

NETWORK CONSENSUS IN THE WASSERSTEIN METRIC SPACE OF PROBABILITY MEASURES*

ADRIAN N. BISHOP[†] AND ARNAUD DOUCET[‡]

Abstract. Distributed consensus in the Wasserstein metric space of probability measures on the real line is introduced in this work. Convergence of each agent’s measure to a common measure is proven under a weak network connectivity condition. The common measure reached at each agent is one minimizing a weighted sum of its Wasserstein distance to all initial agent measures. This measure is known as the Wasserstein barycenter. Special cases involving Gaussian measures, empirical measures, and time-invariant network topologies are considered, where convergence rates and average-consensus results are given. This work has possible applicability in computer vision, machine learning, clustering, and estimation.

Key words. consensus, metric space, Wasserstein space

AMS subject classifications. 49N99, 90C08, 93C83

DOI. 10.1137/19M1268252

1. Introduction. The problem of distributed (network) consensus concerns a group of agents that seek to reach agreement upon certain state variables of interest by exchanging information across a network. Typically, the agents are connected via a network that changes over time due to link failures, node failures, packet drops, etc. For example, in distributed sensor networks the interaction topology may change over time as individual nodes (or some subset of such) may be mobile or unreliable or communication constraints are also present. All such variations in topology can happen randomly and often the network is disconnected for some time. Studies on the convergence of consensus algorithms (to a common agreed “value” at each agent) are often motivated by such complex time-varying networks.

1.1. Background. The consensus problem has a long history, e.g., [1], which is too broad to cover here. We highlight [2, 3, 4, 5, 6, 7] for further history and background.

Many consensus algorithms have been proposed in the literature. References [2, 4, 6, 8] focus on linear update rules (at each agent) and typically concern average-consensus or consensus about some linear function of all initial agent states in Euclidean space. The average-consensus problem has a natural relationship with distributed linear least squares or distributed (linear) maximum likelihood estimation [9] and distributed Kalman filtering [10, 11, 12, 13]. Alternative consensus algorithms using nonlinear update rules have been proposed and studied in [5, 14, 15, 16]. Here consensus to general functions (e.g., the maximum or minimum etc.) of all initial agent states may be sought as in [17, 18] and even finite-time convergence may be achievable [18, 19]. One may also want to achieve consensus to some time-varying reference signal as in [20, 21].

*Received by the editors June 13, 2019; accepted for publication (in revised form) June 22, 2021; published electronically September 21, 2021. A preliminary draft of this work appeared in a conference proceedings as [33].

<https://doi.org/10.1137/19M1268252>

Funding: This work was supported by AFOSR/AOARD via AOARD-144042.

[†]CSIRO and University of Technology Sydney, Broadway, NSW, 2007, Australia (adrian.bishop@uts.edu.au).

[‡]University of Oxford, Oxford, OX1 3LB, UK (doucet@stats.ox.ac.uk).

We note here that the majority of the literature on consensus concerns agreement in Euclidean space as exemplified by the seminal papers of [2, 3, 4, 5, 6, 7]. However, there are exceptions. The problem of synchronization is closely related to consensus but typically deals with the problem of driving a network of oscillators to a common frequency/phase. This work typically concerns nonlinear manifolds such as the circle. A survey on synchronization is given in [22, 23] while consensus and synchronization are related in [24]. Some other notable exceptions of consensus in non-Euclidean spaces are [25, 26, 27, 28, 29, 30]. In particular, [26, 28] consider general nonlinear consensus on manifolds by embedding such manifolds in a suitable high-dimensional Euclidean space. In particular, this embedding approach is used to perform consensus on the special orthogonal group and on Grassmann manifolds. In [30] a consensus algorithm on the Riemannian manifold of (Gaussian) covariance matrices is introduced under the Fisher metric (related to the Kullback–Leibler divergence). The authors in [25, 27, 29] study consensus in different metric spaces which is more closely related to the present work. For example, the author of [29] develops an analogue of Wolfowitz’s theorem [3] for a class of metric spaces with nonpositive curvature which leads to a notion of consensus in such spaces.

The (distributed) consensus problem (as referenced above) has been widely studied across the fields of control and optimization; e.g., see again the seminal articles [2, 3, 4, 5, 6, 7] published in the control literature, the references therein, and the many (thousands) of citing articles. This extensive and broad interest in distributed consensus protocols stems from interesting applications such as distributed estimation, filtering [10, 11, 12, 13] and distributed information fusion [9], interesting applications in distributed control [3, 4, 5, 7] and in distributed optimization [26, 30], and many related topics too broad to discuss here; see also the survey articles [23, 28, 31].

1.2. Contributions. The main contributions of this paper are a novel algorithm and convergence results for distributed consensus in the space of probability measures with time-varying interaction networks. We introduce a well-studied metric known as the Wasserstein distance which allows us to consider an important set of probability measures as a metric space [32]. The proposed consensus algorithm is based on iteratively updating each agent’s probability measure by finding a measure that minimizes the weighted sum of its Wasserstein distances to the agent’s own previous measure plus all neighbor agents’ measures. We show that convergence of the individual agents’ measures to a common probability measure is guaranteed under a weak network connectivity condition. The common measure that is achieved asymptotically at each agent is the one that is closest simultaneously to all initial agent measures in the sense of the Wasserstein distance. Focus in this work is on probability measures over the real line.

This work has potential applicability in the field of computer vision and image processing, distributed computation, clustering and data aggregation, distributed estimation, filtering and information fusion, distributed optimization in metric spaces, and machine learning more broadly, among other fields discussed later. Applications and related topics, particularly relevant in the Wasserstein domain, are discussed later. However, the main focus of this article is on the general Wasserstein consensus idea itself and its convergence.

This paper extends [33] with the addition of results detailing convergence rates and network properties in which particular consensus values may be achieved.

1.3. Paper organization. The main contribution is given in section 2 where a consensus algorithm in the space of probability measures is introduced, and its con-

vergence is studied. In section 3 we establish for specific scenarios—initial Gaussian measures; for empirical measures; and time-invariant networks—the exponential convergence of this consensus algorithm and various computational aspects. In section 4 we discuss potential applications. Concluding remarks are given in section 5.

1.4. Notation and conventions. Consider a group of agents indexed in $\mathcal{V} = \{1, \dots, n\}$ and a set of (possibly) time-varying undirected links $\mathcal{E}(t) \subset \mathcal{V} \times \mathcal{V}$ defining a network graph $\mathcal{G}(t)(\mathcal{V}, \mathcal{E}(t))$. The neighbor set at agent i is denoted by $\mathcal{N}_i(t) = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}(t)\}$. Time is indexed using \mathbb{N} .

The graph adjacency matrix $\mathbf{A}(t) \in \mathbb{R}^{n \times n}$ obeys $\mathbf{A}(t) = \mathbf{A}(t)^\top = [a_{ij}(t)]$, where $a_{ij}(t) = 1 \Leftrightarrow (i, j) \in \mathcal{E}(t)$ and $a_{ij}(t) = 0$ otherwise. Implicit throughout is that $a_{ii}(t) = 1$ for all i and t and thus $i \in \mathcal{N}_i(t)$ for all t . A weighted adjacency matrix is denoted by $\mathbf{W}(t) = [w_{ij}(t)] \in \mathbb{R}^{n \times n}$ with $a_{ij}(t) = 1 \Leftrightarrow w_{ij}(t) > 0$ and $w_{ij} = 0$ otherwise. We require $\sum_{j \in \mathcal{N}_i(t)} w_{ij}(t) = 1$ and assume that $w_{ii}(t) > 0$ for all i and t so that $w_{ij}(t) \in [0, 1)$ whenever $i \neq j$ and for all t .

The adjacency matrix $\mathbf{A}(t)$ defines $\mathcal{G}(t)(\mathcal{V}, \mathcal{E}(t))$ and vice versa because $a_{ij}(t) = 1 \Leftrightarrow (i, j) \in \mathcal{E}(t)$ and $a_{ij}(t) = 0 \Leftrightarrow (i, j) \notin \mathcal{E}(t)$. The weighted adjacency matrix defines $\mathcal{G}(t)(\mathcal{V}, \mathcal{E}(t))$ because $w_{ij}(t) > 0 \Leftrightarrow (i, j) \in \mathcal{E}(t)$ and $w_{ij}(t) = 0 \Leftrightarrow (i, j) \notin \mathcal{E}(t)$, but $\mathcal{G}(t)(\mathcal{V}, \mathcal{E}(t))$ alone defines only the sparsity pattern of $\mathbf{W}(t)$.

Consider the sequence of graphs $\mathcal{G}(t_k), \mathcal{G}(t_{k+1}), \dots, \mathcal{G}(t_{k+T})$ on the same vertex set \mathcal{V} . The union of this sequence is denoted by $\mathfrak{G}(t_k, t_{k+T})(\mathcal{V}, \cup_{t \in [t_k, t_{k+T}]} \mathcal{E}(t))$, i.e., $\mathfrak{G}(t_k, t_{k+T})$ is just a graph on the vertex set \mathcal{V} with edges $\cup_{t \in [t_k, t_{k+T}]} \mathcal{E}(t)$. The sequence is said to be jointly connected if \mathfrak{G} is connected.

