

# Estimating Multi-chirp Parameters using Curvature-guided Langevin Monte Carlo

Sattwik Basu, Debottam Dutta, Yu-lin Wei, Romit Roy Choudhury  
University of Illinois at Urbana-Champaign

**Abstract**—This paper considers the problem of estimating chirp parameters from a noisy mixture of chirps. While a rich body of work exists in this area, challenges remain when extending these techniques to chirps of higher order polynomials. We formulate this as a non-convex optimization problem and propose a modified Langevin Monte Carlo (LMC) sampler that exploits the average curvature of the objective function to reliably find the minimizer. Results show that our *Curvature-guided LMC* algorithm is robust and succeeds even in low SNR regimes, making it viable for practical applications.

**Index Terms**—Chirps, Optimization, Langevin Monte Carlo (LMC), Gaussian smoothing, Curvature

## I. INTRODUCTION

Chirps are a fundamental class of non-stationary signals that underlie many applications. These signals are characterized by complex exponentials with a time-varying instantaneous phase (IP) and amplitude (IP), often modeled as quadratic or logarithmic functions. Of interest is the more complex scenario, where the IP and IA of a finite-duration chirp are  $P^{\text{th}}$  and  $A^{\text{th}}$  order polynomials, respectively. Equation (1) models such a chirp and Figure 1 visualizes it for  $A=3$  and  $P=4$ . Estimating parameters  $\{\rho_a\}_{a=0}^A, \{\varphi_p\}_{p=1}^P$  of such signals becomes difficult when multiple chirps mix at low SNR regimes. This problem has been the focus of several studies in the past [1]–[5] and continues to be the subject of active research. Our goal in this paper is to improve the state of the art towards estimating mixtures of higher dimensional chirps at lower SNRs.

$$y(n) = \sum_{a=0}^A \rho_a \left(\frac{n}{f_s}\right)^a \exp \left[ 2\pi j \sum_{p=1}^P \varphi_p \left(\frac{n}{f_s}\right)^p \right] \quad (1)$$




Fig. 1. Real part of a damped chirp with 4-th order polynomial IP.

Chirp parameter estimation already underlies applications such as wireless radars, audio processing, biomedical sensing and astronomy. In radars, for instance, chirp mixtures are encountered when a transmitted chirp echoes back over a multipath channel from multiple moving targets. The instantaneous phase (IP) continuously varies with time due to the relative motion between the reflectors and the receiver; estimating these parameters correctly reveals information about the position and movements of the targets. Another emerging application pertains to cardiac sounds, where the sound of blood flowing through the heart valves has been shown to be a mixture of 2 higher-order polynomial chirps [6], [7]. The vast majority of doctors find it difficult to manually separate these 2 sounds when they listen through a

stethoscope [8]. Success in chirp parameter estimation could enable electronic stethoscopes that isolate the individual valve sounds in software. This could obviously be valuable to downstream diagnosis and disease localization.

**Past work** on chirp parameter estimation has explored various lines of attack. Nonlinear transformations, like Higher-Order Ambiguity Functions (HAF, ML-HAFs, PHAFs) [1], [3], iteratively transform chirps into sinusoids whose frequencies correspond to the desired chirp parameters. Despite their computational simplicity, HAFs struggle with identifiability issues in multi-component chirps with time-varying amplitudes, and show performance degradation in low-SNR conditions as errors in higher order parameters propagate down to the lower orders. Another body of work uses a maximum a posteriori (MAP)/maximum likelihood (MLE) framework to decouple the amplitude and phase parameter estimation by solving a nonlinear least squares (NLS) problem [9]–[11]. For example, [12] proposes a MCMC approach to solve the NLS problem using importance sampling and a global optimization theorem [13] closely linked with simulated tempering [14] for chirp mixtures with  $P=2$ . The computational burden in this approach increases for higher order chirps because of the need to jointly sample parameters from a discrete grid in parameter space. A more recent approach [15] takes the variational inference route to obtain an approximate posterior over the parameters of  $P=2$  chirp mixtures. A common limitation in most of these works is the performance degradation with increasing parameters ( $P \times A \times N_c$ , where  $N_c$  is the number of chirps in the mixture). We focus on addressing these limitations in our work.

**This paper** tackles mixtures of higher-order chirps ( $P \geq 3$ ,  $A \geq 2$ ,  $N_c = 2$ ) where both amplitude and IP are polynomials. We formulate this as a non-convex optimization problem whose objective function — partially visualized in Figure 2 — contains a deep global minimum, numerous local minima clustered around the global minima, and relatively flat regions elsewhere. We adopt a sampling approach to this optimization problem using a Langevin Monte Carlo (LMC) sampler. While the stochastic behavior of LMC is designed to move out of local minima, we observe that the outcomes were unreliable in our case. Past work in optimization [16] has addressed this with an additional gradient smoothing approach that adds two benefits: (1) helps LMC overcome the local minima by smoothing the local landscape, and (2) improves the gradient magnitude at far away regions so that the LMC sampler can progress towards the global minima,  $\varphi^*$ . Unfortunately, this (Gaussian) smoothing needs careful tuning; excessive smoothing can obscure the global minima while too less smoothing retains the original problem. In response to this, we propose a curvature-based smoothing method (utilizing the trace of the Hessian) to adaptively tune the Gaussian smoothing parameter throughout the whole optimization process. The end result is that our *Curvature-guided LMC* optimizer (CG-LMC) is able to progress towards the global minima  $\varphi^*$  even from far-away regions, and as the optimizer comes near  $\varphi^*$ , it is able to tide over the surrounding stationary points to reliably converge near  $\varphi^*$ .

