written as a maximum of minima of constituents. Using the first identity in [12, Lemma 1], namely

$$\max(a, \min(b, c)) = \max(a, b) + \max(a, c) - \max(a, b, c),$$

we can rewrite this expression as a linear combination of maxima of constituents:

$$f(x) = \sum_{i=1}^p \sigma_i \max(g_{i1}(x), \dots, g_{ik_i}(x)),$$

where $k_i \in \mathbb{N}$, $\sigma_i \in \{-1, 1\}$, for all $i \in \{1, \dots, p\}$. If for all $i \in \{1, \dots, p\}$ the number of constituents k_i is at most $n + 1$, we are done.

Otherwise, for every maximum $\max(g_{i1}(x), \dots, g_{ik_i}(x))$ with k_i more than $n+2$ constituents, by Lemma 2.3, there exists a linear combination of maxima with at most $k_i - 1$ constituents:

$$\max(g_{i1}(x), \dots, g_{ik_i}(x)) = \sum_{j=1}^s \widehat{\sigma}_j \max(g_{j1}(x), \dots, g_{jk_j}(x)),$$

where $k_j \leq k_i - 1$, $\widehat{\sigma}_j \in \{-1, 1\}$, for all $j \in \{1, \dots, s\}$. After a finite number of replacements of maxima with more than $n+1$ constituents, we achieve an integral linear combination of maxima of at most $n+1$ constituents. \square

Remark 2.5 Theorem 4.1 in [5] provides a proof that any convex piecewise linear function f can be written as an integral linear combination of maxima of at most $n+1$ constituents. The proof in this section shows a slightly stronger statement: any (convex and non-convex) piecewise linear function f can be written as an integral linear combination of maxima of at most $n+1$ constituents whose derivatives are affinely independent. We will see in Section 5 that the maximum of affine-linear functions that are affinely independent cannot be expressed as a linear combination of maxima of fewer affine-linear functions.

Example 2.6 Let $g_1(x_1, x_2) = c_1x_1 + c_2x_2$ for some constants c_1, c_2 (not both zero) and let $g_i = g_1 + i$ for $i = 2, 3, 4$. Then clearly $\max(g_1, g_2, g_3, g_4) = g_4$. However, this simplest-possible answer cannot be found by our algorithm (see Section 3). If we choose $\alpha = (-1, 0, 0, 1)$, then we get $S = \{g_1, g_2, g_3\}$, $T = \{g_4\}$ and hence $\max(g_1, g_2, g_3, g_4) = \max(g_1, g_2, g_4) + \max(g_1, g_3, g_4) + \max(g_2, g_3, g_4) - \max(g_1, g_4) - \max(g_2, g_4) - \max(g_3, g_4) + g_4$. In contrast, if we choose $\alpha = (-3, 1, 1, 1)$, then $S = \{g_1\}$, $T = \{g_2, g_3, g_4\}$ and we obtain $\max(g_1, g_2, g_3, g_4) = \max(g_2, g_3, g_4)$ as the final result. However, this weakness can easily be cured by ignoring the condition $|R| \geq n+2$ in Lemma 2.1: now the existence of a vector α is not guaranteed any more, but if it exists, we perform the corresponding decomposition, otherwise, we leave that term unchanged.

3 Reducing the Height

In this section, we give an algorithm for writing the maximum of any number of affine-linear functions from $\mathbb{R}^n \rightarrow \mathbb{R}$ as a linear combination of maxima of at most $n+1$ of these functions.

Algorithm 2 implements the method *reduceMax* that takes the size of the input space n and a maximum function $f(x) := \max(h_1(x), \dots, h_k(x))$ with more than $n+1$ constituents, i.e. $R := \{h_1, \dots, h_k\}$ and $|R| > n+1$. It returns a linear combination of maxima:

$$g := \sum_{i=1}^l c_i \cdot \max(h_1^i, \dots, h_{k_i}^i),$$

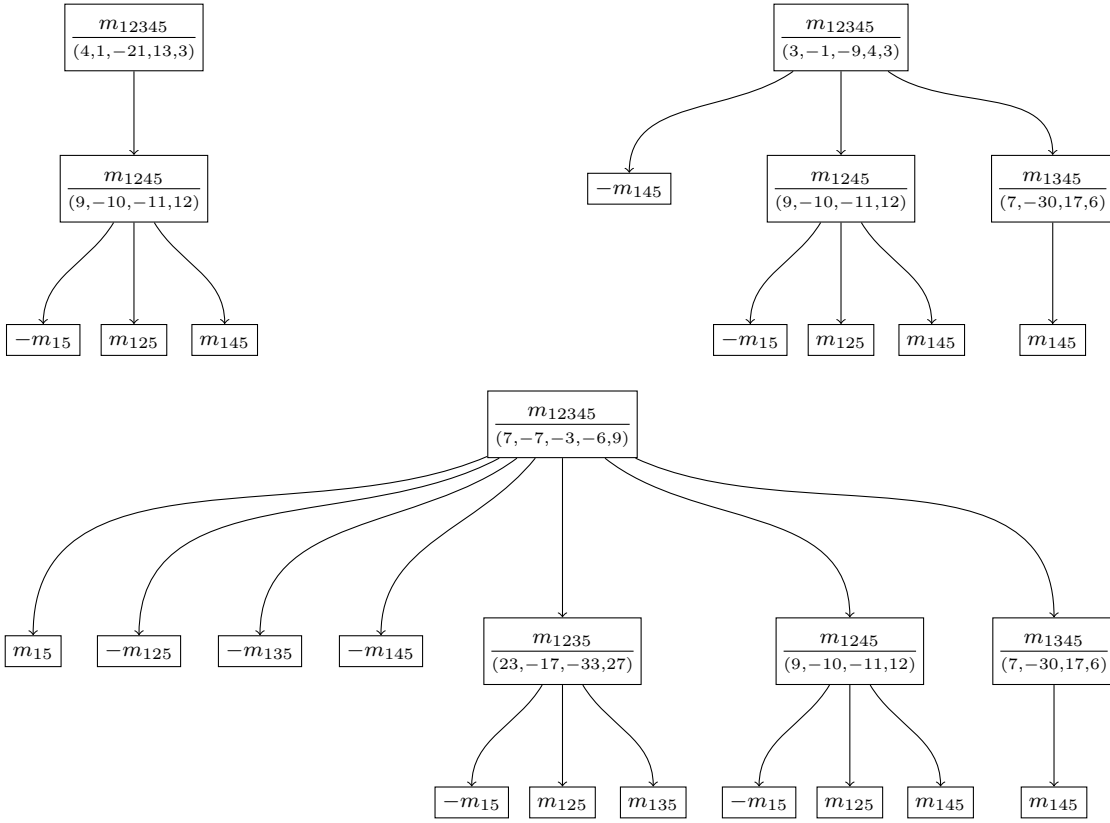


Fig. 2.1 Application of our Algorithm 2 “reduceMax” to the input function $\max(g_1, g_2, g_3, g_4, g_5)$ where $g_1(x_1, x_2) = 3x_1 - 4x_2 + 1$, $g_2(x_1, x_2) = -3x_1 - x_2 - 2$, $g_3(x_1, x_2) = 2x_1 + x_2 - 1$, $g_4(x_1, x_2) = 3x_1 + 2x_2 + 2$, $g_5(x_1, x_2) = -2x_1 + 4x_2 + 3$. We use the abbreviation $m_{ij\dots}$ for $\max(g_i, g_j, \dots)$ and for those maxima that get expanded, we also display the vector α that determines the sets S and T . We show three different executions of the algorithm, which however all yield the same final result (combining equal terms is not shown explicitly here).

where $\{h_1^i, \dots, h_{k_i}^i\} \subset R$ and $k_i \leq n+1$ for all $i \in \{1, \dots, l\}$. The combination is constructed in the following way.

Firstly, one extracts the linear constituents h_1, \dots, h_k from the input maximum f and forms the set R . Then the set R is split into two disjoint subsets $S, T \subset R$ such that $R = S \uplus T$. The *split* operation is the implementation of Lemma 2.1 and is described in Algorithm 1 (more details in this section below). After splitting the set R into the pair of subsets S, T , the linear combination of maxima g is generated. The linear combination of maxima g has the following form:

$$g(x) = (-1)^{|S|+1} \sum_{M \subset S} (-1)^{|M|} \max_{h \in M \cup T} h(x),$$

where the sum runs over all proper subsets of S , including the empty set. By Lemma 2.3, $g(x)$ is equal to the input function $f(x)$, and every maximum in $g(x)$ contains at most $|R| - 1$ constituents, where the set R contains all constituents of the input function $f(x)$. If any summand of $g(x)$ contains an affinely dependent subset of constituents, one replaces it with the corresponding linear combination of maxima by recursively applying the *reduceMax* method. One repeats this simplification procedure until all the summands contain at most $n + 1$ constituents.

