NOTES ON GRADIENT FLOW AND CONSERVED QUANTITIES

IORDAN GANEV

Contents

These are notes stemming from and supplementing the paper [ZGW⁺23]. For simplicity, we assume throughout that all functions are smooth; in many cases this is not necessary. A prerequisite for these notes is familiarity with flows and integral curves (c.f. [Lee12, Chapter 9]). See the appendix for a refresher and for the notation that we will use.

1. Basics on conserved quantities

1.1. **Formulations.** We consider conserved quantities for a function on an open subset of a vector space. This notion can be formalized in several equivalent ways.

Let $U \subseteq \mathbb{R}^p$ be an open subset of the vector space $\mathbb{R}^p$. Let $\mathcal{L} : U \to \mathbb{R}$ be a smooth function. A smooth function $Q : U \to \mathbb{R}$ is a **conserved quantity** with respect to $\mathcal{L}$ if it satisfies any of the equivalent conditions stated in the following lemma:

Lemma 1.1. *Let $Q : U \to \mathbb{R}$ be a smooth function. The following are equivalent:*

(1) *The function Q is constant on any integral curve for the gradient vector field of $\mathcal{L}$*

(2) *Let $\gamma : (a,b) \to U$ be an integral curve for the gradient vector field of $\mathcal{L}$. Then $d(Q \circ \gamma)_t = 0$ for all $t \in (a,b)$.*

(3) *For any $v \in U$, we have $\langle \nabla_v Q, \nabla_v \mathcal{L} \rangle = 0$.*

(4) *Let $\theta : \mathcal{D} \to U$ be a flow for the gradient vector field of $\mathcal{L}$, where $\mathcal{D} \subseteq \mathbb{R} \times U$ is a flow domain. The following diagram commutes:*

$$\begin{array}{ccc} \mathcal{D} & \xrightarrow{\theta} & U \\ {\scriptstyle \text{proj}}\downarrow & & \downarrow{\scriptstyle Q} \\ U & \xrightarrow{Q} & \mathbb{R} \end{array}$$

In other words, $Q(\theta(t,v)) = Q(v)$ for every $t \in \mathbb{R}$ and $v \in U$, so it is invariant for the partial action of $\mathbb{R}$ on U.

Note that we invoke the Euclidean inner product[1] in the third statement.

Sketch of proof. The implications $(1) \Leftrightarrow (2)$ are immediate from definitions. To show the equivalence of the third and second statements, let $v \in U$ and let $\gamma : (a,b) \to U$ be an integral curve for the gradient vector field such that $0 \in (a,b)$ and $\gamma(0) = v$ (such a curve always exists). Then, using the chain rule, we have:

$$\langle \nabla_v Q, \nabla_v \mathcal{L} \rangle = dQ_{\gamma(0)} \circ \nabla_v \mathcal{L} = dQ_{\gamma(0)} \circ d\gamma_0 = d(Q \circ \gamma)_0$$

Finally, the implications $(1) \Leftrightarrow (4)$ follow from the fact that, for a flow θ, the parameterized curve $t \mapsto \theta(t,v)$ is an integral curve for the gradient vector field of $\mathcal{L}$. $\square$

Remark 1.2. Characterizations (1), (2), and (4) immediately generalize to conserved quantities valued in any smooth manifold, that is, $Q : U \to M$. Characterization (3) also generalizes if we replace the condition by $dQ_v(\nabla_v \mathcal{L}) = 0$. (Note that $dQ_v : T_v U \to T_{Q(v)} M$ and $\nabla_v \mathcal{L} \in T_v U$.) One can generalize the entire lemma further by replacing U with a Riemannian manifold, using the Riemannian metric to define the gradient.

1.2. **Limit points.** The following lemma demonstrates that the value of any conserved quantity along an integral curve is determined by its value at the limit of the curve (if it exists).

Lemma 1.3. *Suppose $\gamma : \mathbb{R} \to U$ is an integral curve for the gradient vector field of $\mathcal{L}$, and suppose $v = \lim_{t\to\infty} \gamma(t)$. For any conserved quantity Q, we have that:*

$$Q(\gamma(t)) = Q(v) \qquad \forall t \in \mathbb{R}$$

The result follows from definitions and the continuity of Q. Note that an integral curve cannot always be defined on all of $\mathbb{R}$, and, even if it can, the limit $\lim_{t\to\infty} \gamma(t)$ may not exist. If the limit does exist, the gradient of $\mathcal{L}$ vanishes there, that is, the limit point is a critical point for $\mathcal{L}$.

Conversely, suppose $\mathcal{L}$ has a critical point at $v \in U$. Let S be the subset of points in $\mathbb{R}^p$ that limit to this point under gradient flow. (Equivalently, S is the union of all integral curves whose limit is v.) Then any conserved quantity must be constant on S. We explore an example of this in the next section.

[1] On a Riemannian manifold, one invokes the Riemannian 2-tensor.

1.3. **Example.** Let $\mathcal{L} : V = \mathbb{R}^p \to \mathbb{R}$ be the function taking $v \in V$ to half the square of its L_2 norm: $\mathcal{L}(v) = \frac{1}{2}|v|^2$. A global flow for the reverse gradient vector field of $\mathcal{L}$ is given by:

$$\theta : \mathbb{R} \times V \to V; \qquad (t, v) \mapsto e^{-t}v$$

To check this, we note that $\nabla_v \mathcal{L} = v$ for every $v \in V$, and so:

$$\left(\theta^{(v)}\right)'(t) = -e^{-t}v = -\nabla_{e^{-t}v}\mathcal{L} = -\nabla_{\theta(t,v)}\mathcal{L}$$

The formula for the flow implies that all integral curves have the origin as their limit point. It follows from Lemma 1.3 that any conserved quantity is determined by its value at the origin, and so the only conserved quantities are the constant functions on the whole of $V = \mathbb{R}^p$.

On the other hand, let us restrict our attention to the submanifold $U = \mathbb{R}^p \setminus \{0\}$ of $V = \mathbb{R}^p$. In this case, no integral curve has a limit point within U. Each flow line within U intersects the unit sphere S^{p-1} at a unique point; indeed, for $v \neq 0$, we have that:

$$|\theta(t, v)| = 1 \quad \Leftrightarrow \quad t = \ln(|v|) \quad \Leftrightarrow \quad \theta(t, v) = \frac{v}{|v|}.$$

In other words, flowing from $v \in U$ for time $\ln(|v|)$ takes us to a point on the unit sphere. Consequently, we can identify conserved quantities on $U = \mathbb{R}^p \setminus \{0\}$ with functions on the sphere S^{p-1}. That is:

$$\{\text{conserved quantities on } U = \mathbb{R}^p \setminus \{0\}\} \longleftrightarrow \{\text{functions on } S^{p-1}\}$$

Slightly more explicitly, any function $f : S^{p-1} \to \mathbb{R}$ determines a conserved quantity on $\mathbb{R}^p \setminus \{0\}$ via:

$$Q_f : U \to \mathbb{R}, \qquad v \mapsto f\left(\frac{v}{|v|}\right),$$

and all conserved quantities arise in this way.