2. Consensus in the Wasserstein space of probability measures. The main contribution of this work is given in this section where we introduce and establish the convergence of a consensus algorithm in the Wasserstein metric space of probability measures.

Suppose the state of agent i is given by a Radon probability measure μ_i defined on the Borel sets of (\mathbb{R}, d) where in this section we restrict $d : \mathbb{R} \times \mathbb{R} \rightarrow [0, \infty)$ to be the usual Euclidean distance. Define the space of all such measures on (\mathbb{R}, d) by $\mathfrak{U}(\mathbb{R})$ and the subset of all such measures with finite p th moment by $\mathfrak{U}_p(\mathbb{R})$ where henceforth we assume that $2 \leq p < \infty$. That is, \mathfrak{U}_p is the collection of probability measures such that $\int_{\mathbb{R}} d(x, x_0)^p d\mu_i(x) < \infty$ for a given, arbitrary, $x_0 \in \mathbb{R}$.

One can associate the Wasserstein metric $\ell_p : \mathfrak{U}_p(\mathbb{R}) \times \mathfrak{U}_p(\mathbb{R}) \rightarrow [0, \infty)$ with \mathfrak{U}_p which is defined by

$$\ell_p(\mu_i, \mu_j) = \left(\inf_{\gamma \in \Gamma(\mu_i, \mu_j)} \int_{\mathbb{R} \times \mathbb{R}} d(x_i, x_j)^p d\gamma(x_i, x_j) \right)^{1/p},$$

where $\Gamma(\mu_i, \mu_j)$ denotes the collection of all probability measures on $\mathbb{R} \times \mathbb{R}$ with marginals μ_i and μ_j on the first and second factors; see [34, 35].

Let us recall some standard results about the Wasserstein metric space $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ when $p \geq 2$; see, e.g., [34, 35, 36, 37].

1. $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is a complete and separable metric space.
2. $\lim_{k \rightarrow \infty} \ell_p(\mu_k, \mu) = 0$ is equivalent to weak convergence and convergence of the first p moments.
3. Given two measures $\mu_i, \mu_j \in \mathfrak{U}_p(\mathbb{R})$, then $\ell_p(\mu_i, \mu_j) = \ell_p(\mu_i, \mu) + \ell_p(\mu_j, \mu)$ for some $\mu \in \mathfrak{U}_p(\mathbb{R})$.
4. More generally, there exists a continuously parameterized constant speed path $\mu_s \in \mathfrak{U}_p(\mathbb{R})$, $s \in [0, 1]$, such that for $\mu_i, \mu_j \in \mathfrak{U}_p$ we have $\mu_{s=0} = \mu_i$ and

$\mu_{s=1} = \mu_j$ and $\ell_p(\mu_i, \mu_j) = \ell_p(\mu_i, \mu_s) + \ell_p(\mu_j, \mu_s) \forall s \in [0, 1]$. The measure μ_s is known as the interpolant measure [38].

5. The interpolant measure defines a geodesic, and consequently, (\mathfrak{U}_p, ℓ_p) is geodesic.
6. $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ has vanishing curvature in the sense of Alexandrov (a subset of $\text{CAT}(0)$); see Proposition 4.1 in [36].
7. $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is simply connected; see [36, 37].

All metrics are continuous, and we recall that a constant speed geodesic in $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is a curve $\mu_s : \mathbb{I} \rightarrow \mathfrak{U}_p$ parameterized on some interval $\mathbb{I} \subset \mathbb{R}$ that satisfies $\ell_p(\mu_{s_i}, \mu_{s_j}) = v|s_i - s_j|$ for some constant $v > 0$ and for all $s_i, s_j \in \mathbb{I}$.

Suppose the measure at agent i is updated by

$$(1) \quad \mu_i(t+1) = \operatorname{argmin}_{\eta \in \mathfrak{U}_p(\mathbb{R})} \sum_{j \in \mathcal{N}_i(t)} w_{ij}(t) \ell_p(\eta, \mu_j(t))^p$$

for all $i \in \mathcal{V}$ where we recall that we assume $\sum_{j \in \mathcal{N}_i(t)} w_{ij}(t) = 1$ and $w_{ii}(t) > 0$ so that consequently $w_{ij}(t) \in [0, 1)$ whenever $i \neq j$. This operation is well defined as discussed below.

Application of the update rule (1) to each agent $i \in \mathcal{V}$ corresponds to the proposed nonlinear (distributed) consensus algorithm. Note that we consider only undirected (or bidirectional) network communications in this work for simplicity. Directed communication may be studied as in, e.g., [5, 6], with additional conditions needed for convergence in that setting [5].

2.1. Main result. We state here our main result.

THEOREM 1. *Consider a group of agents \mathcal{V} and network $\mathcal{G}(t)(\mathcal{V}, \mathcal{E}(t))$ where each agent i has initial state $\mu_i(0) \in \mathfrak{U}_p(\mathbb{R})$ and updates its state $\mu_i(t) \in \mathfrak{U}_p(\mathbb{R})$ according to (1). If for all $t_0 \in \mathbb{N}$ the graph union $\mathfrak{G}(t_0, \infty)$ is connected, then there exists $\mu^* \in \mathfrak{U}_p(\mathbb{R})$ such that*

$$\lim_{t \rightarrow \infty} \ell_p(\mu_i(t), \mu^*) = 0$$

for any $i \in \mathcal{V}$.

The proof of Theorem 1 is given in the next subsection following the provision of a number of supporting results.

2.2. Proof of the main result. Before proceeding to the general proof, we note briefly here that the Wasserstein metric $\ell_p : \mathfrak{U}_p(\mathbb{R}) \times \mathfrak{U}_p(\mathbb{R}) \rightarrow [0, \infty)$ may be written as

$$\ell_p(\mu_i, \mu_j) = \left(\int_0^1 |F_i^-(x) - F_j^-(x)|^p dx \right)^{1/p},$$

where $F_i^-(x) : [0, 1] \rightarrow \mathbb{R}$ is the inverse cumulative distribution function for $\mu_i(t) \in \mathfrak{U}_p(\mathbb{R})$, defined in more detailed later. This form for $\ell_p(\mu_i, \mu_j)$ is given by Major in [39]. It follows with $p = 2$ that the solution to (1) at any $i \in \mathcal{V}$ and any $t \in \mathbb{N}$ has an inverse cumulative distribution function given by

$$F_i^-(t+1)(x) = \sum_{j \in \mathcal{N}_i} w_{ij} F_j^-(t)(x)$$

for all $x \in [0, 1]$, as proven more formally later. From this formulation one can derive a kind of convergence result for consensus with the update rule (1); e.g., see Proposition

2 later, its proof, and, e.g., (5) later and the surrounding discussion there. This formulation of the metric in terms of the inverse cumulative distribution functions conveys intuition on the space $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$, for example, implying immediately that $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is isometric to a convex subset of the Banach space $(L_p([0, 1]), d)$ and thus inherits many familiar properties (which is rather unique to the case of measures on \mathbb{R}).

For generality we proceed with the proof using more metric space/geometric arguments that in many places may carry over to the more general case involving measures on \mathbb{R}^n (albeit we do not explore those cases here). Where ideas do not carry over we highlight that in this more general geometric formalism so as to distinguish where a generalized proof may need to be extended.

The proof proceeds now via a series of supporting lemmas. Note that a subset $\mathfrak{X} \subset \mathfrak{U}_p(\mathbb{R})$ is convex if every geodesic segment whose endpoints are in \mathfrak{X} lies entirely in \mathfrak{X} . The (closed) convex hull $\text{co}(\mathfrak{Y})$ of a subset $\mathfrak{Y} \subset \mathfrak{U}_p$ is the intersection of all (closed) convex subsets of \mathfrak{U}_p that contain \mathfrak{Y} .

LEMMA 1. *If $\mu_i(t) \in \mathfrak{U}_p(\mathbb{R})$, then the operation (1) is well defined in the sense that it admits a solution and this solution is unique whenever (at least) one $\mu_j(t) \in \mathfrak{U}_p$, $j \in \mathcal{N}_i(t)$ does not give support to small sets.¹*

Recall that $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is CAT(0) (indeed, it has the stronger property of vanishing curvature, e.g., owing to its relationship with the Banach space $(L_p([0, 1]), d)$), in addition to being uniquely geodesic, complete, and separable, i.e., $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is a Hadamard space. This lemma then follows from the fact that $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is Hadamard, and Fréchet averages such as those defined by operations of the form (1) are well defined in such spaces; see page 334 in [43]. The existence and uniqueness of solutions to (1) is also discussed in [40] more generally; see also [41, 42]. It is worth noting in passing the related work in [25, 29] which deals with similar consensus topics in CAT(0) spaces, and [27] which deals with consensus in a general class of convex metric spaces.

