KERNELS OF CANONICAL ALGEBRAIC CURVATURE TENSORS
CSUSB REU, 2014

ADAM WILLIAMS

ABSTRACT. In this paper, we generalize a result on the possible dimensions of the kernel of a linear combination of a particular type of canonical algebraic curvature tensors. We then introduce a new framework for viewing canonical algebraic curvature tensors, using the wedge product, which allows us to give shorter and more transparent proofs of some basic facts about these tensors.

1. Background

1.1. Model Spaces. Given a pseudo-Riemannian manifold (M, g), together with its Levi-Civita connection ∇, the Riemann curvature tensor R is a $C^\infty(M, \mathbb{R})$-multilinear map on vector fields, defined by the formula

$$R(X, Y, Z, W) = g(\nabla_X \nabla_Y Z - \nabla_Y \nabla_X Z - \nabla_{[X,Y]} Z, W)$$

for smooth vector fields X, Y, Z, W on M. Such a tensor field associates to each $p \in M$ a tensor $R_p \in \otimes^4 T_p^* M$ satisfying the following symmetries:

(1) $R_p(x, y, z, w) = -R_p(y, x, z, w)$

(2) $R_p(x, y, z, w) = R_p(z, w, x, y)$

(3) $R_p(x, y, z, w) + R_p(y, z, x, w) + R_p(z, x, y, w) = 0.$

An *algebraic curvature tensor* on a vector space V is a tensor $R \in \otimes^4 V^*$ which satisfies properties (1)-(3). A triple $\mathfrak{M}_0 = (V, \langle \cdot, \cdot \rangle, R)$, where V is a finite-dimensional real vector space, $\langle \cdot, \cdot \rangle$ is a non-degenerate inner product on V, and R is an algebraic curvature tensor on V is called a *model space* (or a *zero* model space, to distinguish it from a model space which is also equipped with tensors that mimic the symmetries of covariant derivatives of the Riemann curvature tensor). A *weak* model space $\mathfrak{M}_0^w = (V, R)$ lacks an inner product.

Immediately, we can see that each point p on a pseudo-Riemannian manifold (M, g) gives a model space $(T_p M, g_p, R_p)$. Gilkey [2] has shown that (1)-(3) are a "universal" list of the symmetries of the Riemann curvature tensor, in the following sense: Given any model space $\mathfrak{M}_0 = (V, \langle \cdot, \cdot \rangle, R)$, there exists a pseudo-Riemannian manifold (M, g) and a point $p \in M$ at which the model space $(T_p M, g_p, R_p)$ is isomorphic to $\mathfrak{M}_0$ (that is, there is a vector space isomorphism $\ell : T_p M \cong V$ which also preserves the inner product and algebraic curvature tensor).

Some common definitions involving model spaces include the following:

Definition 1.1. *Let R be an algebraic curvature tensor on V. The* kernel *of R is defined as*

$$\ker R = \{x \in V : R(x, y, z, w) = 0 \textit{ for all } y, z, w \in V\}.$$

Remark 1.1. *Using the symmetries of R, it is easy to show (see [3]) that this definition of kernel is* not *biased towards the first slot, in the sense that*

$$\begin{aligned}\ker R &= \{y \in V : R(x, y, z, w) = 0 \textit{ for all } x, z, w \in V\} \\ &= \{z \in V : R(x, y, z, w) = 0 \textit{ for all } x, y, w \in V\} \\ &= \{w \in V : R(x, y, z, w) = 0 \textit{ for all } x, y, z \in V\}.\end{aligned}$$

Definition 1.2. *Let $\mathfrak{M}_0^w = (V, R)$ be a weak model space. We say that $\mathfrak{M}_0^w$* decomposes *over two subspaces $V_1, V_2 \subset V$ if*

(1) $V = V_1 \oplus V_2$, and
(2) For all $x_1, y_1, z_1, w_1 \in V_1$ and all $x_2, y_2, z_2, w_2 \in V_2$,

$$R(x_1 + x_2, y_1 + y_2, z_1 + z_2, w_1 + w_2) = R(x_1, y_1, z_1, w_1) + R(x_2, y_2, z_2, w_2).$$

We denote this criterion by writing $R = R_1 \oplus R_2$).

Remark 1.2. *It is also straightforward to show, using the symmetries of R (see [3]) that the following are equivalent:*

(a) $R = R_1 \oplus R_2$.
(b) $R(v_1, v_2, v_3, v_4) = 0$ whenever $v_i \in V_1$ and $v_j \in V_2$ for some $i, j \in \{1, 2, 3, 4\}$.
(c) $R(x_1, x_2, y, z) = 0$ whenever $x_1 \in V_1$ and $x_2 \in V_2$.

