

Polarization of multi-agent gradient flows over manifolds with application to opinion dynamics

La Mi, Jorge Gonçalves, and Johan Markdahl

Abstract—Multi-agent systems are known to exhibit stable emergent behaviors, including polarization, over \mathbb{R}^n or highly symmetric nonlinear spaces. In this article, we eschew linearity and symmetry of the underlying spaces, and study the stability of polarized equilibria of multi-agent gradient flows evolving on general hypermanifolds. The agents attract or repel each other according to the partition of the communication graph that is connected but otherwise arbitrary. The manifolds are outfitted with geometric features styled “dimples” and “pimples” that characterize the absence of flatness. The signs of inter-agent couplings together with these geometric features give rise to stable polarization under various sufficient conditions. We propose tangible interpretation of the system in the context of opinion dynamics, and highlight throughout the text its versatility in modeling diverse aspects of the polarization phenomenon.

Index Terms—polarization, agents and autonomous systems, network analysis and control, nonlinear systems, stability of NL systems

I. INTRODUCTION

LET there be a collection of simple agents traversing a complex terrain. Their goal is to approach the weighted average positions of their neighbors, which they are programmed to either attract or repel. This is a typical distributed process [1], and is akin to many aspects of social opinion formation, as numerous observers in the systems and control community have identified [2,3]. Accordingly, insights on multi-agent systems, especially on the stability and convergence properties of consensus, have been applied to a large number of opinion dynamics models to study social agreement [4]–[6]. Continuing this tradition, we investigate multi-agent systems that evolve over nonlinear spaces for another type of emergent behavior, namely, polarization. Through this angle, we seek to understand the social phenomenon of opinion polarization, which has become increasingly conspicuous in the backdrop of perceived deepening of social discords [7].

The specific empirical observations concerning the polarization phenomenon we are concerned with are the following. At the individual level, the personal belief system of an agent is simpler and cleaner compared to the richness of the external environment, composed of ideas, events, and issues in the public sphere that are often unrelated to each other. At the macroscopic level, polarization not only arises in a straightforward fashion when two parties holding diametric views are squarely opposed to each other. It might also happen among groups sharing significant common grounds with ostensibly minor differences (e.g., one iota of difference between the Homoousians and Homoiousians [8, Chp. 21]). We show how our results capture these patterns that have not been addressed to date, thus rendering our polarization problem over nonlinear space a suitable phenomenological model for opinion dynamics.

Supported by the Luxembourg National Research Fund (FNR OPEN O19/13904037-SYBION).

Luxembourg Centre for Systems Biomedicine, University of Luxembourg, Belval, L-4367, Luxembourg. Jorge Gonçalves is also with the Department of Plant Sciences, University of Cambridge, Cambridge CB2 3EA, UK. E-mails: la.mi@uni.lu, jorge.goncalves@uni.lu, markdahl@kth.se

More concretely, we study a model of multi-agent gradient flow system confined to manifolds embedded in the Euclidean space. Gradient descent flow, being a sufficiently simple optimizing process, is amenable to rigorous stability analysis and thus widely adopted by many agent-based models as coordinating protocols for robot swarms [9]. The restriction to nonlinear spaces in social dynamics analysis [10], however, is less common. We are motivated by the observation that the variety of issues or situations that an individual confronts is necessarily more complex and nuanced than the set of a few principles, or core beliefs, used to navigate those situations. The external events, the environment, or the expressed opinions of other agents may be naturally presented in the ambient Euclidean space; the core beliefs of an individual are encoded in the local parametrization of a lower dimensional manifold. By the local parametrization of a manifold, we mean the inverse mapping of the chart φ that maps each point on the manifold onto the ambience; it may differ from point to point on the manifold. The analytical investigation into this conceptual interpretation is realized by outfitting the general manifolds with special geometric features styled “dimples” and “pimples”. The interplay between these geometric features and the cooperative/antagonistic interactions among agents then gives rise to different routes to polarization.

There are a few polarization studies on manifolds in the literature, where the n -sphere has received the most attention. Gaitonde et al [11] proved almost sure convergence for a class of Markov processes on the hypersphere. Hong and Strogatz found traveling wave polarization in a variant of the Kuramoto model over the unit circle with conformist and contrarian oscillators [12]. The bifurcation points between different steady states were solved exactly by a series of reduction techniques applied to the mean field approximation. A higher dimensional Kuramoto model analyzed by Ha et al [13] also features positive and negative couplings between agents. They derived stability conditions on initial conditions, relative strengths of the two types of couplings, and frequency matrices that govern the self dynamics. More elaborate state-dependent interaction rules inspired by neuroscience are considered by Crnkčić and Jaćimović over a 3-sphere through a quaternion formulation [14]. The antipodal configuration is asymptotically stable if agents attract or repulse each other when they are respectively close or far. For ring graphs over the 2-sphere, Song et al [15] obtained asymptotically stable polarization with even number of agents. Moreover, the result is almost global if the graph is undirected. For more general manifolds, a recent work by Aydogdu et al [16] explored geodesic and chordal interactions between agents on general Riemannian manifolds, and established the existence of various equilibria and orbits but without stability analysis.

In view of these related works, our contribution is that we provide rigorous stability analysis of polarized equilibria for the multi-agent gradient descent system with arbitrary connected network topology over more general manifolds. When reduced to the hypersphere case, our results generalize and complement the existing ones obtained in the literature [13,15]. Furthermore, we draw on our unique interpretation for embedded lower dimensional manifolds in the context of opinion dynamics, to address aspects of the opinion polarization hitherto unaccounted for.

II. SETUP

A. Geometric features of the manifold

Consider a closed and connected n -dimensional hypermanifold embedded in the Euclidean ambient space

$$\mathcal{H}^n = \{y \in \mathbb{R}^{n+1} \mid c(y) = 0\}$$

implicitly defined by a smooth C^2 function $c : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$, with the usual requirement that the gradient in \mathbb{R}^{n+1} satisfies $\nabla c(x) \neq 0$ for every $x \in \mathcal{H}^n$. A closed manifold is a compact manifold without boundary; closedness and connectedness allow us to quote the Jordan-Brouwer Separation Theorem which says that the hypermanifold \mathcal{H}^n separates its complement $\mathbb{R}^{n+1} - \mathcal{H}^n$ into two disjoint sets [17]. The implicit function c is positive in one of the two disjoint sets, and negative in the other. Without loss of generality, we identify the former with the unbounded set outside \mathcal{H}^n , and the latter with the bounded set inside \mathcal{H}^n . The unit normal $n(x) = \nabla c(x) / \|\nabla c(x)\|$ is outward-pointing, i.e., pointing towards the unbounded set.

The hypermanifold \mathcal{H}^n is equipped with special features: dimples and pimples. To define them, introduce a height function $h_x : \mathcal{H}^n \rightarrow \mathbb{R}$ with respect to a fixed x

$$h_x(y) := \langle n(x), y \rangle, \quad \forall y \in \mathcal{H}^n.$$

The height function gives the altitude of a point y along the axis spanned by $n(x)$. For the example of a 2-sphere, we can take the north pole $(0, 0, 1)$ as the fixed point. Then the normal at the north pole is $(0, 0, 1)$. Consequently, the height function of any point on the sphere gives its z -coordinate. For notational convenience, if the fixed x carries a subscript, e.g. x_i , then h_{x_i} is shortened as h_i . Similarly, $n(x_i)$ is often shortened to n_i . Now we are ready to introduce the definitions for a dimple and a pimple.