The convex hull of the set of measures $\{\mu_i\}$, $i \in \tilde{\mathcal{V}} \subseteq \mathcal{V}$, is defined by

$$\text{co}(\{\mu_i\}) = \left\{ \underset{\eta \in \mathfrak{U}_p(\mathbb{R})}{\text{argmin}} \sum_{i \in \tilde{\mathcal{V}}} w_i \ell_p(\eta, \mu_i)^p \mid w_i \geq 0, \sum_i w_i = 1 \right\}.$$

LEMMA 2. *Consider a collection $\{\mu_i\}$, $i \in \tilde{\mathcal{V}} \subseteq \mathcal{V}$, of distinct measures in $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$. The convex hull of $\{\mu_i\}$ is $\text{co}(\{\mu_i\}) \subset \mathfrak{U}_p(\mathbb{R})$ and is isometric to an l -sided convex polygon in \mathbb{R}^2 with $2 \leq l \leq |\{\mu_i\}|$.*

Before proceeding with the proof we point to [44] for background on comparison triangles and Alexandrov curvature of metric spaces. We also note that in a general geodesic CAT(0) space, i.e., some arbitrary geodesic space with nonpositive curvature, the preceding lemma is not true and the convex hull of a “geodesic triangle” defined

¹A small set is defined [40] as a set of Hausdorff dimension 0. This condition only plays a role in uniqueness and it is generally unnecessary [40, 41]. However, this requirement does exclude empirical measures on \mathbb{R} which arise in numerous applications relevant to this work (as discussed later). Luckily, it is generically true (i.e., excluding particular, nongeneric, arrangements) that (1) has a unique solution even in such cases; see [41, 42]. Note that if all inputs are discrete, we allow for both common and uncommon supports. Going forward we will not repeatedly call on the need for (at least) one initial measure to exclude support on small sets, and later results may be read as implicitly assuming uniqueness (or implicitly assuming exclusion of sole support on small sets).

by three points in such spaces may be of dimension greater than two;² see Chapter II.2 in [44].

Proof. Lemma 2 is a simple consequence of the vanishing curvature property of $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$. We elaborate for completeness. $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ has vanishing curvature in the sense of Alexandrov (see Proposition 4.1 in [36]) which formally means that for any triangle of points $\{\mu_i\}$, $i \in \{i_1, i_2, i_3\}$, and any point on the geodesic $\mu_s \in \mathfrak{U}_p(\mathbb{R})$, $s \in [0, 1]$ such that, for example, $\mu_{s=0} = \mu_{i_1}$ and $\mu_{s=1} = \mu_{i_2}$, then the ℓ_p distance between μ_{i_3} and μ_s , $s \in [0, 1]$, is the same as the corresponding Euclidean distance in a comparison triangle in \mathbb{R}^2 . Consider also any pair of points μ_j and μ_k with μ_j on the geodesic connecting μ_{i_1} and μ_{i_2} and μ_k on the geodesic connecting μ_{i_1} and μ_{i_3} with $\{\mu_j, \mu_k\} \cap \{\mu_i\} = \emptyset$, $i \in \{i_1, i_2, i_3\}$. Then vanishing curvature also implies $\angle_{\mu_{i_1}}(\mu_j, \mu_k)$ is equal to the usual interior Euclidean angle at the corresponding vertex in the comparison triangle in \mathbb{R}^2 . Here the angle $\angle_{\mu_{i_1}}(\mu_j, \mu_k)$ is the Alexandrov angle in arbitrary metric spaces; see Chapter II.1 in [44]. It now follows that the convex hull of any triangle of points $\{\mu_i\}$ in $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ is isometric to a triangle in \mathbb{R}^2 ; e.g., see Proposition 2.9 (Flat Triangle lemma) in [44]. Now define $\mathfrak{C} = \{\Delta_j\}$ to be the collection of geodesic triangles in $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ defined by every combination of three points in $\{\mu_i\}$, $i \in \tilde{\mathcal{V}} \subseteq \mathcal{V}$. Clearly, $\text{co}(\{\mu_i\}) = \cup_j \Delta_j$. Consider also the corresponding collection $\mathfrak{C}^+ = \{\Delta_j^+\}$ of comparison triangles in \mathbb{R}^2 . The Flat Triangle lemma implies that this collection can be arranged in \mathbb{R}^2 such that each angle $\angle_{\mu_i}(\mu_j, \mu_k)$ and each distance $\ell_p(\mu_i, \mu_j)$ for all $i, j, k \in \tilde{\mathcal{V}}$ in $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$ equals exactly the corresponding angle or distance in the comparison configuration of points in \mathbb{R}^2 . Obviously, the convex hull of the comparison configuration is an l -sided convex polygon in \mathbb{R}^2 with $2 \leq l \leq |\{\mu_i\}|$ and equal to $\cup_j \Delta_j^+$. Define the following map:

$$(2) \quad f_{\ell_p, d} : \text{co}(\{\mu_i\}) \rightarrow \mathbb{R}^2, \quad i \in \tilde{\mathcal{V}},$$

so the restriction

$$\begin{aligned} f_{\ell_p, d}(\Delta_j) &= f_{\ell_p, d}(\text{co}(\{\mu_{j_1}, \mu_{j_2}, \mu_{j_3}\})) \\ &= \text{co}(\{f_{\ell_p, d}(\mu_{j_1}), f_{\ell_p, d}(\mu_{j_2}), f_{\ell_p, d}(\mu_{j_3})\}) = \Delta_j^+ \end{aligned}$$

$\forall j \in \mathfrak{C} = \{\Delta_j\}$ is an isometry. Then

$$f_{\ell_p, d}(\text{co}(\{\mu_i\})) = f_{\ell_p, d}(\cup_j \Delta_j) = \cup_j f_{\ell_p, d}(\Delta_j) = \cup_j \Delta_j^+$$

from the Flat Triangle lemma and the property of vanishing curvature. For any two points in $\text{co}(\{\mu_i\})$ there exists a $\Delta_j \in \mathfrak{C}$ that contains them and the restriction $f_{\ell_p, d}(\Delta_j)$ is an isometry to a convex subset of $\cup_j \Delta_j^+$. Thus, $f_{\ell_p, d}$ is an isometry, which completes the proof. \square

LEMMA 3. Consider the convex hull $\text{co}(\{\mu_j(t)\})$, with $j \in \mathcal{N}_i(t)$ at time t . If agent i applies (1), it follows that $\mu_i(t+1)$ is strictly within the convex hull $\text{co}(\{\mu_j(t)\})$ whenever $|\{\mu_j(t)\}| \geq 2$ and two agent states are distinct and $w_{ij}(t) \in (0, 1)$.

²Our Euclidean intuition is generally wrong when it suggests the existence of a two-dimensional convex hull for a triangle defined by three points and the geodesics connecting them (albeit this is hard to visualize of course).

Proof. It is enough to consider two agents³ $i, j \in \mathcal{V}$ with (1) then given by

$$\mu_i(t+1) = \operatorname{argmin}_{\eta \in \mathfrak{U}_p(\mathbb{R})} w_{ii}(t) (\ell_p(\eta, \mu_i(t))^p - \ell_p(\eta, \mu_j(t))^p) + \ell_p(\eta, \mu_j(t))^p$$

and to note that η must lie on a geodesic $\mu_s : \mathbb{I} \rightarrow \mathfrak{U}_p(\mathbb{R})$. The proof relies on showing that $\mu_i(t+1) \notin \{\mu_i(t), \mu_j(t)\}$ when $w_{ii}, w_{ij} \in (0, 1)$. The first term

$$w_{ii}(t) (\ell_p(\eta, \mu_i(t))^p - \ell_p(\eta, \mu_j(t))^p)$$

is strictly negative at $\eta = \mu_i(t)$ and strictly increasing as η moves from $\mu_i(t)$ to $\mu_j(t)$ and conversely $\ell_p(\eta, \mu_j(t))^p$ is strictly positive at $\eta = \mu_i(t)$ and strictly decreasing to zero as η moves from $\mu_i(t)$ to $\mu_j(t)$. Then for any $w_{ii} \in (0, 1)$ and because ℓ_p is continuous, it follows that there exists some μ_ε on μ_s with $\varepsilon > 0$ such that

$$\begin{aligned} w_{ii}(t) (\ell_p(\eta, \mu_i(t))^p - \ell_p(\eta, \mu_j(t))^p) &< 0, \\ |w_{ii}(t) (\ell_p(\eta, \mu_i(t))^p - \ell_p(\eta, \mu_j(t))^p)| &< \ell_p(\eta, \mu_j(t))^p \end{aligned}$$

on $\eta \in \mu_s, s \in [0, \varepsilon]$. Consequently, $\mu_i(t+1)$ is strictly decreasing on $\eta \in \mu_s, s \in [0, \varepsilon]$. Hence for any $w_{i1} \in (0, 1)$ the point $\mu_i(t)$ cannot be a minimum. The same argument applies to $\mu_j(t)$. \square

The following is a simple consequence of the preceding result.