The termination of Algorithm 2 follows from the fact that after every call on the input maximum with k constituents, *reduceMax* returns a finite combination of maxima where each maximum contains at most $k - 1$ constituents.

Algorithm 2 uses the method *split* for dividing a set of constituents R into two disjoint sets S, T . The method *split* is described in Algorithm 1 and it is an implementation of Lemma 2.1. The algorithm splits the input set of constituents R based on the sign of the vector $\alpha \in \mathbb{R}^{|R|}$ that is a solution of the system of equalities:

$$\sum_{i=1}^{|R|} \alpha_i h'_i = 0, \quad \sum_{i=1}^{|R|} \alpha_i = 0,$$

where $h'_i := \left(\frac{\partial h_i}{\partial x_1}, \dots, \frac{\partial h_i}{\partial x_n} \right)^T$ and $h_i \in R$ for all $i \in \{1, \dots, |R|\}$. Due to the fact, that all constituents are linear, it implies that the system of equations is linear:

$$W\alpha = 0,$$

where $\alpha \in \mathbb{R}^{|R|}$ and $W \in \mathbb{R}^{(n+1) \times |R|}$ such that:

$$W := \begin{bmatrix} h'_{11} & h'_{21} & \dots & h'_{|R|1} \\ \vdots & & & \vdots \\ h'_{1n} & h'_{2n} & \dots & h'_{|R|n} \\ 1 & 1 & \dots & 1 \end{bmatrix}.$$

Solving the given system of linear equations is equivalent to finding the null space $\ker(W)$. If the null space is trivial, one does not need to divide the input set R because Lemma 2.1 is not applicable. Otherwise, the vector α can be picked as any vector from $\ker(W)$. By iterating through the entries of the vector α , depending on the sign of the entry α_i , the corresponding constituent h_i is assigned either to T or S . Note that the condition $k \geq n + 2$ in Lemma 2.1 ensures the existence of a non-trivial null space $\ker(W)$.

Example 3.1 Let us take the following function:

$$g(x_1, x_2) := \max(x_1, x_1 + x_2, x_2 + \max(x_1 + x_2 - 7, x_1 + 6x_2 + 4) \\ + 3 \min(4x_2, x_1 - 9, x_1 - x_2)).$$

After applying the function expansion by the rules explained in [12] on g , one receives a linear combination of maxima with 49 summands, where 40 of them contain 5 or more constituents. After applying Algorithm 2 on the expanded version of g , the given sum transforms into a new one with only 7 summands:

$$\begin{aligned}\hat{g}(x_1, x_2) := & \max(x_1, x_1 + 19x_2 + 4) \\ & - \max(x_1 + x_2, 4x_1 + 4x_2 + 4, 4x_1 + 7x_2 - 23) \\ & - \max(x_1 + 19x_2 + 4, 4x_1 + 4x_2 + 4) \\ & + \max(x_1, x_1 + x_2, 4x_1 + 7x_2 - 23) \\ & - \max(x_1, x_1 + 19x_2 + 4, 4x_1 + 7x_2 - 23) \\ & + \max(x_1 + x_2, 4x_1 + 4x_2 + 4) \\ & + \max(x_1 + 19x_2 + 4, 4x_1 + 4x_2 + 4, 4x_1 + 7x_2 - 23).\end{aligned}$$

The function \hat{g} contains maxima with at most 3 constituents, in accordance with Theorem 2.4. However, an example can be found where the number of maxima increases after applying Algorithm 2, compared to the output after applying the rules from [12]. For instance, it holds for the function:

$$\begin{aligned}f(x_1, x_2) := & \max(6x_1 + 5x_2 - 3, 8x_2 - 2, -3x_1 - 5x_2 - 4, \\ & \max(12x_1 - 4x_2 + 1, -7x_1 + 8x_2 + 12) + 3x_1 - 10, \\ & \min(-3x_1 + 4x_2 - 5, 8x_1 + 2)).\end{aligned}$$

