

Convergence of a Neural Network for Sparse Approximation using the Nonsmooth Łojasiewicz Inequality

Aurèle Balavoine, Christopher J. Rozell, Justin Romberg

Abstract—Sparse approximation is an optimization program that produces state-of-the-art results in many applications in signal processing and engineering. To deploy this approach in real-time, it is necessary to develop faster solvers than are currently available in digital. The Locally Competitive Algorithm (LCA) is a dynamical system designed to solve the class of sparse approximation problems in continuous time. But before implementing this network in analog VLSI, it is essential to provide performance guarantees. This paper presents new results on the convergence of the LCA neural network. Using recently-developed methods that make use of the Łojasiewicz inequality for nonsmooth functions, we prove that the output and state trajectories converge to a single fixed point. This improves on previous results by guaranteeing convergence to a singleton even when the optimization program has infinitely many and non-isolated solution points.

I. INTRODUCTION

SPARSE approximation decomposes a signal $y \in \mathbb{R}^M$ in an overcomplete dictionary $\Phi \in \mathbb{R}^{M \times N}$, with $M \ll N$, by constraining the approximation coefficients $a \in \mathbb{R}^N$ to be sparse (i.e. to have only a few non-zero entries). The Locally Competitive Algorithm (LCA) introduced in [1] is a neural network designed to solve this problem and is defined by the following differential equation:

$$\begin{aligned} \tau \dot{u}(t) &= -u(t) - (\Phi^T \Phi - I) a(t) + \Phi^T y \\ a(t) &= T_\lambda(u(t)) \end{aligned} \quad (1)$$

The LCA network (whose architecture is shown in Fig. 1) takes as input the vector $\Phi^T y$. The components of the input drive the vector $u(t)$, which contains the state variables $u_n(t)$ for $n = 1, \dots, N$. The columns of the matrix Φ can be viewed as elements of the dictionary. We assume that they have unit norm and denote them by $\Phi_n \in \mathbb{R}^M$ for $n = 1, \dots, N$. The outputs of the system are the $a_n(t)$ for $n = 1, \dots, N$. They are generated by the nonlinear activation function $T_\lambda(\cdot)$, which is applied entry-wise to the state vector $u(t)$. Each output generates a feedback into each state. The strength of the feedback depends on the level of output activity and on the strength of the inner product between two dictionary elements. The values of these inner products are represented by the interconnection matrix $W = \Phi^T \Phi - I$. This structure ensures that two non-zero coefficients do not carry the same information about the signal. The time constant τ depends on the physical solver; for our analysis, we take $\tau = 1$ without loss of generality.




Fig. 1. Block diagram of the LCA network. The state $u(t)$ is driven by the projection of the signal y onto each of the N dictionary elements Φ_n . The state produces the output $a(t)$ through the activation function $T_\lambda(\cdot)$. The output is then weighted by the interconnection matrix W and fed back.

The LCA is a type of Hopfield-style network [2], and as such a Lyapunov function can be designed for it. An appropriate Lyapunov function for (1) is the typical objective function used to solve sparse approximation problems:

$$V(a) = \frac{1}{2} \|y - \Phi a\|_2^2 + \lambda \sum_{n=1}^N C(a_n). \quad (2)$$

The first term is the mean-squared error of the approximation, while the second term $C(\cdot)$ is a cost penalty on the solution that encourages sparsity. The parameter λ is a tradeoff between these two objectives. The most famous sparse approximation program is ℓ_1 -minimization, also known as Basis Pursuit Denoising. This optimization program plays an important role in signal processing, and in particular in Compressed Sensing, since it allows to recover a sparse signal from many fewer measurements than traditional approaches [3]. For this program, the cost function $C(\cdot)$ is the absolute value, resulting in an ℓ_1 -norm penalty on a .

Theoretical guarantees on system performance (e.g., convergence, convergence speed, etc.) are an important counterpart to ongoing work implementing the LCA in analog circuitry [7], [8]. To this end, in a previous paper [4], we showed that if the activation function $T_\lambda(\cdot)$ and the cost function $C(\cdot)$ satisfy a certain relationship, the fixed points of (1) correspond to the critical points of (2). In addition, using a Lyapunov approach and under certain conditions on the activation function, we showed that the outputs of the network converge to the set of fixed points. When the solution is unique or when the objective function $V(\cdot)$ is strictly convex, this implies convergence to a singleton. However, this result is insufficient to prove that the outputs converge to a single point when solutions of (2) are

This work was partially supported by NSF grant CCF-0905346. Authors are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0250
Email: {aurele.balavoine, crozell, jrom} @gatech.edu.

not isolated. Under some additional conditions on the problem parameters (such as the eigenvalues of certain submatrices of $\Phi^T \Phi$), we also proved that the LCA converges to a single fixed point with exponential rate of convergence [4].

Recently, several papers have developed a new technique based on the Łojasiewicz inequality [5]. Using this inequality, the output of certain networks can be shown to converge to a singleton even when the fixed points are not isolated. However, the specifics of the LCA network prevent us from applying these results directly. In particular, the activation function is zero on some interval and may be unbounded. In addition, the interconnection matrix W may be singular.

The main contribution of this work is to apply a variation of the Łojasiewicz inequality for nonsmooth functions [6] to show two results. First, without assuming that the critical points of (2) are isolated, we show that the output $a(t)$ of the network converges to a single fixed point when starting from any initial point, i.e. $a(t)$ is *globally asymptotically convergent*.

Theorem 1. *Under conditions (3)-(7), the output $a(t)$ of (1) is globally asymptotically convergent, i.e. $\exists a^* \in \mathbb{R}^N$ such that $a(t) \rightarrow a^*$, as $t \rightarrow +\infty$.*

Second, we prove the even stronger result that the state $u(t)$ also converges to a single fixed point.

Theorem 2. *Under conditions (3)-(7), the state $u(t)$ of (1) is globally asymptotically convergent, i.e. $\exists u^* \in \mathbb{R}^N$ such that $u(t) \rightarrow u^*$, as $t \rightarrow +\infty$.*

The necessary hypotheses (3)-(7) on the network are given in Section II, along with a survey of previous works that use the Łojasiewicz inequality and their limitations. Section III gives a summary of the necessary mathematical notions along with several lemmas. Finally, Section IV provides the proofs of the two main results.