COROLLARY 1. *Consider the convex hull $\operatorname{co}(\{\mu_i(0)\})$ of all initial agent states in $(\mathfrak{U}_p(\mathbb{R}), \ell_p)$. If each agent applies (1), it follows that $\operatorname{co}(\{\mu_i(t)\}) \subseteq \operatorname{co}(\{\mu_i(0)\})$ for all t .*

The next result concerns an importance special case of the main result.

LEMMA 4. *Suppose $\mathcal{G}(\mathcal{V}, \mathcal{E})$ is time-invariant and connected. Suppose the state of each agent is $\mu_i(t) \in \mathfrak{U}_p(\mathbb{R})$ and that each agent applies (1). Then there exists $\mu^* \in \mathfrak{U}_p(\mathbb{R})$ such that for any $i \in \mathcal{V}$ it holds that*

$$\lim_{t \rightarrow \infty} \ell_p(\mu_i(t), \mu^*) = 0.$$

Proof. It almost goes without saying that $\ell_2(\mu_i(t), \mu_j(t))^p = 0 \forall i, j \in \mathcal{V}$ with a constant $\mu_i(t)$ in $\mathfrak{U}_p(\mathbb{R})$ an equilibrium state of (1). Consider a Lyapunov-like function $\nu(\mu) : \mathfrak{U}_p^n(\mathbb{R}) \rightarrow \mathbb{R}$ given by

$$(3) \quad \nu(\mu) = \sup_{\eta, \chi \in \{\mu_i(t)\}_{i \in \mathcal{V}}} \ell_p(\eta, \chi)^p$$

and note that $\nu(\mu) \geq 0$ with $\nu(\mu) = 0$ if and only if $\mu_i = \mu_j$ for all $i, j \in \mathcal{V}$.⁴ By Corollary 1 it follows that $\nu(\mu)$ is nonincreasing along trajectories of (1). It suffices to show $\nu(\mu(t+n-1)) < \nu(\mu(t))$ for each t . First, pick a $t_0 \geq 0$ and note $\operatorname{co}(\{\mu_i(t_0)\}) \subseteq \operatorname{co}(\{\mu_i(0)\})$ and $f_{\ell_p, d}(\operatorname{co}(\{\mu_i(t_0)\})) \subseteq f_{\ell_p, d}(\operatorname{co}(\{\mu_i(0)\}))$ from Corollary 1 and where $f_{\ell_p, d}$ is an isometry given by (2). Without loss of generality, via

³This is because Fréchet averages in Euclidean space (on a set of input points) are associative, and can be found iteratively by computing the average initially for a pair of points (in a larger set of inputs), and then computing the average between this result and the next point (in the input set), and so on (adjusting the weights defining the average each time); see [45]. Owing to Lemma 2, this associativity property still holds here.

⁴We abuse notation here slightly. We use the shorthand μ as the argument in $\nu(\mu)$ to represent the collection of all agent measures $\{\mu_i(t)\}_{i \in \mathcal{V}}$ with $|\mathcal{V}| = n$.

Lemma 2, suppose that $f_{\ell_p, d}(\text{co}(\{\mu_i(t_0)\}))$ is an l -sided polygon in \mathbb{R}^2 with $2 \leq l \leq |\mathcal{V}|$ on the collection of vertices $\{\mathbf{x}_j(t_0)\}$, $j \in \{1, \dots, l\}$, with $\mathbf{x}_j(t_0) \in \mathbb{R}^2$. If we chose a t_0 such that $l = 1$, then we would be done. Define the following set-valued function:

$$(4) \quad h_j(t) = \{i \in \mathcal{V} : f_{\ell_p, d}(\mu_i(t)) = \mathbf{x}_j(t_0)\} \quad \forall j \in \{1, \dots, l\}$$

for each time $t \geq t_0$. It is immediate from Lemma 3 that $h_j(t+1) \subseteq h_j(t)$ for all $j \in \{1, \dots, l\}$; i.e., more generally, no agent state $f_{\ell_p, d}(\mu_i(t))$ which is not on the boundary of the l -sided polygon at time t can ever reach this same boundary at $t+1$ as a consequence of Lemma 3. Note that $|h_j(t_0)| \leq n-1$ for all $j \in \{1, \dots, l\}$ with $l \geq 2$ at t_0 . Recall the neighbor set at agent i is given by $\mathcal{N}_i(t)$. Because the network is connected, for each $k \in h_j(t_0)$ the neighbor set obeys $\mathcal{N}_k(t_0) \neq \emptyset$ for each $j \in \{1, \dots, l\}$. Then by Lemma 3 it follows that $h_j(t_0+1) \subset h_j(t_0)$ since at least one $k \in h_j(t_0)$ must be connected to an agent outside $h_j(t_0)$ and this agent's state must change $\mu_k(t_0) \neq \mu_k(t_0+1)$ as a consequence of Lemma 3 such that $f_{\ell_p, d}(\mu_k(t_0+1)) \neq \mathbf{x}_j$. At the next time t_0+1 it holds again that for each $k \in h_j(t_0+1)$ (assuming $h_j(t_0+1) \neq \emptyset$) the neighbor set obeys $\mathcal{N}_k(t_0+1) \neq \emptyset$ for each $j \in \{1, \dots, l\}$. Then by application of Lemma 3 it follows again that $h_j(t_0+2) \subset h_j(t_0+1) \subset h_j(t_0)$. Thus, $h_j(t+1) \subset h_j(t)$ is a strictly decreasing set-valued function unless $h_j(t) = \emptyset$. By at most time t_0+n-1 it follows that $h_j(t_0+n-1) = \emptyset$ and the argument can reset by redefining t_0 . It follows that $f_{\ell_p, d}(\text{co}(\{\mu_i(t_0+n-1)\})) \subset f_{\ell_p, d}(\text{co}(\{\mu_i(t_0)\}))$ for all $t_0 \geq 0$. Following the proof of Lemma 2 we know $\text{co}(\{\mu_i(t_0+n-1)\}) \subset \text{co}(\{\mu_i(t_0)\})$, and thus because we chose t_0 arbitrarily $\nu(\mu(t+n-1)) < \nu(\mu(t))$ for each $t \in \mathbb{N}$ unless $\mu_i(t+n-1) = \mu_i(t) \forall i$, as desired. The existence of a strictly decreasing Lyapunov function completes the proof. \square

The preceding lemma specializes this theorem to the case where the network topology is connected and time-invariant (but otherwise arbitrary). This lemma is of interest on its own in many applications in which the topology is static. Proof of this lemma, given Lemmas 1–3, follows roughly the analysis of [5] on nonlinear consensus in the usual Euclidean metric space.

We are now ready to prove Theorem 1.

Proof of Theorem 1. The proof here relies on extending the previous lemma to the case where $\mathcal{G}(t)(\mathcal{V}, \mathcal{E}(t))$ is time-varying and for all $t_0 \in \mathbb{N}$ the graph union $\mathfrak{G}(t_0, \infty)$ is connected. Recall the same Lyapunov function (3) as used in the proof of Lemma 4 (we assume familiarity with the proof of Lemma 4 going forward).

We note that it suffices to show that there is a countably infinite number of finite time intervals $t \in [t_0^q, \hat{t}_0^q]$, $q \in \mathbb{N}$, such that $\nu(\mu(t_0^q + \hat{t}_0^q)) < \nu(\mu(t_0^q))$.

Pick $t_0^q \geq 0$, $q \in \mathbb{N}$, so $f_{\ell_p, d}(\text{co}(\{\mu_i(t_0^q)\}))$ is an l -sided polygon in \mathbb{R}^2 with $2 \leq l \leq |\mathcal{V}|$ on the collection of vertices $\{\mathbf{x}_j(t_0^q)\}$, $j \in \{1, \dots, l\}$, with $\mathbf{x}_j(t_0^q) \in \mathbb{R}^2$. Recall (4). Then define a sequence of times $\{t_{s(j)}^q\}$, $s(j) \in \mathbb{N}$, each greater than t_0^q for each $j \in \{1, \dots, l\}$ with $l \geq 2$. The connectivity condition implies the existence of such a sequence for each j with the property that if $h_j(t_{s(j)}^q) \neq \emptyset$, there exists a $k \in h_j(t_{s(j)}^q)$ that is connected to an agent outside $h_j(t_{s(j)}^q)$. Then, this agent's state must change $\mu_k(t_{s(j)}^q) \neq \mu_k(t_{s(j)}^q+1)$ as a consequence of Lemma 3 and $f_{\ell_p, d}(\mu_k(t_{s(j)}^q+1)) \neq \mathbf{x}_j(t_0^q)$. Then $h_j(t_{s(j)}^q+1) \subset h_j(t_{s(j)}^q)$ for all $j \in \{1, \dots, l\}$ unless obviously $h_j(t_{s(j)}^q) = \emptyset$. As in the proof of Lemma 4 it holds that $s(j) \geq n-1$ implies $h_j(t_{s(j)+1}^q) = \emptyset$ for all j . Let $\hat{t}_0^q = \min\{t \in \mathbb{N} : t > t_0^q, s(j) \geq n-1 \forall j\}$ and note then that the interval $t \in [t_0^q, \hat{t}_0^q]$ is finite owing to the connectivity condition. Moreover, as in the proof of Lemma 4 one can then show that $\nu(\mu(\hat{t}_0^q)) < \nu(\mu(t_0^q))$. Restart the argument by picking t_0^{q+1} to be

equal or sufficiently close to \hat{t}_0^q and note that the connectivity condition then implies the number of such (finite) intervals $t \in [t_0^q, \hat{t}_0^q]$ is countably infinite on $q \in \mathbb{N}$.

