BRONOWSKI'S CONJECTURE, IDENTIFIABILITY, AND NEUROVARIETIES

ALEX MASSARENTI

ABSTRACT. These notes are based on a seminar delivered at the workshop *Geometry of Neural Networks and Tensors in Nancy (2026)*. They give a compact introduction to the geometry of secant varieties, contact loci, and identifiability, with emphasis on Veronese varieties and polynomial neural networks. We first recall defectivity, weak defectivity, tangential weak defectivity, and identifiability, and explain their relations through Terracini's lemma [Ter11]. We then discuss Veronese varieties, including the Alexander–Hirschowitz theorem [AH95], the exceptional nonidentifiable subgeneric Waring cases of Chiantini–Ottaviani–Vannieuwenhoven [COV17], the generic-rank identifiability theorem of Galuppi–Mella [GM19], and the canonical decompositions of Hilbert and Sylvester [Hil88, Syl04, IK99]. The second part is devoted to Bronowski's conjecture in the modern form considered by Ciliberto–Russo [CR06], to the counterexamples and corrected criterion of Massarenti–Mella [MM24], and to the passage from secant varieties to neurovarieties in [MM25].

CONTENTS

INTRODUCTION

The purpose of these notes is to explain how classical projective geometry organizes modern identifiability questions for tensors, homogeneous polynomials, and polynomial neural networks. They were written for the seminar *Bronowski's Conjecture, Identifiability, and Neurovarieties*, delivered at the workshop *Geometry of Neural Networks and Tensors in Nancy (2026)*.

The central geometric object is a projective variety $X \subset \mathbb{P}^N$ together with its secant varieties. A general point of $\mathbb{S}\mathrm{ec}_h(X)$ is represented as a point in the span of h points of X. The variety X is h-identifiable when this representation is unique. In tensor language, X is usually a variety of rank-one tensors, and identifiability is uniqueness of a tensor decomposition. In the symmetric case, X is a Veronese variety, and the problem becomes uniqueness of Waring decompositions of homogeneous polynomials; this is the setting of the classical theory of power sums and inverse systems [IK99].

The first tool is Terracini's lemma [Ter11]. It translates the dimension of a secant variety into the dimension of a span of tangent spaces. Thus defectivity is a failure of tangent spaces to impose the expected number of independent conditions. For Veronese varieties this has a concrete interpolation meaning: hyperplanes containing tangent spaces to the Veronese correspond to hypersurfaces in $\mathbb{P}^n$ with assigned double points. This gives the bridge between secant varieties, linear systems, and the Alexander–Hirschowitz theorem [AH95]. In particular, the exceptional defective Veronese cases are special failures of interpolation by double points.

Dimension, however, is only the first layer of the story. Even when a secant variety has the expected dimension and the abstract secant map is generically finite, the general fiber may have degree greater than one. This is where contact loci enter. The theory of weakly defective varieties developed by Chiantini–Ciliberto shows that a general tangent hyperplane may have a positive-dimensional contact locus and that this phenomenon is a fundamental obstruction to identifiability [CC02, CC06]. The refined tangential version studies the locus where the span of several tangent spaces is tangent to X; structural results on these tangential contact loci are used in many tensor-identifiability arguments [BBC18]. Non-identifiability forces such tangential contact in the

Date: June 10, 2026.

generically finite range, but the converse is false: a contact locus is the place where another decomposition may occur, not automatic proof that it does occur.

For Veronese varieties, the resulting picture is especially sharp. The Alexander–Hirschowitz theorem classifies defectivity [AH95]. The theorem of Chiantini–Ottaviani–Vannieuwenhoven gives the corresponding subgeneric identifiability statement: outside three exceptional cases, a general symmetric tensor of subgeneric rank is identifiable [COV17]. These exceptional cases have a common geometric explanation through elliptic normal curves and secant order two, in the sense of the secant-order theory of Chiantini–Ciliberto [CC06, COV17]. At the generic rank, the theorem of Galuppi–Mella shows that identifiability is much rarer and is governed by Cremona transformations [GM19]. This recovers the classical canonical forms of binary forms of odd degree due to Sylvester, the pentahedral theorem for cubic surfaces, and Hilbert's theorem on ternary quintics [Syl04, Hil88, IK99]. The corresponding reconstruction algorithms are those of Laface–Massarenti–Rischter [LMR23].

The next theme is Bronowski's conjecture. In the form considered by Massarenti–Mella, the conjecture predicts that, for $X \subset \mathbb{P}^{hn+h-1}$, h-identifiability should be equivalent to the birationality of the general $(h-1)$-tangential projection [MM24]. Ciliberto–Russo revisited this circle of ideas in a modern framework involving varieties with minimal secant degree [CR06]. The implication from identifiability to birationality is known, while the converse is subtle. Massarenti–Mella constructed counterexamples using secant varieties of rational normal curves and showed that birationality of the tangential projection alone is not enough [MM24]. Their examples have degenerate Gauss map, and this is the hidden tangential pathology. The corrected criterion of Massarenti–Mella says that, under nondegeneracy of the Gauss map and nondefectivity one step higher, X is h-identifiable [MM24].