After repeating two transformations, we receive two linear combinations with the number of maxima 3 and 5, respectively, with the final form:

$$\begin{aligned}\hat{f}(x_1, x_2) := & - \max(-4x_1 + 8x_2 + 2, 15x_1 - 4x_2 - 9) \\ & + \max(8x_2 - 2, -4x_1 + 8x_2 + 2, 6x_1 + 5x_2 - 3) \\ & + \max(-4x_1 + 8x_2 + 2, -3x_1 - 5x_2 - 4, 15x_1 - 4x_2 - 9) \\ & + \max(-4x_1 + 8x_2 + 2, 6x_1 + 5x_2 - 3, 15x_1 - 4x_2 - 9) \\ & - \max(-4x_1 + 8x_2 + 2, 6x_1 + 5x_2 - 3).\end{aligned}$$

Although Algorithm 2 can either reduce the number of summands in the final linear combination of maxima or increase it, the final combination seems to be invariant, not depending on the outcome of Algorithm 1. More precisely, we conjecture that the output of Algorithm 2 is independent of how the set of constituents is split by Algorithm 1, as illustrated in Figure 2.1.

4 Duality and Convex Polyhedra

Let n be a positive integer. Instead of all piecewise linear functions on \mathbb{R}^n , we consider in this section the subset of convex and positively homogeneous piecewise linear functions of degree 1, i.e., all functions $f: \mathbb{R}^n \rightarrow \mathbb{R}$ that are convex and satisfy $f(\lambda x) = \lambda f(x)$ for all $x \in \mathbb{R}^n$ and $\lambda \geq 0$. We denote

this subset by \mathcal{F}_n . There is a useful bijective correspondence $\tau: \mathcal{F}_n \rightarrow \mathcal{P}_n$, where we define \mathcal{P}_n as the set of all non-empty convex polytopes in \mathbb{R}^n . For each $f \in \mathcal{F}_n$, we define $\tau(f)$ as the subset of all vectors $h \in \mathbb{R}^n$ such that $\langle h, x \rangle \leq f(x)$ for all $x \in \mathbb{R}^n$. Conversely, if $P \in \mathcal{P}_n$, then $\tau^{-1}(P)$ is the support function $x \mapsto \sup\{\langle h, x \rangle \mid h \in P\}$. By compactness, the supremum is a maximum.

For two polytopes $P, Q \in \mathcal{P}_n$, the Minkowski sum $P + Q$ is defined as the convex polytope $\{a + b \mid a \in P, b \in Q\}$. Let $d \in \mathbb{R}^n \setminus \{0\}$ be a vector (d stands for “direction”). For any polytope $P \in \mathcal{P}_n$, we define the face

$$S_d(P) = \{x \in P \mid \langle x, d \rangle = \max_{y \in P} \langle y, d \rangle\}.$$

Proposition 4.1 *Let $P, Q \in \mathcal{P}_n$ be polytopes. Let $d \in \mathbb{R}^n \setminus \{0\}$ be a direction vector. Then*

a) *The map τ is an isomorphism of semigroups:*

$$\tau^{-1}(P) + \tau^{-1}(Q) = \tau^{-1}(P + Q).$$

b) *Taking faces is an endomorphism of semigroups:*

$$S_d(P + Q) = S_d(P) + S_d(Q).$$

Proof See [4, Lemma 2.1.4]. □

By Proposition 2.7 [7], the function $\tau^{-1}(P+Q)$ is a piecewise linear function with linear regions corresponding to the vertices of the polytope $P + Q$, as illustrated in Figure 4.1.

Proposition 4.2 *Minkowski addition is cancellable: if $A + C = B + C$, then $A = B$.*

Proof This is well-known, and we can prove it easily by translation to functions. Assume $A + C = B + C$. Then

$$\tau^{-1}(A) + \tau^{-1}(C) = \tau^{-1}(A + C) = \tau^{-1}(B + C) = \tau^{-1}(B) + \tau^{-1}(C),$$

hence $\tau^{-1}(A) = \tau^{-1}(B)$ and therefore $A = B$. □

The face $S_d(P)$ is contained in the hyperplane $H_{d,c} = \{x \mid \langle x, d \rangle = c\}$, where $c := \max_{y \in P} \langle y, d \rangle$. For the induction proof in the next section, we need to identify $H_{d,c}$ with \mathbb{R}^{n-1} . To this end, we translate the hyperplane to $H_{d,0}$ by a translation vector cv that satisfies $\langle d, v \rangle = -1$, and then we apply an isomorphism $\phi: H_{d,0} \rightarrow \mathbb{R}^{n-1}$. The face $\phi(S_d(P) + cv)$ is denoted by $F_d(P)$. It depends not only on d , but also on the choice of v and ϕ ; but we may choose v_d and ϕ_d for every d once and for all, subject to the condition $v_{-d} = -v_d$ and $\phi_d = \phi_{-d}$.