This example also exhibits symmetry, namely, the function $\mathcal{L}$ is invariant for the action of the group $G = O(p)$ of orthogonal transformations of $V = \mathbb{R}^p$. Furthermore, the group action commutes with the flow in the following sense:

$$\theta(t, g \cdot v) = e^{-t}(g \cdot v) = g \cdot \left(e^{-t}v\right) = g \cdot \theta(t, v)$$

for any $t \in \mathbb{R}$, $v \in V$, and $g \in O(p)$ (in fact, for any $g \in \mathrm{GL}_p(\mathbb{R})$). Thus, we have an action of $O(p) \times \mathbb{R}$ on $V = \mathbb{R}^p$. The orbit of a point $v \in V$ is defined as:

$$\mathcal{O}_v = \left\{g \cdot (e^{-t}v) \;\;:\;\; g \in O(p),\ t \in \mathbb{R}\right\}$$

It is straightforward to verify that:

$$\mathcal{O}_v = \begin{cases} U = \mathbb{R}^p \setminus \{0\} & \text{if } v \neq 0 \\ \{0\} & \text{if } v = 0 \end{cases}$$

Thus, all nonzero points belong to the same orbit under the action of $O(p) \times \mathbb{R}$.

We make the following cursory observations. They are not immediately relevant in this simple example, but useful when connecting to the more general setting.

- The action of $O(p)$ on $V = \mathbb{R}^p$ restricts to a transitive action on S^{p-1}, i.e. there is a single $O(p)$-orbit on S^{p-1}. The stabilizer of $v_0 = (1, 0, \dots, 0)$ can be identified with the orthogonal group $O(p-1)$ (embedded in $O(p)$ in the lower right). We conclude that there is a diffeomorphism
$$O(p)/O(p-1) \xrightarrow{\sim} S^{p-1}.$$
- The flow θ restricts to a diffeomorphism:
$$\alpha : S^{p-1} \times \mathbb{R} \longrightarrow U = \mathbb{R}^p \setminus \{0\}, \qquad (v, t) \mapsto e^{-t} v$$
(This is a restatement of the fact that an integral curve intersects the unit circle in a unique point.) This diffeomorphism is equivariant for the action of $O(p) \times \mathbb{R}$, where $O(p)$ acts transitively on S^{p-1} and $\mathbb{R}$ acts on itself by addition.
- The composition
$$\mathcal{L} \circ \alpha : S^{p-1} \times \mathbb{R} \longrightarrow \mathbb{R}$$
is simply the projection map onto the second factor.
- By Lemma 1.1, $Q : U \to \mathbb{R}$ is a conserved quantity if and only if it is invariant for the action of $\mathbb{R}$. The equivariance of ϕ implies that this happens if and only if the composition $Q \circ \alpha : S^{p-1} \times \mathbb{R} \longrightarrow \mathbb{R}$ is invariant for the action of $\mathbb{R}$, which in turn holds if and only if $Q \circ \alpha$ only depends on the first factor, that is, there is a function $f : S^{p-1} \to \mathbb{R}$ making the following diagram commute:
$$\begin{array}{ccc} S^{p-1} \times \mathbb{R} & \xrightarrow{Q \circ \alpha} & \mathbb{R} \\ {\scriptstyle \mathrm{proj}_1} \downarrow & \nearrow_{f} & \\ S^{p-1} & & \end{array}$$
where the left vertical is the projection map $(v, t) \mapsto v$.

1.4. **Example.** Consider the function $\mathcal{L} : v \mapsto |v|$. The gradient at $v \neq 0$ is $\frac{v}{|v|}$, and the function is not smooth at 0. A maximal flow for the reverse gradient vector field appears to be given by:
$$\theta : \mathcal{D} \to V, \qquad (t, v) \mapsto \begin{cases} \left(1 - \frac{t}{|v|}\right) v & \text{if } v \neq 0 \\ 0 & \text{if } v = 0 \end{cases}$$
where $\mathcal{D} = \{(t, v) \in \mathbb{R} \times \mathbb{R}^p \setminus \{0\} \ : \ t < |v|\} \cup (\mathbb{R} \times \{0\})$. This flow does not extend to a global flow.

To verify that θ is a flow for the reverse gradient vector field, note that, for $v \neq 0$, we have: $\left(\theta^{(v)}\right)'(t) = \frac{-v}{|v|} = \frac{-\theta(t,v)}{|\theta(t,v)|} = -\nabla_{\theta(t,v)} \mathcal{L}$. Note also that $\theta(0, v) = v$, and

$\theta(s,\theta(t,v)) = \theta(s+t,v)$. The last identity follows from the computation:

$$\theta(s,\theta(t,v)) = \left(1 - \frac{s}{|\theta(t,v)|}\right)\theta(t,v) = \theta(t,v) - \frac{sv}{|v|} = v - \frac{(s+t)v}{|v|}$$

As in the previous example, each flow line intersects the unit sphere S^{p-1} at a unique point; indeed, for $v \neq 0$, we have that:

$$|\theta(t,v)| = 1 \quad \Leftrightarrow \quad t = |v| - 1 \quad \Leftrightarrow \quad \theta(t,v) = \frac{v}{|v|}.$$

Hence, conserved quantities on $U = \mathbb{R}^p \setminus \{0\}$ are given by functions on the unit sphere S^{p-1}. Also, the function $\mathcal{L}$ is invariant for the action of the orthogonal group, and this action commutes with gradient flow:

$$\theta(t, g\cdot v) = g\cdot v - t\frac{g\cdot v}{|g\cdot v|} = g\cdot\left(v - \frac{tv}{|v|}\right) = g\cdot\theta(t,v)$$

2. Invariant group actions

Let $\mathcal{L} : V = \mathbb{R}^p \to \mathbb{R}$ be a smooth function. The differential $d\mathcal{L}_v$ of $\mathcal{L}$ at $v \in V$ is row vector, while the gradient $\nabla_v\mathcal{L}$ of $\mathcal{L}$ at v is a column vector[2]:

$$d\mathcal{L}_v = \begin{bmatrix} \left.\frac{\partial\mathcal{L}}{\partial x_1}\right|_v & \cdots & \left.\frac{\partial\mathcal{L}}{\partial x_p}\right|_v \end{bmatrix} \qquad\qquad \nabla_v\mathcal{L} = \begin{bmatrix} \left.\frac{\partial\mathcal{L}}{\partial x_1}\right|_v \\ \vdots \\ \left.\frac{\partial\mathcal{L}}{\partial x_p}\right|_v \end{bmatrix}$$

Hence $\nabla_v\mathcal{L}$ is the transpose of $d\mathcal{L}_v$, that is: $\nabla_v\mathcal{L} = (d\mathcal{L}_v)^T$.