**Experiments** with synthetic chirps at relatively low signal-to-noise

(SNR) demonstrate low estimation error and high reliability. Our CG-LMC method outperforms two baselines, namely a classical Langevin Monte Carlo optimizer (LMC) [17] and a Noise-Annealed Langevin Monte Carlo optimizer (NA-LMC) [18].

## II. PROBLEM FORMULATION

### A. Multi-component Chirps

The signal model for a  $N$  sample discrete time multi-component chirp  $y(n)$  sampled at a rate  $f_s$  is shown.

$$y(n) = \sum_{c=1}^{N_c} A_c(n) \exp \left[ 2\pi j \sum_{p=1}^P \varphi_{c,p} \left( \frac{n}{f_s} \right)^p \right] + w(n) \quad (2)$$

where

$$A_c(n) = \sum_{a=0}^{A_c} \rho_{c,a} \left( \frac{n}{f_s} \right)^a \quad \text{and} \quad w(n) \sim \mathcal{N}_{\mathbb{C}}(0, \sigma^2)$$

for  $n = 0, \dots, N-1$ . Equation (2) represents a superimposition of  $N_c$  chirps polluted by Gaussian noise  $w(n)$ . Each chirp is assumed to have a  $P$ -th order phase polynomial and  $A_c$ -th order amplitude polynomial. Estimating the parameters of the combined chirp entails estimating all the parameters shown in (3).

$$\boldsymbol{\varphi} = [\varphi_1 \dots \varphi_{N_c}]^T, \boldsymbol{\rho} = [\rho_1 \dots \rho_{N_c}]^T \quad (3)$$

where  $\boldsymbol{\varphi}_c = [\varphi_{c,1} \dots \varphi_{c,P}]$  and  $\boldsymbol{\rho}_c = [\rho_{c,0} \dots \rho_{c,A_c}]$  represent the phases and amplitudes of the  $c$ -th chirp, respectively.

We can rewrite the measured signal model in matrix-vector form to setup a nonlinear regression approach as shown

$$\mathbf{y} = \mathbf{H}(\boldsymbol{\varphi})\mathbf{b}(\boldsymbol{\rho}) + \mathbf{w} \quad (4)$$

Here,  $\mathbf{y} \in \mathbb{C}^{N \times 1}$  represents the measured signal, and the basis matrix,  $\mathbf{H}(\boldsymbol{\varphi}) \in \mathbb{C}^{N \times \sum_{c=1}^{N_c} (A_c+1)}$  is comprised of the sub-matrices  $\mathbf{H}(\boldsymbol{\varphi}_c) \in \mathbb{C}^{N \times (A_c+1)}$  for  $c = 1, \dots, N_c$  such that

$$\mathbf{H}(\boldsymbol{\varphi}) = \left[ \mathbf{H}(\boldsymbol{\varphi}_1) \dots \mathbf{H}(\boldsymbol{\varphi}_{N_c}) \right] \quad (5)$$

The columns of the  $c$ -th sub-matrix expressed compactly are  $\mathbf{H}(\boldsymbol{\varphi}_c) = [\mathbf{h}_{c,0} \dots \mathbf{h}_{c,A_c}]$  where  $\mathbf{h}_{c,a} \in \mathbb{C}^{N \times 1}$  is

$$\mathbf{h}_{c,a} = \begin{bmatrix} \left( \frac{0}{f_s} \right)^a \exp \left[ 2\pi j \sum_{p=1}^P \varphi_{c,p} \left( \frac{0}{f_s} \right)^p \right] \\ \vdots \\ \left( \frac{N-1}{f_s} \right)^a \exp \left[ 2\pi j \sum_{p=1}^P \varphi_{c,p} \left( \frac{N-1}{f_s} \right)^p \right] \end{bmatrix} \quad (6)$$

for  $a = 0, \dots, A_c$ . The vector  $\mathbf{b}(\boldsymbol{\rho}) \in \mathbb{C}^{\sum_{c=1}^{N_c} (A_c+1) \times 1}$  consists of amplitude parameters such that

$$\mathbf{b}(\boldsymbol{\rho}) = \left[ \mathbf{b}(\boldsymbol{\rho}_1) \dots \mathbf{b}(\boldsymbol{\rho}_{N_c}) \right]^T \quad (7)$$

with  $\mathbf{b}(\boldsymbol{\rho}_c) = [\rho_{c,0} \dots \rho_{c,A_c}]$  for  $c = 1, \dots, N_c$ . From here on, we denote  $\mathbf{H}(\boldsymbol{\varphi})$  and  $\mathbf{b}(\boldsymbol{\rho})$  as  $\mathbf{H}$  and  $\mathbf{b}$  for the sake of brevity.




Fig. 2. Slice of the 4-D objective function for two chirp mixture with  $P = 2$ . Parameters  $\varphi_{12}$ ,  $\varphi_{22}$  are varied while others are fixed at optimum values.