We thus have a strictly decreasing Lyapunov function $\nu(\mu(\hat{t}_0^q)) < \nu(\mu(t_0^q))$ on the sequence of finite intervals $t \in [t_0^q, \hat{t}_0^q]$, $q \in \mathbb{N}$, and this completes the proof. \square

3. Special cases and convergence details. First, given Theorem 1, it is worth noting the following result.

PROPOSITION 1. *Consider a group of agents \mathcal{V} and network $\mathcal{G}(t)(\mathcal{V}, \mathcal{E}(t))$, where each agent i has initial state $\mu_i(0) \in \mathfrak{U}_p(\mathbb{R})$ and updates its state $\mu_i(t)$ according to (1). Suppose for all $t_0 \in \mathbb{N}$ the graph union $\mathfrak{G}(t_0, \infty)$ is connected so that Theorem 1 applies and there exists $\mu^* \in \mathfrak{U}_p(\mathbb{R})$ such that $\lim_{t \rightarrow \infty} \ell_p(\mu_i(t), \mu^*) = 0$ for any $i \in \mathcal{V}$. Then there exists some symmetric weight matrix $\overline{\mathbf{W}} = [\overline{w}_{ij}] \in \mathbb{R}^{n \times n}$ with $\overline{w}_{ij} \in (0, 1)$ and $\sum_{j \in \mathcal{V}} \overline{w}_{ij} = 1$ for all i such that for any $i \in \mathcal{V}$*

$$\mu^* = \operatorname{argmin}_{\eta \in \mathfrak{U}_p(\mathbb{R})} \sum_{j \in \mathcal{V}} \overline{w}_{ij} \ell_p(\eta, \mu_j(0))^p,$$

where we emphasize that $\overline{\mathbf{W}}$ is not (generally) the same as $\mathbf{W}(t)$ but it is solely dependent on the sequence $\mathbf{W}(t)$, $t \in \mathbb{N}$, and (possibly) the initial measures $\{\mu_i(0)\}_{i \in \mathcal{V}}$.

Proof of this proposition is straightforward given the actual convergence result stated in Theorem 1. This result states that the common measure which all agent states converge to is within the convex hull of all initial agent measures in \mathfrak{U}_p .

An interesting open problem is how one can design the evolution of $\mathbf{W}(t)$, $t \in \mathbb{N}$, such that for a set of measures $\{\mu_i(0)\}_{i \in \mathcal{V}}$ the final weighting matrix $\overline{\mathbf{W}}$ specifies a limit μ^* , i.e., $\lim_{t \rightarrow \infty} \ell_p(\mu_i(t), \mu^*) = 0$ for any $i \in \mathcal{V}$, that is optimal, or desired, in some sense (e.g., minimum variance over all possible $\overline{\mathbf{W}}$ given $\mu_i(0) \in \mathfrak{U}_p(\mathbb{R})$, $i \in \mathcal{V}$). One example of “average” consensus is given later.

In the remainder of this section we consider convergence to particular limits of interest, e.g., to a limit equally close to all agents’ initial measures. We also consider convergence speeds and we consider computational aspects of the update protocol for given classes of input measures.

3.1. General convergence speeds and average Wasserstein consensus.

In this subsection we consider only undirected, connected, and time-invariant network graphs $\mathcal{G}(\mathcal{V}, \mathcal{E})$. We consider only measures with finite second moment, and we work solely in the Wasserstein metric space denoted by $(\mathfrak{U}_2(\mathbb{R}), \ell_2)$.

The first result considers the convergence speed of the entire group of agents under the protocol (1).

PROPOSITION 2. *Consider a group of agents \mathcal{V} and a connected time-invariant network $\mathcal{G}(\mathcal{V}, \mathcal{E})$ where each agent i updates its state $\mu_i(t) \in \mathfrak{U}_2(\mathbb{R})$ according to (1). Then there exists $\mu^* \in \mathfrak{U}_2(\mathbb{R})$ such that for any $i \in \mathcal{V}$,*

$$\lim_{t \rightarrow \infty} \ell_2(\mu_i(t), \mu^*) = 0$$

at an exponential rate.

Proof. If $\mu_i(0) \in \mathfrak{U}_2(\mathbb{R})$, then the solution to (1) at any $i \in \mathcal{V}$ and any $t \in \mathbb{N}$ can be written in the form

$$\mu_i(t+1)(M) = \left(\sum_{j \in \mathcal{N}_i} w_{ij} T_j^i(t) \right) \# \mu_i(t)(M)$$

for all Borel sets M on (\mathbb{R}, d) ; see [40]. Here $(T_j^i(t))\#\mu_i(t)$ denotes the push forward of $\mu_i(t)$ to $\mu_j(t)$ through the nondecreasing measurable map $T_j^i(t) : \mathbb{R} \rightarrow \mathbb{R}$ such that $(T_j^i(t))\#\mu_i(t) = \mu_j(t)$. Obviously, we have $(T_i^i(t))\#\mu_i(t) = \mu_i(t)$. For any measure $\psi(t)$ dominated by the Lebesgue measure on \mathbb{R} it follows that

$$\left(\sum_{j \in \mathcal{N}_i} w_{ij} T_j^i(t)\right) \# \mu_i(t)(M) = \left(\sum_{j \in \mathcal{N}_i} w_{ij} T_j(t)\right) \# \psi(t)(M)$$

where $(T_j(t))\#\psi(t)$ denotes the push forward of $\psi(t)$ to $\mu_j(t)$; see [46]. Write the cumulative distribution function for each $\mu_i(t) \in \mathfrak{U}_2(\mathbb{R})$ by $F_i(x) : \mathbb{R} \rightarrow [0, 1]$ and $F_i(x) = \mu_i((-\infty, x])$. Define its inverse $F_i^-(x) : [0, 1] \rightarrow \mathbb{R}$ by

$$F_i^-(x) = \inf_y \{y \in \mathbb{R} : F_i(y) \geq x\}$$

for all $x \in [0, 1]$. One can show [40] that $T_j^i(t) = F_j^- \circ F_i$ or if $\psi(t)$ is uniform on $[0, 1]$, then $T_j^i(t) = F_j^-$ and

$$\mu_i(t+1)(M) = \left(\sum_{j \in \mathcal{N}_i} w_{ij} F_j^-\right) \# \psi(t)(M)$$

with $\psi(t)$ uniform on $[0, 1]$. It follows directly that the solution to (1) at any $i \in \mathcal{V}$ and any $t \in \mathbb{N}$ has an inverse cumulative distribution function given by

$$(5) \quad F_i^-(t+1)(x) = \sum_{j \in \mathcal{N}_i} w_{ij} F_j^-(t)(x)$$

for all $x \in [0, 1]$.

Now one can stack these functions so $F^-(t+1)(\mathbf{x}) = \mathbf{W}F^-(t)(\mathbf{x}) \forall \mathbf{x} \in [0, 1]^n$. From the assumed network connectivity condition and the weighting assumptions we conclude that \mathbf{W} is row-stochastic and primitive with a distinct maximum eigenvalue of 1. The remaining $n-1$ eigenvalues have an absolute value strictly less than 1. The convergence rate of $F^-(t)(\mathbf{x})$ is determined by the convergence rate of \mathbf{W}^t to the rank one matrix $\mathbf{1}\mathbf{u}^\top$ associated with the maximum eigenvalue. Writing

$$\mathbf{W}^t = \sum_{i=1}^n \lambda_i^t \mathbf{v}_i \mathbf{u}_i^\top = \mathbf{1}\mathbf{u}^\top + \sum_{i=2}^n \lambda_i^t \mathbf{v}_i \mathbf{u}_i^\top,$$

where λ_i is the i th eigenvalue of \mathbf{W} , it then follows that $\|\mathbf{W}^t - \mathbf{1}\mathbf{u}^\top\| = \|\sum_{i=2}^n \lambda_i^t \mathbf{v}_i \mathbf{u}_i^\top\|$ vanishes exponentially at a rate dominated by the absolute value of the second largest eigenvalue (which is strictly less than 1) and the proof is complete. \square

Note that a time-invariant network model is certainly not necessary for exponential convergence but we do not consider further generalization in this work. It is also important to note that the time-varying network connectivity condition allowed in Theorem 1 is also certainly too weak to ensure exponential convergence in general. Indeed, Theorem 1 does not even require the network to be jointly connected until some arbitrary finite future time.