Definition 2.1: If for some $x \in \mathcal{H}^n$, $y = x$ is a strict local minimizer of $h_x(y)$ in a sufficiently small neighborhood $\mathcal{I}_x = \{y \in \mathcal{H}^n \mid \|y - x\| < \epsilon\}$, then \mathcal{I}_x is referred to as a dimple, and x the bottom of the dimple. Similarly, if for some $x \in \mathcal{H}^n$, $y = x$ is a strict local maximizer of $h_x(y)$ in a sufficiently small neighborhood \mathcal{I}_x , then \mathcal{I}_x is referred to as a pimple, and x the bottom of the pimple.

Fig. 1 illustrates the concepts of dimples and pimples in \mathbb{R}^3 by three common fruits.

Remark 2.1: The dimple or pimple \mathcal{I}_x does not necessarily contain only one bottom x . It may be that \mathcal{I}_x has a single set of bottoms covering all or part of \mathcal{I}_x , or there are multiple disjoint sets of bottoms within \mathcal{I}_x . We require \mathcal{I}_x to be sufficiently small in Definition 2.1 to exclude the latter case, by shrinking the radius ϵ of the neighborhood around x . However, it is not always possible to completely avoid disjoint sets of bottoms when, e.g., the embedding of the hypermanifold is not analytic.

B. Multi-agent networks

Evolving on the hypermanifold is a homogeneous multi-agent system with N agents, associated with an undirected, connected, and weighted graph structure $\mathcal{G} = (\mathcal{V}, \mathcal{E}, A)$. The adjacency matrix $A = [a_{ij}]$ is symmetrical and has non-negative entries. The vertices \mathcal{V} are divided into two groups $\mathcal{V}_u = \{1, 2, \dots, M\}$ and $\mathcal{V}_l = \{M+1, \dots, N\}$ for $1 < M < N$. The edge set \mathcal{E} is partitioned into intragroup and intergroup sets $\mathcal{E}_+ = \{\{i, j\} \in \mathcal{E} \mid i, j \in \mathcal{V}_u \text{ or } i, j \in \mathcal{V}_l\}$ and $\mathcal{E}_- = \{\{i, j\} \in \mathcal{E} \mid i \in \mathcal{V}_l, j \in \mathcal{V}_u\}$.

Such a partition is introduced to enforce different coupling rules over edges in \mathcal{E}_+ and \mathcal{E}_- . The couplings are positive over all edges in \mathcal{E}_+ , whereas those over \mathcal{E}_- can be either all positive or all negative. This ‘‘edge coloring’’ equivalently generates a *structurally balanced* graph [18] if we allow the graph to be signed such that edge weights

on elements in \mathcal{E}_+ are all positive, and edge weights on elements in \mathcal{E}_- are either all positive or all negative. In fact, doing so would not affect any of our results conceptually, as such, no loss of generality is incurred with the non-negativity requirement on A .

Remark 2.2: It is noted in [18, Lem. 1] that there exists a gauge transformation that brings a structurally balanced signed graph to a nonnegative one. This operation is akin to simultaneously reshuffling group membership and assigning to the affected agents opposite positions in the Euclidean space. It is not applicable in our general case (except the sphere case of special interest in §III-B), because no symmetry is assumed in the underlying nonlinear space, as detailed in §II-A. The symmetry assumption on A , however, is essential, as we deal with gradient flows, see §II-C.

C. Gradient flow dynamics

Let us denote the states of the agents individually by x_i and collectively by $\chi := (x_i)_{i=1}^N$. The agents evolve according to a simple rule of gradient descent flow in continuous time. Given a disagreement function $V : \mathcal{H}^n \rightarrow \mathbb{R}$, the dynamics of each agent is

$$\dot{x}_i = -\text{grad}_i V(\chi) = -P_i (\nabla_i V(\chi)), \quad (1)$$

where grad_i is the intrinsic gradient on the tangent space $T_{x_i} \mathcal{H}^n$ at the point x_i , $P_i = I - n_i n_i'$ is a positive semidefinite projection matrix [19] on $T_{x_i} \mathcal{H}^n$, and $\nabla_i V(\chi)$ is the standard gradient in the Euclidean space of the disagreement function.

As mentioned in §II-B, we divide the N agents into two groups so that members within the same group are attracted to each other, whereas members in different groups can be made to either oppose or attract each other. To model the situation with attractive intragroup coupling and repulsive intergroup coupling on \mathcal{H}^n , we use the disagreement function

$$V_-(\chi) := \frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_+} a_{ij} \|x_j - x_i\|^2 - \frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij} \|x_j - x_i\|^2. \quad (2)$$

Another appealing option in the literature is to retain the sum of square structure [18] by changing the second term

$$\frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_+} a_{ij} \|x_j - x_i\|^2 + \frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij} \|x_j + x_i\|^2. \quad (3)$$

However, (3) is not suitable for nonlinear spaces. Take a two particle system evolving on the unit circle as an example, every antipodal formation minimizes (3) only if the unit circle is centered at the origin. Right shift by one unit, and consensus at the origin becomes the new minimizer. To be coordinate agnostic, we choose (2) to arrive at

$$\begin{aligned} \dot{x}_i &= P_i \left(\sum_{j \in \mathcal{V}_u} a_{ij} (x_j - x_i) - \sum_{j \in \mathcal{V}_l} a_{ij} (x_j - x_i) \right), \quad i \in \mathcal{V}_u \\ \dot{x}_i &= P_i \left(\sum_{j \in \mathcal{V}_l} a_{ij} (x_j - x_i) - \sum_{j \in \mathcal{V}_u} a_{ij} (x_j - x_i) \right), \quad i \in \mathcal{V}_l. \end{aligned} \quad (4)$$

The sphere model with homogeneous frequency matrices [13, (2.3)] is a special case of (4).

For the situation with purely attractive coupling, we simply change the sign of the second term in (2):

$$V_+(\chi) = \frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}} a_{ij} \|x_j - x_i\|^2, \quad (5)$$

following which the state equation for agent i becomes

$$\dot{x}_i = (I - n_i n_i') \sum_{j \in \mathcal{V}} a_{ij} (x_j - x_i), \quad i \in \mathcal{V}. \quad (6)$$




Fig. 1. Three common fruits illustrating Definition 2.1: a lemon with a pair of pimples (left), an apple with a pair of dimples (middle), and a peach with one pimple and one dimple (right).

Or more compactly, (6) can be written as $\dot{\chi} = -PL \otimes I_{\chi}$, where $P = \text{diag}(P_i)$ and L is the familiar graph Laplacian. It is readily recognizable as a nonlinear higher-dimensional version of the Abel-son model [20], where the nonlinearity comes from the underlying nonlinear hypermanifold. Moreover, the sphere version of (6) is a higher-dimensional Kuramoto model with homogeneous frequencies [21, 22].

D. The assemblage

Assembling the aforementioned ingredients in §II-B and §II-C, we have a multi-agent gradient flow system with repulsive (4) or attractive (6) intergroup interactions evolving on a hypermanifold \mathcal{H}^n . The hypermanifold is equipped with a pair of dimples or pimples as illustrated in Fig. 1, each containing one of the two groups of agents \mathcal{V}_u and \mathcal{V}_l . We are interested in possible polarization arising in this setting as a result of the interplay between the graph couplings and the geometry of the underlying nonlinear space.