### B. MAP estimation

Since  $\mathbf{w}$  is assumed to be white Gaussian noise, the likelihood function can be expressed as

$$p(\mathbf{y}; \boldsymbol{\varphi}, \boldsymbol{\rho}) = \frac{1}{(2\pi)^{N/2} \sigma^N} \exp \left( -\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{H}\mathbf{b}\|_2^2 \right) \quad (8)$$

To obtain MAP estimates of vectors  $\boldsymbol{\varphi}$  and  $\boldsymbol{\rho}$  with uniform priors, we see that

$$\operatorname{argmax}_{\boldsymbol{\varphi}, \boldsymbol{\rho}} \ln p(\boldsymbol{\varphi}, \boldsymbol{\rho} | \mathbf{y}) = \operatorname{argmax}_{\boldsymbol{\varphi}, \boldsymbol{\rho}} \ln p(\mathbf{y} | \boldsymbol{\varphi}, \boldsymbol{\rho}) + \ln p(\boldsymbol{\varphi}, \boldsymbol{\rho}) \quad (9)$$

$$\therefore [\boldsymbol{\varphi}, \boldsymbol{\rho}]_{MAP} = \operatorname{argmin}_{\boldsymbol{\varphi}, \boldsymbol{\rho}} \|\mathbf{y} - \mathbf{H}\mathbf{b}\|_2^2 \quad (10)$$

### C. Nonlinear Least Squares

Minimizing Equation (10) is complicated due to the unknown nonlinear dependency on  $\boldsymbol{\varphi}$ . However, noting that the NLS criterion is quadratic in  $\mathbf{b}$ , for any given vector  $\boldsymbol{\varphi}$ ,  $\hat{\mathbf{b}} = \mathbf{G}^{-1} \mathbf{H}^* \mathbf{y}$  where,  $\mathbf{G} = (\mathbf{H}^* \mathbf{H} + \gamma \mathbf{I})$ . Here,  $\gamma$  is a regularization parameter and  $*$  represents the Hermitian transpose of a complex matrix.

Substituting  $\hat{\mathbf{b}}$  back into Equation (10) yields a modified cost function  $J: \mathbb{R}^{N_c P} \rightarrow \mathbb{R}$  that depends only on  $\boldsymbol{\varphi}$ .

$$J(\boldsymbol{\varphi}) = \left\| \mathbf{y} - \mathbf{H}\hat{\mathbf{b}} \right\|_2^2 = \mathbf{y}^* \mathbf{P}_{\mathbf{H}}^{\perp} \mathbf{y} \quad (11)$$

Minimizing this objective function gives us the optimal phase vector  $\boldsymbol{\varphi}^*$ . That is,

$$\boldsymbol{\varphi}^* = \operatorname{argmin}_{\boldsymbol{\varphi}} J(\boldsymbol{\varphi}) = \operatorname{argmin}_{\boldsymbol{\varphi}} \mathbf{y}^* \mathbf{P}_{\mathbf{H}}^{\perp} \mathbf{y} \quad (12)$$

Once nonlinear parameters  $\boldsymbol{\varphi}$  are estimated,  $\boldsymbol{\rho}$  can be computed via least squares. However, optimizing  $J(\boldsymbol{\varphi})$  in Equation (12) is not trivial as it is highly non-convex (Fig. 2). The complexity grows with increasing dimensions of the parameter space.

## III. CURVATURE-GUIDED LANGEVIN MONTE CARLO (CG-LMC)

### A. Background

Recent works have demonstrated that optimization problems can be addressed by sampling from a distribution whose mode is concentrated around the global optimum. MCMC algorithms such as Importance Sampling and Metropolis-Hastings [19], [20] are well-known for this purpose. However these algorithms are problematic in high-dimensional settings due to high rejection rates and difficulties in designing proposal distributions [21]. Moreover, the distribution associated with the non-convex objective functions are multi-modal adding further complexity to the optimization. In these scenarios, Langevin Monte Carlo (LMC) offers attractive benefits [22].

LMC has found increasing interest in the nonconvex optimization literature [17], [23]–[25] due to its resemblance with Stochastic

Gradient Descent (SGD) [26]–[28]. The core idea is to simulate an overdamped Langevin Diffusion equation shown in Equation (13)

$$d\mathbf{X}(t) = -\nabla F(\mathbf{X}(t)) dt + \sqrt{2\beta^{-1}} d\mathbf{B}(t), t \geq 0 \quad (13)$$

where the distribution of  $\mathbf{X}(t)$  as  $t \rightarrow \infty$  converges to a stationary distribution  $\pi(\cdot) \propto \exp(-\beta F(\cdot))$ . Here,  $F$  is called the potential function,  $\mathbf{B}(t)$  is a Brownian process in  $\mathbb{R}^d$ , and  $\beta > 0$  is an inverse temperature parameter. For sufficiently large  $\beta$ , drawing samples from  $\pi(\cdot)$  produces values close to the minimizer with high probability.