The final part connects this secant viewpoint to polynomial neural networks. For networks with polynomial activations, the output coordinates are homogeneous forms, and the Zariski closure of the coefficient image is a neurovariety [MM25]. In the shallow single-output case this is exactly a secant variety of a Veronese variety. In deeper architectures, one obtains a constrained secant construction over projected iterated Veronese varieties [MM25]. A Terracini-type formula decomposes the tangent space to a neurovariety into the usual final secant contribution and additional level-wise normal spaces coming from the intermediate layers [MM25].

This viewpoint gives an Alexander–Hirschowitz type theorem for neurovarieties [MM25]. Under suitable room conditions and nondefectivity of the last Veronese block, the single-output neurovariety has the expected dimension. In the multi-output case, the condition $n_L \geq 2$ plays a decisive role: a single output may lie on finitely many admissible secant spaces, but a general tuple of at least two outputs determines the admissible secant space uniquely [MM25]. This is the geometric mechanism behind global identifiability of the network, up to the standard permutation and scaling symmetries.

The guiding principle throughout the notes is that identifiability is controlled by tangent and secant geometry. Defectivity detects dimensional failure, contact loci detect the possible location of non-uniqueness, Bronowski-type projections test identifiability through birational geometry, and neurovarieties show that the same mechanisms persist in the geometry of polynomial neural networks.

1. Secants, contact loci, and identifiability

We work over a field k of characteristic zero. Let $X \subset \mathbb{P}^N$ be irreducible, nondegenerate, and of dimension n.

The h-secant variety $\mathbb{S}\mathrm{ec}_h(X)$ is the Zariski closure of the union of the $(h-1)$-planes $\langle x_1, \dots, x_h \rangle$ spanned by h independent points of X. Its expected dimension is

$$\mathrm{expdim}(\mathbb{S}\mathrm{ec}_h(X)) = \min\{h(n+1)-1, N\}.$$

The variety X is h-defective if $\dim \mathbb{S}\mathrm{ec}_h(X) < \mathrm{expdim}(\mathbb{S}\mathrm{ec}_h(X))$.

The abstract h-secant variety is

$$\mathcal{S}_h(X) = \overline{\{(p, x_1 + \cdots + x_h) \in \mathbb{P}^N \times X^{(h)} \mid p \in \langle x_1, \dots, x_h \rangle\}},$$

where $X^{(h)}$ is the h-th symmetric product of X. It comes with the natural map

$$\pi_h^X : \mathcal{S}_h(X) \longrightarrow \mathbb{S}\mathrm{ec}_h(X), \qquad (p, x_1 + \cdots + x_h) \longmapsto p.$$

The general fiber of π_h^X parametrizes the decompositions of a general point $p \in \mathbb{S}\mathrm{ec}_h(X)$ as a point in the span of h points of X. When π_h^X is generically finite, its degree is the h-secant degree of X, and X is h-identifiable if this degree is one.

Remark 1.1. For tensors, X is the variety of rank-one tensors and h-identifiability is uniqueness of a tensor decomposition of rank h. For Veronese varieties it is uniqueness of a Waring decomposition.

Lemma 1.2 (Terracini)**.** *Let $x_1, \dots, x_h \in X$ be general smooth points, and let*

$$p \in \langle x_1, \dots, x_h \rangle$$

be general. Then

$$T_p \operatorname{\mathbb{S}ec}_h(X) = \langle T_{x_1} X, \dots, T_{x_h} X \rangle.$$

Thus X is h-defective precisely when the span of h general embedded tangent spaces has dimension smaller than $\operatorname{expdim}(\mathbb{S}\mathrm{ec}_h(X))$.

Definition 1.3. Let $A = \{x_1, \dots, x_h\} \subset X_{\mathrm{reg}}$ be general and set

$$M_A = \langle T_{x_1} X, \dots, T_{x_h} X \rangle.$$

The h-tangential projection is the rational map

$$\tau_h^X : X \dashrightarrow X_h$$

obtained by projecting from M_A. The h-tangential contact locus is

$$\Gamma_h(A) = \overline{\{x \in X_{\mathrm{reg}} \mid T_x X \subseteq M_A\}},$$

where one usually keeps only the union of the components through the points x_i. The variety X is h-tangentially weakly defective if $\dim \Gamma_h(A) > 0$.

Let $H \subset \mathbb{P}^N$ be a general hyperplane containing M_A. The h-contact locus of H is the union of the components through the x_i of $\operatorname{Sing}(H \cap X)$. The variety X is h-weakly defective if this locus has positive dimension.

Since every hyperplane containing M_A is tangent to X at every point of $\Gamma_h(A)$, one has

$$h\text{-tangentially weakly defective} \Rightarrow h\text{-weakly defective}.$$

The converse is false. For instance, let

$$X = \mathbb{P}^1 \times \mathbb{P}^m \subset \mathbb{P}^{2m+1}$$

be the Segre variety, with $m \geq 2$. At a general point $x = ([u], [v])$, a general hyperplane tangent to X has contact locus

$$\{[u]\} \times \mathbb{P}^{m-1},$$

so X is 1-weakly defective. However, $T_y X \subseteq T_x X$ only for $y = x$, hence X is not 1-tangentially weakly defective.