Definition 2.2 (polarization): The agents are said to be polarized if $x_i = x_j$ for all $\{i, j\} \in \mathcal{E}_+$ and $x_i \neq x_j$ for all $\{i, j\} \in \mathcal{E}_-$.

This definition is less in common with the notion of bipartite consensus over linear Euclidean space in [18] meaning all agents are equal up to signs, or with the strong/weak polarization [11] and complete polarization [13] over the hypersphere which are antipodal. Those definitions are convenient for highly symmetric spaces, whereas we work with more general hypermanifolds without symmetry and cannot define polarization this way. However, our definition enjoys more flexibility, as we can model opposition over seemingly trivial differences (from the external point of view) with Definition 2.2 plus $\|x_i - x_j\| < \epsilon$ for $\{i, j\} \in \mathcal{E}_-$. In this sense, the partial polarization defined in [13] is more closely related, where the two groups on the sphere need not be antipodal and is possible in their case due to the presence of frequency terms.

Definition 2.2 characterizes a polarized configuration without specifying whether the states are in equilibrium, limit-cycle, or more complex non-stationary modes. We focus on polarization as an equilibrium (set) and its stability properties, because equilibrium is the only possible attractor as shown by the following proposition. Let u_i denote the aggregate attraction for x_i

$$u_i = \sum_{j \in \mathcal{V}} a_{ij} (x_j - x_i). \quad (7)$$

Proposition 2.1: The system (6) converges to an equilibrium set on $(\mathcal{H}^n)^N$. At the equilibrium, u_i defined in (7) is either zero or parallel to n_i for all $i \in \mathcal{V}$.

Proof: For $V_+(\chi)$ in (5),

$$\begin{aligned} \dot{V}_+ &= \sum_{\{i,j\} \in \mathcal{E}} a_{ij} (\langle x_j - x_i, \dot{x}_j \rangle - \langle x_j - x_i, \dot{x}_i \rangle) \\ &= \sum_{i \in \mathcal{V}} \langle \dot{x}_i, \sum_{j \in \mathcal{V}} a_{ij} (x_i - x_j) \rangle + \sum_{j \in \mathcal{V}} \langle \dot{x}_j, \sum_{i \in \mathcal{V}} a_{ij} (x_j - x_i) \rangle \\ &= 2 \sum_{i \in \mathcal{V}} \langle \dot{x}_i, \sum_{j \in \mathcal{V}} a_{ij} (x_i - x_j) \rangle \\ &= -2 \sum_{i \in \mathcal{V}} \langle P_i u_i, u_i \rangle = -2 \sum_{i \in \mathcal{V}} \|P_i u_i\|^2, \end{aligned}$$

where the last equality comes from $P_i^2 = P_i$. Thus, we have $\dot{V}_+ \leq 0$. By LaSalle's invariance principle, system (6) converges to the set $\{\chi \mid \dot{V}_+ = 0\}$. To achieve $\dot{V}_+ = 0$, either $u_i = 0$ or $u_i \in \ker P_i$. In the latter case, as $\ker P_i = \text{span}(n_i)$, we have u_i parallel to n_i . Inspecting (6), we see that both cases lead to $\dot{x}_i = 0$ for all $i \in \mathcal{V}$. Therefore, the system (6) converges to an equilibrium set on $(\mathcal{H}^n)^N$. ■

As noted in §II-B, we may use signed graphs for intergroup antagonism, then (2) and (4) would respectively reduce in form to (5) and (6). And the same conclusion as Prop. 2.1 can be reached via identical arguments for system (4) derived from V_- .

Remark 2.3: Proposition 2.1 dictates convergence to a set, rather than to a point of equilibrium. Hence, it does not exclude non-stationary solutions. An example of a gradient system displaying non-trivial convergence behavior is the Mexican hat, which has trajectories converging to the unit circle [23, Introduction].

The first statement in Prop. 2.1 can be proved using more standard arguments for gradient flows on manifolds [24, App. C.12]. Nevertheless, our derivation, analogous to that of [25, Prop. 11], makes explicit the second statement which points to the necessary condition for polarization to occur. Namely, the normals at the two distinct points forming the polarized configuration must be aligned. This observation motivates the sufficient conditions for stable polarization that we are about to see in the next section.

We collect a few previous results and associated definitions that will pave the way for the development of our main results. The definitions of concepts such as stability and a local minimizer of a real function are well known. Here, we clarify their meanings when a set rather than a point is in question, which is perhaps less standard. Introduce the *Hausdorff distance* between two sets $\mathcal{Y}, \mathcal{Z} \subset \mathbb{R}^n$,

$$d_H(\mathcal{Y}, \mathcal{Z}) := \max\left\{\sup_{y \in \mathcal{Y}} \inf_{z \in \mathcal{Z}} \|y - z\|, \sup_{z \in \mathcal{Z}} \inf_{y \in \mathcal{Y}} \|y - z\|\right\}.$$

Definition 2.3 (stability): A set of equilibria \mathcal{S} is Lyapunov stable if, for each $\epsilon > 0$, there is $\delta = \delta(\epsilon)$ such that $d_H(x, \mathcal{S})|_{t=0} < \delta$

implies $d_H(x, \mathcal{S})|_t < \epsilon$ for all $t \geq 0$; is asymptotically stable if it is stable and δ can be chosen such that $d_H(x, \mathcal{S})|_{t=0} < \delta$ implies $\lim_{t \rightarrow \infty} d_H(x, \mathcal{S}) = 0$.

Definition 2.4 (local minimizer): A set $\mathcal{S} \subset \mathcal{M}$ is said to be a *local minimizer* of a real function $f: \mathcal{M} \rightarrow \mathbb{R}$ from a metric space (\mathcal{M}, d_H) if for some $\epsilon > 0$ there is an open neighborhood $\mathcal{N}(\mathcal{S}) = \{x \in \mathcal{M} | d_H(x, \mathcal{S}) < \epsilon\}$ such that $f|_{\mathcal{S}} \leq f(x)$ for all $x \in \mathcal{N}(\mathcal{S})$. Moreover, if the inequality is strict for all $x \in \mathcal{N}(\mathcal{S}) \setminus \mathcal{S}$, then \mathcal{S} is said to be a *strict local minimizer*.

Definition 2.5 (isolated critical): A set $\mathcal{S} \subset \mathcal{M}$ of critical points of a real function $f: \mathcal{M} \rightarrow \mathbb{R}$ from a metric space (\mathcal{M}, d_H) is said to be *isolated critical* if for some $\epsilon > 0$ there is an open neighborhood $\mathcal{N}(\mathcal{S}) = \{x \in \mathcal{M} | d_H(x, \mathcal{S}) < \epsilon\}$ such that $\mathcal{N}(\mathcal{S}) \setminus \mathcal{S}$ is void of critical points.

Proposition 2.2 (Prop. 3 [26]): Let \mathcal{M} be closed and take any $V: \mathcal{M} \rightarrow \mathbb{R}$ that is C^2 . Let \mathcal{S} be a compact set of local minimizers of V . If \mathcal{S} is a strict local minimizer, then \mathcal{S} is a Lyapunov stable equilibrium set of $\dot{x} = -\text{grad } V$. If \mathcal{S} is also isolated critical, then it is asymptotically stable.