For simulating Equation (13) the Euler-Maruyama [29] scheme is used to obtain the discrete-time Markov chain (LMC)

$$\mathbf{X}^{(k+\eta)} = \mathbf{X}^{(k)} - \eta \nabla F(\mathbf{X}^{(k)}) + \sqrt{2\eta\beta^{-1}} \boldsymbol{\xi}^{(k)} \quad (14)$$

where  $\eta > 0$  is the stepsize,  $k$  the iteration index and  $\boldsymbol{\xi}_k \sim \mathcal{N}(0, \mathbf{I}_d)$ . Following each step, the Metropolis-Hastings algorithm is used to either accept or reject the proposed sample  $\mathbf{X}^{(k+\eta)}$  [30]. The interplay between gradient and the stochastic perturbation helps in overcoming local stationary points to improve convergence. As  $k$  increases, the iterates  $\mathbf{X}^{(k+\eta)}$ , converge on the minimizer of  $F$ .

### B. LMC for optimizing $J(\boldsymbol{\varphi})$

To apply LMC to optimize  $J(\boldsymbol{\varphi})$  in Equation (12), we define  $\pi(\boldsymbol{\varphi}) \propto \exp(-\beta J(\boldsymbol{\varphi}))$  as the stationary distribution. Observe that  $\pi(\boldsymbol{\varphi})$  is the marginalized posterior  $p(\boldsymbol{\varphi}|\mathbf{y})$  in Equation (9) representing the decoupling of  $\boldsymbol{\varphi}$  from  $\boldsymbol{\rho}$ . Since  $J(\boldsymbol{\varphi})$  is non-convex,  $\pi(\boldsymbol{\varphi})$  is multi-modal with the largest mode concentrating around  $\boldsymbol{\varphi}^*$ . To draw samples from this distribution, we use the Markov chain

$$\boldsymbol{\varphi}^{(k+1)} = \boldsymbol{\varphi}^{(k)} - \eta \nabla_{\boldsymbol{\varphi}} J(\boldsymbol{\varphi}^{(k)}) + \sqrt{2\eta\beta^{-1}} \boldsymbol{\xi}^{(k)} \quad (15)$$

where  $\boldsymbol{\xi}^{(k)} \sim \mathcal{N}(0, \mathbf{I}_{N_c P})$  and  $\boldsymbol{\varphi}^{(0)} \sim \pi^{(0)}$ , an initial distribution. The gradient elements of  $\nabla_{\boldsymbol{\varphi}} J \in \mathbb{R}^{N_c P \times 1}$  (derivation omitted in interest of space) contains partial derivatives  $\nabla_{\varphi_{c,p}} J$  where

$$\nabla_{\varphi_{c,p}} J = -2\text{Re} \left[ \mathbf{y}^* \mathbf{P}_H^\perp \nabla_{\varphi_{c,p}} \mathbf{H} \mathbf{G}^{-1} \mathbf{H}^* \mathbf{y} \right] \quad (16)$$

for  $c = 1 \dots N_c$  and  $p = 1 \dots P$ . The term  $\nabla_{\varphi_{c,p}} \mathbf{H}$  in Equation (16) contains element wise partial derivatives of  $\mathbf{H}$  with respect to each parameter  $\varphi_{c,p}$  for  $c = 1 \dots N_c$  and  $p = 1 \dots P$ .

$$\begin{aligned} \nabla_{\varphi_{c,p}} \mathbf{H}(\boldsymbol{\varphi}) &= \left[ \nabla_{\varphi_{c,p}} \mathbf{H}(\boldsymbol{\varphi}_1) \dots \nabla_{\varphi_{c,p}} \mathbf{H}(\boldsymbol{\varphi}_{N_c}) \right] \\ &= \left[ \mathbf{0} \dots \nabla_{\varphi_{c,p}} \mathbf{H}(\boldsymbol{\varphi}_c) \dots \mathbf{0} \right] \\ &= \left[ \mathbf{0} \dots 2\pi j \mathbf{n}_p \odot [\mathbf{h}_{c,0} \dots \mathbf{h}_{c,A_c}] \dots \mathbf{0} \right] \end{aligned} \quad (17)$$

where  $\mathbf{n}_p = \left[ \left( \frac{0}{f_s} \right)^p \dots \left( \frac{N-1}{f_s} \right)^p \right]^T$  and  $\odot$  denotes element-wise multiplication.

Although LMC offers promise, it is not always able to “jump out” of a local minima, despite its stochastic nature [31]. Such issues become pronounced in higher dimensional chirps where the objective function has many complex ripples. Smoothing techniques have been proposed [18] to help the LMC sampler combat such cases.