Proposition 1.4 (Structure of the tangential contact locus)**.** *Assume that X is not h-defective. Then the general tangential contact locus $\Gamma_h(A)$ is equidimensional. Moreover, if $\Gamma_h(A)$ is reducible, then it has exactly h irreducible components of the same dimension, and after reordering*

$$\Gamma_h(A) = \Gamma_1 \cup \cdots \cup \Gamma_h, \qquad x_i \in \Gamma_i.$$

In this case the general point x_i lies on only one component. Following the standard terminology, $\Gamma_h(A)$ is of type I if it is irreducible, and of type II if it is reducible.

This dichotomy is useful in applications. Type I means that the h general points lie on a common positive-dimensional tangential contact locus. Type II means that the contact is split into h congruent pieces, one through each marked point. These structural facts are due to the theory of weakly defective varieties and contact loci developed by Chiantini–Ciliberto and refined in later work [CC02, CC06, BBC18].

Remark 1.5 (Terracini and linear systems for Veronese varieties)**.** Let $V_d^n = \nu_d(\mathbb{P}^n) \subset \mathbb{P}^N$, with $N = \binom{n+d}{d} - 1$, and let

$$A = \{p_1, \dots, p_h\} \subset \mathbb{P}^n$$

be general. Hyperplanes in $\mathbb{P}^N$ containing

$$\langle T_{\nu_d(p_1)} V_d^n, \dots, T_{\nu_d(p_h)} V_d^n \rangle$$

correspond to degree d hypersurfaces in $\mathbb{P}^n$ singular at the points of A, that is, to

$$H^0(\mathbb{P}^n, \mathcal{I}_{2A}(d)).$$

Indeed, a hyperplane in $\mathbb{P}^N = \mathbb{P}(H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(d))^\vee)$ is a degree d form F, and containing $T_{\nu_d(p)} V_d^n$ is the same as imposing

$$F(p) = 0, \qquad \frac{\partial F}{\partial x_0}(p) = \cdots = \frac{\partial F}{\partial x_n}(p) = 0.$$

Thus

$$\dim \langle T_{\nu_d(p_1)} V_d^n, \dots, T_{\nu_d(p_h)} V_d^n \rangle = N - h^0(\mathbb{P}^n, \mathcal{I}_{2A}(d)).$$

Therefore V_d^n is h-defective precisely when

$$h^0(\mathbb{P}^n, \mathcal{I}_{2A}(d)) > \max\left\{0, \binom{n+d}{d} - h(n+1)\right\}.$$

This is the standard translation between secants of Veronese varieties and interpolation by hypersurfaces with assigned double points.

Example 1.6 (The quadratic Veronese surface)**.** Consider

$$V_2^2 = \nu_2(\mathbb{P}^2) \subset \mathbb{P}^5.$$

For $h = 2$, the expected dimension of $\mathbb{S}\mathrm{ec}_2(V_2^2)$ is

$$\min\{2(2+1) - 1, 5\} = 5.$$

Let $p, q \in \mathbb{P}^2$ be two general points and let $L = \langle p, q \rangle$ be the line through them. The double line $2L$ is a conic singular along L, hence singular at both p and q. Thus

$$H^0(\mathbb{P}^2, \mathcal{I}_{2p+2q}(2)) = \langle L^2 \rangle.$$

Two general double points do not impose independent conditions on plane conics. By Terracini,

$$\dim \langle T_{\nu_2(p)} V_2^2, T_{\nu_2(q)} V_2^2 \rangle = 4,$$

and $\mathbb{S}\mathrm{ec}_2(V_2^2)$ is defective.

The same defect is visible projectively. The image $\nu_2(L)$ is a plane conic

$$C = \nu_2(L) \subset \langle \nu_2(L) \rangle \simeq \mathbb{P}^2.$$

The tangent line to C at $\nu_2(p)$ is contained in $T_{\nu_2(p)} V_2^2$, and the tangent line to C at $\nu_2(q)$ is contained in $T_{\nu_2(q)} V_2^2$. Since two tangent lines to a plane conic meet, the two tangent planes to V_2^2 meet in one point. In coordinates, if $p = [u]$ and $q = [v]$, then

$$T_{[u^2]} V_2^2 = \mathbb{P}(u \cdot \mathbb{C}^3), \qquad T_{[v^2]} V_2^2 = \mathbb{P}(v \cdot \mathbb{C}^3),$$

and

$$T_{[u^2]} V_2^2 \cap T_{[v^2]} V_2^2 = \mathbb{P}(u \cdot v).$$

This intersection lowers the dimension of the tangent span by one.