III. MAIN RESULTS

In this section, we present and discuss our main results concerning the stability properties of polarized equilibria. They arise in different combinations of attractive/repulsive interactions with dimple/pimple geometric features, best exemplified by the fruits portrayed in Fig. 1. Despite the symmetrical shapes that the fruits assume in the figure, we emphasize that our general results do not require any spatial symmetry of the hypermanifold embedding.

The subsections are headed by various cases of geometric features. We are motivated to study these cases because, as pointed out in the abstract, nonlinear hypermanifolds are typically free of flatness; dimples and pimples are the norm. So although our results are only sufficient, the cases we study are representative of general polarizing conditions.

A. Pimple pairs

Consider a pair of pimples on the hypermanifold, one containing the group of agents \mathcal{V}_u and the other \mathcal{V}_l . The assumption is then

Assumption 1: The sets \mathcal{I}_u and \mathcal{I}_l are a pair of pimples, and $x_i \in \mathcal{I}_u$ for all $i \in \mathcal{V}_u$, $x_i \in \mathcal{I}_l$ for all $i \in \mathcal{V}_l$.

1) Lemon: To begin, we may picture the system (4) evolving on a lemon-like manifold in Fig. 1 (left). For the dynamics, we are interested in (4) with attractive intragroup coupling and repulsive intergroup coupling corresponding to the disagreement function (2).

Denote a *closed* ball centered at a point x with radius r as $\mathcal{B}_r(x)$. Let $x_o = \frac{1}{2}(x_u + x_l)$ denote the midpoint between x_u and x_l , and $r_o = \frac{1}{2}\|x_u - x_l\|$ the half distance between them. The following results concern the set

$$\mathcal{C}_{\text{lem}} := \{\chi \in (\mathcal{H}^n)^N \mid x_i = x \forall i \in \mathcal{V}_u, x_i = y \forall i \in \mathcal{V}_l, (x, y) \in Y\}, \quad (8)$$

where $Y := \{(x, y) \in \mathcal{I}_u \times \mathcal{I}_l \mid \|x - y\| = 2r_o\}$. This set is non-empty as $\mathcal{C}_{\text{lem}} \supset \chi^*$, where

$$\chi^* := \{\chi \in (\mathcal{H}^n)^N \mid x_i = x_u \forall i \in \mathcal{V}_u, x_i = x_l \forall i \in \mathcal{V}_l\}. \quad (9)$$

It may contain other elements when, for instance, the pair of pimples belongs to a sphere or a sphere-like hypermanifold that fully or partially overlaps with the boundary of the closed ball $\mathcal{B}_{r_o}(x_o)$.

Proposition 3.1: For system (4) under Assumption 1, if there exists a pair of distinct pimple bottoms $x_u \in \mathcal{I}_u$ and $x_l \in \mathcal{I}_l$ such that \mathcal{I}_u and \mathcal{I}_l are entirely contained in $\mathcal{B}_{r_o}(x_o)$, then a strict

local minimum of V_- is $V_-^* := -2r_o^2 \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij}$, and the corresponding strict local minimizer is a compact set of polarized configurations \mathcal{C}_{lem} defined in (8).

Proof: We shall minimize V_- , which is composed of disagreement terms $\|x_j - x_i\|$ among agents. Under Assumption 1, the bounds on the disagreement terms differ for agents belonging to the same or opposing groups. For all $\{i, j\} \in \mathcal{E}_-$, assume without loss of generality that $i \in \mathcal{V}_u$ and $j \in \mathcal{V}_l$. The term $\|x_j - x_i\|^2$ in (2) is upper bounded by

$$\|x_j - x_i\| \leq 2r_o = \|x_u - x_l\|.$$

The inequality is because both pimples are entirely contained in $\mathcal{B}_{r_o}(x_o)$. The upper bound is achieved when $x_i = x$ and $x_j = y$ for every pair $(x \in \mathcal{I}_u, y \in \mathcal{I}_l)$ such that $\|x - y\| = 2r_o$, an example of which is $x = x_u$ and $y = x_l$.

For all $\{i, j\} \in \mathcal{E}_+$,

$$\|x_j - x_i\|^2 \geq 0,$$

where the equality is achieved when $x_i = x_j$, of which a special case is $x_i = x_j = x_u$ for $i, j \in \mathcal{V}_u$ and $x_i = x_j = x_l$ for $i, j \in \mathcal{V}_l$. Therefore,

$$V_-(\chi) \geq -\frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij} \|x_j - x_i\|^2 \geq -2r_o^2 \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij}.$$

The minimum is achieved only when $x_i = x, i \in \mathcal{V}_u$ and $x_i = y, i \in \mathcal{V}_l$ for every pair of $(x \in \mathcal{I}_u, y \in \mathcal{I}_l)$ such that $\|x - y\| = 2r_o$. ■

Corollary 3.2: For system (4) under Assumption 1, if the two pimples \mathcal{I}_u and \mathcal{I}_l satisfy the conditions given in Prop. 3.1, then \mathcal{C}_{lem} defined in (8) is a Lyapunov stable set of polarized equilibria.

Proof: This is a direct application of Prop. 2.2 to Prop. 3.1. ■

Next, we derive an asymptotic stability result for when \mathcal{I}_u and \mathcal{I}_l live on nice manifolds.

Theorem 3.3: For system (4) under Assumption 1, if the two pimples \mathcal{I}_u and \mathcal{I}_l satisfy the conditions given in Prop. 3.1, and in addition, there is a neighborhood $\mathcal{N}_a(\mathcal{C}_{\text{lem}})$ on $(\mathcal{H}^n)^N$ that belongs to an analytic manifold, then \mathcal{C}_{lem} defined in (8) is an asymptotically stable set of polarized equilibria.

Proof: The first part of the proof is to show that pointwise, there are no equilibrium points other than those in \mathcal{C}_{lem} in the neighborhood of each $\chi_p \in \mathcal{C}_{\text{lem}}$. We label quantities by a subscript “p” to stress that they vary with each different $\chi_p \in \mathcal{C}_{\text{lem}}$. Following a variant [27, Sec. 9] of the Łojasiewicz inequality valid on analytic Riemannian manifolds, the analytic function $V_-(\chi)$ in (2) behaves in the following way in a neighborhood $\mathcal{N}_1(\chi_p) \subset \mathcal{N}_a(\mathcal{C}_{\text{lem}})$ of every polarized equilibrium $\chi_p \in \mathcal{C}_{\text{lem}}$:

$$|V_-(\chi) - V_-(\chi_p)|^{\alpha_p} \leq \kappa_p \|\text{grad } V_-(\chi)\|,$$

for $\alpha_p < 1$ and $\kappa_p > 0$. Suppose that $\chi \in \mathcal{N}_1(\chi_p)$ is an element in the set of equilibria \mathcal{Q} such that $\mathcal{Q} \cap \mathcal{C}_{\text{lem}} = \emptyset$, then $\text{grad } V_-(\chi) = 0$, c.f. (1). Consequently, $V_-(\chi) = V_-(\chi_p)$. However, Prop. 3.1 says that the local minimum V_-^* is achieved only if $\chi \in \mathcal{C}_{\text{lem}}$, a contradiction. Therefore, $\mathcal{Q} \cap \mathcal{N}_1(\chi_p) = \emptyset, \forall \chi_p \in \mathcal{C}_{\text{lem}}$.