Definition 1.3. *A model space $\mathfrak{M}_0 = (V, \langle\cdot,\cdot\rangle, R)$* decomposes *over $V_1, V_2 \subset V$ if the weak model space $\mathfrak{M}_0^w = (V, R)$ decomposes over these subspaces* and *$V_1 \perp V_2$ (that is, $\langle v_1, v_2\rangle = 0$ for all $v_1 \in V_1, v_2 \in V_2$).*

1.2. **Canonical Algebraic Curvature Tensors.** Given a symmetric bilinear form $\phi \in S^2(V^*)$, we can define an algebraic curvature tensor R_ϕ on V via

$$R_\phi(x, y, z, w) = \phi(x, w)\phi(y, z) - \phi(x, z)\phi(y, w).$$

Fiedler [1] and Gilkey [2] have shown that, letting $\mathcal{A}(V)$ denote the vector space of all algebraic curvature tensors on V,

$$\mathcal{A}(V) = \operatorname{span}\{R_\phi : \phi \in S^2(V^*)\}.$$

2. A Result on the Kernel of Linear Combination of Canonical Algebraic Curvature Tensors

Given a weak model space $\mathfrak{M}_0^w = (V, R)$, the *structure group* $G_{\mathfrak{M}_0^w}$ of $\mathfrak{M}_0^w$ is defined to be

$$G_{\mathfrak{M}_0^w} = \{A \in GL(V) \mid A^*R = R\},$$

where A^*R is the *pullback* of R by A, defined by $A^*R(x, y, z, w) = R(Ax, Ay, Az, Aw)$. It is easy to verify that $G_{\mathfrak{M}_0^w}$ is in fact a group.

The following analysis reduces the study of the structure group of an algebraic curvature tensor R to more basic questions: Given $A \in G_{\mathfrak{M}_0^w}$, we know that $\ker R$ is an invariant subspace of V under A, because

$$R(Ax, y, z, w) = A^*R(x, A^{-1}y, A^{-1}z, A^{-1}w) = 0$$

for any $x, y, z, w \in V$. And since A is invertible, the restriction $A|_{\ker R}$ is in fact an isomorphism. So if we take a basis $\beta = \{\epsilon_1, \ldots, \epsilon_k, \eta_1, \ldots, \eta_l\}$ for V so that $\eta_1, \ldots, \eta_k$ is a basis for $\ker R$, then the matrix of A with respect to β has the form

$$A_\beta = \begin{bmatrix} \bar{A} & 0 \\ C & N \end{bmatrix},$$

where Dunn, Franks, and Palmer [3] have shown that:

(1) $\bar{A}$ is in $G_{\bar{\mathfrak{M}}_0^w}$, where $\bar{\mathfrak{M}}_0^w = (\bar{V}, \bar{R})$ is the model space with $\bar{V} = V/\ker R$, and $\bar{R}$ satisfies $\bar{R}(\pi x, \pi y, \pi z, \pi w) = R(x, y, z, w)$ for all $x, y, z, w \in V$ (where $\pi : v \mapsto v + \ker R$ is the natural projection).
(2) $N \in GL(\ker R)$.

(3) There are no restrictions on C.

So, understanding the structure group of R amounts in large part to understanding its kernel. Hence the next theorem investigates the possible dimensions of $\ker(R_\phi \pm R_\psi)$, for $\phi, \psi \in S^2(V^*)$ in order to start gaining an understanding of $\ker R$ and thus of $G_{\mathfrak{M}_0^w}$ in general. First, a lemma:

Lemma 2.1. *Up to a sign change, every linear combination $R = aR_\phi + bR_\psi$ of canonical algebraic curvature tensors, where $a, b \neq 0$, can be written as $R_{\tilde{\phi}} + \delta R_{\tilde{\psi}}$ for some $\tilde{\phi}, \tilde{\psi} \in S^2(V^*)$ and $\delta \in \{-1, 1\}$.*

Proof. Up to a sign change of R, we may assume $a > 0$. Set

$$\begin{aligned} \tilde{\phi} &= \sqrt{a}\phi \\ \tilde{\psi} &= \sqrt{|b|}\psi \\ \delta &= b/|b|. \end{aligned}$$

It is then immediate from the definition that $R_{\tilde{\phi}} + \delta R_{\tilde{\psi}} = aR_\phi + bR_\psi$.