2.1. First approach toward invariance.

Suppose $\rho : G \to \mathrm{GL}(V)$ is an action of a Lie group G on V such that $\mathcal{L}$ is G-invariant. We write simply $g\cdot v$ for $\rho(g)(v)$, and g for $\rho(g)$.

Lemma 2.1. *For any $v \in V$ and $g \in G$, we have:*

$$\nabla_v\mathcal{L} = g^T \cdot (\nabla_{g\cdot v}\mathcal{L})$$

Proof. The proof is a computation:

$$\begin{aligned}\nabla_v\mathcal{L} &= (d_v\mathcal{L})^T = (d(\mathcal{L}\circ g)_v)^T = (d\mathcal{L}_{g\cdot v}\circ dg_v)^T = (d\mathcal{L}_{g\cdot v}\circ g)^T = g^T\cdot(d\mathcal{L}_{g\cdot v})^T \\ &= g^T\cdot(\nabla\mathcal{L}_{g\cdot v})\end{aligned}$$

The second equality relies on the hypothesis that $\mathcal{L}\circ g = \mathcal{L}$, the third on the chain rule, and the fourth on the fact that $dg_v = g$ since g is a linear map. □

[2]Following the usual conventions, we regard column vectors as elements of V and row vectors as elements of the dual vector space V^*.

One can perform the same computation in coordinates:

$$(\nabla_v \mathcal{L})_i = (d\mathcal{L}_v)^i = \left.\frac{\partial \mathcal{L}}{\partial x_i}\right|_v = \left.\frac{\partial (\mathcal{L} \circ g)}{\partial x_i}\right|_v = \left.\frac{\partial \mathcal{L}}{\partial x_j}\right|_{gv} \left.\frac{\partial g_j}{\partial x_i}\right|_v$$
$$= \left(\nabla_{gv}\mathcal{L}\right)_j g^i_j = (g^T)^j_i \left(\nabla_{gv}\mathcal{L}\right)_j = \left(g^T \cdot \nabla_{gv}\mathcal{L}\right)_i$$

2.2. **Alternative approach.** Suppose $\rho : G \to \mathrm{GL}(V)$ is an action of a Lie group G on V. Then G acts on the vector space of smooth functions $C^\infty(V)$ on V by precomposition:

$$(g \triangleright F)(v) = F(g^{-1} \cdot v)$$

for $g \in G$, $F \in C^\infty(V)$, and $v \in V$. Using the same methods as above, one can show that:

$$g^T \cdot \left(\nabla_{g\cdot v}\left(g \triangleright F\right)\right) = \nabla_v F$$

If F is G-invariant, we recover Lemma 2.1 above.

2.3. **Examples.** One can easily verify the result of Lemma 2.1 in the case of the loss functions $v \mapsto \frac{1}{2}|v|^2$ and $v \mapsto |v|$, both of which are invariant with respect to the orthogonal group $O(p)$.

3. Orthogonal transformations and conserved quantities

We now discuss loss functions invariant for an orthogonal group action.

3.1. **Commuting actions.** Let $\mathcal{L} : V = \mathbb{R}^p \to \mathbb{R}$ be a smooth function. Suppose a global flow exists for the reverse gradient vector field of $\mathcal{L}$. Thus, there is an action of the additive group of $\mathbb{R}$ on V:

$$\theta : \mathbb{R} \times V \to V$$

such that $\frac{d}{dt}[t \mapsto \theta(t, v)]_0 = -\nabla_v \mathcal{L}$ for all $v \in V$ (this equality takes place in the tangent space $T_v V$ of v in V). We expect this discussion to generalize to the case where a global flow does not exist.

Suppose $\rho : G \to O(V)$ is an action of a Lie group G on V by orthogonal transformations such that $\mathcal{L}$ is G-invariant. We write simply $g \cdot v$ for $\rho(g)(v)$.

Lemma 3.1. *For all $t \in \mathbb{R}$, $g \in G$, and $v \in V$, we have:*

$$\theta(t, g \cdot v) = g \cdot \theta(t, v)$$

Sketch of proof. The fact that g is an orthogonal transformation implies that $g^T = g^{-1}$, and hence the result of Lemma 2.1 becomes $g \cdot \nabla_v \mathcal{L} = \nabla_{g\cdot v}\mathcal{L}$. Moreover, $t \mapsto g \cdot \theta(t, v)$ is an integral curve starting at $g \cdot v$. Since $t \mapsto \theta(t, g \cdot v)$ is the unique maximal such integral curve, the result follows. □

Corollary 3.2. *There is a well-defined action of $G \times \mathbb{R}$ on V given by:*

$$\begin{aligned} a : G \times \mathbb{R} \times V &\longrightarrow V \\ (g, t, v) &\longmapsto g \cdot \theta(t, v) \end{aligned}$$

3.2. **Conserved quantities on orbits.** For $v \in V$, let $G_v = \{g \in G \mid g \cdot v = v\}$ be the stabilizer of v in G, which is a closed subgroup of G. Furthermore, let $\mathcal{O}_v$ be the $G \times \mathbb{R}$-orbit of v, that is, the image under a of $G \times \mathbb{R} \times \{v\}$. Then $\mathcal{O}_v$ is a smooth manifold (though its closure may not be smooth).

Proposition 3.3. *For $v \in V$, there is a surjective map:*

$$\phi : \mathcal{O}_v \longrightarrow G/G_v.$$

A smooth function $Q : \mathcal{O}_v \to \mathbb{R}$ is a conserved quantity for the restriction of $\mathcal{L}$ to $\mathcal{O}_v$ if and only if $Q = f \circ \phi$ for a smooth function $f : G/G_v \to \mathbb{R}$.

In other words, there is a natural identification of conserved quantities on the orbit $\mathcal{O}_v$ with functions on the manifold G/G_v.

$$\left\{ \begin{matrix} \text{conserved quantities} \\ \text{on the orbit } \mathcal{O}_v \end{matrix} \right\} \longleftrightarrow \left\{ \begin{matrix} \text{functions on} \\ G/G_v \end{matrix} \right\}.$$

Sketch of proof of Proposition 3.3. We begin by considering the stabilizer $H_v = \mathrm{Stab}_{G \times \mathbb{R}}(v)$ of v in $G \times \mathbb{R}$. From basic group theory, there is a diffeomorphism $\alpha_v : (G \times \mathbb{R})/H_v \to \mathcal{O}_v$ making the following diagram commute:

$$\begin{array}{ccc} G \times \mathbb{R} & \xrightarrow{\quad a_v \quad} & V \\ \downarrow & & \uparrow \\ (G \times \mathbb{R})/H_v & \xrightarrow{\quad \alpha_v \quad} & \mathcal{O}_v \end{array}$$

where the top horizontal map is $(g, t) \mapsto a(g, t, v)$, the left vertical map is the quotient map, and the right vertical map is the inclusion. Note that α_v is $G \times \mathbb{R}$-equivariant.