### C. Gaussian-smoothed Langevin Monte Carlo

Given the landscape of our objective function  $J(\boldsymbol{\varphi})$ , the initial points located far from the global minimum have negligible gradients while those initialized nearby risk being trapped in local stationary points. An effective solution is to add additional noise during Langevin updates to perturb the gradients [18]. Adding noise has an implicit gradient smoothing effect as explained in [16], [32]. To see this, we write a modified LMC as

$$\boldsymbol{\varphi}^{(k+1)} = \boldsymbol{\varphi}^{(k)} - \eta \nabla_{\boldsymbol{\varphi}} J(\boldsymbol{\varphi}^{(k)}) + \sqrt{2\eta\beta^{-1}} \boldsymbol{\xi}^{(k)} + \sigma \boldsymbol{\epsilon}^{(k)} \quad (18)$$

where  $\boldsymbol{\epsilon}^{(k)} \sim \pi(\boldsymbol{\epsilon}^{(k)}) = \mathcal{N}(0, \mathbf{I}_{N_c P})$  is another Gaussian perturbation vector independent from  $\boldsymbol{\xi}^{(k)}$  with variance adjusted by  $\sigma$ . Now, by setting  $\boldsymbol{\theta}^{(k+1)} = \boldsymbol{\varphi}^{(k+1)} - \sigma \boldsymbol{\epsilon}^{(k)}$  followed by taking the expectation with respect to  $\boldsymbol{\epsilon}^{(k-1)}$ , we get

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\epsilon}^{(k-1)}} \left[ \boldsymbol{\theta}^{(k+1)} \right] &= \boldsymbol{\theta}^{(k)} - \eta \mathbb{E}_{\boldsymbol{\epsilon}^{(k-1)}} \left[ \nabla_{\boldsymbol{\varphi}} J(\boldsymbol{\theta}^{(k)} + \sigma \boldsymbol{\epsilon}^{(k-1)}) \right] \\ &\quad + \sqrt{2\eta\beta^{-1}} \boldsymbol{\xi}^{(k)} \end{aligned} \quad (19)$$

The expectation term in the R.H.S of Equation (19) is a convolution between  $\nabla_{\boldsymbol{\varphi}} J(\boldsymbol{\varphi}_k)$  and a Gaussian pdf. This is referred to as the Gaussian-smoothed gradient  $\nabla_{\boldsymbol{\varphi}} J_\sigma$  [32]. As a result of smoothing, the iterates  $\boldsymbol{\theta}^{(k+1)}$ , on average, move along a smoothed gradient field minimizing the traps due to local minima and saddle points.

### D. Need for Curvature-guided Gaussian-smoothing

It is crucial to note that the choice of  $\sigma$  determines the extent of smoothing, which in turn guides convergence. To show this, we focus on the expectation term in Equation (19).

$$\nabla_{\boldsymbol{\varphi}} J_\sigma = \int \nabla_{\boldsymbol{\varphi}} J(\boldsymbol{\theta}^{(k)} + \sigma \boldsymbol{\epsilon}^{(k-1)}) \pi(\boldsymbol{\epsilon}^{(k-1)}) d\boldsymbol{\epsilon}^{(k-1)} \quad (20)$$

$$= \frac{1}{\sigma} \int \nabla_{\boldsymbol{\varphi}} J(\boldsymbol{\varphi}^{(k)}) \pi\left(\frac{\boldsymbol{\varphi}^{(k)} - \boldsymbol{\theta}^{(k)}}{\sigma}\right) d\boldsymbol{\varphi}^{(k)} \quad (21)$$

Equation 21 represents an inner product between the gradient  $\nabla_{\boldsymbol{\varphi}} J$  and a kernel  $\pi(\cdot)$  that is scaled by  $\sigma$  and translated by  $\boldsymbol{\theta}^{(k)}$ . A large  $\sigma$  widens the kernel  $\pi$ , smoothing over wide regions of the parameter space, which can obscure the global minimum. A smaller  $\sigma$  sharpens the kernel smoothing only over a small local neighborhood. Thus, for reliable convergence,  $\sigma$  needs to be adaptive; large when far away from the global minima so that the flat landscape can tilt towards the global minima (see Fig. 2). When arriving closer to the global minima,  $\sigma$  needs to shrink so that only the local minima is smoothed without destroying the natural curvature of the landscape.

Without knowledge of the global minima, adapting  $\sigma$  is difficult. Past work [18] have empirically scheduled  $\sigma$  to start as a large value and decrease with iterations. We propose to leverage the average curvature of the objective function to regulate this adaptation. We obtain curvature information from  $\text{tr} \left\{ \nabla_{\boldsymbol{\varphi}}^2 J \right\}$ , i.e., the trace of the Hessian of  $J$ , and update  $\sigma$  as:

$$\sigma^{(k+1)} = \max \left[ \sigma_{\min}, \sigma^{(k)} - \mu_\sigma \left| \text{tr} \left\{ \nabla_{\boldsymbol{\varphi}}^2 J(\boldsymbol{\varphi}^{(k)}) \right\} \right| \right] \quad (22)$$

where  $\mu_\sigma$  is the step size,  $k$  is the iteration index and  $\sigma_{\min}$  is a user defined minimum value. The Hessian,  $\nabla_{\boldsymbol{\varphi}}^2 J$  in (22) is approximated using Stein’s Lemma [33] with the same perturbation vector  $\boldsymbol{\epsilon}^{(k)}$  that is used to compute  $\nabla_{\boldsymbol{\varphi}} J_\sigma$ . In addition, we require the magnitude of average curvature and not the concavity of  $J$ , hence we use the absolute value of the trace. We noticed that a recent work [34] has used the trace as well but in the context of Gaussian Homotopy.

**Initialization.** Our algorithm begins with multiple randomized, diffused initializations, priming the LMC algorithm using shorter-length signals before transitioning to the full-length measured signal  $\mathbf{y}$ . These shorter signals correspond to objective functions with broader global minima [35], providing a more reliable starting point for the lower order phase parameters. This approach eliminates any dependency on additional initialization algorithms used in prior works.