□

This theorem is a generalization of a result of Strieby [4]:

Theorem 2.2. *Let V be an n-dimensional vector space ($n \geq 3$), ϕ a non-degenerate inner product on V with signature (p, q), and ψ a symmetric bilinear form on V. Suppose further that $\beta = \{e_1, \ldots, e_n\}$ is an orthonormal basis for V which also diagonalizes ψ. That is, the matrices $\phi_\beta = (\phi(e_i, e_j))_{ij}$ and $\psi_\beta = (\psi(e_i, e_j))_{ij}$ are given by*

$$\phi_\beta = \begin{pmatrix} \epsilon_1 & & & & & 0 \\ & \ddots & & & & \\ & & \epsilon_p & & & \\ & & & \epsilon_{p+1} & & \\ & & & & \ddots & \\ 0 & & & & & \epsilon_{p+q} \end{pmatrix}, \qquad \psi_\beta = \begin{pmatrix} \lambda_1 & & 0 \\ & \ddots & \\ 0 & & \lambda_n \end{pmatrix}$$

(where $\epsilon_1, \ldots, \epsilon_p = +1$ and $\epsilon_{p+1}, \ldots, \epsilon_{p+q} = -1$). Let $R = R_\phi + \delta R_\psi$, where $\delta \in \{-1, 1\}$. Furthermore assume that $\ker R \neq 0$; that is, that there exists $v = \sum c_r e_r \in \ker R$ and an index l with $c_l = 1$. Then

(1) *There exists a non-zero real number λ such that the diagonal entries of ψ_β have the following form:*

$$\lambda_i = \begin{cases} \epsilon_i \lambda & i \neq l \\ -\delta/\lambda & i = l \end{cases}.$$

(2) $R = 0$ *only if* $\delta = -1$, $\lambda = \pm 1$, *and* $\phi = \lambda\psi$.

(3) *If* $R \neq 0$, *then* $\ker R = \operatorname{span}\{e_l\}$.

In particular, given such ϕ, ψ, *and* R, *the only possible dimensions of* $\ker R$ *are* 0, 1, *and* n.

Proof. Recall that $R_\phi(x, y, z, w) = \phi(x, w)\phi(y, z) - \phi(x, z)\phi(y, w)$, and similarly for R_ψ. We observe that for any $i \neq l$ we have, since $v \in \ker R$,

$$\begin{aligned} 0 &= R(v, e_i, e_i, e_l) \\ &= \phi(v, e_l)\phi(e_i, e_i) + \delta\psi(v, e_l)\psi(e_i, e_i) \\ &= \sum c_i \phi(e_i, e_l)\phi(e_i, e_i) + \delta \sum c_i \psi(e_i, e_l)\psi(e_i, e_i) \\ &= \epsilon_l \epsilon_i + \delta \lambda_l \lambda_i. \end{aligned}$$

So

$$\lambda_l \lambda_i = -\epsilon_l \epsilon_i \delta. \tag{1}$$

Now, for any $j \leq i$ and $i, j \neq l$, we either have $\epsilon_i = \epsilon_j$ or $\epsilon_i = -\epsilon_j$. In the first case, we have the two equations

$$\begin{aligned} \lambda_l \lambda_i &= -\epsilon_l \epsilon_i \delta \\ \lambda_l \lambda_j &= -\epsilon_l \epsilon_i \delta, \end{aligned}$$

which can be subtracted to yield

$$\lambda_l(\lambda_i - \lambda_j) = 0.$$

But since $\lambda_l \lambda_i = -\epsilon_l \epsilon_i \delta \neq 0$, we know that $\lambda_l \neq 0$, and so $\lambda_i = \lambda_j$ for $\epsilon_i = \epsilon_j$. Similarly, if $\epsilon_i = -\epsilon_j$, then we have

$$\begin{aligned} \lambda_l \lambda_i &= -\epsilon_l \epsilon_i \delta \\ \lambda_l \lambda_j &= \epsilon_l \epsilon_i \delta, \end{aligned}$$

which yields

$$\lambda_l(\lambda_i + \lambda_j) = 0,$$

meaning that $\lambda_i = -\lambda_j$ for $\epsilon_i = -\epsilon_j$. Hence if we set $\lambda = \lambda_i$ for $i = 1, \dots, p$ (or $\lambda = -\lambda_i$ for $i = p + 1, \dots, p + q$), we have that

$$\lambda_i = \begin{cases} \epsilon_i \lambda & i \neq l \\ -\epsilon_l \delta/\lambda & i = l \end{cases},$$

which completes the proof of the first claim.