Having shown that every point in \mathcal{C}_{lem} is isolated from \mathcal{Q} , we proceed to demonstrate that no sequence in \mathcal{Q} can approach \mathcal{C}_{lem} arbitrarily close. Suppose on the contrary that there is such a sequence $\{\chi_i\}_{i=1}^\infty \in \mathcal{Q}$, then $\inf_{\chi \in \mathcal{Q}} V_-(\chi) = V_-^*$ and $\lim_{i \rightarrow \infty} V_-(\chi_i) = V_-^*$. The sequence is bounded, since \mathcal{I}_u and \mathcal{I}_l are bounded sets. By the Bolzano-Weierstrass theorem, the sequence $\{\chi_i\}_{i=1}^\infty$ has a subsequence that converges to some point χ_q such that $V(\chi_q) = V_-^*$, which in turn implies that $\chi_q \in \mathcal{C}_{\text{lem}}$. This is to say that this subsequence in \mathcal{Q} must converge to a point $\chi_q \in \mathcal{C}_{\text{lem}}$, contradicting the fact that $\mathcal{Q} \cap \mathcal{N}_1(\chi_q) = \emptyset$.

Thus, we have shown that \mathcal{C}_{lem} is isolated critical, and therefore is asymptotically stable by Prop. 2.2. ■

Remark 3.1: The additional requirement on the analyticity of the manifold in Thm 3.3 is a local one. In fact, the sole purpose of introducing the neighborhood $\mathcal{N}_a(\bullet)$ is to emphasize this local nature. We do not require the whole manifold to be analytic for \mathcal{C}_{lem} or χ^* to be asymptotically stable. For instance, $\mathcal{N}_a(\mathcal{C}_{\text{lem}})$ may be a subset of $(\mathcal{H}^n)^N \cap \mathcal{M}^N$, where \mathcal{H}^n is the hypermanifold on which the agents inhabit, whereas \mathcal{M} is an analytic manifold.

The condition proposed in Prop. 3.1 essentially seeks to ensure that $2r_o$ is the largest possible distance between all possible pairs of points $\{x, y\} \in \mathcal{I}_u \times \mathcal{I}_l$ in the neighborhood, whereby V_-^* is a strict local minimizer of V_- . However, the conservatism introduced by this approach can be immediately identified. Even if one pimple, say \mathcal{I}_u , is outside $\mathcal{B}_{r_o}(x_o)$, if the other pimple \mathcal{I}_l is very “steep”, it may still be the case that $2r_o$ is the largest possible distance between all possible pairs of points $\{x, y\} \in \mathcal{I}_u \times \mathcal{I}_l$. Nonetheless, if both pimples are outside $\mathcal{B}_{r_o}(x_o)$, then instability can be established for χ^* . For brevity, we only show the loss of asymptotic stability in the following.

Proposition 3.4: For system (4) under Assumption 1, if the two pimples \mathcal{I}_u and \mathcal{I}_l are entirely outside $\mathcal{B}_{r_o}(x_o)$ except for the bottoms x_u and x_l , then the equilibrium χ^* defined in (9) is not asymptotically stable.

Proof: That χ^* is an equilibrium can be checked by substituting (9) into (4). To show instability, we find a perturbation that causes the trajectory not to settle into χ^* . We choose a perturbation $\tilde{\chi}$ that leaves the configuration polarized, that is $x_i = \tilde{x}_u \neq x_u$ for all $i \in \mathcal{V}_u$ and $x_i = \tilde{x}_l \neq x_l$ for all $i \in \mathcal{V}_l$, such that $\tilde{x}_u - \tilde{x}_l$ passes through x_o . As both pimples are outside $\mathcal{B}_{r_o}(x_o)$ except for the isolated points x_u and x_l , $\|\tilde{x}_u - \tilde{x}_l\|$ is guaranteed to be greater than $\|x_u - x_l\|$. Consequently, the disagreement function is upper-bounded:

$$\begin{aligned} V_-(\tilde{\chi}) &= -\frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij} \|\tilde{x}_u - \tilde{x}_l\| \\ &< -\frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij} \|x_u - x_l\| = V_-(\chi^*). \end{aligned}$$

Since the disagreement function V_- decreases along the solutions of $\dot{\chi} = -\text{grad } V_-(\chi)$ [28], which in this case is (4), the perturbed configuration $\tilde{\chi}$ does not return to χ^* . ■

2) Apple donut: Stable polarization in a pair of pimples also exists in the system (6) with attractive intergroup coupling, when the normals of the two pimples point toward each other, see the apple donut in Fig. 2.

Observe that the graph \mathcal{G} is connected and both \mathcal{V}_u and \mathcal{V}_l are non-empty, so is \mathcal{E}_- . Consequently, under polarization as per Def. 2.2, u_i defined in (7) is non-zero, and hence must be parallel to n_i according to Prop. 2.1. This observation motivates the conditions in the ensuing results.

Proposition 3.5: For system (6) under Assumption 1, if there exists a pair of distinct pimple bottoms $x_u \in \mathcal{I}_u$ and $x_l \in \mathcal{I}_l$ such that

- 1) $x_u - x_l$ is parallel to $n(x_u)$, and
- 2) $h_u(x_l)$ is a local minimum satisfying $h_u(x_l) > h_u(x_u)$,

then a strict local minimum of V_+ is $V_+^* := 2r_o^2 \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij}$, and the corresponding strict local minimizer is χ^* defined in (9).

Proof: For all $\{i, j\} \in \mathcal{E}_-$, assume without loss of generality that $i \in \mathcal{V}_u$ and $j \in \mathcal{V}_l$. The term $\|x_j - x_i\|^2$ in (5) is lower bounded by

$$\begin{aligned} \|x_j - x_i\|^2 &\geq \langle x_j - x_i, n(x_u) \rangle^2 = (h_u(x_j) - h_u(x_i))^2 \\ &\geq (h_u(x_u) - h_u(x_l))^2 = 4r_o^2. \end{aligned} \quad (10)$$

The first lower bound is achieved only when $x_j - x_i$ is parallel to $n(x_u)$. For the second lower bound, observe that x_u is the bottom of the pimple \mathcal{I}_u . Consequently, $h_u(x_u)$ is by Definition 2.1 a strict local maximum achieved only by x_u . Combining this observation with condition 2 in the proposition, we conclude that the second lower bound is achieved only when $h_u(x_i) = h_u(x_u)$, implying $x_i = x_u$, and $h_u(x_j) = h_u(x_l)$. With condition 1 in the proposition, the only configuration to achieve both lower bounds is thus x_i and x_j being in their respective bottoms x_u and x_l .

For all $\{i, j\} \in \mathcal{E}_+$, we again have $\|x_j - x_i\|^2 \geq 0$, with the equality achieved when $x_i = x_j$. Therefore,

$$V_+(\chi) \geq \frac{1}{2} \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij} \|x_j - x_i\|^2 \geq 2r_o^2 \sum_{\{i,j\} \in \mathcal{E}_-} a_{ij}.$$

The minimum is achieved only when $x_i = x_u, i \in \mathcal{V}_u$ and $x_i = x_l, i \in \mathcal{V}_l$. ■

Corollary 3.6: For system (6) under Assumption 1, if the two pimples \mathcal{I}_u and \mathcal{I}_l satisfy the properties given in Prop. 3.5, then χ^* defined in (9) is a Lyapunov stable polarized equilibrium.