Now write coordinates on $\mathbb{P}^5 = \mathbb{P}(\mathrm{Sym}^2 \mathbb{C}^3)$ as

$$[z_{00} : z_{01} : z_{02} : z_{11} : z_{12} : z_{22}]$$

and set

$$M_z = \begin{pmatrix} z_{00} & z_{01} & z_{02} \\ z_{01} & z_{11} & z_{12} \\ z_{02} & z_{12} & z_{22} \end{pmatrix}.$$

The Veronese surface is the locus of symmetric matrices of rank one, while its second secant variety is the locus of symmetric matrices of rank at most two. Hence

$$\mathbb{S}\mathrm{ec}_2(V_2^2) = \{[z] \in \mathbb{P}^5 \mid \det(M_z) = 0\}.$$

Explicitly,

$$z_{00} z_{11} z_{22} + 2 z_{01} z_{02} z_{12} - z_{00} z_{12}^2 - z_{11} z_{02}^2 - z_{22} z_{01}^2 = 0.$$

Thus $\mathbb{S}\mathrm{ec}_2(V_2^2) \subset \mathbb{P}^5$ is a cubic hypersurface.

Proposition 1.7 (Contact loci and identifiability)**.** *Assume that the abstract h-secant map π_h^X is generically finite. If X is not h-weakly defective, then X is h-identifiable. In particular, if X is not h-tangentially weakly defective, then X is h-identifiable.*

Here is the geometric reason for the tangential statement. Assume that π_h^X is generically finite but not birational. Fix a general decomposition

$$A = \{x_1, \dots, x_h\} \subset X$$

and let $p \in \langle A \rangle$ be general. A second decomposition

$$B_p = \{y_1(p), \dots, y_h(p)\} \subset X$$

gives

$$p \in \langle B_p \rangle.$$

Terracini's lemma applied to both decompositions gives

$$\langle T_{x_1}X, \dots, T_{x_h}X\rangle = \langle T_{y_1(p)}X, \dots, T_{y_h(p)}X\rangle.$$

As p moves in the general $(h-1)$-plane $\langle A\rangle$, the second decompositions cannot remain fixed. Hence the points $y_j(p)$ sweep a positive-dimensional locus whose tangent spaces are contained in the fixed span $\langle T_{x_1}X, \dots, T_{x_h}X\rangle$. Thus non-identifiability forces a positive-dimensional tangential contact locus.

The main implications can be summarized as follows:

$$h\text{-defective} \Rightarrow h\text{-tangentially weakly defective} \Rightarrow h\text{-weakly defective},$$

and, in the generically finite range,

$$\text{not } h\text{-weakly defective} \Rightarrow \text{not } h\text{-tangentially weakly defective} \Rightarrow h\text{-identifiable} \Rightarrow \text{not } h\text{-defective}.$$

Remark 1.8 (Non-reversibility)**.** None of the relevant converses holds in general.

(i) The converse of

$$h\text{-defective} \Rightarrow h\text{-tangentially weakly defective}$$

is false. The Veronese surface $V_6^2 \subset \mathbb{P}^{27}$ is not 9-defective, but it is 9-tangentially weakly defective. The general ternary sextic of rank 9 has exactly two Waring decompositions, and the tangential contact locus is the elliptic curve induced by the unique plane cubic through the nine points of a decomposition [COV17].

(ii) The converse of

$$h\text{-tangentially weakly defective} \Rightarrow h\text{-weakly defective}$$

is false, as shown above by $X = \mathbb{P}^1 \times \mathbb{P}^m \subset \mathbb{P}^{2m+1}$ for $h = 1$ and $m \geq 2$.

(iii) The converse of

$$\text{not } h\text{-tangentially weakly defective} \Rightarrow h\text{-identifiable}$$

is false. The Grassmannian $\mathrm{Gr}(\mathbb{P}^2, \mathbb{P}^7)$ in its Plucker embedding is 3-tangentially weakly defective, but the general point of $\mathbb{S}\mathrm{ec}_3(\mathrm{Gr}(\mathbb{P}^2, \mathbb{P}^7))$ is identifiable [BV18].

(iv) The converse of

$$h\text{-identifiable} \Rightarrow \text{not } h\text{-defective}$$

would say that nondefectivity implies identifiability, and is false. Again $V_6^2 \subset \mathbb{P}^{27}$ with $h = 9$ is nondefective, but a general point of $\mathbb{S}\mathrm{ec}_9(V_6^2)$ has two decompositions [COV17].

Thus defectivity is a dimensional failure, tangential weak defectivity is a tangential contact phenomenon, and weak defectivity is a hyperplane-contact phenomenon. Non-identifiability forces tangential contact in the generically finite range, but contact does not automatically produce non-identifiability.

2. Veronese varieties

Let $V_d^n = \nu_d(\mathbb{P}^n) \subset \mathbb{P}^N$, where $N = \binom{n+d}{d} - 1$. A point of $\mathbb{S}\mathrm{ec}_h(V_d^n)$ corresponds to a degree d form with a Waring expression

$$F = \lambda_1 \ell_1^d + \dots + \lambda_h \ell_h^d. \tag{2.1}$$

Here $\ell_i \in H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))$ and $\lambda_i \in \mathbb{C}^*$. Thus h-identifiability of V_d^n means uniqueness of (2.1), up to permutation of the summands and rescaling of the linear forms.