II. BACKGROUND

A. Hypotheses

We proved in [4] that the fixed points of the network (1) coincide with the critical points of the objective function (2) if the activation function $T_\lambda(\cdot)$ and the cost function $C(\cdot)$ satisfy

$$u_n - a_n = u_n - T_\lambda(u_n) \in \lambda \partial C(a_n). \quad (3)$$

In addition, we showed that the objective function (2) is decreasing along the network trajectories if the activation function has the form

$$a_n(t) = T_\lambda(u_n(t)) = \begin{cases} 0, & |u_n(t)| \leq \lambda \\ f(u_n(t)), & |u_n(t)| > \lambda \end{cases}, \quad (4)$$

where the function $f(\cdot)$ is a real-valued function defined on $\mathcal{D} = (-\infty, -\lambda] \cup [\lambda, +\infty)$, is continuous on \mathcal{D} , differentiable on the interior of \mathcal{D} , and satisfies the following properties:

$$f(-u_n) = -f(u_n), \quad \forall u_n \in \mathcal{D} \quad \text{and} \quad f(\lambda) = 0 \quad (5a)$$

$$f'(u_n) > 0, \quad \forall u_n \in \text{int}(\mathcal{D}) \quad (5b)$$

$$f(u_n) \leq u_n, \quad \forall u_n \in \mathcal{D} \quad \text{s.t.} \quad u_n \geq 0. \quad (5c)$$




Fig. 2. Example of activation functions satisfying conditions (5)-(7). The plain dark line corresponds to the soft-threshold function for ℓ_1 -minimization.

Two activation functions satisfying these conditions are shown in Fig. 2. Conditions (4) and (5) ensure that the state and output trajectories are continuous for all time and that the cost function $C(\cdot)$ is strictly increasing with the absolute value of the outputs.

Under conditions (3)-(5), $V(\cdot)$ is a Lyapunov function for the network. Using this fact, we showed in [4] that the output trajectories of (1) converge to the set of fixed points satisfying $\{a \in \mathbb{R}^N \text{ s.t. } \dot{a}(t) = 0\}$. When the critical points of (2) are isolated, we proved that the output and state variables both converge to a singleton. This is the case, for instance, when the objective function $V(\cdot)$ is strictly convex. However, when the points in this set are not isolated, Lyapunov theory is insufficient to guarantee that the output converges to a singleton.

For the purpose of this paper, we add the following requirements:

$$f(\cdot) \text{ is subanalytic on } \mathcal{D} \quad \text{and} \quad (6)$$

$$\forall U > \lambda, \exists \alpha_U > 0 \text{ s.t. } \forall u_n \in [\lambda, U] \quad f'(u_n) \leq \alpha_U. \quad (7)$$

The additional condition (6) ensures that the cost $C(\cdot)$ and thus the objective $V(a)$ are subanalytic. This notion concerns geometric properties of the graph of a function and will be reviewed in detail in Section III and will allow us to apply the Łojasiewicz inequality for nonsmooth functions. Condition (7) ensures that $f'(\cdot)$ remains bounded on bounded intervals. If $f'(\cdot)$ is continuous, then (7) holds immediately. However, (6) and (7) are more general conditions, and $f'(\cdot)$ need not be continuous. The activation functions satisfying these conditions correspond to a large class of cost functions that are often used in practice [9]. In particular, the soft-thresholding function (black plain line in Fig. 2) satisfies all of the requirements and leads to the ℓ_1 -minimization program.

B. The Łojasiewicz inequality

Since techniques based on Lyapunov functions only guarantee convergence to a set of fixed points, recent papers have developed a new technique based on the Łojasiewicz inequality

[5] to overcome this limitation. This inequality states that for all $\bar{x} \in \mathbb{R}^N$, there exists $\nu \in [0, 1)$, $C > 0$ and $\Delta > 0$ such that the gradient of a real-analytic function $F : \mathbb{R}^N \rightarrow \mathbb{R}$ satisfies:

$$|F(x) - F(\bar{x})|^\nu \leq C \|\nabla F(x)\| \quad \forall x \in B_\Delta(\bar{x}).$$

Using this inequality, the trajectories of certain networks can be shown to be finite, ensuring convergence to a singleton even when the fixed points are not isolated. In [10], a general approach is taken where the network's equation has the form

$$\begin{aligned} \dot{u}(t) &= -Du(t) - \nabla F(a(t)) \\ a(t) &= T(u(t)) \end{aligned}$$

In [10], the functions $F(\cdot)$ and $T(\cdot)$ are assumed to be analytic (which implies the existence of derivatives of any order), and the activation function $T(\cdot)$ is required to be bounded and strictly increasing. Recently, an extension of the Łojasiewicz inequality was developed for nonsmooth functions in [6]. In addition, the authors of [6] show how a network satisfying the differential inclusion $\dot{u}(t) \in -\partial F(u(t))$ has finite-length trajectories if $F(\cdot)$ is subanalytic and either lower semicontinuous convex or lower- C^2 . The notation $\partial F(x)$ represents the subgradient of a function $F(\cdot)$ at x (see Section III for a more precise definition). Using this result, the study in [10] was extended to the nonsmooth case in [11] for networks satisfying

$$\dot{u}(t) \in -\partial F(u(t)),$$

where the function $F(\cdot)$ is subanalytic and the network solves a quadratic program with linear constraints. Finally, the paper [12] also makes use of the nonsmooth Łojasiewicz inequality to prove that a network of the form

$$\begin{aligned} \dot{u}(t) &\in -Du(t) - \partial V(a(t)) + \theta \\ a(t) &= T(u(t)) \end{aligned}$$

converges to a singleton. In [12], the activation function must be bounded and the diagonal matrix D must have strictly positive entries.

We will show in Section III that the LCA network (1) satisfies the differential inclusion

$$\dot{u}(t) \in -\partial V(a(t)).$$

Despite its similarities to previous studies, the specifics of the LCA network prevent us from applying the results in the works cited above. Specifically, to meet the requirements of sparse approximation, the activation function $T_\lambda(\cdot)$ forces many outputs to be zero by having a value of zero on $[-\lambda, \lambda]$. In addition, the activation function is generally nonsmooth and unbounded. This last property discourages outputs from growing away from zero. Finally, the interconnection matrix W may have both positive and negative eigenvalues, as well as a significantly large nullspace. These characteristics may lead the objective function $V(a(t))$ to remain constant while the state vector $u(t)$ is still evolving. Nevertheless, these previous studies have inspired the authors to apply similar techniques based on the nonsmooth Łojasiewicz inequality to

prove convergence of the LCA network in the case of non-isolated fixed points.