COROLLARY 2. *Consider a group of agents \mathcal{V} and a connected time-invariant network $\mathcal{G}(\mathcal{V}, \mathcal{E})$. Suppose that \mathbf{W} is doubly stochastic and that each agent i updates its state $\mu_i(t) \in \mathfrak{U}_2(\mathbb{R})$ according to (1). For any $i \in \mathcal{V}$, we have $\lim_{t \rightarrow \infty} \ell_2(\mu_i(t), \mu^*) = 0$ at an exponential rate where*

$$\mu^* = \operatorname{argmin}_{\eta \in \mathfrak{U}_2(\mathbb{R})} \frac{1}{n} \sum_{j \in \mathcal{V}} \ell_2(\eta, \mu_j(0))^2.$$

Proof. This result follows again because linear consensus in \mathbb{R} over a time-invariant network with a doubly stochastic weighting matrix leads asymptotically to “average” consensus [9]. Looking at (5) we see that (nonlinear) consensus via (1) is related to (linear) consensus in the space of inverse cumulative distribution functions. Moving from the limiting inverse cumulative distribution function to a probability measure does not change the limiting $1/n$ averaging coefficient. \square

This corollary provides sufficient conditions⁵ under which exponential convergence to a measure is achieved and where the consensus measure achieved asymptotically at each agent is an average distance to all initial measures. In this case, as per Proposition 1, we have $\bar{\mathbf{W}} = \frac{1}{n} \mathbf{1}\mathbf{1}^\top$. We note that other consensus measures may be more desirable, e.g., one may want to reach an agreement on that measure with the smallest variance within the convex hull of all initial measures.

3.2. Convergence with Gaussian measures. In this subsection we consider the flow of operation (1) when $\mu_i(0)$ is a Gaussian measure. As in the preceding subsection, we consider only the case $(\mathfrak{U}_2(\mathbb{R}), \ell_2)$.

Suppose $\mu_i(t) \in \mathfrak{U}_2(\mathbb{R})$ for all $i \in \mathcal{V}$ admits a Gaussian density of the form $\mathcal{N}(p_i, P_i)$. Then it follows that $\mu_i(t + 1) \in \mathfrak{U}_2(\mathbb{R})$ is a Gaussian measure [38, 40] of density $\mathcal{N}(q, Q)$ where the updated mean and variance are given by

$$q = \sum_{j \in \mathcal{N}_i(t)} w_{ij}(t) p_j,$$

$$Q = \left(\sum_{j \in \mathcal{N}_i(t)} w_{ij}(t) P_j^{1/2} \right)^2.$$

One may find this result in the scalar case by studying (5) and noting in this case that

$$F_i^-(t)(x) = p_i + P_i^{1/2} \sqrt{2} \operatorname{erf}^{-1}(2x - 1),$$

where $\operatorname{erf}(\cdot)$ is the standard error function. The (weighted) averaging in (5) taken pointwise in x is then just an averaging over the input means and standard deviations.

The following corollary then specializes those results in the preceding subsection to Gaussian measures.

COROLLARY 3. *Consider a group of agents \mathcal{V} and a connected time-invariant network $\mathcal{G}(\mathcal{V}, \mathcal{E})$. Assume \mathbf{W} is doubly stochastic and $\mu_i(0) \in \mathfrak{U}_2(\mathbb{R})$ admits a Gaussian density $\mathcal{N}(p_i(0), P_i(0))$. Then $\mu_i(t + 1) \in \mathfrak{U}_2(\mathbb{R})$ in (1) admits a Gaussian density $\mathcal{N}(p_i(t + 1), P_i(t + 1))$, where*

$$p_i(t + 1) = \sum_{j \in \mathcal{N}_i} w_{ij} p_j(t),$$

$$P_i^{1/2}(t + 1) = \sum_{j \in \mathcal{N}_i} w_{ij} P_j^{1/2}(t).$$

Moreover, we have for any $i \in \mathcal{V}$ $\lim_{t \rightarrow \infty} \ell_2(\mu_i(t), \mu^*) = 0$ exponentially fast where μ^* satisfies

$$\mu^* = \operatorname{argmin}_{\eta \in \mathfrak{U}_2(\mathbb{R})} \frac{1}{n} \sum_{j \in \mathcal{V}} \ell_2(\eta, \mu_j(0))^2.$$

This corollary collapses to a (classical) scalar average consensus algorithm [7] on the mean and standard deviation at each iteration. More general results achieving

⁵A time-invariant network topology and a doubly stochastic weighting matrix. The time-invariance constraint can be relaxed—it is just sufficient—but we do not consider generalization here.

average consensus in this Gaussian setting that accommodate time-varying networks, finite-time convergence, etc. [4, 6, 7, 19] may be substituted.

Although the update (1) is linear (in mean and standard deviation) and closed in the event of Gaussian input measures, this is not generally true. The consensus problem (1) is, in general, inherently nonlinear.

3.3. Computational aspects of the update protocol. In the case of Gaussian input measures, we have shown that the updating step of our consensus algorithm can be performed in closed form and that the resulting algorithm resembles a particular case of standard linear consensus in \mathbb{R} ; e.g., see [9].

Consider the important scenario where all the input initial measures are empirical measures,⁶

$$\mu_i^N(0)(dx) = \frac{1}{N} \sum_{j=1}^N \delta_{x_j^i(0)}(dx),$$

where $\delta_y(dx)$ denotes the delta-Dirac measure located at y . In this case, the minimization in (1) can be solved exactly. First, we define the order statistics,

$$x_{1,*}^i(t) \leq x_{2,*}^i(t) \leq \dots \leq x_{N,*}^i(t)$$

corresponding to $\{x_j^i(t)\}$, $1 \leq j \leq N$. Then, with notation as in (1) we define

$$\bar{x}_{k,*}^i(t) = \sum_{j \in \mathcal{N}_i(t)} w_{ij}(t) x_{k,*}^j(t)$$

for all $1 \leq k \leq N$. Then,

$$\mu_i^N(t+1)(dx) = \frac{1}{N} \sum_{j=1}^N \delta_{\bar{x}_{j,*}^i(t)}(dx).$$

This computation involves only sorting and averaging of numbers in \mathbb{R} .

We note as an aside that for empirical measures in \mathbb{R}^n , for any integer $n \geq 1$, the minimization in (1) can still be solved exactly via a finite-dimensional linear program [47, 48] and the resulting measure is again an empirical distribution.⁷ However, the computational requirements of this linear program (in \mathbb{R}^n , when $n \geq 2$) may explode quickly with the number of input measures and the number of atoms of these measures; see [47, 48, 49, 50]. Numerous fast approximation methods have been derived for computing the barycenter update (1) itself [42, 45, 51, 52]. The details of these algorithms are beyond the scope of this work, but it follows that the update (1) itself is thus computable with empirical measures in higher dimensions \mathbb{R}^n . Albeit not the focus of this article, we note in passing that convergence of the distributed protocol (1) with empirical measures on \mathbb{R}^n follows readily from arguments on the convergence of Euclidean (linear) consensus [5, 49].

Consider now, more generally, arbitrary input measures on \mathbb{R} . The optimization problem (1) typically does not admit a closed-form solution. However, it is convex [40, 42] and thus numerical methods/approximations are feasible and already exist

⁶Interestingly, if each input measure is defined by a single Dirac (in $(\mathcal{M}_2(\mathbb{R}), \ell_2)$, i.e., with $p = 2$), then the classical (linear) consensus algorithm in \mathbb{R} is recovered, e.g., as in [5, 9]. Of course, typically one is interested in more general empirical input measures.

⁷As discussed previously, for special, certainly nongeneric, arrangements of discrete measures the minimization in (1) may not have a unique solution (though a solution always exists) [42].

in a number of cases; see [41, 42, 52]. The issue of convergence under approximate update steps is unclear and beyond the scope of this work; but practically one would intuit that close approximation leads to close convergence. While convexity of the minimization problem is advantageous in general, for measures on \mathbb{R} there are yet further virtues. The update in (5) is typically computable in closed form and thus we “almost” have a general closed-form expression for (1). The relationship between (1) and the inverse cumulative distribution in (5) is the basis for the solution with empirical measures on \mathbb{R} given in the preceding paragraph. This relationship has also been explored in [40, 41, 42] with further example computations, and as a lead into more general computational results.