We consider first the case $\nabla_v \mathcal{L} = 0$. Then $\theta(t, v) = v$ for all $t \in \mathbb{R}$, and so $(g, t) \in G \times \mathbb{R}$ stabilizes v if and only if $g \in G_v$. In other words, $H_v = G_v \times \mathbb{R}$ and so $(G \times \mathbb{R})/H_v = G/G_v$, which is the G-orbit of v. We set $\phi = \alpha_v^{-1}$. Since there is no flow from $g \cdot v$ for any $g \in G$, any function on $\mathcal{O}_v \simeq G/G_v$ is a conserved quantity.

Henceforth assume $\nabla_v \mathcal{L} \neq 0$. We claim that, in this case, $H_v = G_v \times \{0\}$. To see this, first observe that the fact that $v \mapsto \theta(t, v)$ is a diffeomorphism of V for any $t \in \mathbb{R}$ implies that $\nabla_{\theta(t,v)} \mathcal{L} \neq 0$ for all $t \in \mathbb{R}$. Since the flow θ generates the reverse gradient vector field of $\mathcal{L}$, for every $t > 0$, we have that $\mathcal{L}(\theta(t, v)) < \mathcal{L}(v)$. Consequently, using the fact that $\mathcal{L}$ is G-invariant, we also have:

$$\mathcal{L}(a(g, v, t)) = \mathcal{L}(g \cdot \theta(t, v)) = \mathcal{L}(\theta(t, v)) < \mathcal{L}(v) \qquad \text{for } t > 0$$

Similarly, $\mathcal{L}(a(g,v,t)) > \mathcal{L}(v)$ for $t < 0$. It follows that $(g,t) \in G \times \mathbb{R}$ stabilizes v if and only if $t = 0$ and $g \in G_v$.

We are now ready to define $\phi : \mathcal{O}_v \to G/G_v$. Since $(G \times \mathbb{R})/H_v = G/G_v \times \mathbb{R}$, we set $\phi = \text{proj}_1 \circ \alpha_v^{-1}$. That is, ϕ is the following composition:

$$\phi : \mathcal{O}_v \xrightarrow{\alpha_v^{-1}} (G \times \mathbb{R})/H_v = G/G_v \times \mathbb{R} \xrightarrow{\text{proj}_1} G/G_v$$

where proj_1 is the projection onto the first coordinate.

We now turn our attention to conserved quantities. Suppose $f : G/G_v \to \mathbb{R}$ is a smooth function and set $Q = f \circ \phi$. To show that Q is a conserved quantity, is suffices to show that Q is $\mathbb{R}$-invariant. Let $w \in \mathcal{O}_v$ be arbitrary, so that $w = \alpha_v([g],t)$ for some $g \in G$ and $t \in \mathbb{R}$, where $[g]$ is the image of g in G/G_v. It is straightforward to verify that $Q(w) = f([g])$. On the other hand, for any $s \in \mathbb{R}$, we have:

$$\theta(s,w) = g \cdot \theta(s+t,v) = a_v(g,s+t) = \alpha_v([g],s+t)$$

Therefore,

$$Q(\theta(s,w)) = Q(\alpha_v([g],s+t)) = f \circ \text{proj}_1 \circ \alpha_v^{-1} \circ \alpha_v([g],s+t) = f([g]) = Q(w).$$

We conclude that $Q = f \circ \phi$ is a conserved quantity.

Finally, suppose $Q : \mathcal{O}_v \to \mathbb{R}$ is a conserved quantity. Then Q is $\mathbb{R}$-invariant, and so the composition

$$G/G_v \times \mathbb{R} \xrightarrow{\alpha_v} \mathcal{O}_v \xrightarrow{Q} \mathbb{R}$$

is $\mathbb{R}$-invariant, since α_v is $\mathbb{R}$-equivariant. Hence it factors through the projection map $\text{proj}_1 : G/G_v \times \mathbb{R} \to G/G_v$. In other words, there is a function $f : G/G_v \to \mathbb{R}$ such that $Q \circ \alpha_v = f \circ \text{proj}_1$. The result follows. □

3.3. **Example: orthogonal change-of-basis.** We explore the example of orthogonal change-of-basis symmetries on the parameter space of an MLP. In the case of radial activations, such symmetries leave the feedforward function fixed. The results in this section are exploratory and not particularly conclusive.

Suppose an MLP has L layers and widths given by the tuple $\mathbf{n} = (n_0, n_1, n_2, \ldots, n_{L-1}, n_L)$. In other words, the width of the i-th layer is n_i, with n_0 and n_L being the widths of the input and output spaces. The parameter space of such an MLP is defined as the vector space of all possible choices of trainable parameters. Hence, it is given by a a product of matrix spaces:

$$\mathsf{Param}(\mathbf{n}) = \left(\mathbb{R}^{n_1 \times n_0} \times \mathbb{R}^{n_2 \times n_1} \times \cdots \times \mathbb{R}^{n_L \times n_{L-1}}\right) \times \left(\mathbb{R}^{n_1} \times \mathbb{R}^{n_2} \times \cdots \times \mathbb{R}^{n_L}\right)$$

We denote an element therein as a pair of tuples $(\mathbf{W}, \mathbf{b})$ where $\mathbf{W} = (W_i \in \mathbb{R}^{n_i \times n_{i-1}})_{i=1}^L$ are the weights and $\mathbf{b} = (b_i \in \mathbb{R}^{n_i})_{i=1}^L$ are the biases. To describe certain symmetries of the parameter space, consider the following product of orthogonal groups, with sizes corresponding to the widths of the hidden layers:

$$O(\mathbf{n}^{\text{hid}}) = \mathcal{O}(n_1) \times O(n_2) \times \cdots \times O(n_{L-1})$$

There is an action of $O(\mathbf{n}^{\text{hid}})$ on $\mathsf{Param}(\mathbf{n})$ by change-of-basis. Explicitly, the element $\mathbf{Q} = (Q_i)_{i=1}^{L-1} \in O(\mathbf{n}^{\text{hid}})$ transforms $(\mathbf{W}, \mathbf{b}) \in \mathsf{Param}(\mathbf{n})$ as:

$$\mathbf{W} \mapsto \mathbf{Q} \cdot \mathbf{W} := \left(Q_i \circ W_i \circ Q_{i-1}^{-1} \right)_{i=1}^{L}, \qquad \mathbf{b} \mapsto \mathbf{Q} \cdot \mathbf{b} := (Q_i \circ b_i)_{i=1}^{L}, \tag{3.1}$$

where $Q_0 = \text{id}_{n_0}$ and $Q_L = \text{id}_{n_L}$. We write $\mathbf{Q} \cdot (\mathbf{W}, \mathbf{b})$ for $(\mathbf{Q} \cdot \mathbf{W}, \mathbf{Q} \cdot \mathbf{b})$.