Algorithm 1: split (Lemma 2.1)

Input: $\{h_1, \dots, h_k\}$ – set of linear constituents, where $h_i: \mathbb{R}^n \rightarrow \mathbb{R}, \forall i \in \{1, \dots, k\}$
Output: (S, T) – disjoint decomposition of R

- 1 $W := \left(\begin{pmatrix} h'_1 \\ 1 \end{pmatrix}, \dots, \begin{pmatrix} h'_k \\ 1 \end{pmatrix} \right)$
- 2 $K := \ker(W)$
- 3 **if** K is zero **then**
 └ **Return:** $\emptyset, \{h_1, \dots, h_k\}$
- 4 $\alpha :=$ pick a non-zero vector from K
- 5 $v :=$ constant terms of h_1, \dots, h_k
- 6 $S := \emptyset$
- 7 $T := \emptyset$
- 8 $c := v^T \alpha$
- 9 **for** $i \in \{1, \dots, k\}$ **do**
- 10 └ **if** $\alpha_i > 0$ **then**
- 11 └ $T := T \cup \{h_i\}$
- 12 └ **else**
- 13 └ $S := S \cup \{h_i\}$
- 14 **if** $c > 0$ **then**
 └ **Return:** S, T

Return: T, S

Algorithm 2: reduceMax (Lemma 2.3)

Input: $g := \max(h_1, \dots, h_k)$ – max function, where $g: \mathbb{R}^n \rightarrow \mathbb{R}$.
Output: Linear combination of max functions.

- 1 $R := \{h_1, \dots, h_k\}$
- 2 $(S, T) := \text{split}(R)$
- 3 $A := \max(T)$
- 4 **if** $S = \emptyset$ **then**
 └ **Return:** A
- 5 **if** $|T| > 1$ **then**
 └ $A := \text{reduceMax}(A)$
- 7 **for** all proper subsets M of S **do**
- 8 └ $P := M \cup T$
- 9 └ $A := A + (-1)^{|M|} \text{reduceMax}(\max(P))$

Return: $(-1)^{|S|+1} A$

5 A Lower Bound for the Height

Recall that an n -simplex is a polytope which is the convex hull of $n + 1$ points that are affinely independent. The faces of simplices are again simplices.

A polytope $P \in \mathcal{P}_n$ is called a *zero volume* polytope if and only if it has no interior points; this is the case if and only if it is contained in a hyperplane. If P has zero volume and $d \in \mathbb{R}^n \setminus \{0\}$ is a direction vector, then one of the following two cases holds:

- either $S_d(P) = S_{-d}(P) = P$,
- or both faces $F_d(P)$ and $F_{-d}(P)$ have zero volume as polytopes in \mathcal{P}_{n-1} .