Now we are ready to analyze R. First we note that $R_{ijkm} = 0$ unless there are exactly two distinct indices, and also entries of the form R_{iijj} are zero. Thus the only curvature entries we need to consider are those of the form R_{ijji} for $i \neq j$, $i, j \neq l$; and those of the form R_{illi} for $i \neq l$. For the latter, we observe that

$$\begin{aligned} R_{illi} &= R(e_i, e_l, e_l, e_i) \\ &= \phi(e_i, e_i)\phi(e_l, e_l) + \delta\psi(e_i, e_i)\psi(e_l, e_l) \\ &= \epsilon_i\epsilon_l + \delta\lambda_i\lambda_l \\ &= \epsilon_i\epsilon_l + \delta\epsilon_i\lambda\frac{(-\epsilon_l)\delta}{\lambda} \\ &= \epsilon_i\epsilon_l - \epsilon_i\epsilon_l \\ &= 0. \end{aligned}$$

And for the former, we have

$$\begin{aligned} R_{ijji} &= \phi(e_i, e_i)\phi(e_j, e_j) + \delta\psi(e_i, e_i)\psi(e_j, e_j) \\ &= \epsilon_i\epsilon_j + \delta\lambda_i\lambda_j \\ &= \epsilon_i\epsilon_j + \delta\epsilon_i\lambda\epsilon_j\lambda \\ &= \epsilon_i\epsilon_j\left(1 + \delta\lambda^2\right). \end{aligned}$$

So assume $R = 0$. Then $R_{ijj} = \epsilon_i\epsilon_j\left(1 + \delta\lambda^2\right) = 0$ for all $i \neq j$, $i, j \neq l$. Since $\epsilon_i, \epsilon_j \neq 0$, we have

$$1 + \delta\lambda^2 = 0.$$

So, $\delta = -1$ and $\lambda = \pm 1$. It thus follows from claim (1) that if $\lambda = 1$ then $\psi = \phi$, and if $\lambda = -1$ then $\psi = -\phi$, which proves the second claim.

So assume $R \neq 0$. By the preceding analysis, there exist i, j with $i \neq j$, $i, j \neq l$ such that $R_{ijji} \neq 0$. And since the values of such R_{ijji} differ only by a sign, we know that in fact $R_{ijji} \neq 0$ for *every* $i \neq j$ and $i, j \neq l$. Using this, we can show that $\ker R = \langle e_l \rangle$. Let $w = \sum b_r e_r \in \ker R$. Then for any $i \neq l$, we can choose $j \neq, i, l$ (since $\dim V \geq 3$), and we have

$$\begin{aligned} 0 &= R\left(\sum b_r e_r, e_j, e_j, e_i\right) \\ &= \sum b_r R_{rjji} \\ &= b_i R_{ijji}. \end{aligned}$$

Since $R_{ijji} \neq 0$, we conclude that $b_i = 0$, and since $i \neq l$ was arbitrary, this means that $w = b_l e_l \in \langle e_l \rangle$, which proves the third claim.

□

3. Canonical Algebraic Curvature Tensors as Wedge Products

In this section, we introduce a superficially different definition of the object R_ϕ for $\phi \in S^2(V^*)$, in order to give simpler and more transparent proofs of some basic facts about these canonical algebraic curvature tensors. First, we recall two standard definitions:

Definition 3.1. *Given a symmetric bilinear form $\phi \in S^2(V)$, the* kernel *of ϕ is defined as*

$$\ker\phi = \{x \in V : \phi(x, y) = 0 \textit{ for all } y \in V\},$$

and the rank *of ϕ is defined as* $\operatorname{rk}\phi = \dim V - \ker\phi$.

Lemma 3.1. *Let $\phi \in S^2(V^*)$. For each $x \in V$, ϕ induces a linear map $\phi_x : V \to \mathbb{R}$ defined by*

$$\phi_x(y) = \phi(x, y)$$

for all $y \in V$. Then the map $\Phi : x \mapsto \phi_x$ from V to its dual V^* is linear, and

$$\ker \Phi = \ker \phi$$

and

$$\operatorname{rk} \Phi = \operatorname{rk} \phi.$$

Proof. Observe that

$$\begin{aligned} \ker \Phi &= \{x \in V : \phi_x = 0\} \\ &= \{x \in V : \phi_x(y) = 0 \text{ for all } y \in V\} \\ &= \{x \in V : \phi(x, y) = 0 \text{ for all } y \in V\} \\ &= \ker \phi. \end{aligned}$$

The fact that $\operatorname{rk} \Phi = \operatorname{rk} \phi$ then follows from the rank-nullity theorem.