The expected dimension of the h-th secant variety is

$$\operatorname{expdim} \mathbb{S}\mathrm{ec}_h(V_d^n) = \min\{h(n+1) - 1, N\}. \tag{2.2}$$

The Alexander–Hirschowitz theorem [AH95] says that, for $d \geq 3$, $\mathbb{S}\mathrm{ec}_h(V_d^n)$ has dimension (2.2), except in the following defective cases:

n	d	h
2	4	5
3	4	9
4	4	14
4	3	7

In these four cases $\mathbb{S}\mathrm{ec}_h(V_d^n)$ has dimension one less than expected. Notice that the exceptional cubic case is $(4, 3, 7)$: the statement that a general cubic in five variables has rank 8 is the corresponding statement about generic rank, since $\mathbb{S}\mathrm{ec}_7(V_3^4)$ is defective and does not fill $\mathbb{P}^{34}$.

For quadrics the situation is governed by symmetric matrices. The variety V_2^n parametrizes rank-one symmetric matrices, and $\mathbb{S}\mathrm{ec}_h(V_2^n)$ parametrizes symmetric matrices of rank at most h. Hence, for $2 \le h \le n$, its dimension is smaller than the naive secant expectation. Thus the full defective list, including quadrics, is

n	d	h
$n \ge 2$	2	$2 \le h \le n$
2	4	5
3	4	9
4	4	14
4	3	7

Therefore, outside this list, the secant varieties of Veronese varieties have the expected dimension. In particular, if $h(n+1)-1 \le N$, then outside the defective cases the abstract secant map is generically finite. However, generically finite does not mean birational. Thus nondefectivity is only the first numerical condition for identifiability, and it does not imply identifiability. The theorem of Chiantini, Ottaviani, and Vannieuwenhoven [COV17] gives the sharp subgeneric identifiability statement for Veronese varieties: a general symmetric tensor of subgeneric rank is identifiable, except for the three nonidentifiable cases

$$(n,d,h) = (2,6,9), \quad (3,4,8), \quad (5,3,9),$$

where the general tensor has exactly two decompositions.

Example 2.3 (The defective case $(n,d,h) = (2,4,5)$)**.** Consider the Veronese surface

$$V_4^2 = \nu_4(\mathbb{P}^2) \subset \mathbb{P}^{14}.$$

The expected dimension of its fifth secant variety is

$$\min\{5(\dim V_4^2 + 1) - 1, 14\} = \min\{14, 14\} = 14.$$

Thus one expects $\mathbb{S}\mathrm{ec}_5(V_4^2) = \mathbb{P}^{14}$, that is one expects a general ternary quartic to be a sum of five fourth powers of linear forms. The Alexander–Hirschowitz theorem says that this expectation fails: $\mathbb{S}\mathrm{ec}_5(V_4^2)$ is a hypersurface in $\mathbb{P}^{14}$ [AH95]. Hence the general ternary quartic has rank 6, not 5.

Let $A = \{p_1, \dots, p_5\} \subset \mathbb{P}^2$ be a general set of five points. By Terracini's lemma,

$$T_z \, \mathbb{S}\mathrm{ec}_5(V_4^2) = \langle T_{\nu_4(p_1)} V_4^2, \dots, T_{\nu_4(p_5)} V_4^2 \rangle$$

for $z \in \langle \nu_4(p_1), \dots, \nu_4(p_5) \rangle$ general. Therefore $\mathbb{S}\mathrm{ec}_5(V_4^2)$ is defective precisely when the span of these five tangent planes has dimension smaller than 14.

We translate this into a linear system on $\mathbb{P}^2$. Hyperplanes in $\mathbb{P}^{14}$ containing $T_{\nu_4(p_i)} V_4^2$ correspond to plane quartics singular at p_i. Hence hyperplanes containing all five tangent planes correspond to the linear system

$$|\mathcal{I}_{2A}(4)|.$$

A double point in the plane imposes 3 linear conditions on quartics. Since $h^0(\mathbb{P}^2, \mathcal{O}_{\mathbb{P}^2}(4)) = 15$, five general double points should impose 15 independent conditions. The expected vector-space dimension of $H^0(\mathbb{P}^2, \mathcal{I}_{2A}(4))$ is therefore 0.

However, five general points in $\mathbb{P}^2$ lie on a unique conic Q. The double conic $2Q$ is a quartic singular along Q, hence in particular singular at the five points of A. Thus

$$2Q \in H^0(\mathbb{P}^2, \mathcal{I}_{2A}(4)),$$

and the expected vanishing fails. In fact, for A general, this is the only quartic singular at the five points, so

$$H^0(\mathbb{P}^2, \mathcal{I}_{2A}(4)) = \langle Q^2 \rangle.$$

Geometrically, the defect is caused by the positive-dimensional contact locus. The unique hyperplane containing the five tangent planes corresponds to the double conic $2Q$, and its contact locus on V_4^2 contains $\nu_4(Q)$. Thus the five prescribed tangent points do not impose independent tangent conditions since they are forced to lie on the same conic. This is the simplest geometric picture behind the defect:

$$\text{five double points on quartics} \quad \rightsquigarrow \quad \text{a double conic.}$$

In terms of Waring decompositions, this says that the family of sums

$$F = \lambda_1 \ell_1^4 + \cdots + \lambda_5 \ell_5^4$$

does not fill the whole space of ternary quartics. The missing dimension is detected by the unique quartic Q^2 apolar to the tangent data. Therefore the general ternary quartic cannot be written as a sum of five fourth powers; one needs six summands.