Proof: This is a direct application of Prop. 2.2 to Prop. 3.5. ■

Theorem 3.7: For system (6) under Assumption 1, if the two pimples \mathcal{I}_u and \mathcal{I}_l satisfy the properties given in Prop. 3.5, and in addition, there is a neighborhood $\mathcal{N}_a(\chi^*)$ on $(\mathcal{H}^n)^N$ that belongs to an analytic manifold, then χ^* defined in (9) is an asymptotically stable polarized equilibrium.

Theorem 3.7 concerns the asymptotic stability of an equilibrium point as the singleton set χ^* , rather than an equilibrium set as in (8). The line of reasoning is consequently simplified to the first paragraph of the proof of Theorem 3.3 and is thus omitted.

B. Sphere: a special case of the lemon

In the special case of $\mathcal{H}^n = \mathcal{S}^n$ when the hypermanifold is the n -sphere, for every $x \in \mathcal{S}^n$ and its opposite pole $y = -x$, \mathcal{I}_x and \mathcal{I}_y form a pair of pimples satisfying the conditions in Theorem 3.3. Therefore, we have a corollary for the following subset of \mathcal{C}_{lem} , which is a polarization set specialized on the n -sphere

$$\begin{aligned} \mathcal{C}_{\text{sph}} &:= \{\chi \in (\mathcal{S}^n)^N \mid \\ &\quad x_i = x \ \forall i \in \mathcal{V}_u, x_i = y \ \forall i \in \mathcal{V}_l, \|x - y\| = 2\} \\ &= \{\chi \in (\mathcal{S}^n)^N \mid \\ &\quad x_i = x \ \forall i \in \mathcal{V}_u, x_i = -x \ \forall i \in \mathcal{V}_l, x \in \mathcal{S}^n\}. \end{aligned} \quad (11)$$

Corollary 3.8: For system (4) evolving on \mathcal{S}^n , \mathcal{C}_{sph} given by (11) constitutes an asymptotically stable set of polarized equilibria.

Furthermore, for n -spheres excepting the circle \mathcal{S}^1 , a stronger result of almost global asymptotic stability can be obtained by exploiting the spherical symmetry. By almost global asymptotic stability, we mean

Definition 3.1 (almost global asymptotic stability): A set of equilibria $\mathcal{D} \subset (\mathcal{H}^n)^N$ is almost globally attractive if for all initial conditions except a measure-zero set, it holds that $\lim_{t \rightarrow \infty} \chi(t) \in \mathcal{D}$. Moreover, if \mathcal{D} is stable, \mathcal{D} is almost globally asymptotically stable.

The measure-zero set is with respect to the Lebesgue measure.

Theorem 3.9: For system (4) evolving on \mathcal{S}^n with $n \geq 2$, the polarization set \mathcal{C}_{sph} given by (11) is almost globally asymptotically stable. Moreover, every trajectory converges to some point in \mathcal{C}_{sph} at a locally exponential rate.

Proof: There is no loss of generality in assuming that the origin is located at the n -sphere center, as (2) is coordinate agnostic. Apply a coordinate transformation $y_i = x_i$ for all $i \in \mathcal{V}_u$ and $y_i = -x_i$ for all $i \in \mathcal{V}_l$. System (4) becomes

$$\dot{y}_i = (I - y_i y_i') \sum_{j \in \mathcal{V}} a_{ij} y_j, \quad \forall i \in \mathcal{V}, \quad (12)$$




Fig. 2. A lemon donut with a pair of dimples (left), an apple donut with a pair of pimples (middle), and a peach donut with one dimple and one pimple (right).

where we have used the facts that $n_i = x_i$ and $(I - x_i x_i')x_i = 0$ that are valid on the unit n -spheres. System (12) is a special case (constant a_{ij}) of a family of consensus protocols considered in [25, Thm. 13], which guarantees the almost global asymptotic stability of the consensus set

$$\mathcal{C} = \{(y_i)_{i=1}^N \in (\mathcal{S}^n)^N \mid y_i = y_j, \forall i, j \in \mathcal{V}\}.$$

This consensus set \mathcal{C} maps to the polarization set \mathcal{C}_{sph} by reversing the bijective coordinate transform. Therefore, applying [25, Thm. 13] to system (12) and then reversing the coordinate transformation, we obtain the first conclusion. For the second statement, notice that the adjacency matrix A is constant and nonnegative. Then, [25, Thm. 13] again applies. ■

Remark 3.2: Theorem 3.9 together with Corollary 3.8 complements and generalizes several existing results in the literature. For instance, [13, Thm. 3.4] provides regions of attraction with exponential stability under a condition on relative strengths of attractive and repulsive gains, assuming all-to-all networks. The almost global asymptotic stability result for polarization in [15, Thm. 4.5] is only established on the 2-sphere, and is restricted to cycle graphs with even number of agents. In contrast, we are able to conclude without extra conditions that the polarized equilibria set is almost globally asymptotically stable on \mathcal{S}^n with $n \geq 2$, and locally asymptotically stable on \mathcal{S}^n for all $n \in \mathbb{N}$.

C. Dimple pairs

There is an apparent symmetry between a pair of pimples and a pair of dimples. Indeed, through identical arguments we can obtain mirroring conditions for stable polarization within a pair of dimples. In this subsection, we operate under the following assumption:

Assumption 2: The sets \mathcal{I}_u and \mathcal{I}_l are a pair of dimples, and $x_i \in \mathcal{I}_u$ for all $i \in \mathcal{V}_u$, $x_i \in \mathcal{I}_l$ for all $i \in \mathcal{V}_l$.

If the manifold resembles the apple in Fig. 1, we have something akin to Theorem 3.7 for the apple donut.

Proposition 3.10: For system (6) under Assumption 2, if there exists a pair of distinct dimple bottoms $x_u \in \mathcal{I}_u$ and $x_l \in \mathcal{I}_l$ such that $x_u - x_l$ is parallel to $n(x_u)$, and $h_u(x_l)$ is a local maximum satisfying $h_u(x_l) < h_u(x_u)$, then χ^* defined in (9) is Lyapunov stable. Furthermore, if there is a neighborhood $\mathcal{N}_a(\chi^*)$ on $(\mathcal{H}^n)^N$ that belongs to an analytic manifold, then χ^* is asymptotically stable.

Proof: See proofs of Cor. 3.6 and Thm. 3.7. ■

On the other hand, if the manifold resembles the lemon donut in Fig. 2, we have a result analogous to Thm. 3.3.

Proposition 3.11: For system (4) under Assumption 2, if there exists a pair of distinct dimple bottoms $x_u \in \mathcal{I}_u$ and $x_l \in \mathcal{I}_l$ such that \mathcal{I}_u and \mathcal{I}_l are entirely contained in $\mathcal{B}_{r_o}(x_o)$, then \mathcal{C}_{lem} defined

in (8) is Lyapunov stable. Furthermore, if there is a neighborhood $\mathcal{N}_a(\mathcal{C}_{\text{lem}})$ on $(\mathcal{H}^n)^N$ that belongs to an analytic manifold, then \mathcal{C}_{lem} is asymptotically stable.

Proof: Identical to the proofs of Cor. 3.2 and Thm. 3.3. ■

By now it is clear that the various cases we have classified in §III-A and §III-C hinge on aligned normals that point either toward or away from each other. Orientability is thus essential. But this is implied by the hypermanifold being compact (see §II-A).