□

Recall some common definitions: Given vector spaces V and W, a map $M : V^k \to W$ is *multilinear* if the map $v \mapsto M(v_1, \dots, v_{i-1}, v, v_{i+1}, \dots, v_k)$ is linear for every $i = 1, \dots, k$ and every $(v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_k) \in V^{k-1}$. A multilinear map A is called *alternating* if for every permutation $\sigma \in S_k$ and every $v_1, \dots, v_k \in V$, we have

$$A(v_{\sigma(1)}, \dots, v_{\sigma(k)}) = \operatorname{sgn}(\sigma) A(v_1, \dots, v_k)$$

(in the case of real vector spaces, this is equivalent to the condition that $A(v_1, \dots, v_k) = 0$ whenever $v_i = v_j$ for some $i \neq j$).

Now we recall the linear algebraic construction of the *wedge product* of vectors: Given a finite-dimensional real vector space V, one can construct a certain real associative algebra containing V, known as the *exterior algebra* of V and denoted $\Lambda(V)$. The multiplication operation $(\alpha, \beta) \mapsto \alpha \wedge \beta$ on $\Lambda(V)$ is called the wedge product (or *exterior* product), and can be described intuitively as the "most general alternating product of vectors." The most useful formal interpretation of this intuition for our purposes is the following: The vector space $\Lambda(V)$ can be decomposed as a direct sum

$$\Lambda(V) = \bigoplus_{k=0}^{\infty} \Lambda^k(V),$$

where $\Lambda^k(V)$ is a $\binom{\dim V}{k}$-dimensional vector space spanned by vectors of the form $v_1 \wedge \cdots \wedge v_k$ for $v_1, \dots, v_k \in V$, and satisfying the following universal property: For every real vector space W and every alternating, multilinear map $A : V^k \to W$, there exists a unique linear map $\ell : \Lambda^k(V) \to W$ such that the following diagram commutes,

$$\begin{tikzcd} \Lambda^k(V) \arrow[dr, "\ell"] & \\ V^k \arrow[u, "\wedge^k"] \arrow[r, "A"'] & W \end{tikzcd}$$

where $\wedge^k$ denotes the map $(v_1, \dots, v_k) \mapsto v_1 \wedge \cdots \wedge v_k$.

The following two results are well-known facts about the exterior product of vectors, but are proved here for the sake of completeness:

Lemma 3.2. *Let V be an n-dimensional vector space and $\Lambda(V)$ its exterior algebra. Then for every natural number k and every $v_1, \ldots, v_k \in V$, $v_1 \wedge \cdots \wedge v_k = 0$ if and only if $v_1, \ldots, v_k$ are linearly dependent.*

Proof. Suppose that $v_1, \ldots, v_k$ are dependent. Then for some index j and scalars a_i we have

$$v_j = \sum_{i \neq j} a_i v_i.$$

We then apply the alternating and multilinear properties of the wedge product to conclude

$$\begin{aligned} v_1 \wedge \cdots \wedge v_k &= v_1 \wedge \cdots \wedge v_{j-1} \wedge \left(\sum_{i \neq j} a_i v_i \right) \wedge v_{j+1} \wedge \cdots \wedge v_k \\ &= \sum_{i \neq j} a_i v_1 \wedge \cdots \wedge v_{j-1} \wedge v_i \wedge v_{j+1} \wedge \cdots \wedge v_k \\ &= 0. \end{aligned}$$

For the second direction, we prove the contrapositive: Suppose $v_1, \ldots, v_k$ are linearly independent. Then we claim that there exists an alternating, multilinear form $A : V^k \to \mathbb{R}$ such that $A(v_1, \ldots, v_k) = 1$. To construct one such form, we can extend $v_1, \ldots, v_k$ to a basis $v_1, \ldots, v_n$ for V and define A on this basis by

$$A(v_{i_1}, \ldots, v_{i_k}) = \begin{cases} \operatorname{sgn}(\sigma) & i_1 = \sigma(1), \ldots, i_k = \sigma(k) \text{ for some } \sigma \in S_k \\ 0 & \text{else} \end{cases}$$

and extending mutlilinearly to the rest of V^k. Then $A : V^k \to \mathbb{R}$ is alternating, with $A(v_1, \ldots, v_k) = 1$.