Remark 2.4 (Comparison with the other Alexander–Hirschowitz exceptions)**.** The other defective Veronese cases have the same broad nature, but not always the same literal geometry. In every case, Terracini's lemma reduces defectivity to the failure of a general set of double points to impose independent conditions on forms of degree d. Thus the common mechanism is the presence of an unexpected hypersurface singular at the assigned points, often coming from a low-degree variety through them [AH95].

For $(n,d,h) = (3,4,9)$, nine general points in $\mathbb{P}^3$ lie on a unique quadric surface Q. Then $2Q$ is a quartic surface singular along Q, hence singular at the nine points. This is directly parallel to the plane case $(2,4,5)$:

$$\text{nine double points on quartic surfaces} \quad \rightsquigarrow \quad \text{a double quadric surface.}$$

For $(n,d,h) = (4,4,14)$, fourteen general points in $\mathbb{P}^4$ lie on a special low-degree surface configuration producing an unexpected quartic singular at the points. The defect is again a failure of independent conditions for double points, but the geometry is no longer just the square of a unique hypersurface through the points.

For $(n,d,h) = (4,3,7)$, the defective secant is $\mathbb{S}\mathrm{ec}_7(V_3^4)$. This is the exceptional cubic case. Its geometry is more subtle than the double-conic and double-quadric pictures, since the unexpected cubic singular at seven general points is not simply the double of a hypersurface. This is one reason why the cubic part of the Alexander–Hirschowitz theorem requires a separate and more delicate argument.

Thus the answer is: yes, the philosophy is the same, since all these defects come from special linear systems with assigned double points; no, the concrete geometry is not always the same. The cases $(2,4,5)$ and $(3,4,9)$ are the cleanest ones, since the unexpected form is visibly a double conic or a double quadric.

The precise subgeneric statement for Veronese varieties is due to Chiantini, Ottaviani, and Vannieuwenhoven [COV17]. They prove that a general tensor of subgeneric rank is identifiable, with exactly three exceptions:

$$(n,d,h) = (2,6,9), \quad (3,4,8), \quad (5,3,9).$$

In each of these three cases there are exactly two Waring decompositions.

Remark 2.5 (Weak defectivity and the three exceptions)**.** The proof of [COV17] is formulated in terms of weak defectivity and tangential weak defectivity, following [CC02, BCO14]. Let $A = \{p_1, \ldots, p_h\} \subset \mathbb{P}^n$ be general. Hyperplanes tangent to V_d^n at $\nu_d(A)$ correspond, via apolarity, to degree d hypersurfaces in $\mathbb{P}^n$ singular at A. In the three exceptional cases the common singular locus is larger than A: it is a unique elliptic normal curve through A.

More precisely, for $(n,d,h) = (2,6,9)$ the unique sextic plane curve singular at 9 general points is twice the unique plane cubic through them. For $(n,d,h) = (3,4,8)$ the common singular locus is the elliptic normal quartic which is the complete intersection of the pencil of quadrics through the 8 points. For $(n,d,h) = (5,3,9)$ it is the unique elliptic normal sextic through the 9 points, classically related to Coble's construction and Gale transform [Cob22, Dol04, COV17].

The common mechanism is the following. Let $E \subset \mathbb{P}^n$ be the unique elliptic normal curve through the h points of a general decomposition. Then $\nu_d(E)$ is an elliptic normal curve of degree $d(n+1)$. In the three exceptional cases one has $d(n+1) = 2h$, hence

$$\nu_d(E) \subset \mathbb{P}^{2h-1}$$

is an elliptic normal curve of degree $2h$. A general point of $\mathbb{P}^{2h-1}$ has exactly two decompositions as a point in an h-secant $(h-1)$-plane to such a curve [CC06]. This is the geometric source of the two Waring decompositions.

Example 2.6 (The ternary sextic case $(2,6,9)$)**.** Let $F \in \mathbb{C}[x_0, x_1, x_2]_6$ be a general ternary sextic of rank 9, and assume that one decomposition is

$$F = \lambda_1 \ell_1^6 + \cdots + \lambda_9 \ell_9^6.$$

Let $A = \{[\ell_1], \ldots, [\ell_9]\} \subset \mathbb{P}^{2*}$. Since A is general, there is a unique plane cubic $E \subset \mathbb{P}^{2*}$ through A. Under the 6-Veronese embedding, E is mapped to

$$C = \nu_6(E) \subset \mathbb{P}^{17}.$$

Indeed $\mathcal{O}_E(6)$ has degree 18 and $h^0(E, \mathcal{O}_E(6)) = 18$, hence C is an elliptic normal curve of degree $18 = 2 \cdot 9$.

The span $\langle \nu_6(A) \rangle$ is a P^8 which is 9-secant to C. The point F belongs to this P^8. Since the 9-secant order of an elliptic normal curve of degree 18 in $\mathbb{P}^{17}$ is 2, a general such F belongs to exactly one other 9-secant P^8 of C. Thus there is a second divisor

$$B = \{[m_1], \ldots, [m_9]\} \subset E$$

such that

$$F = \mu_1 m_1^6 + \cdots + \mu_9 m_9^6.$$

This explains geometrically why the general ternary sextic of rank 9 is not identifiable.