TABLE I  
COMPARISON BETWEEN CG-LMC, LMC, AND NA-LMC SHOWING THE MEAN  $\pm$  STDEV OF THE ESTIMATED PARAMETER FOR DIFFERENT SNRS.

Algorithm Parameters	CG-LMC		LMC		NA-LMC	
	3 dB	12 dB	3 dB	12 dB	3 dB	12 dB
$\varphi_{1,1} = 10$	10.17 $\pm$ 0.55	9.89 $\pm$ 0.59	10.05 $\pm$ 0.97	9.69 $\pm$ 2.11	9.93 $\pm$ 1.58	9.67 $\pm$ 0.81
$\varphi_{1,2} = 40$	40.85 $\pm$ 1.88	44.91 $\pm$ 0.86	39.98 $\pm$ 8.71	43.67 $\pm$ 13.59	41.20 $\pm$ 10.93	41.58 $\pm$ 6.58
$\varphi_{1,3} = -70$	-68.96 $\pm$ 3.91	-76.32 $\pm$ 1.81	-80.43 $\pm$ 24.51	-83.22 $\pm$ 5.07	-77.04 $\pm$ 5.63	-77.44 $\pm$ 2.74
$\varphi_{1,4} = 110$	109.84 $\pm$ 2.40	104.66 $\pm$ 5.34	134.39 $\pm$ 21.57	125.64 $\pm$ 10.21	126.36 $\pm$ 24.37	118.78 $\pm$ 23.59
$\varphi_{2,1} = 50$	50.08 $\pm$ 0.55	49.98 $\pm$ 0.21	50.55 $\pm$ 0.92	50.21 $\pm$ 0.91	50.27 $\pm$ 1.20	51.08 $\pm$ 1.18
$\varphi_{2,2} = 60$	62.58 $\pm$ 1.27	60.68 $\pm$ 2.17	60.33 $\pm$ 4.74	61.85 $\pm$ 4.91	57.34 $\pm$ 9.02	57.77 $\pm$ 9.89
$\varphi_{2,3} = -90$	-104.46 $\pm$ 7.05	-93.23 $\pm$ 1.28	-102.16 $\pm$ 7.19	-106.05 $\pm$ 16.29	-85.59 $\pm$ 18.26	-79.32 $\pm$ 22.61
$\varphi_{2,4} = 105$	112.26 $\pm$ 3.91	107.11 $\pm$ 12.06	129.99 $\pm$ 17.15	130 $\pm$ 26.03	106.80 $\pm$ 8.20	119 $\pm$ 21.02

#### IV. EXPERIMENTS AND RESULTS

We compare **CG-LMC** against two baselines, **LMC** and **NA-LMC**, all implemented in MATLAB. We run on 5 different chirp mixtures ( $N_c = 2$ ) and all chirps are parameterized by  $A = 3$  and  $P = 4$ . The chirp signals are 1.0s long, sampled at  $f_s = 1000\text{Hz}$ , and their  $\varphi$  and  $\rho$  values are selected arbitrarily. We ensure that the IF never exceeds  $f_s/2$ . Gaussian noise is added to achieve 3dB and 12dB SNRs. We report the accuracy of the estimated chirp parameters (mean and std. dev. over 5 runs); we also show the convergence behavior by tracking different initializations near and far from the global minima. Our code is available at <https://github.com/basusattwik/ChirpEstimation>

**Table I** reports the phase parameter estimates for one specific chirp with detailed comparisons. At 3dB SNR, the mean of the phase parameter estimates from CG-LMC are closer to the true value with smaller standard deviations, than those of LMC or NA-LMC, except in parameters  $\varphi_{2,3}, \varphi_{2,4}$ . This demonstrates the benefit of curvature-guided Gaussian smoothing as the low SNR regimes have more jagged objective functions; LMC in contrast tends to exhibit higher error and NA-LMC tends to exhibit higher std. dev. At a higher SNR of 12dB, CG-LMC shows improved estimates with smaller variance compared to both LMC and NA-LMC. Of course, all methods show improvements as expected.

**Fig. 3** sheds light on the internal behavior of CG-LMC for the same experiment as above. We track a total of 9 CG-LMC sample runs with the first 3 initialized close to the global minimum (blue curves), 3 moderately far (orange curves), and 3 very far (green curves) from the global minimum. For comparison, 3 sample runs of NA-LMC are similarly initialized, near, far and very far away. Fig. 3(a) shows how CG-LMC is able to reliably converge to the global minima regardless of their initialization. Of course, far away initializations tend to take longer compared to nearby ones. In contrast, NA-LMC struggles to converge when the initialization is far away, leading to high standard deviation.