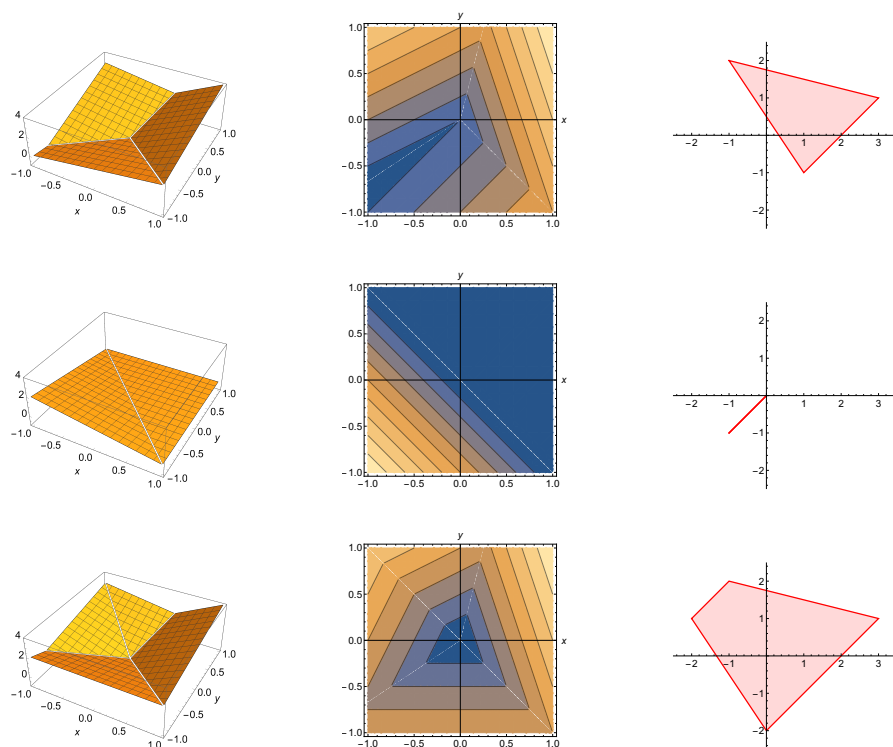


Fig. 4.1 Two positively homogeneous piecewise linear functions, $f(x, y) = \max(x - y, 3x + y, -x + 2y)$ (first row) and $g(x, y) = \max(0, -x - y)$ (second row), as well as their sum $h = f + g$ (third row). Each row shows the function $\mathbb{R}^2 \rightarrow \mathbb{R}$ as a three-dimensional plot (first column) and as a contour plot (second column). The third column shows the corresponding polygon $\tau(f)$ etc. in \mathcal{P}_n . One clearly sees how $h(x, y) = \max(-2y, -2x + y, 3x + y, -x + 2y)$ produces the Minkowski sum $\tau(h) = \tau(f) + \tau(g)$ of polygons.




Fig. 5.1 The polytope P is a zero-summand since it can be written as a Minkowski sum of three line segments.

We say that a polytope $P \in \mathcal{P}_n$ is a *zero-summand* if and only if there are convex polytopes $P_1, \dots, P_r, Q_1, \dots, Q_s \in \mathcal{P}_n$ of zero volume such that $P + P_1 + \dots + P_r = Q_1 + \dots + Q_s$, as illustrated in Figure 5.1.

Lemma 5.1 *Let $P \in \mathcal{P}_n$ be a zero-summand. Let $d \in \mathbb{R}^n \setminus \{0\}$ be a direction vector such that $F_{-d}(P) \in \mathcal{P}_{n-1}$ has zero volume. Then $F_d(P)$ is also a zero-summand.*

Proof Since $P \in \mathcal{P}_n$ is a zero summand, there exist convex polytopes $P_i, Q_j \in \mathcal{P}_n$, where $i \in \{1, \dots, r\}$ and $j \in \{1, \dots, s\}$ of zero volume such that $P + P_1 + \dots + P_r = Q_1 + \dots + Q_s$.

We may assume, without loss of generality, that the zero volume polytopes P_i or Q_j that have a face $F_d(P_i) = F_{-d}(P_i)$ of nonzero volume – or similarly with Q_j – are P_1, \dots, P_k and Q_1, \dots, Q_l , for some $k \leq r$ and $l \leq s$. Set $A := P + P_1 + \dots + P_r$; therefore also $A = Q_1 + \dots + Q_s$. We apply F_d to A and neglect all polytopes of zero volume. This yields

$$F_d(A) = F_d(P) + F_d(P_1) + \dots + F_d(P_k)$$

plus polytopes of zero volume, and

$$F_d(A) = F_d(Q_1) + \dots + F_d(Q_l)$$

plus polytopes of zero volume. Now we apply F_{-d} to A and neglect all polytopes of zero volume, taking into account the equations $F_d(P_i) = F_{-d}(P_i)$ for $i = 1, \dots, k$ and $F_d(Q_j) = F_{-d}(Q_j)$ for $j = 1, \dots, l$, and the assumption that $F_{-d}(P) \in \mathcal{P}_{n-1}$ has zero volume. This yields

$$F_{-d}(A) = F_d(P_1) + \dots + F_d(P_k)$$

plus polytopes of zero volume, and

$$F_{-d}(A) = F_d(Q_1) + \dots + F_d(Q_l)$$

plus polytopes of zero volume. Summing up, we get that $F_d(A) + F_{-d}(A)$ is equal to both sides of the equation

$$\begin{aligned} & F_d(P) + F_d(P_1) + \dots + F_d(P_k) + F_d(Q_1) + \dots + F_d(Q_l) \\ &= F_d(P_1) + \dots + F_d(P_k) + F_d(Q_1) + \dots + F_d(Q_l) \end{aligned}$$

modulo polytopes of zero volume. By Proposition 4.2, it follows that $F_d(P)$ is a zero summand. \square