The second decomposition can also be described by liaison. If B is the second decomposition, then $A\cap B=\emptyset$ and $A\cup B$ is a complete intersection of type $(3,6)$ in $\mathbb{P}^{2*}$ [CO21]. Thus A and B are residual to each other in the intersection of the cubic E with a plane sextic S:

$$A\cup B=E\cap S.$$

This description is special to the plane case. The other two exceptional cases have the same elliptic-normal-curve mechanism, but not the same literal plane liaison construction.

Remark 2.7. In all three cases the two decompositions are controlled by an elliptic normal curve $E\subset\mathbb{P}^n$ through the points of a decomposition, and $\nu_d(E)$ is an elliptic normal curve of degree $2h$ in $\mathbb{P}^{2h-1}$. Hence the two decompositions come from the secant order 2 of this curve. However, the concrete residual construction $A\cup B=E\cap S$ with E a plane cubic and S a plane sextic is specific to $(2,6,9)$. For $(3,4,8)$ the curve is the complete intersection of two quadrics in $\mathbb{P}^3$, while for $(5,3,9)$ it is the Coble elliptic normal sextic in $\mathbb{P}^5$.

Algorithm 2.8 (Laface–Massarenti–Rischter algorithm in the case $(2,6,9)$)**.** Let $F\in\mathbb{C}[x_0,x_1,x_2]_6$ be a general ternary sextic of rank 9. For $s\geq 0$, let $H^s_{\partial F}$ be the projective span of the order s partial derivatives of F.

The algorithm in [LMR23, Algorithm 5.2] specializes as follows.

(1) Compute $H^3_{\partial F}\subset\mathbb{P}(\mathbb{C}[x_0,x_1,x_2]_3)$. For a general rank 9 ternary sextic the middle catalecticant has one-dimensional kernel, hence $H^3_{\partial F}$ is a hyperplane in $\mathbb{P}^9$.

(2) Intersect this hyperplane with the cubic Veronese surface:
$$C_{6,2,9}=H^3_{\partial F}\cap V^2_3\subset\mathbb{P}^9.$$
Pulling back by $\nu_3:\mathbb{P}^{2*}\to V^2_3$ gives the unique plane cubic $E\subset\mathbb{P}^{2*}$ through the points of the decompositions.

(3) Embed E by quartics:
$$W=\nu_4(E)\subset V^2_4\subset\mathbb{P}^{14}.$$
Then project from $H^2_{\partial F}$:
$$\pi_{H^2_{\partial F}}:\mathbb{P}^{14}\dashrightarrow\mathbb{P}^8,\qquad \overline{W}=\pi_{H^2_{\partial F}}(W)\subset\mathbb{P}^8.$$

(4) Compute the scheme $Z\subset\mathbb{P}^8$ cut out by the linear and quadratic equations in the ideal of $\overline{W}$. These are the equations of degree $c<3$.

(5) Find the linear irreducible components $\overline{R}\simeq\mathbb{P}^2$ of Z. Each decomposition gives such a plane, since the projected span of the quartic powers
$$\overline{H}_\ell=\pi_{H^2_{\partial F}}\langle\ell_1^4,\dots,\ell_9^4\rangle$$
is a plane meeting $\overline{W}$ in the 9 projected points.

(6) For each such component $\overline{R}$, take its inverse image $R\subset\mathbb{P}^{14}$ under $\pi_{H^2_{\partial F}}$ and compute
$$R\cap V^2_4.$$
If this intersection consists of 9 reduced points $\ell_1^4,\dots,\ell_9^4$, recover the linear forms ℓ_i.

(7) Finally solve the linear system
$$F=\lambda_1\ell_1^6+\cdots+\lambda_9\ell_9^6$$
in the unknowns λ_i. The compatible components give the Waring decompositions of F. For a general ternary sextic of rank 9, the output consists of the two decompositions predicted by [COV17].

Remark 2.9 (How to find the second decomposition once one is known)**.** Assume that one decomposition $A=\{[\ell_1],\dots,[\ell_9]\}$ of a general ternary sextic F is already known. Chiantini and Ottaviani describe a direct liaison procedure [CO21, Remark 3.9]. Let $R=\mathbb{C}[x_0,x_1,x_2]$. The ideal of A has Hilbert–Burch resolution

$$0\longrightarrow R(-5)^3\xrightarrow{M}R(-4)^3\oplus R(-3)\longrightarrow I_A\longrightarrow 0,$$

where the degree matrix of M is

$$\begin{pmatrix}1&1&1\\1&1&1\\1&1&1\\2&2&2\end{pmatrix}.$$

The unique cubic through A is obtained from the kernel of the middle catalecticant C^3_F. To construct the residual set B, erase the bottom row of M and add a column of three quadrics Q_1,Q_2,Q_3. The maximal minors of the resulting matrix M' generate the ideal of a set B linked to A in a complete intersection of type $(3,6)$.