**Fig. 3(b) and (c)** further show the smoothing parameter  $\sigma$  and the trace of the Hessian for CG-LMC and NA-LMC. Observe that NA-LMC is not adaptive, so  $\sigma$  reduces over iterations as a static step function (hence, a single line in Figure 3(b)). While CG-LMC and NA-LMC are both reducing  $\sigma$ , observe that CG-LMC's green curves wait far longer before reducing  $\sigma$ . This allows CG-LMC to smoothen over large areas when the iterates are far from the global minima, allowing them to move in the correct direction. This can be validated from Fig. 3(c), where the curvature is low for the green curves over the initial iterations. Once the curvature begins to increase, the value of  $\sigma$  begins to fall. On the other hand, the sample runs that were initialized near and moderately far (blue and red), show a rapid decrease in  $\sigma$  owing to the already large curvature near the global minimum. In contrast, since NA-LMC is "blind" to the landscape of the objective function, the method of static smoothing may or may not guide the optimization to the global minima. This negatively impacts the reliability (or standard deviation) of estimation.

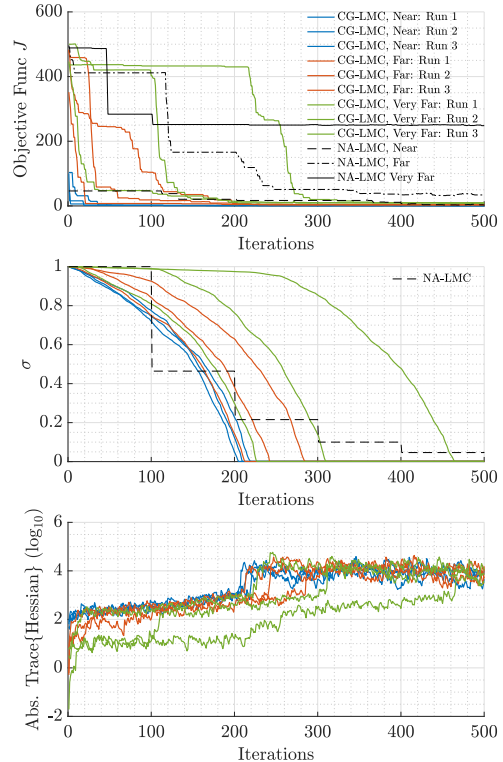


Fig. 3. (a) 9 sample runs in CG-LMC and 3 in NA-LMC, (b) Comparing  $\sigma$  adaptation in CG-LMC vs. the fixed schedule in NA-LMC, (c) Visualizing  $\text{tr}\{\nabla_{\varphi}^2 J\}$  over different CG-LMC runs.

**Additional chirp mixtures** were designed at 3dB SNR using different  $\varphi$  and  $\rho$  values. The mean absolute errors (ME) in the estimated phase parameters were [1.03, 3.55, 5.12, 11.24, 3.59, 4.19, 7.75, 7.87] over 5 runs while NA-LMC showed [0.77, 5.73, 13.33, 24.30, 4.67, 23.50, 12.26, 14.06]. For CG-LMC, errors are close to those shown in Table I demonstrating a generalization of performance across a range of parameters. In addition the low values of ME in CG-LMC indicates robust estimation at low SNRs.

#### V. CONCLUSIONS AND FUTURE WORK

We propose CG-LMC, an optimization algorithm for estimating chirp parameters, each with polynomial phase and amplitude. Although we build on existing methods to optimize the chirp-centric objective function, we show that leveraging the intrinsic curvature of the function can be valuable, especially in terms of making the algorithm relatively robust to initialization. In future, tests need to be performed to compare the estimators with the Cramér-Rao bounds. Perhaps source separation methods could be explored to separate chirps (if possible), and then use CG-LMC to estimate parameters on a lower dimensional space. We leave these to future research.