Corollary 5.2 *An n -simplex in \mathcal{P}_n is not a zero summand.*

Proof Induction on n : if $n = 1$, then the zero volume polytopes are single points, and therefore, the zero summands are also single points. But a 1-simplex is a line segment of positive length and not a single point.

If $\Delta \in \mathcal{P}_n$ is an n -simplex for some $n > 1$, then there is a direction vector d such that $F_{-d}(\Delta)$ is a point and $F_d(\Delta)$ is an $(n-1)$ -simplex in \mathcal{P}_{n-1} . By the induction hypothesis, $F_d(\Delta)$ is not a zero summand. Also, $F_{-d}(\Delta)$ has zero volume. By Lemma 5.1, applied in contraposition, it follows that Δ is not a zero summand. \square

For the rest of the paper, a linear function $g: \mathbb{R}^n \rightarrow \mathbb{R}$ is a function of the form $g(x) := a^T x$, where $a \in \mathbb{R}^n$.

Lemma 5.3 *Let g_0, \dots, g_n be linear functions, whose derivatives are affinely independent. Then the positively homogeneous function $\max(g_0, \dots, g_n)$ is not a linear combination of maxima of at most n linear functions.*

Proof Assume, indirectly, that $\max(g_0, \dots, g_n) = \alpha_1 f_1 + \dots + \alpha_k f_k$, where f_1, \dots, f_k are maxima of at most n linear functions. Without loss of generality, we may assume that the α_i are either 1 or -1 . Let us assume that $\alpha_1 = \dots = \alpha_l = 1$ and $\alpha_{l+1} = \dots = \alpha_k = -1$. Hence $\max(g_0, \dots, g_n) + f_{l+1} + \dots + f_k = f_1 + \dots + f_l$. By Proposition 4.1, we obtain

$$\tau(\max(g_0, \dots, g_n)) + P_{l+1} + \dots + P_k = P_1 + \dots + P_l.$$

For $i = 1, \dots, k$, the function f_i is in \mathcal{F}_n , and $P_i := \tau(f_i)$ is a zero volume polytope because it is the convex hull of at most n points. This shows that $\tau(\max(g_0, \dots, g_n))$ is a zero summand. But this contradicts Corollary 5.2, because $\tau(\max(g_0, \dots, g_n))$ is an n -simplex. \square

Theorem 5.4 *Let $f: \mathbb{R}^n \rightarrow \mathbb{R}$ be a piecewise linear function such that for every $(x_1, \dots, x_n) \in \mathbb{R}^n$ holds:*

$$(x_1, \dots, x_n) \mapsto \max(0, x_1, \dots, x_n).$$

The function f is not a linear combination of maxima of less than $n + 1$ affine-linear functions.

Proof By Lemma 5.3, the function f is not a linear combination of maxima of less than $n + 1$ linear functions: Assume, indirectly, that there are integers s, k_1, \dots, k_s with $k_i \leq n$ for $i = 1, \dots, s$, real numbers $\alpha_1, \dots, \alpha_s$ and linear functions $g_{i,j}: \mathbb{R}^n \rightarrow \mathbb{R}$, $i = 1, \dots, s$, $j = 1, \dots, k_i$ such that

$$m_i := \max_{j=1, \dots, k_i} (g_{i,j}) \text{ and } f = \sum_{i=1}^s \alpha_i m_i.$$

We will then construct a representation of f as a linear combination of maxima of less than $n + 1$ linear functions, giving a contradiction.

For $i = 1, \dots, s$, we proceed as follows. We define $c_i := m_i(\mathbf{0})$. We may assume without loss of generality that there exists $r_i \leq k_i$ such that $g_{i,j}(\mathbf{0}) = c_i$ if $j \leq r_i$ and $g_{i,j}(\mathbf{0}) < c_i$ if $j > r_i$. For $j = 1, \dots, r_i$, we set $h_{i,j} := g_{i,j} - c_i$. Then $h_{i,j}(\mathbf{0}) = 0$, which implies that the functions $h_{i,j}$ are all linear. Now we set

$$n_i := \max_{j=1, \dots, r_i} (h_{i,j}) \text{ and } e := \sum_{i=1}^s \alpha_i n_i.$$

Then e is a linear combination of maxima of less than $n + 1$ linear functions. We will prove that $e = f$, which will finish the indirect proof.