Now, by the universal property of the exterior power $\Lambda^k(V)$, there exists a linear map $\ell : V^k \to \Lambda^k(V)$ such that $\ell(x_1 \wedge \cdots \wedge x_k) = A(x_1, \ldots, x_k)$ for all $(x_1, \ldots, x_k) \in V^k$. In particular, $\ell(v_1 \wedge \cdots \wedge v_k) = 1 \neq 0$, and since ℓ is linear, this implies $v_1 \wedge \cdots \wedge v_k \neq 0$, which is what we wanted to show.

□

Lemma 3.3. *Let V_1 and V_2 be vector subspaces of V, and suppose that for all $v_1 \in V_1$ and all $v_1 \in V_2$, $v_1 \wedge v_2 = 0$. Then either $V_1 = 0$ or $V_2 = 0$ or V_1 and V_2 are the same one-dimensional subspace of V.*

Proof. Assume $V_1 \neq 0$ and $V_2 \neq 0$. Then there exists $v_1 \neq 0$ in V_1. Then, for every non-zero $v_2 \in V_2$, $v_1 \wedge v_2 = 0$, so v_1 and v_2 are dependent by the previous Lemma. So $v_2 \in \operatorname{span}\{v_1\}$. Since v_2 was an arbitrary non-zero vector in V_2, we conclude $V_2 \subset \operatorname{span}\{v_1\}$. Since $V_2 \neq 0$ by assumption, this means $V_2 = \operatorname{span}\{v_1\}$.

Note that we took an arbitrary non-zero $v_1 \in V_1$ and showed that $V_2 = \operatorname{span}\{v_1\}$. Now, in particular, v_1 is a non-zero vector in V_2, and so we can apply a symmetric argument to conclude that $V_1 = \operatorname{span}\{v_1\}$ as well, which proves the claim.

□

Now we are ready to give the 'new' definition of R_ϕ for $\phi \in S^2(V^*)$:

Definition 3.2. *For a symmetric, bilinear form on a vector space V, define $R_\phi : V^2 \to \Lambda^2(V^*)$ by*

$$R_\phi(x, y) = \phi_y \wedge \phi_x.$$

In order to see why this new definition of R_ϕ is equivalent to the old one, we make use of the canonical isomorphism

$$\Lambda^k(V^*) \cong \mathrm{Alt}_k(V, \mathbb{R}),$$

where $\mathrm{Alt}_k(V, \mathbb{R})$ is the vector space of alternating k-multilinear forms $A : V^k \to \mathbb{R}$. This isomorphism is given by letting $\alpha_1 \wedge \cdots \wedge \alpha_k \in \Lambda^k(V^*)$ act on V^k via

$$(\alpha_1 \wedge \cdots \wedge \alpha_k)(v_1, \ldots, v_k) = \sum_{\sigma \in S_k} \mathrm{sgn}(\sigma) \alpha_1(v_{\sigma_1}) \cdots \alpha_k(v_{\sigma(k)}).$$

Hence we can view $R_\phi(x, y)$ as an alternating bilinear form on V, and by the above formula, for any $z, w \in V$ we have

$$R_\phi(x, y)(z, w) = \phi(x, w)\phi(y, z) - \phi(x, z)\phi(y, w).$$

Lemma 3.4. *The following definitions*

(a) $\ker R_\phi = \{x \in V : R_\phi(x, y) = 0$ *for all* $y \in V\}$.
(b) Let $\mathfrak{M}_0^w = (V, R_\phi)$ be a weak model space with algebraic curvature tensor R_ϕ. Then we say $\mathfrak{M}_0^w$ decomposes *into two subspaces $V_1, V_2 \subset V$ if $V = V_1 \oplus V_2$ and for every $v_1 \in V_1$ and $v_2 \in V_2$, $R_\phi(v_1, v_2) = 0$.*
(c) Let $\mathfrak{M}_0 = (V, \langle \cdot, \cdot \rangle, R_\phi)$ be a model space with algebraic curvature tensor R_ϕ. Then we say $\mathfrak{M}_0$ decomposes *into two subspaces $V_1, V_2 \subset V$ if $\mathfrak{M}_0^w = (V, R_\phi)$ decomposes into V_1 and V_2 and V_1 is perpendicular to V_2 with respect to $\langle \cdot, \cdot \rangle$.*

agree with the old definitions, in the sense that

(1) $\ker R_\phi = \{x \in V : R_\phi(x, y)(z, w) = 0$ *for all* $y, z, w \in V\}$.
(2) $R_\phi(v_1, v_2) = 0$ *whenever* $v_1 \in V_1$, $v_2 \in V_2$ *if and only if* $R_\phi = (R_\phi)_1 \oplus (R_\phi)_2$.