Remark 3.3: The asymptotic stability results in Thms. 3.3 and 3.7 (Props. 3.10 and 3.11) rely on an analyticity assumption of the manifold. Although a smooth manifold may be topologically equivalent to an analytic one, it is worth stressing that our results concern manifolds with special geometric features, see Definition 2.1, and are thus not amenable to such topological equivalence.

IV. NUMERICAL EXAMPLES

We illustrate our main results of §III with simple numerical examples, and explore other possible routes to polarization not covered by the main results.

A. Illustrations of the main results

We demonstrate the stable polarization processes on lemon, sphere, and apple manifolds. There are $N = 7$ agents divided into two groups $\mathcal{V}_u = \{1, 2, 3\}$ and $\mathcal{V}_l = \{4, 5, 6, 7\}$. They are connected with uniform weights over all edges by a network A_1 portrayed in Fig. 3.




Fig. 3. Symmetric (A_1) and asymmetric (A_2) network topologies.

The agents update their beliefs according to the rules specified by (4) or (6). They are randomly initialized in a pair of pimples or dimples. Note that one of the agents in \mathcal{V}_l are initialized very close to the final state on all the fruit-like manifolds. In the case of the sphere, they are randomly initialized on the entire 2-sphere, since every pair of antipodal points can serve as the pimple bottoms, and the radii ϵ of their neighborhoods may be very large up to 2. These specifications conform to the conditions in Theorem 3.3 and Proposition 3.10, and the trajectories depicted in Fig. 4 converge to polarized equilibria as expected. The accompanying convergence rates are shown in Fig. 5 in log scale. The predominantly linear trends




Fig. 4. Two groups of agents with network A_1 form stable polarization on various manifolds. From left to right: system (4) on a lemon, (4) on a sphere, (6) on an apple, and mixed intergroup coupling on a peach. Initial and final states correspond to small and large dots on the trajectories. Some parts outside the geometric features are intentionally left blank in part to recall the fact that they do not play a role in the local stability results, and in part to highlight the trajectories.




Fig. 5. Rates of convergence to the polarized equilibrium x_p in log scale for the cases in Fig. 4: lemon (yellow solid), sphere (yellow dotted), apple (green solid), and peach (pink solid).

are evidence of exponential convergence, and in the sphere case is also locally guaranteed by Theorem 3.9.

The example on the lemon in Fig. 4 visually demonstrate that our polarization model, is able to capture the phenomenon of radicalization, where opinions become more extreme. Many opinion models in the literature are contractive even when explicit mechanisms to foster social cleavage is present (antagonistic interactions [18], bounded confidence [29], and stubborn agents [30]), unless nonlinear factors are incorporated such as biased assimilation [31] and saturation [32].

The example on the apple in Fig. 4 suggest that even when two parties try to reach a common ground by making compromises, thereby getting closer to each other in the Euclidean space, the unbridgeable core belief landscape presents an obstruction to the eventual concord. Especially interesting is when the Euclidean distance between the polarized points are small as seen from the extrinsic point of view, it may represent an irreconcilable long lasting hostility between two groups with minor differences.

For the peach with a pimple and a dimple in Fig. 1, we are unable to find a polarization condition for the gradient flow dynamics of either (4) or (6); we only provide the example in Fig. 4 (right) to show the possibility. In Fig. 4 (right), the agents in \mathcal{V}_u are repulsed by those in \mathcal{V}_l , and the agents in \mathcal{V}_l are attracted to those in \mathcal{V}_u . Such dynamics can be derived from the following disagreement functions

$$V_{\mp}(\chi) = \begin{cases} V_{-}(\chi) & i \in \mathcal{V}_u \\ V_{+}(\chi) & i \in \mathcal{V}_l. \end{cases}$$

Since asymmetric intergroup interactions are not generated by a single gradient flow, the polarized equilibrium we numerically found is not covered by our theoretical approach. In fact, such systems do not necessarily converge to an equilibrium set, (not just because the sum-of-squares in the proof of Prop. 2.1 is lost) as we have found simulated counterexamples on the sphere with network A_2 and

on a peach having a “flatter” pimple with a bipartite network. The non-convergent peach case and the convergent case in Fig. 4 (right) suggest that depending on the flatness of the geometric features, one of the components in V_{\mp} may play a dominant role. Such a characterization of the geometric features goes beyond the simple height-function-based definitions of dimples and pimples.

B. Weakly anomalous intergroup coupling

It is conceivable that for system (4) on a lemon, introducing weakly attractive intergroup coupling to a small number of edges in \mathcal{E}_{-} does not necessarily compromise polarization or its stability properties. In turn, the tolerance for such deviation from purely repulsive intergroup coupling may be viewed as a measure of robustness. Along this line and based on the example in §IV-A, we switch the sign of interaction and change the weight on each edge in \mathcal{E}_{-} successively one at a time, to see to what degree can polarization be maintained given different manifolds and initial conditions. Each experiment is allowed sufficient time to run for agents starting at different initial conditions to evolve. The final configuration is considered polarized if the order parameters $\|\rho_u\|$ and $\|\rho_l\|$ are equal to 1.000 within the final integration time 200, where $\rho_u = (x_1 + \dots + x_M)/(M\|x_u\|)$ and similarly for ρ_l . Complete phase synchronization also results in $\|\rho_u\| = \|\rho_l\| = 1$, but is ruled out by visual inspection.

Table I summarizes the findings. The maximal weights for each edge in \mathcal{E}_{-} when changed to attractive coupling are the same for 7 different initial conditions. The tests are run on the lemon and the sphere with graph structures A_1 and A_2 in Fig. 3.

TABLE I
EDGE ROBUSTNESS

network	manifold	a_{24}	a_{27}	a_{34}	a_{35}
A_1	lemon	0.83	0.099	0.83	0.099
	sphere	0.41	0.23	0.41	0.23
A_2	lemon	0.2	0.94	0.9	N/A
	sphere	0.56	0.82	0.35	N/A

For A_1 , $a_{24} = a_{34}$ and $a_{27} = a_{35}$ because the network topology has a symmetric structure. The initial conditions appear to have no effect on the attainable weights. This may be because no new equilibrium points are created within the dimples by the change of sign on the link, and the only option is to converge to the known equilibrium point, irrespective of the initial conditions. We also observe that the edge $\{2, 4\}$ is more robust to treason than $\{2, 7\}$.

The asymmetric network A_2 is obtained by switching the link $\{6, 7\}$ to $\{4, 7\}$, leaving the group membership unchanged. The N/A

values mean that the system fail to completely polarize no matter how small the weight is, indicating zero tolerance. The $\{3, 5\}$ link is especially sensitive because agent 6 is completely dependent on agent 5, and is thus swept along once agent 5 switches allegiance. It is difficult to quantitatively characterize the actual robustness measure, as the outcome heavily depends on the geometry of the underlying manifold, as well as the network topology.

V. CONCLUSIONS

We have established sufficient conditions for asymptotic stability of polarized equilibria of arbitrarily connected multi-agent gradient flow systems over nonlinear spaces. These sufficient conditions are tailored to four scenarios likely to be found on nonflat hypermanifolds, arising from the combination of dimple/pimple pairs and attractive/repulsive intergroup couplings. In particular, the hypersphere as a special case of general manifolds provides effortless generalization of previously known results. Its highly symmetric nature further allows us to prove almost global asymptotic stability except for the unit circle. To strengthen the proposed opinion dynamics interpretation of the low dimensional core belief space embedded in the high dimensional external space, a natural next step is to work with manifolds of dimensions much lower than the ambient space, whereby the normal space dimension exceeds 1. Although this can be partially addressed by the introduction of products of hypermanifolds, where the combined dimension of all normal spaces can be arbitrarily large.