III. PRELIMINARIES

The keys to applying the Łojasiewicz inequality are the subanalyticity of the objective function $V(a)$ and the properties of its subgradient. We adopt the definition of subgradient used in [13] and [14]. For consistency, in this section we review some definitions from nonsmooth analysis and the Łojasiewicz inequality for nonsmooth functions. As we define these notions, we also apply them to the LCA network and develop several lemmas that will be useful in the proofs of the main theorems.

A. Subgradient of the objective

The usual *one-sided directional derivative* of a function $F : \mathbb{R}^N \rightarrow \mathbb{R}$ at $x \in \mathbb{R}^N$ in the direction $v \in \mathbb{R}^N$ is

$$F'(x; v) = \lim_{t \downarrow 0} \frac{F(x + tv) - F(x)}{t}.$$

Since some nonsmooth functions may fail to admit one-sided derivatives, this definition is relaxed to the following notion of *generalized directional derivative*:

$$F^\circ(x; v) = \limsup_{\substack{y \rightarrow x \\ t \downarrow 0}} \frac{F(y + tv) - F(y)}{t}.$$

With this definition, the existence of directional derivatives of $F(\cdot)$ at x are not necessary. For instance, the quantity $F^\circ(x; v)$ is well-defined when $F(\cdot)$ is only Lipschitz near x . Using this quantity, the *subgradient* of $F(\cdot)$ at $x \in \mathbb{R}^N$ is defined as the set

$$\partial F(x) = \{ \xi \in \mathbb{R}^N \text{ s.t. } F^\circ(x; v) \geq \xi^T v, \quad \forall v \in \mathbb{R}^N \}.$$

When $F(\cdot)$ is smooth at x , $\partial F(x)$ is a singleton that coincides with the classic notion of gradient $\partial F(x) = \{ \nabla F(x) \}$.

Note that for all $u_n \in \mathcal{D}$, $f(u_n)$ is well-defined and Eq. (3) implies that $\partial C(a_n)$ reduces to a single point:

$$\partial C(a_n) = \frac{dC(a_n)}{da_n} = u_n - f(u_n).$$

In other words, $C(\cdot)$ is differentiable on $\mathbb{R} \setminus \{0\}$. As a consequence, $V(\cdot)$ is also differentiable almost everywhere (a.e.), except at points $a \in \mathbb{R}^N$ such that $a_n = 0$ for some $n = 1, \dots, N$. Because $V(\cdot)$ is differentiable a.e., its subgradient simplifies to the following definition [13]:

$$\partial V(a) = co \left\{ \lim_{i \rightarrow \infty} \nabla V(b_i) : b_i \rightarrow a, b_i \notin \mathcal{S}, b_i \notin \Omega_V \right\},$$

where co is the convex hull, Ω_V is the set of points where $V(\cdot)$ fails to be differentiable, and \mathcal{S} is any set of Lebesgue measure 0 in \mathbb{R}^N . In other words, $\partial V(a)$ is the smallest convex set containing the limit points of the gradients of the function along any sequence of points $\{b_i\}$ approaching a while avoiding $\Omega_V \cup \mathcal{S}$. When $F'(x; v)$ exists and $F'(x; v) = F^\circ(x; v)$ for all $v \in \mathbb{R}^N$, the function $F(\cdot)$ is said to be *regular* at x [13, Def. 2.3.4]. $C(\cdot)$ admits left and right derivatives and is clearly regular on \mathbb{R} . This implies that $V(\cdot)$ is regular

on \mathbb{R}^N . Consequently, proposition 2.3.3 of [13] holds with equality and yields

$$\partial V(a(t)) = -\Phi^T y + \Phi^T \Phi a(t) + \lambda \partial C(a(t)), \quad (8)$$

where $\partial C(a(t)) = [\partial C(a_1(t)), \dots, \partial C(a_N(t))]^T$. Finally, condition (3) yields

$$-\dot{u}(t) \in \partial V(a(t)). \quad (9)$$

In order to compute the time derivative of $V(\cdot)$ along the LCA trajectories, we use the following result [13, Th. 2.3.9 (iii)], which is a generalization of the chain rule for regular functions.

Theorem (Chain Rule). *Assume $F(x) : \mathbb{R}^N \rightarrow \mathbb{R}$ is regular in \mathbb{R}^N and $x(t) : [0, +\infty) \rightarrow \mathbb{R}^N$ is locally absolutely continuous on $[0, +\infty)$. Then, $F(x(t))$ is regular on \mathbb{R}^N . In addition, its time derivative $\dot{F}(x(t))$ exists for almost all $t \geq 0$ and satisfies*

$$\dot{F}(x(t)) = \frac{d}{dt} F(x(t)) = \zeta^T \dot{x}(t) \quad \forall \zeta \in \partial F(x(t)). \quad (10)$$

Using this result, we can compute the time derivative of $V(\cdot)$ along the LCA trajectories. Noting that $V(a(t))$ only depends on the outputs $a_n(t)$ that are non-zero, it is useful to define the set $\Gamma(t) = \{n \in \{1, \dots, N\} \text{ s.t. } |u_n(t)| > \lambda\}$. We call this set of indices the *active set* and the corresponding $u_n(t)$ and $a_n(t)$ the *active nodes*. The active set is time-dependent, since state variables may cross the threshold λ in either direction. However, for readability purposes, we remove the explicit dependence on time in the notation and write it as Γ . The notations $a_\Gamma(t)$ and $u_\Gamma(t)$ refer to the vectors $a(t)$ and $u(t)$ where entries outside of Γ are set to zero. In particular, $a(t) = a_\Gamma(t)$ since $a_n(t) = 0$ for $n \in \Gamma^c$. The following lemma gives two characterizations of the time derivative of $V(\cdot)$ in terms of the active nodes.

Lemma 1. *The time derivative of $V(a(t))$ satisfies the two equalities*

$$\dot{V}(a(t)) = - \sum_{n \in \Gamma} f'(u_n(t)) |\dot{u}_n(t)|^2 \quad (11)$$

$$\dot{V}(a(t)) = - \sum_{n \in \Gamma} \frac{1}{f'(u_n(t))} |\dot{a}_n(t)|^2 \quad (12)$$

for almost all $t \geq 0$.