Proof. The equivalence is immediate from the definitions. □

We use these facts to give a new proof of a result of Gilkey [2]:

Theorem 3.5. *Let $\phi \in S^2(V^*)$, and assume* $\mathrm{rk}\, \phi \geq 2$. *Then*

(1) If the model space $\mathfrak{M}_0^w = (V, R_\phi)$ decomposes over V_1 and V_2, then either $V_1 \subset \ker \phi$ or $V_2 \subset \ker \phi$.
(2) If ϕ is non-degenerate, then the model space $\mathfrak{M}_0^w = (V, R_\phi)$ is indecomposable.
(3) If $\langle \cdot, \cdot \rangle$ is an inner product on V and $\ker \phi$ is totally isotropic, then the model space $\mathfrak{M}_0 = (V, \langle \cdot, \cdot \rangle, R_\phi)$ is indecomposable.
(4) $\ker R_\phi = \ker \phi$.

Proof. Proof of (1): Suppose $\mathfrak{M}_0^w$ decomposes into V_1 and V_2. Then whenever $v_1 \in V_1$ and $v_2 \in V_2$, we have $R_\phi(v_1, v_2) = \phi_{v_2} \wedge \phi_{v_1} = 0$. That is for any $\alpha_1 \in \Phi[V_1] = \{\phi_{v_1} \mid v_1 \in V_1\}$, $\alpha_2 \in \Phi[V_2]$, we have $\alpha_1 \wedge \alpha_2 = 0$. By Lemma 3.3, this implies that either $\Phi[V_1] = 0$ or $\Phi[V_2] = 0$, or $\Phi[V_1]$ and $\Phi[V_2]$ are the same 1-dimensional subspace of V^*. The latter case is a contradiction: since $V_1 + V_2 = V$ by assumption and Φ is linear,

$$\dim \Phi[V] = \dim \Phi[V_1 + V_2] = \dim(\Phi[V_1] + \Phi[V_2]) = \dim \Phi[V_1] = \dim \Phi[V_2] = 1,$$

which contradicts the fact that $\operatorname{rk}\phi = \operatorname{rk}\Phi \geq 2$. So either $\Phi[V_1] = 0$ or $\Phi[V_2] = 0$, in which case $V_1 \subset \ker\phi$ or $V_2 \subset \ker\phi$ (since $\ker\phi = \ker\Phi$).

Proof of (2): This follows immediately from (1), since ϕ is non-degenerate if and only if $\ker\phi = 0$.

Proof of (3): This follows from (1) as well: First we claim that if $V = V_1 \perp V_2$ is a non-trivial orthogonal decomposition of V with respect to $\langle\cdot,\cdot\rangle$, then it cannot be the case that V_1 or V_2 is totally isotropic. To see this, suppose V_1 was totally isotropic. Since the decomposition is nontrivial, there is a non-zero vector $v_1 \in V_1$. Since $V_1 \perp V_2$, we know that $v_1 \perp v_2$ for all $v_2 \in V_2$. And since V_1 is totally isotropic, $v_1 \perp v_1'$ for all $v_1' \in V_1$. Since any $v \in V$ can be written as $v_1' + v_2$ for $v_1' \in V$, $v_2' \in V_2$, we have $\langle v_1, v\rangle = \langle v_1, v_1'\rangle + \langle v_1, v_2\rangle = 0$ for all $v \in V$. So the inner product is degenerate, which contradicts our assumption. So we can conclude that neither V_1 nor V_2 is totally isotropic in the decomposition. Hence $V_1, V_2 \not\subset \ker\phi$, so $R_\phi \neq R_1 \oplus R_2$.

Proof of (4): If $x \in \ker\phi$, then $\phi_x = 0$, which immediately implies $R_\phi(x, y) = 0$. So let $x \in \ker R_\phi$. Then $\phi_y \wedge \phi_x = 0$ for every $y \in V$. Since $\operatorname{rk}\phi \geq 2$, there exist $y, z \in V$ with ϕ_y and ϕ_z linearly independent. Since $\phi_y \wedge \phi_x = 0$, ϕ_x and ϕ_y are dependent by Lemma 3.2. Hence, since ϕ_y is non-zero, ϕ_x is a scalar multiple of ϕ_y: There exists $a \in \mathbb{R}$ with $\phi_x = a\phi_y$. Similarly, ϕ_x and ϕ_z are dependent and $\phi_z \neq 0$, so there is $b \in \mathbb{R}$ with $\phi_x = b\phi_z$. So

$$\phi_x - \phi_x = 0,$$

which means

$$a\phi_y - b\phi_z = 0.$$

Since ϕ_y are ϕ_z are independent, this implies $a = b = 0$, and so $\phi_x = 0$, which means $x \in \ker\phi$.