4. Discussion and applications. The output of each iteration of operation (1) is known in the literature as the Wasserstein barycenter. Similarly, the limit $\mu^* \in \mathcal{U}_p(\mathbb{R})$ to which all agents converge upon repeated iteration of operation (1) is also a Wasserstein barycenter (on a fully connected graph of all agent’s initial measures). In other words, this work studies the convergence properties of a consensus algorithm concerned with distributed (iterative) computation of the Wasserstein barycenter over a (possibly) time-varying, arbitrary, network topology. We consider undirected networks in this work for simplicity; but directed networks may be studied as in, e.g., [5, 6], with additional conditions needed for convergence in that setting [5].

While this is the first such study in this direction, potential applications/uses for the Wasserstein barycenter (itself) have been considered previously in a number of fields [42, 45, 48, 49, 50, 51, 51, 52, 53, 54], and this list is by no means exhaustive.

Arguably the most popular domain in which the Wasserstein barycenter has found applications is in computer vision and image/video processing [42, 48, 51]. We do not consider specifics here, but the interested reader may consult [42], where numerous examples and a detailed discussion are given. It is noted [42] that state-of-the-art advancements in a number of related problems have arisen via the use of Wasserstein barycenters. Importantly, both Gaussian and discrete measures find applicability through the Wasserstein barycenter in computer vision and image/video processing; again, see [42].

Applications in machine learning and Bayesian statistics have also made use of the Wasserstein barycenter [41, 52, 53], and it is envisioned that this technology (and the related optimal transportation problem) will find wider adoption in this field. In this setting, distributed (or even parallel) computation of the Wasserstein barycenter is likely important; e.g., distributed Bayesian computation on large data sets is the subject of [53].

Related work on consensus in spaces of probability distributions has been considered in the field of distributed estimation and information fusion. Suppose each agent starts with a (posterior) probability measure associated to some common underlying event of interest. Then one may like to combine all these measures (which amount to each agents’ estimate and/or belief of the underlying event) into a common probability measure that captures all the agents’ beliefs; this is often called opinion pooling. Related work in [55] considers the application of consensus to the problem of distributed Bayesian computations. In [56, 57, 58] the consensus algorithm from [55] is further extended and applied in distributed estimation and filtering. The Bayesian ideas in [55, 56, 57, 58] are related to so-called log-linear opinion pools which are related to the barycenter defined with respect to a Kullback–Leibler divergence (in an analogous fashion to the Wasserstein barycenter) [58]. A Monte Carlo approximation of the consensus algorithm from [55] was studied in [57]. The log linear opinion pool

for distributed information fusion was extended in [59, 60] to the general barycenter of a Bregman divergence (of which the Kullback–Leibler-based barycenter is a special case). Monte Carlo approximations were also considered in [59, 60]. While Bregman and Kullback–Leibler divergences have Bayesian-type interpretations, it is also possible to consider information fusion and distributed estimation in the context of Wasserstein barycenters as in [41, 49, 50]. As shown herein, and in [50], information fusion with the Wasserstein barycenter has advantageous computability properties in the space of empirical distributions. Further study of the Wasserstein barycenter in the context of information fusion and estimation is an ongoing topic.

We finally highlight that the “consensus” terminology throughout this article is used in the sense of network consensus and agreement as in [1, 2, 3, 4, 5, 6, 7]. This differs from the topic of consensus (or ensemble) clustering or consensus aggregation [61, 62] which may also make use of the Wasserstein barycenter [63, 64] for distributional clustering, etc. The latter topic may be an application of networked-type consensus as considered herein.

5. Concluding remarks. Distributed consensus in the Wasserstein metric space of probability measures was introduced in this paper. It is shown that convergence of the individual agents’ measures to a common measure value is guaranteed if a relatively weak network connectivity condition is satisfied. The measure that is achieved asymptotically at each agent is the measure that minimizes a weighted sum of its Wasserstein distances to these initial measures and is known as the Wasserstein barycenter in the literature.

Finally, we note that following [5], it would be straightforward to consider an extension to the case in which the network topology is directed and one expects analogous results (concerning connectivity) to apply in the Wasserstein space considered herein. For brevity, and notational simplicity, we do not explore this scenario further. Moreover, one may seek analogous results in the Wasserstein metric space of measures defined on the Borel sets of (\mathbb{R}^m, d) for some $m \geq 2$. We conjecture that similar results hold in this case. However, while many of the lemmas used herein carry over immediately, this generalization is not immediate. Indeed, the Wasserstein metric space in such cases is positively curved, so it does not resemble Euclidean space and it is not CAT(0).

REFERENCES

- [1] M.H. DEGROOT, *Reaching a consensus*, J. Amer. Statist. Assoc., 69 (1974), pp. 118–121.
- [2] J.N. TSITSIKLIS, D.P. BERTSEKAS, AND M. ATHANS, *Distributed asynchronous deterministic and stochastic gradient optimization algorithms*, IEEE Trans. Automat. Control, 31 (1986), pp. 803–812.
- [3] A. JADBABAIE, J. LIN, AND A.S. MORSE, *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Trans. Automat. Control, 48 (2003), pp. 988–1001.
- [4] R. OLFATI-SABER AND R.M. MURRAY, *Consensus problems in networks of agents with switching topology and time-delays*, IEEE Trans. Automat. Control, 49 (2004), pp. 1520–1533.
- [5] L. MOREAU, *Stability of multiagent systems with time-dependent communication links*, IEEE Trans. Automat. Control, 50 (2005), pp. 169–182.
- [6] W. REN AND R.W. BEARD, *Consensus seeking in multiagent systems under dynamically changing interaction topologies*, IEEE Trans. Automat. Control, 50 (2005), pp. 655–661.
- [7] R. OLFATI-SABER, J.A. FAX, AND R.M. MURRAY, *Consensus and cooperation in networked multi-agent systems*, Proc. IEEE, 95 (2007), pp. 215–223.
- [8] M. CAO, A.S. MORSE, AND B.D.O. ANDERSON, *Reaching a consensus in a dynamically changing environment: A graphical approach*, SIAM J. Control Optim., 47 (2008), pp. 575–600, <https://doi.org/10.1137/060657005>.

- [9] L. XIAO, S. BOYD, AND S. LALL, *A scheme for robust distributed sensor fusion based on average consensus*, in Proceedings of the 4th International Symposium on Information Processing in Sensor Networks, Los Angeles, IEEE, 2005, pp. 63–70.
- [10] D.P. SPANOS AND R.M. MURRAY, *Distributed sensor fusion using dynamic consensus*, in Proceedings of the 16th IFAC World Congress, Prague, Czech Republic, 2005.
- [11] R. OLFATI-SABER, *Distributed Kalman filtering for sensor networks*, in Proceedings of the 46th IEEE Conference on Decision and Control, New Orleans, 2007, pp. 5492–5498.
- [12] R. CARLI, A. CHIUSO, L. SCHENATO, AND S. ZAMPIERI, *Distributed Kalman filtering based on consensus strategies*, IEEE J. Sel. Areas Commun., 26 (2008), pp. 622–633.
- [13] F.S. CATTIVELLI AND A.H. SAYED, *Diffusion strategies for distributed Kalman filtering and smoothing*, IEEE Trans. Automat. Control, 55 (2010), pp. 2069–2084.
- [14] Q. HUI AND W.M. HADDAD, *Distributed nonlinear control algorithms for network consensus*, Automatica J. IFAC, 44 (2008), pp. 2375–2381.
- [15] W. YU, G. CHEN, AND M. CAO, *Consensus in directed networks of agents with nonlinear dynamics*, IEEE Trans. Automat. Control, 56 (2011), pp. 1436–1441.
- [16] A. AJORLOU, A. MOMENI, AND A.G. AGHDAM, *Sufficient conditions for the convergence of a class of nonlinear distributed consensus algorithms*, Automatica J. IFAC, 47 (2011), pp. 625–629.
- [17] J. CORTÉS, *Distributed algorithms for reaching consensus on general functions*, Automatica J. IFAC, 44 (2008), pp. 726–737.
- [18] X. WANG AND Y. HONG, *Distributed finite-time χ -consensus algorithms for multi-agent systems with variable coupling topology*, J. Syst. Sci. Complex., 23 (2010), pp. 209–218.
- [19] J. CORTÉS, *Finite-time convergent gradient flows with applications to network consensus*, Automatica J. IFAC, 42 (2006), pp. 1993–2000.
- [20] M. ZHU AND S. MARTÍNEZ, *Discrete-time dynamic average consensus*, Automatica J. IFAC, 46 (2010), pp. 322–329.
- [21] Y. HONG, J. HU, AND L. GAO, *Tracking control for multi-agent consensus with an active leader and variable topology*, Automatica J. IFAC, 42 (2006), pp. 1177–1182.
- [22] S.H. STROGATZ, *From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators*, Phys. D, 143 (2000), pp. 1–20.
- [23] F. DÖRFLER AND F. BULLO, *Synchronization in complex networks of phase oscillators: A survey*, Automatica J. IFAC, 50 (2014), pp. 1539–1564.
- [24] Z. LI, Z. DUAN, G. CHEN, AND L. HUANG, *Consensus of multiagent systems and synchronization of complex networks: A unified viewpoint*, IEEE Trans. Circuits Syst. I. Regul. Pap., 57 (2010), pp. 213–224.
- [25] S. EMRE TUNA AND R. SEPULCHRE, *Consensus under general convexity*, in Proceedings of the 46th IEEE Conference on Decision and Control, New Orleans, IEEE, 2007, pp. 294–299.
- [26] A. SARLETTE AND R. SEPULCHRE, *Consensus optimization on manifolds*, SIAM J. Control Optim., 48 (2009), pp. 56–76, <https://doi.org/10.1137/060673400>.
- [27] I. MATEI AND J.S. BARAS, *The asymptotic consensus problem on convex metric spaces*, IEEE Trans. Automat. Control, 60 (2015), pp. 907–921.
- [28] R. SEPULCHRE, *Consensus on nonlinear spaces*, Annu. Rev. Control, 35 (2011), pp. 56–64.
- [29] P. GROHS, *Wolfowitz’s Theorem and Convergence of Consensus Algorithms in Hadamard Spaces*, Technical Report SAM Report 2012-27, ETH, Zurich, Switzerland, 2012.
- [30] S. BONNABEL, *Stochastic gradient descent on Riemannian manifolds*, IEEE Trans. Automat. Control, 58 (2013), pp. 2217–2229.
- [31] F. GARIN AND L. SCHENATO, *A survey on distributed estimation and control applications using linear consensus algorithms*, in Networked Control Systems, Lect. Notes Control Inf. Sci. 406, Springer, Berlin, 2010, pp. 75–107.
- [32] C.R. GIVENS AND R.M. SHORTT, *A class of Wasserstein metrics for probability distributions*, Michigan. Math. J., 31 (1984), pp. 231–240.
- [33] A.N. BISHOP AND A. DOUCET, *Distributed Nonlinear Consensus in the Space of Probability Measures*, in Proceedings of the 19th IFAC World Congress, Cape Town, South Africa, 2014.
- [34] C. VILLANI, *Topics in Optimal Transportation*, American Mathematical Society, Providence, RI, 2003.
- [35] L. AMBROSIO, N. GIGLI, AND G. SAVARÉ, *Gradient Flows in Metric Spaces and in the Space of Probability Measures*, Birkhäuser Verlag, Basel, 2005.
- [36] B. KLOECKNER, *A geometric study of Wasserstein spaces: Euclidean spaces*, Ann. Sc. Norm. Super. Pisa Cl. Sci. (5), 9 (2010), pp. 297–323.