Proof. Since $V(a(t))$ is regular, we can choose any element in $\partial V(a(t))$ to compute the time derivative of $V(\cdot)$ along the trajectories of the neural network. In particular, we can pick $-\dot{u}(t) \in \partial V(a(t))$. Moreover, active nodes satisfy $a_n(t) = f(u_n(t))$, and using the usual chain rule we get $\dot{a}_n(t) =$

$f'(u_n(t))\dot{u}_n(t)$. As a consequence, (10) yields

$$\begin{aligned} \dot{V}(a(t)) &= -\dot{u}(t)^T \dot{a}(t) \\ &= - \sum_{n=1}^N \dot{u}_n(t) \dot{a}_n(t) \\ &= - \sum_{n \in \Gamma} f'(u_n(t)) |\dot{u}_n(t)|^2 \\ &= - \sum_{n \in \Gamma} \frac{1}{f'(u_n(t))} |\dot{a}_n(t)|^2, \end{aligned}$$

which completes the proof. \square

By condition (5b), $f'(u_n) > 0$ for all $n \in \Gamma$, so this lemma demonstrates that the objective function is strictly decreasing on non-stationary output trajectories.

B. Properties of the network

The following lemma gives some properties of the cost function that will be useful in the proof of the main results.

Lemma 2. *Without loss of generality, assume that $C(0) = 0$. Then, conditions (3) and (5) on the activation function yield the following properties:*

$$C(a_n) \geq 0 \quad \text{and} \quad C(a_n) = C(-a_n) \quad \forall a_n \in \mathbb{R}^N \quad (13)$$

$$\text{sign}(u_n) = \text{sign}(a_n) \quad \forall u_n \in \mathcal{D} \quad (14)$$

$$|a_n|^2 \leq u_n a_n \leq |u_n|^2 \quad \forall u_n \in \mathbb{R}. \quad (15)$$

Proof. Since $f(\cdot)$ is continuous and strictly increasing on \mathcal{D} , $f^{-1}(\cdot)$ is well-defined and strictly increasing on $f(\mathcal{D})$. In addition, property (5a) of $f(\cdot)$ implies that $f^{-1}(\cdot)$ satisfies $f^{-1}(-a_n) = -f^{-1}(a_n)$ and $f^{-1}(\lambda) = 0$. As a consequence, $C(\cdot)$ is continuous on \mathbb{R} , and for all $a_n > 0$ we see that $dC(a_n)$ exists and is equal to

$$dC(a_n) = u_n - a_n = u_n - f(u_n) = f^{-1}(a_n) - a_n.$$

This quantity is positive by (5c). This proves that $C(a_n) \geq C(0) = 0$ for all $a_n > 0$. Moreover, for all $a_n > 0$, the following holds:

$$\begin{aligned} C(-a_n) &= \int_0^{-a_n} dC(s) \\ &= \int_0^{-a_n} (f^{-1}(s) - s) ds \\ &= \int_0^{a_n} (f^{-1}(-s) + s) (-ds) \\ &= \int_0^{a_n} (f^{-1}(s) - s) ds = C(a_n). \end{aligned}$$

So $C(-a_n) = C(a_n)$ for all $a_n \in \mathbb{R}$ and (13) holds.

By (5), for all $u_n > \lambda$, we have

$$\begin{aligned} 0 &< a_n = f(u_n) \leq u_n \\ -u_n &\leq -f(u_n) = f(-u_n) = -a_n < 0, \end{aligned}$$

which proves (14), i.e. $\text{sign}(a_n) = \text{sign}(u_n)$. This also shows that $a_n u_n = \text{sign}(a_n) |a_n| \text{sign}(u_n) |u_n| = |a_n| |u_n|$ and $|a_n| \leq |u_n|$ for all $u_n \in \mathcal{D}$. This last inequality can be

extended to \mathbb{R} , since for $|u_n| \leq \lambda$, $a_n = 0$. Finally, for all $u_n \in \mathbb{R}$, we obtain:

$$|a_n|^2 \leq |a_n| |u_n| = a_n u_n \leq |u_n|,$$

which proves (15). \square

Note that we could choose any value for $C(0)$. In all cases, the objective function $V(\cdot)$ will be lower bounded by $\lambda NC(0)$, and a lower-bound on $V(\cdot)$ is all that is required in the proofs. Taking $C(0) = 0$ simplifies the lower bound to $V(\cdot) \geq 0$ on all of \mathbb{R}^N . Using these properties, the following lemma states that the objective function is also upper-bounded for all time, and so are the output and state variables.

Lemma 3. *For all $t \geq 0$, we have $V(a(t)) \leq V(a(0))$. In addition, the output $a(t)$ and state variables $u(t)$ of the system (1) are bounded for all $t \geq 0$.*

Proof. From (12) and (5b), we have that $\dot{V}(a(t)) \leq 0$ for almost all $t \geq 0$. As a consequence, $V(a(t))$ is non-increasing and for all $t > 0$ we have:

$$V(a(t)) - V(a(0)) = \int_0^t \dot{V}(a(s)) ds.$$

Since $0 < t$ and $\dot{V}(a(t)) \leq 0$ for almost all $s \in (0, t)$, by the positivity of the integral we see that $V(a(t)) \leq V(a(0))$ for all $t \geq 0$.