The reduction of a widths vector $\mathbf{n}$ is a tuple $\mathbf{n}^{\text{red}} = (n_0^{\text{red}}, n_1^{\text{red}}, \ldots, n_L^{\text{red}})$ defined recursively by setting $n_0^{\text{red}} = n_0$, then $n_i^{\text{red}} = \min(n_i\ ,\ n_{i-1}^{\text{red}} + 1)$ for $i = 1, \ldots, L-1$, and finally $n_L^{\text{red}} = n_L$.

Proposition 3.4. *Let $(\mathbf{W}, \mathbf{b})$ be a point in the parameter space such that each $n_i \times (1 + n_{i-1})$ matrix $[b_i\ W_i]$ has full rank. The stabilizer in $O(\mathbf{n}^{\text{hid}})$ of this point is given by:*

$$O(\mathbf{n}^{\text{hid}})_{(\mathbf{W},\mathbf{b})} = O(n_1 - n_1^{\text{red}}) \times O(n_2 - n_2^{\text{red}}) \times \cdots \times O(n_{L-1} - n_{L-1}^{\text{red}})$$

In particular, if $\mathbf{n} = \mathbf{n}^{\text{red}}$, then the stabilizer of $(\mathbf{W}, \mathbf{b})$ is trivial.

To clarify, for each i, we use the following embedding of $O(n_i - n_i^{\text{red}})$ into $O(n_i)$:

$$O(n_i - n_i^{\text{red}}) \hookrightarrow O(n_i), \qquad P \mapsto \begin{bmatrix} \text{id}_{n_i^{\text{red}}} & 0 \\ 0 & P \end{bmatrix}$$

The zeros abbreviate zero matrices of the appropriate size. Additionally, if $n_i = n_i^{\text{red}}$, then $O(n_i - n_i^{\text{red}})$ is the trivial group.

Proof Idea for Proposition 3.4. Using a generalization of a QR decomposition, one can reduce to the case where the $n_i \times (1 + n_{i-1})$ matrix $[b_i\ W_i]$ has the following form:

$$[b_i\ W_i] = \begin{bmatrix} T_i^{(1)} & T_i^{(2)} \\ 0 & T_i^{(4)} \end{bmatrix}$$

where $T_i^{(1)}$ is a full-rank matrix of size $n_i^{\text{red}} \times (1 + n_{i-1}^{\text{red}})$, $T_i^{(2)}$ is a matrix of size $n_i^{\text{red}} \times (n_{i-1} - n_{i-1}^{\text{red}})$ and $T_i^{(4)}$ is a matrix of size $(n_i - n_i^{\text{red}}) \times (n_{i-1} - n_{i-1}^{\text{red}})$. The definition of n_i^{red} implies that either $T_i^{(1)}$ is square and invertible, or $T_i^{(1)}$ is surjective. In both cases, there exists a matrix S_i of size $(1 + n_{i-1}^{\text{red}}) \times n_i^{\text{red}}$ such that $T_i^{(1)} \circ S_i = \text{id}_{n_i^{\text{red}}}$.

Let $\mathbf{Q} = (Q_1, \ldots, Q_{L-1})$ belong to the stabilizer in $O(\mathbf{n}^{\text{hid}})$ of $(\mathbf{W}, \mathbf{b})$. We proceed by induction on i to show that Q_i belongs to the image of $O(n_i - n_i^{\text{red}})$ for all $i = 1, \ldots, L-1$. Write

$$Q_i = \begin{bmatrix} Q_i^{(1)} & Q_i^{(2)} \\ Q_i^{(3)} & Q_i^{(4)} \end{bmatrix}$$

where $Q_i^{(1)}$ is of size $n_i^{\text{red}} \times n_i^{\text{red}}$, $Q_i^{(4)}$ is of size $(n_i - n_i^{\text{red}}) \times (n_i - n_i^{\text{red}})$, and so forth. The base step can technically be taken to be $i = 0$. For the induction step, take $i > 0$ and suppose $Q_{i-1} \in O(n_{i-1} - n_{i-1}^{\text{red}})$, so that $Q_{i-1} = \begin{bmatrix} \text{id}_{n_{i-1}^{\text{red}}} & 0 \\ 0 & P_{i-1} \end{bmatrix}$ for a matrix $P_{i-1} \in O(n_{i-1} - n_{i-1}^{\text{red}})$. Then, writing $\mathbf{T}$ instead of $(\mathbf{W}, \mathbf{b})$, we have:

$$
\begin{aligned}
(\mathbf{Q} \cdot \mathbf{T})_i = \mathbf{T}_i \Leftrightarrow Q_i \circ T_i \circ \begin{bmatrix} 1 & 0 \\ 0 & Q_{i-1}^{-1} \end{bmatrix} = T_i & \\
\Leftrightarrow \begin{bmatrix} Q_i^{(1)} & Q_i^{(2)} \\ Q_i^{(3)} & Q_i^{(4)} \end{bmatrix} \begin{bmatrix} T_i^{(1)} & T_i^{(2)} \\ 0 & T_i^{(4)} \end{bmatrix} &= \begin{bmatrix} T_i^{(1)} & T_i^{(2)} \\ 0 & T_i^{(4)} \end{bmatrix} \begin{bmatrix} \text{id}_{1+n_{i-1}^{\text{red}}} & 0 \\ 0 & P_{i-1} \end{bmatrix} \\
\Rightarrow \begin{cases} Q_i^{(1)} \circ T_i^{(1)} = T_i^{(1)} \\ Q_i^{(3)} \circ T_i^{(1)} = 0 \end{cases} \quad \Rightarrow \begin{cases} Q_i^{(1)} = \text{id}_{n_i^{\text{red}}} \\ Q_i^{(3)} = 0 \end{cases} &
\end{aligned}
$$

In the computation above, we use the induction hypothesis, and the matrix S_i such that $T_i^{(1)} \circ S_i = \text{id}_{n_i^{\text{red}}}$. The orthogonality of Q_i now implies that $Q_i^{(2)} = 0$ and $Q_i^{(4)} \in O(n_i - n_i^{\text{red}})$. Thus $Q = \begin{bmatrix} \text{id}_{n_i^{\text{red}}} & 0 \\ 0 & Q_i^{(4)} \end{bmatrix}$. □

We discuss some consequences of Proposition 3.4. Let $G = O(\mathbf{n}^{\text{hid}})$ and let $(\mathbf{W}, \mathbf{b})$ be full-rank in the sense that $[b_i \; W_i]$ is a full rank matrix for all i. Furthermore, take a widths vector $\mathbf{n}$ with $\mathbf{n} = \mathbf{n}^{\text{red}}$. Then conserved quantities on the $G \times \mathbb{R}$-orbit of $(\mathbf{W}, \mathbf{b})$ are given by functions on G. However, this identification is not straightforward, one reason being that the '$\mathbf{Q}$' of the generalized QR decomposition changes with gradient flow. In the examples of $x \mapsto |x|$ and $x \mapsto \frac{1}{2}|x|^2$, every flow passed through the unit sphere in a unique point. There does not seem such a straightforward criterion in the present setting.