- [37] J. BERTRAND AND B. KLOECKNER, *A geometric study of Wasserstein spaces: Hadamard spaces*, J. Topol. Anal., 4 (2012), pp. 515–542.
- [38] R.J. MCCANN, *A convexity principle for interacting gases*, Adv. Math., 128 (1997), pp. 153–179.
- [39] P. MAJOR, *On the invariance principle for sums of independent identically distributed random variables*, J. Multivariate Anal., 8 (1978), pp. 487–517.
- [40] M. AGUEH AND G. CARLIER, *Barycenters in the Wasserstein space*, SIAM J. Math. Anal., 43 (2011), pp. 904–924, <https://doi.org/10.1137/100805741>.
- [41] J. BIGOT AND T. KLEIN, *Consistent Estimation of a Population Barycenter in the Wasserstein Space*, preprint, <https://arxiv.org/abs/1212.2562v1>, 2012.
- [42] N. BONNEEL, J. RABIN, G. PEYRÉ, AND H. PFISTER, *Sliced and Radon Wasserstein barycenters of measures*, J. Math. Imaging Vision, 51 (2015), pp. 22–45.
- [43] D. BURAGO, Y. BURAGO, AND S. IVANOV, *A Course in Metric Geometry*, American Mathematical Society, Providence, RI, 2001.
- [44] M.R. BRIDSON AND A. HAEFLIGER, *Metric Spaces of Non-Positive Curvature*, Springer, Berlin, 1999.
- [45] E. BOISSARD, T. LE GOUIC, AND J.-M. LOUBES, *Distribution's Template Estimate with Wasserstein Metrics*, preprint, <https://arxiv.org/abs/1111.5927>, 2011.
- [46] N. BONNEEL, J. RABIN, G. PEYRÉ, AND H. PFISTER, *Sliced and Radon Wasserstein barycenters of measures*, J. Math. Imaging Vision, 51 (2015), pp. 22–45.
- [47] W. GANGBO AND A. SWIECH, *Optimal maps for the multidimensional Monge-Kantorovich problem*, Comm. Pure Appl. Math., 51 (1998), pp. 23–45.
- [48] J. RABIN, G. PEYRÉ, J. DELON, AND M. BERNOT, *Wasserstein barycenter and its application to texture mixing*, in Scale Space and Variational Methods in Computer Vision, Springer, Berlin, Heidelberg, 2012, pp. 435–446.
- [49] A.N. BISHOP, *Gossip-based distributed data fusion of empirical probability measures*, in Proceedings of the 2014 IEEE Workshop on Statistical Signal Processing, Gold Coast, Australia, IEEE, 2014, pp. 372–375.
- [50] A.N. BISHOP, *Information fusion via the Wasserstein barycenter in the space of probability measures: Direct fusion of empirical measures and Gaussian fusion with unknown correlation*, in Proceedings of the 17th International Conference on Information Fusion, Salamanca, Spain, 2014, pp. 1–7.
- [51] G. CARLIER, A. OBERMAN, AND E. OUDET, *Numerical Methods for Matching for Teams and Wasserstein Barycenters*, preprint, <https://arxiv.org/abs/1411.3602>, 2014.
- [52] M. CUTURI AND A. DOUCET, *Fast computation of Wasserstein barycenters*, in Proceedings of the 31st International Conference on Machine Learning, Beijing, China, 2014, pp. 685–693.
- [53] S. SRIVASTAVA, V. CEVHER, Q. TRAN-DINH, AND D.B. DUNSON, *WASP: Scalable Bayes via barycenters of subset posteriors*, in Proceedings of the 18th International Conference on Artificial Intelligence and Statistics, San Diego, 2015, pp. 912–920.
- [54] J. ARROYO, G. GONZÁLEZ-RIVERA, C. MATÉ, AND A.M. SAN ROQUE, *Smoothing methods for histogram-valued time series: An application to value-at-risk*, Stat. Anal. Data Min., 4 (2011), pp. 216–228.
- [55] R. OLFATI-SABER, E. FRANCO, E. FRAZZOLI, AND J.S. SHAMMA, *Belief consensus and distributed hypothesis testing in sensor networks*, in Networked Embedded Sensing and Control, Lect. Notes Control Inf. Sci. 331, Springer, Berlin, Heidelberg, 2005, pp. 169–182.
- [56] G. BATTISTELLI AND L. CHISCI, *Kullback–Leibler average, consensus on probability densities, and distributed state estimation with guaranteed stability*, Automatica J. IFAC, 50 (2014), pp. 707–718.
- [57] I.L. MANUEL AND A.N. BISHOP, *Distributed Monte Carlo information fusion and distributed particle filtering*, in Proceedings of the 19th IFAC World Congress, Cape Town, South Africa, 2014, pp. 8681–8688.
- [58] S. BANDYOPADHYAY AND S.-J. CHUNG, *Distributed Bayesian filtering using logarithmic opinion pool for dynamic sensor networks*, Automatica J. IFAC, 97 (2018), pp. 7–17.
- [59] C.N. TAYLOR AND A.N. BISHOP, *Homogeneous functionals and Bayesian data fusion with unknown correlation*, Inf. Fusion, 45 (2019), pp. 179–189.
- [60] C.N. TAYLOR AND A.N. BISHOP, *Distributed power mean fusion*, in Proceedings of the 22nd International Conference on Information Fusion, 2019, pp. 1–8.
- [61] S. VEGA-PONS AND J. RUIZ-SHULCLOPER, *A survey of clustering ensemble algorithms*, Int. J. Pattern Recognit. Artif. Intell., 25 (2011), pp. 337–372.

- [62] J. GHOSH AND A. ACHARYA, *A survey of consensus clustering*, in Handbook of Cluster Analysis, Chapman and Hall/CRC, Boca Raton, FL, 2015, pp. 518–539.
- [63] P.C. ALVAREZ-ESTEBAN, E. DEL BARRIO, J.A. CUESTA-ALBERTOS, AND C. MATRÁN, *Wide consensus aggregation in the Wasserstein space: Application to location-scatter families*, Bernoulli, 24 (2018), pp. 3147–3179.
- [64] I. VERDINELLI AND L. WASSERMAN, *Hybrid Wasserstein distance and fast distribution clustering*, Electron. J. Stat., 13 (2019), pp. 5088–5119.