Next, we show that the state $u(t)$ is bounded. For this result, we begin by showing that both $\|\Phi a(t)\|_2$ and $\|\Phi u(t)\|_2$ are bounded for all $t \geq 0$. Condition (13) guarantees that $C(a_n) \geq 0$ for all $a_n \in \mathbb{R}$, so for all $t \geq 0$ we have:

$$\frac{1}{2} \|y - \Phi a(t)\|_2^2 \leq V(a(t)) \leq V(a(0)).$$

The triangle inequality yields

$$\|\Phi a(t)\|_2 - \|y\|_2 \leq \sqrt{2V(a(0))}.$$

This shows that $\|\Phi a(t)\|_2$ is bounded for all $t \geq 0$. For this reason, there must exist a constant $C_1 \geq 0$ such that, for all $t \geq 0$,

$$\|(I - \Phi \Phi^T) \Phi a(t) + \Phi^T y\|_2 \leq \sigma_1 \|\Phi a(t)\|_2 + \|\Phi^T y\|_2 \leq C_1,$$

where $\sigma_1 \geq 0$ is the largest eigenvalue of the interconnection matrix $W = \Phi \Phi^T - I$. This inequality implies that $\|\Phi u(t)\|_2$ is also bounded for $t \geq 0$. Indeed, using the Cauchy-Schwartz inequality, the time-derivative of $1/2 \|\Phi u(t)\|_2^2$ satisfies

$$\begin{aligned} \frac{d}{dt} \frac{1}{2} \|\Phi u(t)\|_2^2 &= u(t) \Phi^T \dot{\Phi} u(t) \\ &= u^T \Phi^T \Phi (-u(t) + a(t) - \Phi^T \Phi a(t) + \Phi^T y) \\ &\leq -\|\Phi u(t)\|_2^2 + \|\Phi u(t)\|_2 C_1 \\ &\leq -\|\Phi u(t)\|_2 (\|\Phi u(t)\|_2 - C_1). \end{aligned}$$

As a consequence, the set $\{u \in \mathbb{R}^N \text{ s.t. } \|\Phi u\|_2 \leq C_1\}$ is attractive, and by continuity, $\|\Phi u(t)\|_2$ is bounded for all $t \geq 0$. We cannot conclude directly that $\|u(t)\|_2$ is bounded because the matrix Φ may be singular. Any vector u in its nullspace can grow unbounded while $\|\Phi u\|_2$ remains bounded.

However, $u(t)$ can be decomposed into its component $u_1(t)$ that lies in the nullspace of Φ and its component $u_2(t)$ that lies in the range of Φ^T . These two vectors are orthogonal (this comes from the singular value decomposition of Φ), and we will show that each of them is bounded. Since $u_1(t)$ is in the nullspace of Φ , we have $\Phi u(t) = \Phi u_2(t)$. Since $u_2(t)$ is in the range of Φ^T , $\exists x_2(t) \in \mathbb{R}^M$ such that $u_2(t) = \Phi^T x_2(t)$. Using the Cauchy-Schwartz inequality, we find

$$\begin{aligned} \|x_2(t)\|_2 \|\Phi u(t)\|_2 &\geq x_2(t)^T \Phi u(t) \\ &= x_2(t)^T \Phi u_2(t) \\ &= x_2(t)^T \Phi \Phi^T x_2(t) \geq \sigma_2^2 \|x_2(t)\|_2^2, \end{aligned}$$

where $\sigma_2 > 0$ is the smallest singular value of Φ^T restricted to its range (so it is strictly positive). Letting σ_3 be the largest singular value of Φ^T , we obtain

$$\|u_2(t)\|_2 = \|\Phi^T x_2(t)\|_2 \leq \sigma_3 \|x_2(t)\|_2 \leq \sigma_3 \sigma_2^{-2} \|\Phi u(t)\|_2.$$

Since $\|\Phi u(t)\|_2$ is bounded, so is $\|u_2(t)\|_2$. On the other hand, using the fact that $\Phi u_1(t) = 0$, we can compute the time-derivative of $1/2 \|u_1(t)\|_2^2$ as follows:

$$\begin{aligned} \frac{d}{dt} \frac{1}{2} \|u_1(t)\|_2^2 &= u_1(t)^T \dot{u}_1(t) \\ &= u_1(t)^T (-u(t) + a(t) + \Phi^T y - \Phi^T \Phi a(t))_1 \\ &= -u_1(t)^T u_1(t) + u_1(t)^T a_1(t) \leq 0, \end{aligned}$$

where the last inequality follows from (15). We conclude that $\|u_1(t)\|_2$ is also bounded for all $t \geq 0$. This shows that $\|u(t)\|_2 \leq \|u_1(t)\|_2 + \|u_2(t)\|_2$ is bounded $\forall t \geq 0$.

Finally, since $f(\cdot)$ is continuous on \mathcal{D} and $\|u(t)\|_2$ is bounded for all $t \geq 0$, $\|f(u(t))\|_2$ is also bounded for all $t \geq 0$. The elements of the output $a(t)$ are either equal to zero or to $f(u(t))$, which shows that $\|a(t)\|_2$ is also bounded. \square

Lemma 3 states that $u(t)$ is bounded for all $t \geq 0$, which together with conditions (5b) and (7) guarantee that there exists $+\infty > \alpha \geq \beta > 0$ such that $\forall n \in \mathbb{N}$

$$\alpha \geq f'(u_n(t)) \geq \beta \quad \forall t \geq 0. \quad (16)$$

The inequalities hold for all time and the two constants α and β will be used in the proof of the first main result.

C. Subanalyticity of the objective function

Finally, we show that $V(\cdot)$ is subanalytic and state the Łojasiewicz inequality for nonsmooth functions.

A function is subanalytic if its graph obeys certain geometric properties. This notion involves algebraic manipulations of sets defined by real-analytic equations and inequalities. More precisely, a set $A \subset \mathbb{R}^N$ is said to be *semianalytic* if each point $x \in \mathbb{R}^N$ admits a neighborhood \mathcal{N} for which

$$A \cap \mathcal{N} = \bigcup_{i=1}^p \bigcap_{j=1}^q \{x \in \mathcal{N}, f_{ij}(x) = 0, g_{ij}(x) > 0\}$$

where $f_{ij}, g_{ij} : \mathcal{N} \rightarrow \mathbb{R}$ are real-analytic functions for all $1 \leq i \leq p$, $1 \leq j \leq q$, and p and q are some integers. The set A is said to be *subanalytic* if it is locally the projection

of a semianalytic set, i.e. each point $x \in \mathbb{R}^N$ admits a neighborhood \mathcal{N} such that $A \cap \mathcal{N} = \{x \in \mathbb{R}^N, (x, y) \in B\}$, where B is a bounded semianalytic subset of $\mathbb{R}^N \times \mathbb{R}^M$ for some $M \geq 1$. A function $F : \mathbb{R}^N \rightarrow \mathbb{R}$ is said to be *subanalytic* if its graph, $\text{Graf}F = \{(x, y) \text{ s.t. } y = F(x)\}$, is a subanalytic subset of $\mathbb{R}^N \times \mathbb{R}$.