4. The infinitesimal action

This section provides a different but equivalent perspective on the results of Section 5.1 of the paper.

4.1. **The infinitesimal action.**

Let G be a Lie group with Lie algebra $\mathfrak{g} = \text{Lie}(G)$, identified with the tangent space at the identity 1 of G. Suppose $G \times V \to V$ is an action of G on $V = \mathbb{R}^p$. For now, this action is not necessarily linear. For $v \in V$, set:

$$a_v : G \to V; \qquad g \mapsto g^{-1} \cdot v$$

Definition 4.1. The infinitesimal action of G on V is the following map of Lie algebras:

$$\mathfrak{g} \to \Gamma(V, TV); \qquad U \mapsto \overline{U} = [v \mapsto d(a_v)_1(U)]$$

Hence the value $\overline{U}_v$ of the vector field $\overline{U}$ at $v \in V$ is given by first computing the differential of the map $g \mapsto g^{-1}v$ at the identity of G, and then evaluating at $U \in \mathfrak{g}$.

4.2. **Symmetries of the loss function.** Suppose $\mathcal{L} : V = \mathbb{R}^p \to \mathbb{R}$ is a differentiable, G-invariant function, so that $\mathcal{L}(g \cdot v) = \mathcal{L}(v)$ for all $v \in V$ and $g \in G$.

Lemma 4.2. *Fix $v \in V$. The infinitesimal action of any $U \in \mathfrak{g}$ at v is orthogonal to the gradient of $\mathcal{L}$ at v:*

$$\langle \overline{U}_v, \nabla_v \mathcal{L} \rangle = 0$$

Proof. We compute:

$$\langle \overline{U}_v, \nabla_v \mathcal{L} \rangle = d\mathcal{L}_v \left(\overline{U}_v \right) = d\mathcal{L}_v \circ d(a_v)_1(U) = d(\mathcal{L} \circ a_v)_1(U)$$

The fact that $\mathcal{L}$ is G-invariant implies that the composition $\mathcal{L} \circ a_v : G \to \mathbb{R}$ factors through the constant map $g \mapsto v$, so its differential is constantly zero. $\square$

Two relevant commutative diagrams are displayed below. The left encodes the fact that $\mathcal{L}$ is G-invariant. The right is obtained from the left when restricting to $G \times \{v\}$ (and possibly adding an inverse).

$$\begin{array}{ccc} G \times V & \xrightarrow{\text{action}} & V \\ {\scriptstyle\text{project}}\downarrow & & \downarrow{\scriptstyle\mathcal{L}} \\ V & \xrightarrow[\mathcal{L}]{} & \mathbb{R} \end{array} \qquad\qquad \begin{array}{ccc} G & \xrightarrow{a_v} & V \\ {\scriptstyle\text{const}}\downarrow & & \downarrow{\scriptstyle\mathcal{L}} \\ \{v\} & \xrightarrow[\mathcal{L}]{} & \mathbb{R} \end{array}$$

Note that Lemma 4.2 generalizes to Riemannian manifolds. Namely, if M is a Riemannian manifold with an action of G, then the gradient vector field of an invariant function is pointwise orthogonal to the infinitesimal action of $\mathfrak{g}$.

4.3. **Linear actions.** Now suppose the action of G on the vector space $V = \mathbb{R}^p$ is linear, i.e., V is a representation of G. Hence we have a map of groups $\pi : G \to \mathrm{GL}(V)$. Differentiating at the identity, we obtain a representation

$$d\pi_1 : \mathfrak{g} \longrightarrow \mathbb{R}^{p \times p}$$

of the Lie algebra of G, where we use the identifications $\mathfrak{g} = T_1 G$, and $T_{\mathrm{id}_d}\mathrm{GL}(V) = \mathbb{R}^{p \times p}$. We write simply $d\pi$ for $d\pi_1$. This representation is related to the infinitesimal action as follows:

Lemma 4.3. *Fix $v \in V$. For any $U \in \mathfrak{g}$, we have $\overline{U}_v = -d\pi(U)(v)$.*

Sketch of proof. The diagram below on the left commutes by the definition of a_v, while the right one is obtained from the left one by computing the differential at $1 \in G$:

$$\begin{array}{ccccc} G & \xrightarrow{\pi} & \mathrm{GL}(V) & \hookrightarrow & \mathbb{R}^{p\times p} \\ {\scriptstyle\mathrm{inv}}\uparrow & & & & \downarrow{\scriptstyle\mathrm{ev}_v} \\ G & & \xrightarrow{a_v} & & V \end{array} \qquad \begin{array}{ccc} \mathfrak{g} & \xrightarrow{d\pi} & \mathbb{R}^{p\times p} \\ {\scriptstyle -1}\uparrow & & \downarrow{\scriptstyle\mathrm{ev}_v} \\ \mathfrak{g} & \xrightarrow{d(a_v)_1} & V \end{array}$$

Here, $\mathrm{inv} : G \to G$ is the map $g \mapsto g^{-1}$ and ev_v is the evaluation map at v, which is linear. It follows that $\overline{U}_v = d(a_v)_1(U) = -\mathrm{ev}_v(d\pi(U)) = -d\pi(U)(v)$. □

4.4. **Conserved quantities and the infinitesimal action.** For $U \in \mathfrak{g}$, abbreviate $d\pi(U)$ by $\hat{U}$.

Lemma 4.4. *Fix $U \in \mathfrak{g}$. There exists a differentiable function $Q : V \to \mathbb{R}$ such that*

$$\nabla_v Q = \overline{U}_v$$

for all $v \in V$ if and only if $d\pi(U) \in \mathbb{R}^{p\times p}$ is a symmetric matrix.