## REFERENCES

- [1] S. Peleg and B. Friedlander, "The discrete polynomial-phase transform," *IEEE Transactions on Signal Processing*, vol. 43, no. 8, pp. 1901–1914, 1995.
- [2] S. Peleg and B. Friedlander, "Multicomponent signal analysis using the polynomial-phase transform," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 32, no. 1, pp. 378–387, 1996.
- [3] S. Barbarossa, A. Scaglione, and G. Giannakis, "Product high-order ambiguity function for multicomponent polynomial-phase signal modeling," *IEEE Transactions on Signal Processing*, vol. 46, no. 3, pp. 691–708, 1998.
- [4] P. Djuric and S. Kay, "Parameter estimation of chirp signals," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 38, no. 12, pp. 2118–2126, 1990.
- [5] R. Liang and K. Arun, "Parameter estimation for superimposed chirp signals," in *Proceedings ICASSP-92: 1992 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 5, pp. 273–276 vol.5, 1992.
- [6] J. Xu, L. Durand, and P. Pibarot, "Nonlinear transient chirp signal modeling of the aortic and pulmonary components of the second heart sound," *IEEE Transactions on Biomedical Engineering*, vol. 47, no. 10, pp. 1328–1335, 2000.
- [7] J. Xu, L.-G. Durand, and P. Pibarot, "Extraction of the aortic and pulmonary components of the second heart sound using a nonlinear transient chirp signal model," *IEEE Transactions on Biomedical Engineering*, vol. 48, no. 3, pp. 277–283, 2001.
- [8] T. WR, "In defence of auscultation: a glorious future?," *Heart Asia*, vol. 9(1), pp. 44–47, 2017.
- [9] B. Friedlander and J. Francos, "Estimation of amplitude and phase of non-stationary signals," in *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers*, pp. 431–435 vol.1, 1993.
- [10] T. L. Jensen, J. K. Nielsen, J. R. Jensen, M. G. Christensen, and S. H. Jensen, "A fast algorithm for maximum-likelihood estimation of harmonic chirp parameters," *IEEE Transactions on Signal Processing*, vol. 65, no. 19, pp. 5137–5152, 2017.
- [11] D. Kundu and S. Nandi, "On chirp and some related signals analysis: A brief review and some new results," *Sankhya A*, vol. 83, pp. 844–890, 2021.
- [12] S. Saha and S. Kay, "Maximum likelihood parameter estimation of superimposed chirps using monte carlo importance sampling," *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 224–230, 2002.
- [13] M. Pincus, "A closed form solution of certain programming problems," *Operations Research*, vol. 16, no. 3, pp. 690–694, 1968.
- [14] E. Marinari and G. Parisi, "Simulated tempering: a new monte carlo scheme," *EPL*, vol. 19, pp. 451–458, 1992.
- [15] J. Neri, P. Depalle, and R. Badeau, "Damped chirp mixture estimation via nonlinear bayesian regression," in *2021 24th International Conference on Digital Audio Effects (DAFx)*, pp. 65–72, 2021.
- [16] N. Chatterji, J. Diakonikolas, M. I. Jordan, and P. Bartlett, "Langevin monte carlo without smoothness," in *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics* (S. Chappa and R. Calandra, eds.), vol. 108 of *Proceedings of Machine Learning Research*, pp. 1716–1726, PMLR, 26–28 Aug 2020.
- [17] M. Raginsky, A. Rakhlin, and M. Telgarsky, "Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis," in *Proceedings of the 2017 Conference on Learning Theory* (S. Kale and O. Shamir, eds.), vol. 65 of *Proceedings of Machine Learning Research*, pp. 1674–1703, PMLR, 07–10 Jul 2017.
- [18] Y. Song and S. Ermon, *Generative modeling by estimating gradients of the data distribution*. Red Hook, NY, USA: Curran Associates Inc., 2019.
- [19] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*. Springer Texts in Statistics, Springer New York, NY, 2 ed., 2004.
- [20] D. Luengo, L. Martino, M. F. Bugallo, J. Read, and J. Míguez, "A survey of monte carlo methods for parameter estimation," *EURASIP Journal on Advances in Signal Processing*, vol. 2020, no. 25, 2020.
- [21] J. S. Speagle, "A conceptual introduction to markov chain monte carlo methods," 2020.
- [22] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, May 2011.
- [23] Y.-A. Ma, Y. Chen, C. Jin, N. Flammarion, and M. I. Jordan, "Sampling can be faster than optimization," *Proceedings of the National Academy of Sciences*, vol. 116, no. 42, pp. 20881–20885, 2019.
- [24] P. Xu, J. Chen, D. Zou, and Q. Gu, "Global convergence of langevin dynamics based algorithms for nonconvex optimization," in *Advances in Neural Information Processing Systems* (S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds.), vol. 31, Curran Associates, Inc., 2018.
- [25] A. Dalalyan, "Further and stronger analogy between sampling and optimization: Langevin monte carlo and gradient descent," in *Proceedings of the 2017 Conference on Learning Theory* (S. Kale and O. Shamir, eds.), vol. 65 of *Proceedings of Machine Learning Research*, pp. 678–689, PMLR, 07–10 Jul 2017.
- [26] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," 2018.
- [27] M. Welling and Y. W. Teh, "Bayesian learning via stochastic gradient langevin dynamics," in *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, (Madison, WI, USA), p. 681–688, Omnipress, 2011.
- [28] K. P. Murphy, *Probabilistic Machine Learning: An introduction*. MIT Press, 2022.
- [29] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*. Stochastic Modelling and Applied Probability, Springer Berlin, Heidelberg, 1 ed., 1992.
- [30] G. O. Roberts and R. L. Tweedie, "Exponential convergence of Langevin distributions and their discrete approximations," *Bernoulli*, vol. 2, no. 4, pp. 341 – 363, 1996.
- [31] H. Lee, A. Risteski, and R. Ge, "Beyond log-concavity: Provable guarantees for sampling multi-modal distributions using simulated tempering langevin monte carlo," in *Advances in Neural Information Processing Systems* (S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds.), vol. 31, Curran Associates, Inc., 2018.
- [32] Y. Nesterov and V. Spokoiny, "Random gradient-free minimization of convex functions," *Foundations of Computational Mathematics*, vol. 17, no. 3, pp. 527–566, 2017.
- [33] J. Zhu, "Hessian estimation via stein's identity in black-box problems," in *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference* (J. Bruna, J. Hesthaven, and L. Zdeborova, eds.), vol. 145 of *Proceedings of Machine Learning Research*, pp. 1161–1178, PMLR, 16–19 Aug 2022.
- [34] H. Iwakiri, Y. Wang, S. Ito, and A. Takeda, "Single loop gaussian homotopy method for non-convex optimization," *Advances in Neural Information Processing Systems*, vol. 35, pp. 7065–7076, 2022.
- [35] J. Angeby, "Estimating signal parameters using the nonlinear instantaneous least squares approach," *IEEE Transactions on Signal Processing*, vol. 48, no. 10, pp. 2721–2732, 2000.