Let U be a small neighborhood of $\mathbf{0}$ such that for each i , we have

$$\max_{j=1, \dots, k_i} g_{i,j} = \max_{j=1, \dots, r_i} g_{i,j} = \max_{j=1, \dots, r_i} (h_{i,j} + c_i) = \left(\max_{j=1, \dots, r_i} h_{i,j} \right) + c_i$$

inside U . Then we have $m_i = n_i - c_i$ inside U and therefore

$$f = \sum_{i=1, \dots, s} \alpha_i(n_i + c_i) = e + \sum_{i=1}^s c_i.$$

However, $\sum_{i=1}^s c_i = f(\mathbf{0}) = 0$, and therefore $f = e$ inside U . Both functions f and e are positively homogeneous; so, if they coincide in U , then they coincide everywhere. \square

6 Conclusion

It has been shown that any piecewise linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ can be represented as a linear combination of maxima with at most $n + 1$ arguments, where the linear arguments of each maximum are picked from the set of affine-linear parts of the function f . We develop an algorithm for calculating this representation. It is an open question whether the derived representation is invariant under certain choices that can be made inside Algorithm 2. After running a series of experiments, as illustrated in Example 3.1, we conjecture that this is the case. Proving this conjecture could be a possible direction for future research.

By proving that the function $\max(0, x_1, \dots, x_n)$ is not a linear combination of maxima of less than $n + 1$ affine-linear functions, we confirm the optimal representation conjecture formulated by Shuning Wang and Xusheng Sun in [12]. Using these two contributions, we can state that every piecewise linear function f can be expressed as a ReLU neural network with at most $\lceil \log_2(n + 1) \rceil$ layers and $\mathcal{O}(p^{2n^2+3n+1})$ neurons in every hidden layer, where p is the number of constituents in the function f , see Theorem 4.4 [5]. As a direct implication of Theorem 2.4, we get the refinement of the upper bound of the number of neurons for every hidden layer that equals $\mathcal{O}(p^{n+1})$.

References

1. Arora, R., Basu, A., Mianjy, P., Mukherjee, A.: Understanding deep neural networks with rectified linear units. In: ICLR (2018)
2. Chua, L., Deng, A.C.: Canonical piecewise-linear representation. IEEE Trans. Circuits Syst. **35**(1), 101–111 (1988)
3. Goodfellow, I., Bengio, Y., Courville, A.: Deep learning, pp. 164–172. MIT Press (2016)
4. Gritzmann, P., Sturmfels, B.: Minkowski addition of polytopes: computational complexity and applications to Gröbner bases. SIAM J. Discrete Math. **6**(2), 246–269 (1993)
5. Hertrich, C., Basu, A., Di Summa, M., Skutella, M.: Towards lower bounds on the depth of relu neural networks. NeurIPS **34**, 3336–3348 (2021)
6. Kripfganz, A., Schulze, R.: Piecewise affine functions as a difference of two convex functions. Optimization **18**(1), 23–29 (1987)
7. Montúfar, G., Ren, Y., Zhang, L.: Sharp bounds for the number of regions of maxout networks and vertices of minkowski sums. SIAM J. Appl. Algebra Geom. **6**(4), 618–649 (2022)
8. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: ICML, p. 807814. Omnipress (2010)

-
9. Ovchinnikov, S.: Max-min representation of piecewise linear functions. *Beitr. Algebra Geom.* **43**(1), 297–302 (2002)
 10. Schlüter, N., Darup, M.S.: Novel convex decomposition of piecewise affine functions. In: *Proc. 21th IFAC World Congress*. Elsevier (2020)
 11. Tarela, J., Martinez, M.: Region configurations for realizability of lattice piecewise-linear models. *Math. Comput. Model.* **30**(11), 17–27 (1999)
 12. Wang, S., Sun, X.: Generalization of hinging hyperplanes. *IEEE Trans. Inf. Theory* **51**, 4425–4431 (2005)