Proof. Abbreviate the matrix $d\pi(U)$ by $\hat{U}$. Let $x_1, \ldots, x_p$ be the standard coordinates on $V = \mathbb{R}^p$, so that $\overline{U} = \left(-\hat{U}^i_j x_i\right)_{j=1}^{p}$. First suppose such a Q exists, so that $\frac{\partial Q}{\partial x_j} = -\hat{U}^i_j x_i$. Then the i, j entry of the Hessian of Q is:

$$\mathrm{Hess}(Q)^i_j = \frac{\partial Q}{\partial x_i \partial x_j} = \frac{\partial}{\partial x_i}\left(-\hat{U}^\ell_j x_\ell\right) = -\hat{U}^i_j$$

Thus $\mathrm{Hess}(Q) = -\hat{U}$. Since the Hessian is symmetric, it follows that $\hat{U}$ is as well. Conversely, suppose $\hat{U}$ is symmetric. Take Q to be the quadratic form defined by the matrix $-\hat{U}$, that is,

$$Q(v) = -\frac{1}{2} v^T \hat{U} v.$$

In coordinates, $Q = -\frac{1}{2}\sum_{i,j} \hat{U}^i_j x_i x_j$. Using the fact that $\hat{U}$ is symmetric, one verifies that $\frac{\partial Q}{\partial x_j} = -\hat{U}^i_j x_i$, for each $j = 1, \ldots, p$. □

Note that any real symmetric matrix can be diagonalized via orthogonal matrices.

4.5. **General case.** Suppose the transpose map $\mathbb{R}^{p\times p} \to \mathbb{R}^{p\times p}$ taking a matrix A to its transpose A^T lifts to an involution $\tau : \mathfrak{g} \to \mathfrak{g}$, so that $(\hat{U})^T = \widehat{\tau(U)}$ and $\tau^2 = \mathrm{id}_{\mathfrak{g}}$.

Lemma 4.5. *The function $Q_U(v) = -\frac{1}{2} v^T \hat{U} v$ is a conserved quantity for $\mathcal{L}$.*

Proof. The gradient of Q_U is given by:

$$\nabla_v(Q_U) = \left(\frac{\hat{U} + \hat{U}^t}{2}\right) v = \frac{\overline{U}_v + \overline{\tau(U)}_v}{2}$$

By Lemmas 1.1 and 4.2, we see that Q_U is a conserved quantity. □

We can write $Q_U = Q_{U^{\text{sym}}} + Q_{U^{\text{ant}}}$ where $U^{\text{sym}} = \frac{U+\tau(U)}{2}$ is the symmetric part of U and $U^{\text{ant}} = \frac{U-\tau(U)}{2}$ is the antisymmetric part of U. We see that

$$Q_U = Q_{\tau(U)} = Q_{U^{\text{sym}}} \qquad \text{and} \qquad \nabla_v (Q_U) = \nabla_v (Q_{U^{\text{sym}}}) = \overline{U^{\text{sym}}}_v$$

for all $v \in V$. Additionally, $Q_{U^{\text{ant}}} = 0$, so Q_U is depends only on the symmetric part of U. In particular, if $\tau(U) = -\tau(U)$, so that $\hat{U}$ is antisymmetric, then $Q_U = 0$.

We can also regard the above lemma as providing conserved quantities on $\mathbb{R}^p$ valued in the dual of the Lie algebra $\mathfrak{g}^*$ rather than in $\mathbb{R}$:

$$Q : \mathbb{R}^p \to \mathfrak{g}^*, \qquad v \mapsto [U \mapsto Q_U(v)]$$

5. Example: a diagonalizable action

Consider the function $\mathcal{L} : \mathbb{R}^2 \to \mathbb{R}$ given by $\mathcal{L}(x, y) = xy$. The reverse gradient vector field is:

$$\nabla_{(x,y)}\mathcal{L} = -y\frac{\partial}{\partial x} - x\frac{\partial}{\partial y}$$

A flow is given by:

$$\begin{aligned} &\theta : \mathbb{R} \times \mathbb{R}^2 \to \mathbb{R}^2 \\ &\quad (t, x, y) \mapsto (x \cosh(t) - y \sinh(t)\ ,\ y \cosh(t) - x \sinh(t)) \end{aligned}$$

where $\cosh(t) = \frac{e^t+e^{-t}}{2}$ and $\sinh(t) = \frac{e^t-e^{-t}}{2}$. The verification that θ is a flow relies on the fact that $\cosh'(t) = \sinh(t)$ and $\sinh'(t) = \cosh(t)$. We have an explicit formula for the value of $\mathcal{L}$ along a flow line:

$$\mathcal{L}(\theta(t, x, y)) = xy \cosh(2t) - (x^2 + y^2) \sinh(2t).$$

The function $\mathcal{L}$ is invariant for the action of the multiplicative group $G = \mathbb{R}^\times = \mathbb{R} \setminus \{0\}$ by:

$$\begin{aligned} \rho : \mathbb{R}^\times &\to \mathrm{GL}_2(\mathbb{R}) \\ \lambda &\mapsto \begin{bmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{bmatrix} \end{aligned}$$

So $\lambda \cdot (x, y) = (\lambda x\ ,\ \lambda^{-1} y)$. One can immediately see that the flow does not commute with the action of G. where $z = (x, y) \in \mathbb{R}^2$. A conserved quantity in this case is given by:

$$Q : \mathbb{R}^2 \to \mathbb{R}, \qquad (x, y) \mapsto x^2 - y^2$$

The gradient vector field of Q is $2x\frac{\partial}{\partial x} + 2y\frac{\partial}{\partial y}$. It is also easy to verify directly that $Q(\theta(t, x, y)) = Q(x, y)$.

Appendix A. Flows and integral curves

A.1. **Jacobians and differentials.** Let $U \subset \mathbb{R}^n$ be an open subset of Euclidean space $\mathbb{R}^n$, and let $F : U \to \mathbb{R}^m$ be a differentiable function. Let $F_1, \ldots, F_m : U \to \mathbb{R}$ be the components of F, so that $F(u) = (F_1(u), \ldots, F_m(u))$. The Jacobian of F, also know as differential of F, at $u \in U$ is the following matrix of partial derivatives evaluated at u:

$$dF_u = \begin{bmatrix} \frac{\partial F_1}{\partial x_1}\Big|_u & \frac{\partial F_1}{\partial x_2}\Big|_u & \cdots & \frac{\partial F_1}{\partial x_n}\Big|_u \\ \frac{\partial F_2}{\partial x_1}\Big|_u & \frac{\partial F_2}{\partial x_2}\Big|_u & \cdots & \frac{\partial F_2}{\partial x_n}\Big|_u \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial F_m}{\partial x_1}\Big|_u & \frac{\partial F_m}{\partial x_2}\Big|_u & \cdots & \frac{\partial F_m}{\partial x_n}\Big|_u \end{bmatrix}$$

The differential dF_u defines a linear map from $\mathbb{R}^n$ to $\mathbb{R}^m$, that is, an element of $\mathrm{Hom}(\mathbb{R}^n, \mathbb{R}^m)$. Observe that if F itself is linear, then, as matrices, $dF_u = F$ for all points $u \in U$. More generally, let M and N be smooth manifolds, and suppose $F : M \to N$ is a smooth map. The differential of F at $m \in M$ is a linear map between the tangent spaces:

$$dF_m : T_m M \to T_{F(m)} N$$

The map dF_m is computed in local coordinate charts as the Jacobian of partial derivatives. If $G : N \to L$ is another smooth map, then the chain rule implies that:

$$d(G \circ F)_m = dG_{F(m)} \circ dF_m.$$

A.2. **Vector fields.** Let $U \subseteq \mathbb{R}^n$ be an open subset of Euclidean space.