From condition (6), $f(\cdot)$ is subanalytic, and so $V(a)$ is also subanalytic. Indeed, we can write the graph of $V(\cdot)$ as the projection onto the first and last component of the set

$$\begin{aligned} & \left\{ a, v_1, a_2, v_2, v \in \mathbb{R}^{2N+3} \text{ s.t. } \frac{1}{2} \|y - \Phi a\|_2 = v_1, \right. \\ & \quad \lambda C(a_2) = v_2, \quad a = a_2, \quad v = v_1 + v_2 \} \\ & = (\text{Graf}F_1 \times \text{Graf}F_2 \times R) \cap \left\{ a, a_2, v_1, v_2, v \in \mathbb{R}^{2N+3} \right. \\ & \quad \left. \text{s.t. } a = a_2, \quad v = v_1 + v_2 \right\}, \end{aligned}$$

where $F_1(a) = \frac{1}{2} \|y - \Phi a\|_2^2$ and $F_2(a) = \lambda C(a)$ are subanalytic.

The following theorem gives the Łojasiewicz inequality for nonsmooth functions [6, Th 3.1.]. It provides some bound on the decay of the function in terms of its nonsmooth slope. The *nonsmooth slope* of $F(\cdot)$ at $x \in \mathbb{R}^N$ is defined as

$$m(\partial F(x)) = \inf \{ \|\xi\|_2, \xi \in \partial F(x) \},$$

and represents the smallest norm in the set $\partial F(x)$.

Theorem (Nonsmooth Łojasiewicz inequality). *Suppose that a function $F : \mathbb{R}^N \rightarrow \mathbb{R}$ is subanalytic and continuous in \mathbb{R}^N . Then, for any $\bar{x} \in \mathbb{R}^N$, there exist $\nu \in [0, 1)$, $C > 0$, and $\Delta > 0$ such that*

$$|F(x) - F(\bar{x})|^\nu \leq C m(\partial F(x)) \quad \forall x \in B_\Delta(\bar{x}).$$

The subanalyticity of $V(\cdot)$ and the Łojasiewicz inequality rely on geometric properties and do not require smoothness of the function.

IV. PROOF OF THE RESULTS

Having established the lemmas in the previous section, we are now ready to prove the main results. First, using the Łojasiewicz inequality on $V(\cdot)$, we can prove that the output trajectories necessarily converge to a single fixed point.

Proof of Theorem 1. We begin by showing that $V(a(t))$ is convergent. Indeed, by (12), (5b) and (13), we see that $V(a(t))$ is decreasing and $V(a(t)) \geq 0$ for all $t \geq 0$. As a consequence, $V(a(t))$ converges as t goes to infinity. Denote by V^* this limit. On the other hand, by Lemma 3, $a(t)$ is bounded for all $t \geq 0$. By the Bolzano-Weierstrass theorem, there exists a sequence of increasing times $\{t_k\}_{k \in \mathbb{N}}$ such that $\{a(t_k)\}_{k \in \mathbb{N}}$ converges. Let a^* be the limit point of this converging sequence. We will show that the output $a(t)$ converges to a^* with a proof by contradiction.

By the continuity of $V(a(t))$, the limit satisfies $V(a^*) = V^*$. Applying the nonsmooth Łojasiewicz inequality to $V(\cdot)$ at a^* , there exist $\nu \in [0, 1)$, $C > 0$, and $\Delta > 0$ such that

$$|V(a) - V^*|^\nu \leq C m(\partial V(a)), \quad \forall a \in B_\Delta(a^*). \quad (17)$$

We fix $0 < \delta \leq \Delta$. Since $a(t_k)$ converges to a^* , there exists $K \in \mathbb{N}$ such that for all $k \geq K$

$$\|a(t_k) - a^*\|_2 < \frac{\delta}{4}. \quad (18)$$

Since $V(a(t))$ is decreasing and converges to V^* , there exists $T \geq 0$ such that, for all $t \geq T$

$$0 \leq V(a(t)) - V^* \leq \left[\frac{\beta \delta (1 - \nu)}{4C\alpha} \right]^{\frac{1}{1-\nu}}, \quad (19)$$

where C and ν are defined in (17) and α, β in (16). We now define the time index $p = \min \{k \in \mathbb{N} \text{ s.t. } k \geq K \text{ and } T \leq t_k\}$. Time t_p exists, since the sequence of time $\{t_k\}_{k \in \mathbb{N}}$ is increasing and goes to infinity. In addition, t_p is such that it satisfies both (18) and (19). We also define

$$t_q = \sup \{t \geq t_p \text{ s.t. } \forall s \in [t_p, t) \quad \|a(s) - a^*\|_2 < \delta\}.$$

If $t_q = +\infty$, then for all time $t \geq t_p$, $\|a(t) - a^*\|_2 \leq \delta$. Since δ can be chosen arbitrarily small, this proves that the output $a(t)$ converges to the single fixed point a^* . By contradiction, assume that $t_q < +\infty$. This implies that for all time $s \in [t_p, t_q)$, the output trajectory remains within a ball of radius δ around the fixed point, i.e. $a(s) \in B_\delta(a^*)$, but leaves this ball at time t_q , i.e. $\|a(t_q) - a^*\|_2 = \delta$. According to (16), we have $\forall n \in \Gamma$,

$$\|\dot{a}_n(t)\|_2 = \|f'(u_n(t)) \odot \dot{u}_n(t)\|_2 \geq \beta \|\dot{u}_n(t)\|_2,$$

which implies that

$$\begin{aligned} \|\dot{a}(t)\|_2 &= \|\dot{a}_\Gamma(t)\|_2 \geq \beta \|\dot{u}_\Gamma(t)\|_2 \\ &\geq \beta m(\partial V(a_\Gamma(t))) = \beta m(\partial V(a(t))). \end{aligned}$$

Furthermore, from (12) and (16) we see that

$$\dot{V}(a(t)) = - \sum_{n \in \Gamma} \frac{1}{f'(u_n(t))} |\dot{a}_n(t)|^2 \leq -\frac{1}{\alpha} \|\dot{a}(t)\|_2^2.$$

Putting everything together, we get

$$\dot{V}(a(t)) \leq -\frac{1}{\alpha} \|\dot{a}(t)\|_2^2 \leq -\frac{\beta}{\alpha} \|\dot{a}(t)\|_2 m(\partial V(a(t))).$$