□

4. Open Questions

(1) The case where a symmetric form ψ can be orthogonally diagonalized with respect to an arbitrary non-degenerate inner product is very special. However, It is known that, given a *Lorentzian* inner product ϕ (i.e. ϕ has signature $(n-1, 1)$ and a symmetric bilinear form ψ on V, ψ can always be diagonalized outside of a subspace of dimension at most 3 [5]. Can we use this information to understand the possible dimensions of $\ker(R_\phi + \delta R_\psi)$ for Lorentzian ϕ? In particular, can we apply the method used for diagonalized symmetric forms in this project to understand what happens on the "diagonalized" subspace, and then proceed case-by-case through all of the possible forms ψ can take outside of this subspace?

(2) It is a well known fact [2] about canonical algebraic curvature tensors that if ϕ, $\psi \in S^2(V^*)$ and $\operatorname{rk}\phi \geq 3$ and $R_\phi = R_\psi$, then $\phi = \pm\psi$. Gilkey's proof involves choosing a very special basis for V and calculating the values of R_ϕ and R_ψ on certain tuples of basis elements. Can we give a basis-independent proof using wedge products?

(3) Fiedler [1] and Gilkey [2] have also shown that

$$\mathcal{A}(V) = \operatorname{span}\{R_\phi : \phi \in \Lambda^2(V^*)\},$$

where R_ϕ for an anti-symmetric form $\phi \in \Lambda^2(V^*)$ is an algebraic curvature tensor given by

$$R_\phi(x, y, z, w) = \phi(x, w)\phi(y, z) - \phi(x, z)\phi(y, w) - 2\phi(x, y)\phi(z, w)$$

for $x, y, z, w \in V$. Can such a tensor be represented as a wedge product, and can we use that representation to prove the analogous version of Theorem 3.5 for these canonical ACTs?

(4) *Canonical algebraic covariant derivative curvature tensors* model the symmetries of the covariant derivative ∇R of the Riemann curvature tensor, and take the form $\nabla R_{\phi,\psi}$ for $\phi \in S^2(V^*)$, $\psi \in S^3(V^*)$, where

$$\nabla R_{\phi,\psi}(x, y, z, w; s) = \phi(x, w)\psi(y, z, s)+\phi(y, z)\psi(x, w, s)-\phi(x, z)\psi(y, w, s)-\phi(y, w)\psi(x, z, s)$$

for $x, y, z, w, s \in V$. Such a tensor can be represented as a sum of wedge products

$$\nabla R_{\phi,\psi}(x, y, z, w; s) = (\phi_w \wedge \psi_{zs} + \phi_z \wedge \psi_{ws})\,(x, y),$$

where $\psi_{uv} \in V^*$ is the map $t \mapsto \psi(u, v, t)$. Can the wedge product method be used to analyze the kernels of these tensors?

(5) If $R = R_\phi + \delta R_\psi$ and $R = R_1 \oplus R_2$, does this imply that $R_\phi = (R_\phi)_1 \oplus (R_\phi)_2$ and $R_\psi = (R_\psi)_1 \oplus (R_\psi)_2$? IE, if $R_{\phi+\delta R_\psi}$ is decomposable, must R_ϕ and R_ψ be decomposable over the same subspaces?

(6) It follows from Lemma 3.2 that if R_ϕ and R_ψ are linearly dependent, then the vectors ϕ_x, ϕ_y, ψ_x, ψ_y are linearly dependent for every $x, y \in V$ (although the converse is not true). Are wedge products a useful tool for studying linear dependence of sets of ACTs in general?

References

[1] B. Fielder (2003) *Determination of the structure of algebraic curvature tensors by means of Young symmetrizers*, Seminaire Lotharingien de Combinatoire, 48 (2003) Article B48d.

[2] P. Gilkey. *The Geometry of Curvature Homogeneous Pseudo-Riemannian Manifolds*, Imperial College Press (2007).

[3] C. Dunn, C. Franks, J. Palmer *On the structure group of a decomposable model space*, to appear in Contributions in Algebra and Geometry, 2015.

[4] D. Strieby. *The decomposability of full and weak model spaces*, CSUSB 2009 REU, http://www.math.csusb.edu/reu/ds09.pdf.

[5] B. O'Neill. *Semi-Riemannian Geometry with Applications to Relativity*, Academic Press (1983).