Definition A.1 (Simplified)**.** A vector field on U is a smooth map $\xi : U \to \mathbb{R}^n$. We write $\Gamma(U, TU)$ for the set of vector fields on U.

The idea is that a vector field is the assignment of a direction to every point in U. This direction is a vector in $\mathbb{R}^n$. We write $\xi_u \in \mathbb{R}^n$ for the value of ξ at $u \in U$. For reasons that we do not explain here, we will often denote the vector field ξ as:

$$\xi = f_1 \frac{\partial}{\partial x_1} + f_2 \frac{\partial}{\partial x_2} + \cdots + f_n \frac{\partial}{\partial x_n}$$

where f_i are the components of ξ, regarded as a map $U \to \mathbb{R}^n$, so $\xi(u) = (f_1(u), \ldots, f_n(u))$. Similarly, we may write

$$\xi_u = f_1(u) \frac{\partial}{\partial x_1}\Big|_u + f_2(u) \frac{\partial}{\partial x_2}\Big|_u + \cdots + f_n(u) \frac{\partial}{\partial x_n}\Big|_u$$

for the value of ξ at a $u \in U$. More generally, a vector field on a smooth manifold M is a smooth section of the tangent bundle $TM \to M$, that is, a choice of a tangent vector ξ_p in the tangent space $T_p M$ for every point $p \in M$. Vector fields can also be understood in terms of derivations of the algebra $C^\infty(M)$ of global smooth functions on M.

Returning to the context of an open subset $U \subseteq \mathbb{R}^n$ of Euclidean space, we arrive at the definition of an integral curve:

Definition A.2. An **integral curve** for the vector field ξ on U is a parameterized curve $\gamma : (a,b) \to U$ such that $\dot{\gamma}(t) = \xi_{\gamma(t)}$ for all t in the open interval $(a,b) \subseteq \mathbb{R}$.

To be explicit, in coordinates, the curve γ is tuple of curves $(\gamma_1, \ldots, \gamma_n)$, where $\gamma_i : (a,b) \to \mathbb{R}$, and $\dot{\gamma}(s)$ is the differential $d\gamma_s = \left(\frac{d\gamma_1}{dt}\Big|_s, \ldots, \frac{d\gamma_n}{dt}\Big|_s \right)$.

A.3. **Flows.** We first discuss global flows before discussing more general flows. We continue to work over an open subset $U \subseteq \mathbb{R}^n$ of Euclidean space. As made explicit in the following definition, a global flow is an action of the additive group $\mathbb{R}$ on U.

Definition A.3. A **global flow** on U a smooth map

$$\theta : \mathbb{R} \times U \to U$$

such that, for all $u \in U$, we have $\theta(0,u) = u$ and $\theta(s, \theta(t,u)) = \theta(s+t, u)$ for $s, t \in \mathbb{R}$.

Let θ be a flow on U, and, for $u \in U$, consider the parameterized curve $\theta^{(u)} : \mathbb{R} \to U$, $t \mapsto \theta(t,u)$. The vector field generated by the flow θ is[3]: $u \mapsto d(\theta^{(u)})_0$.

We have seen that global flows generate vector fields. However, it turns out that an arbitrary vector field may not be generated by a global flow. Therefore, we consider the following, more general, definitions.

Definition A.4. A **flow domain** is a subset $\mathcal{D} \subseteq \mathbb{R} \times U$ such that, for every $u \in U$, subset

$$D^{(u)} := \{t \in \mathbb{R} \mid (t,u) \in \mathcal{D}\}$$

is an open interval in $\mathbb{R}$ containing zero.

Definition A.5. Let $\xi : U \to \mathbb{R}^n$ be a vector field on U. A **flow** on U generating ξ is a smooth map

$$\theta : \mathcal{D} \to U$$

where $\mathcal{D} \subset \mathbb{R} \times U$ is a flow domain, such that, for all $u \in U$, we have:

- $\theta(0,u) = u$
- For $t \in \mathcal{D}^{(u)}$ and $s \in \mathcal{D}^{(\theta(t,u))}$ such that $s + t \in \mathcal{D}^{(u)}$, we have:

$$\theta(s, \theta(t,u)) = \theta(s+t, u)$$

- The following parameterized curve is an integral curve for ξ:

$$\theta^{(u)} : \mathcal{D}^{(u)} \to U, \qquad t \mapsto \theta(t,u)$$

One can show that every vector field is generated by a flow.

[3]This is the negative of the infinitesimal action of the Lie algebra of $\mathbb{R}$ on U by vector fields.

Theorem A.6. [Lee12, Theorem 9.12] *For every vector field ξ, there exists a unique maximal smooth flow $\theta : \mathcal{D} \to U$ generating ξ. It has the following properties:*

(1) *For every $u \in U$, the curve $\theta^{(u)}$ is the unique maximal integral curve of ξ starting at u.*
(2) *If $s \in \mathcal{D}^{(u)}$, then $\mathcal{D}^{(\theta(s,u))} = \mathcal{D}^{(u)} - s$.*
(3) *For $t \in \mathbb{R}$, consider:*
$$U_t := \{u \in U \mid (t,u) \in \mathcal{D}\}$$
This is an open subset of U, and
$$\theta_t : U_t \to U_{-t}, \qquad u \mapsto \theta(t,u)$$
is a diffeomorphism with inverse θ_{-t}.
(4) *For $(t,u) \in \mathcal{D}$, we have that $d(\theta_t)_u(\xi_u) = \xi_{\theta(t,p)}$.*

The subset U_t consists of all points for which a flow of time t exists. When a global flow exists, $U_t = U$ for all t.

References

[Lee12] John M. Lee, *Introduction to Smooth Manifolds*, Graduate Texts in Mathematics, vol. 218, Springer, 2012.

[ZGW+23] Bo Zhao, Iordan Ganev, Robin Walters, Rose Yu, and Nima Dehmamy, *Symmetries, flat minima, and the conserved quantities of gradient flow*, The Eleventh International Conference on Learning Representations (ICLR), 2023.