By definition of t_q , and since $\delta < \Delta$ for all $t \in (t_p, t_q)$, we have $a(t) \in B_\delta(a^*) \subset B_\Delta(a^*)$, and so (17) yields

$$\begin{aligned} \|\dot{a}(t)\|_2 &\leq \frac{\alpha}{\beta} \frac{-\dot{V}(a(t))}{m(\partial V(a(t)))} \\ &\leq \frac{\alpha C}{\beta} \frac{-\dot{V}(a(t))}{(V(a(t)) - V^*)^\nu}. \quad (\text{since } -\dot{V} \leq 0) \end{aligned}$$

Then we can find an upper bound for the following integral:

$$\begin{aligned}
\|a(t_q) - a(t_p)\|_2 &= \left\| \int_{t_p}^{t_q} \dot{a}(s) ds \right\|_2 \leq \int_{t_p}^{t_q} \|\dot{a}(s)\|_2 ds \\
&\leq -\frac{\alpha C}{\beta} \int_{t_p}^{t_q} \frac{\dot{V}(a(s))}{(V(a(s)) - V^*)^\nu} ds \\
&= -\frac{\alpha C}{\beta} \int_{V(a(t_p))}^{V(a(t_q))} \frac{dV}{(V - V^*)^\nu} \\
&= \frac{\alpha C}{\beta(1-\nu)} \left[(V(a(t_p)) - V^*)^{1-\nu} \right. \\
&\quad \left. - (V(a(t_q)) - V^*)^{1-\nu} \right] \\
&\leq \frac{\alpha C}{\beta(1-\nu)} (V(a(t_p)) - V^*)^{1-\nu} \\
&\leq \frac{\delta}{4}. \quad (\text{from (19)})
\end{aligned}$$

Finally, we see that

$$\begin{aligned}
\delta &= \|a(t_q) - a^*\|_2 \leq \|a(t_q) - a(t_p)\|_2 + \|a(t_p) - a^*\|_2 \\
&\leq \frac{\delta}{4} + \frac{\delta}{4} = \frac{\delta}{2} < \delta.
\end{aligned}$$

We have reached a contradiction, which proves that $t_q = +\infty$. Consequently, for all $t \geq t_p$, we have $\|a(t) - a^*\|_2 \leq \delta$. Since δ can be chosen arbitrarily small, this shows that $\lim_{t \rightarrow +\infty} a(t) = a^*$, and thus the output converges. \square

The second proof of this section shows that the state variables also converge to a single fixed point u^* .

Proof of Theorem 2. From Theorem 1, the output converges to some fixed point $a^* \in \mathbb{R}^N$. The dynamical equation (1) can be written in terms of the distance $\tilde{a}(t) = a(t) - a^*$ of the output from the fixed point:

$$\dot{u}(t) = -u(t) - \Phi^T \Phi a^* + \Phi^T y + a^* - \Phi^T \Phi \tilde{a}(t) + \tilde{a}(t).$$

Defining $u^* = -\Phi^T \Phi a^* + \Phi^T y + a^*$ yields the following equation:

$$\dot{u}(t) = -u(t) + u^* - (\Phi^T \Phi - I) \tilde{a}(t).$$

The solutions of this differential equation have the following form for all $t \geq 0$:

$$u(t) = u^* + e^{-t} (u(0) - u^*) + e^{-t} \int_0^t e^s (\Phi^T \Phi - I) \tilde{a}(s) ds.$$

The second term $e^{-t} (u(0) - u^*)$ obviously converges to zero as t goes to infinity. We will also show that the last term converges to zero, thus proving that $u(t)$ converges to u^* . Denote by $h(t)$ this integral term, and consider its norm:

$$\begin{aligned}
\|h(t)\|_2 &= \left\| e^{-t} \int_0^t e^s (\Phi^T \Phi - I) \tilde{a}(s) ds \right\|_2 \\
&= e^{-t} \int_0^t e^s \|(\Phi^T \Phi - I) \tilde{a}(s)\|_2 ds \\
&\leq e^{-t} \sigma_1 \int_0^t e^s \|\tilde{a}(s)\|_2 ds,
\end{aligned}$$

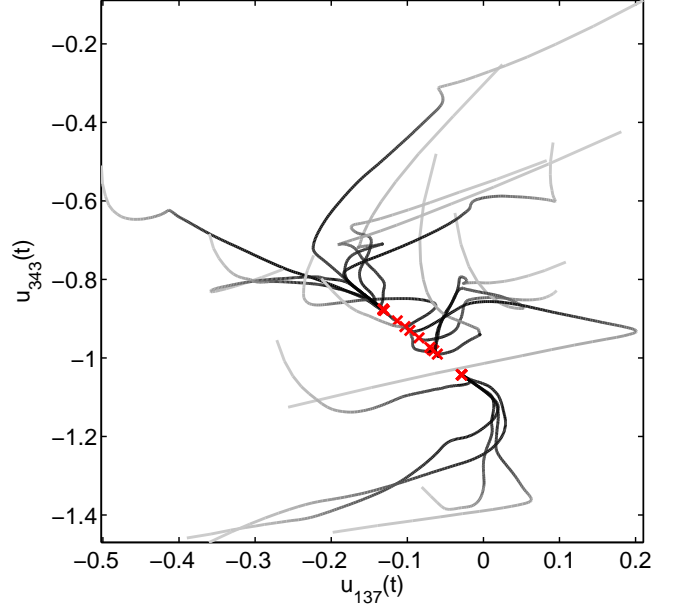


Fig. 3. Convergence of LCA trajectories obtained for 20 different initial points and projected onto the space spanned by two randomly chosen non-zero entries of a^\dagger . The color gradient indicates the time evolution. A red cross indicates the fixed point reached by the system.