In light of the stark contrast between the abundance of opinion dynamics models and the lack of experimental validation thereof [2], we sketch a simplistic way how validation of the model that we studied may be done in practice. For a given subject consisting of multiple components, collect an initial sample of vector valued opinions from different individuals by surveys to form a point cloud. The graph structure and signs of interaction may be obtained through the surveys or inferred from social network interactions. The point cloud is assumed to be sampled from an underlying low dimensional manifold, which can be recovered through widely available manifold embedding techniques. The inference of both the communication network and the manifold can be constrained to conform to the structural balance requirement and Assumptions 1 or 2, respectively. Collect a subsequent sample of opinion point cloud as the ground truth. If the learned manifold satisfies the conditions of either of our main results, one compares the outcome of the distributed gradient flow system with the ground truth. This is of course one of many possible ways to carry out the validation.

ACKNOWLEDGMENT

The authors would like to thank Daniele Proverbio for many helpful suggestions for improving the manuscript.

REFERENCES

- [1] R. Tron, B. Afsari, and R. Vidal, "Riemannian consensus for manifolds with bounded curvature," *IEEE Transactions on Automatic Control*, vol. 58, no. 4, pp. 921–934, 2012.
- [2] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks. Part I," *Annual Reviews in Control*, vol. 43, pp. 65–79, 2017.
- [3] A. V. Proskurnikov and R. Tempo, "A tutorial on modeling and analysis of dynamic social networks. Part II," *Annual Reviews in Control*, vol. 45, pp. 166–190, 2018.
- [4] M. Cao, A. Stewart, and N. E. Leonard, "Convergence in human decision-making dynamics," *Systems & Control Letters*, vol. 59, no. 2, pp. 87–97, 2010.
- [5] A. Mirtabatabaei and F. Bullo, "Opinion dynamics in heterogeneous networks: Convergence conjectures and theorems," *SIAM Journal on Control and Optimization*, vol. 50, no. 5, pp. 2763–2785, 2012.
- [6] W. S. Rossi and P. Frasca, "Opinion dynamics with topological gossiping: Asynchronous updates under limited attention," *IEEE Control Systems Letters*, vol. 4, no. 3, pp. 566–571, 2020.
- [7] M. Gentzkow, J. M. Shapiro, and M. Taddy, "Measuring group differences in high-dimensional choices: method and application to congressional speech," *Econometrica*, vol. 87, no. 4, pp. 1307–1340, 2019.
- [8] E. Gibbon, *The decline and fall of the Roman Empire*, vol. 2. Encyclopaedia Britannica, 1989.
- [9] R. Sepulchre, "Consensus on nonlinear spaces," *Annual reviews in control*, vol. 35, no. 1, pp. 56–64, 2011.
- [10] M. Caponigro, A. C. Lai, and B. Piccoli, "A nonlinear model of opinion formation on the sphere," *Discrete & Continuous Dynamical Systems*, vol. 35, no. 9, p. 4241, 2015.
- [11] J. Gaitonde, J. Kleinberg, and E. Tardos, "Polarization in geometric opinion dynamics," in *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 499–519, Association for Computing Machinery, 2021.
- [12] H. Hong and S. H. Strogatz, "Kuramoto model of coupled oscillators with positive and negative coupling parameters: An example of conformist and contrarian oscillators," *Physical Review Letters*, vol. 106, no. 5, p. 054102, 2011.
- [13] S.-Y. Ha, D. Kim, J. Lee, and S. E. Noh, "Emergence of bicluster aggregation for the swarm sphere model with attractive-repulsive couplings," *SIAM Journal on Applied Dynamical Systems*, vol. 19, no. 2, pp. 1225–1270, 2020.
- [14] A. Crnkčić and V. Jaćimović, "Swarms on the 3-sphere with adaptive synapses: Hebbian and anti-Hebbian learning rule," *Systems & Control Letters*, vol. 122, pp. 32–38, 2018.
- [15] W. Song, J. Markdahl, S. Zhang, X. Hu, and Y. Hong, "Intrinsic reduced attitude formation with ring inter-agent graph," *Automatica*, vol. 85, pp. 193–201, 2017.
- [16] A. Aydogdu, S. McQuade, and N. Duteil, "Opinion dynamics on a general compact Riemannian manifold," *Networks and Heterogeneous Media*, vol. 12, no. 3, pp. 489–523, 2017.
- [17] E. L. Lima, "The Jordan-Brouwer separation theorem for smooth hypersurfaces," *The American Mathematical Monthly*, vol. 95, no. 1, pp. 39–42, 1988.
- [18] C. Altafini, "Consensus problems on networks with antagonistic interactions," *IEEE Transactions on Automatic Control*, vol. 58, no. 4, pp. 935–946, 2012.
- [19] J. Markdahl, "Consensus seeking gradient descent flows on boundaries of convex sets," in *2020 American Control Conference (ACC)*, pp. 830–835, 2020.
- [20] R. P. Abelson, "Mathematical models in social psychology," in *Advances in Experimental Social Psychology*, vol. 3, pp. 1–54, Elsevier, 1967.
- [21] F. Ritort, "Solvable dynamics in a system of interacting random tops," *Physical review letters*, vol. 80, no. 1, p. 6, 1998.
- [22] J. Markdahl, D. Proverbio, L. Mi, and J. Gonçalves, "Almost global convergence to practical synchronization in the generalized Kuramoto model on networks over the n -sphere," *Communications Physics*, vol. 4, no. 1, pp. 1–9, 2021.
- [23] C. Lageman, *Convergence of gradient-like dynamical systems and optimization algorithms*. PhD thesis, Universität Würzburg, 2007.
- [24] U. Helmke and J. B. Moore, *Optimization and dynamical systems*. Springer, 1996.
- [25] J. Markdahl, J. Thunberg, and J. Gonçalves, "Almost global consensus on the n -sphere," *IEEE Transactions on Automatic Control*, vol. 63, no. 6, pp. 1664–1675, 2018.
- [26] J. Markdahl, "Synchronization on Riemannian manifolds: Multiply connected implies multistable," *IEEE Transactions on Automatic Control*, vol. 66, no. 9, pp. 4311–4318, 2021.
- [27] K. Kurdyka, T. Mostowski, and A. Parusinski, "Proof of the gradient conjecture of R. Thom," *Annals of Mathematics*, pp. 763–792, 2000.
- [28] J. Jost, *Riemannian geometry and geometric analysis*. Springer, 2008.
- [29] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *Journal of artificial societies and social simulation*, vol. 5, no. 3, 2002.
- [30] S. E. Parsegov, A. V. Proskurnikov, R. Tempo, and N. E. Friedkin, "Novel multidimensional models of opinion dynamics in social networks," *IEEE Transactions on Automatic Control*, vol. 62, no. 5, pp. 2270–2285, 2016.
- [31] L. Wang, Y. Hong, G. Shi, and C. Altafini, "Signed social networks with biased assimilation," *IEEE Transactions on Automatic Control*, 2021.
- [32] A. Bizyayeva, A. Franci, and N. E. Leonard, "Nonlinear opinion dynamics with tunable sensitivity," *IEEE Transactions on Automatic Control*, 2022.