where $\sigma_1 \geq 0$ is the largest eigenvalue of the interconnection matrix $W = \Phi \Phi^T - I$. To show convergence to zero, we split the integral into two parts. Since $a(t)$ converges to a^* , $\tilde{a}(t)$ converges to 0 as $t \rightarrow +\infty$. Thus, for any $\tilde{\epsilon} > 0$, there exists a time $t_c \geq 0$ such that, $\forall t \geq t_c$, $\|\tilde{a}(t)\|_2 \leq \tilde{\epsilon}$. Moreover, since $\|\tilde{a}(t)\|_2$ is continuous and goes to zero as t goes to infinity, it admits a maximum μ , $\forall t \in \mathbb{R}$. This yields, for all $t \geq 2t_c$

$$\begin{aligned}
\|h(t)\|_2 &\leq e^{-t} \|\Phi^T \Phi - I\|_2 \mu \int_0^{t_c} e^s ds \\
&\quad + e^{-t} \|\Phi^T \Phi - I\|_2 \tilde{\epsilon} \int_{t_c}^t e^s ds \\
&\leq \|\Phi^T \Phi - I\|_2 \mu [e^{-t+t_c} - e^{-t}] \\
&\quad + \|\Phi^T \Phi - I\|_2 \tilde{\epsilon} [1 - e^{-t+t_c}] \\
&\leq \|\Phi^T \Phi - I\|_2 \mu [e^{-t/2} - e^{-t}] + \|\Phi^T \Phi - I\|_2 \tilde{\epsilon}.
\end{aligned}$$

Since the left term converges to 0 and $\tilde{\epsilon}$ can be chosen arbitrarily small, this shows that the trajectory $u(t)$ converges to the trajectory $u^* + e^{-t} (u(0) - u^*)$ as t goes to infinity, and thus, we can conclude that $u(t) \xrightarrow[t \rightarrow +\infty]{} u^*$. \square

V. SIMULATION

To illustrate the theoretical result, we create an example for which there exists a subspace of non-isolated solutions. The matrix Φ has dimension $M = 256$ by $N = 512$ and is generated uniformly at random from a Gaussian distribution (then normalized to have columns with unit norm). A sparse vector a^\dagger is generated by selecting uniformly at random the

location of 5 non-zero entries. Their amplitudes are generated from a normal distribution and normalized to one. The column in Φ corresponding to one of the non-zero entries is replaced by a random linear combination of the other 4 columns and normalized. The measurements are $y = \Phi a^\dagger + n$, where n is a Gaussian noise with standard deviation $\sigma = 0.01$. The network is simulated through a discrete approximation in Matlab with a step size of 0.001, a time constant of $\tau = 0.01$ and the soft-threshold activation function with a threshold set to $\lambda = 0.03$. Figure 3 shows the trajectories for 20 random starting points projected onto the space spanned by two randomly selected non-zero entries of a^\dagger . Despite the solutions being non-isolated and lying on a linear subspace, the system converges and reaches a single fixed point for every starting point.

VI. SUMMARY

The LCA is a neural network defined by a set of differential equations. It was shown in [4] that the fixed points of the network coincide with the set of solutions of the sparse approximation problem with an appropriate cost function. However, the Lyapunov approach used in [4] was insufficient to guarantee convergence of the outputs to a single fixed point in the case where the sparse approximation problem had non-isolated solutions. A technique using the Łojasiewicz inequality has recently been developed that overcomes this limitation. For networks satisfying a certain type of differential inclusion with a subanalytic objective, the trajectories can be shown to converge to a single fixed point even when there are infinitely many and non-isolated solutions. However, the characteristics of the LCA necessitate a more careful treatment. In particular, its activation function is nonsmooth and possibly unbounded, and its interconnection matrix may be singular. In this paper, we were able to apply the Łojasiewicz inequality when the activation function is subanalytic to prove convergence of both the output and state variables without assuming isolated solutions. This improves on the results obtained in [4] and provides further evidence that the LCA is a reliable network to solve this important class of optimization programs, thus supporting pursuing its implementation in analog.

REFERENCES

- [1] C. J. Rozell, D. H. Johnson, R. G. Baraniuk, and B. A. Olshausen, "Sparse coding via thresholding and local competition in neural circuits," *Neural Comput.*, vol. 20, no. 10, pp. 2526–2563, Oct. 2008.
- [2] J. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons," *Proc. Natl. Acad. Sci.*, vol. 81, no. 10, pp. 3088–3092, May 1984, PMID: 6587342.
- [3] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, vol. 43, no. 1, pp. 129–159, Mar. 2001.
- [4] A. Balavoine, J. Romberg, and C. J. Rozell, "Convergence and rate analysis of neural networks for sparse approximation," *IEEE Trans. Neural Netw.*, vol. 23, no. 9, pp. 1377–1389, Sept. 2012.
- [5] S. Łojasiewicz, "Une propriété topologique des sous-ensembles analytiques réels," *Colloques internationaux du C.N.R.S. Les équations aux dérivées partielles*, vol. 117, pp. 87–89, 1963.
- [6] J. Bolte, A. Daniilidis, and A. Lewis, "The Łojasiewicz inequality for nonsmooth subanalytic functions with applications to subgradient dynamical systems," *SIAM J. Optim.*, vol. 17, no. 4, pp. 1205–1223, Jan. 2007.

- [7] S. Shapero, A. S. Charles, C. J. Rozell, and P. Hasler, "Low power sparse approximation on reconfigurable analog hardware," *IEEE J. Emerg. Sel. Topic Circuits Syst.*, vol. 2, no. 3, pp. 530–541, Sept. 2012.
- [8] S. Shapero, C. J. Rozell, and P. Hasler, "Configurable hardware integrate and fire neurons for sparse approximation," *Neural Netw.*, 2013, in press.
- [9] A. S. Charles, P. Garrigues, and C. J. Rozell, "A common network architecture efficiently implements a variety of sparsity-based inference problems," *Neural computation*, vol. 24, no. 12, pp. 3317–3339, Dec. 2012.
- [10] M. Forti and A. Tesi, "Absolute stability of analytic neural networks: an approach based on finite trajectory length," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 51, no. 12, pp. 2460–2469, Dec. 2004.
- [11] M. Forti, P. Nistri, and M. Quincampoix, "Convergence of neural networks for programming problems via a nonsmooth Łojasiewicz inequality," *IEEE Trans. Neural Netw.*, vol. 17, no. 6, pp. 1471–1486, Nov. 2006.
- [12] W. Lu and J. Wang, "Convergence analysis of a class of nonsmooth gradient systems," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 55, no. 11, pp. 3514–3527, Dec. 2008.
- [13] F. H. Clarke, *Optimization and Nonsmooth Analysis*, Society for Industrial Mathematics, Jan. 1987.
- [14] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*, vol. 317, Springer, 2011.