The quadrics Q_1, Q_2, Q_3 are chosen so that the maximal minors of M' are apolar to F. This is a linear condition in the coefficients of the Q_i's. Once such Q_i's are found, the residual scheme B gives the second decomposition of F, and the coefficients are obtained by solving

$$F = \mu_1 m_1^6 + \cdots + \mu_9 m_9^6,$$

where $B = \{[m_1], \ldots, [m_9]\}$.

2.9. **Generic rank: Galuppi–Mella, Hilbert, and Sylvester.** Let $V_d^n = \nu_d(\mathbb{P}^n) \subset \mathbb{P}^N$, with $N = \binom{n+d}{d} - 1$. Denote by $r_{\text{gen}}(n,d)$ the rank of a general degree d form in $n+1$ variables. Generic identifiability can hold only in the perfect case

$$\binom{n+d}{d} = r_{\text{gen}}(n,d)(n+1),$$

since otherwise the abstract secant map has positive-dimensional general fiber. Outside the Alexander–Hirschowitz defective cases, this number is the expected generic rank.

Galuppi and Mella prove that the perfect cases which are actually identifiable are exactly the classical ones [GM19]. More precisely, a general form of degree d in $n+1$ variables is identifiable at the generic rank if and only if

$$(n, d, r_{\text{gen}}) = (1, 2k-1, k), \quad (3,3,5), \quad (2,5,7).$$

The first family is the binary odd degree case, classically treated by Sylvester. The second case is Sylvester's pentahedral theorem: a general cubic surface in $\mathbb{P}^3$ is a sum of five cubes in a unique way. The third case is Hilbert's theorem: a general ternary quintic is a sum of seven fifth powers in a unique way [Syl04, Hil88, IK99].

The birational idea in [GM19] is the following. If a general form F is identifiable at the generic rank r, then the abstract r-secant map of V_d^n is birational. By the tangential projection criterion of Mella [Mel06], the projection of V_d^n from the span of $r-1$ general tangent spaces must be birational. For the Veronese embedding, this tangential projection is described on $\mathbb{P}^n$ by the linear system

$$\mathcal{L}_{n,d}(2^{r-1}) \subset |\mathcal{O}_{\mathbb{P}^n}(d)|$$

of degree d hypersurfaces with $r-1$ general double points. Thus generic identifiability forces $\mathcal{L}_{n,d}(2^{r-1})$ to define a Cremona transformation of $\mathbb{P}^n$. Galuppi–Mella classify exactly when this happens. The Cremona cases are

$$(n, d, r-1) = (1, 2k-1, k-1), \quad (3,3,4), \quad (2,5,6),$$

and they give precisely the three identifiable generic-rank cases above.

Remark 2.10 (Why the result is so restrictive)**.** The numerical perfectness condition only says that the general fiber of the secant map is expected to be finite. Galuppi–Mella show that degree one is much more rigid. The map attached to $\mathcal{L}_{n,d}(2^{r-1})$ must be Cremona, and this almost never occurs. Thus the scarcity of canonical Waring forms at the generic rank is a birational phenomenon.

We now recall the reconstruction algorithms of Laface–Massarenti–Rischter [LMR23]. For a form $F \in S^d V$, let $H^s_{\partial F} \subset \mathbb{P}(S^{d-s}V)$ be the projective span of the order s partial derivatives of F.

Algorithm 2.11 (Sylvester pentahedral reconstruction)**.** Let $F \in \mathbb{C}[x_0, x_1, x_2, x_3]_3$ be a general cubic. Then $r_{\text{gen}} = 5$, and Sylvester's theorem gives a unique expression

$$F = \lambda_1 L_1^3 + \cdots + \lambda_5 L_5^3.$$

The algorithm in [LMR23, Algorithm 3.5] reconstructs the five planes $\Pi_i = \{L_i = 0\} \subset \mathbb{P}^{3*}$ as follows.

First compute the span of first partial derivatives

$$H^1_{\partial F} \subset \mathbb{P}(\mathbb{C}[x_0, x_1, x_2, x_3]_2).$$

For a general cubic, $H^1_{\partial F} \simeq \mathbb{P}^3$. Then compute

$$Z = H^1_{\partial F} \cap \mathbb{S}\mathrm{ec}_2(V_2^3).$$

The scheme Z consists of $\binom{5}{3} = 10$ reduced points. Geometrically, these points correspond to the ten triple intersections

$$\Pi_i \cap \Pi_j \cap \Pi_k, \qquad 1 \le i < j < k \le 5.$$

Indeed, if $\xi \in \mathbb{P}^{3*}$ lies on three of the planes Π_i, then the directional derivative $D_\xi F$ is a linear combination of the two remaining squares, hence $D_\xi F \in \mathbb{S}\mathrm{ec}_2(V_2^3)$.

Now each plane Π_i contains exactly $\binom{4}{2} = 6$ of the ten triple intersection points. Therefore one searches among the subsets of six points of Z those which span a plane. By [LMR23, Lemma 3.4], the five planes obtained in this way are exactly $\Pi_1, \dots, \Pi_5$. Once the linear forms L_i are recovered, the coefficients λ_i are found by solving the linear system

$$F = \lambda_1 L_1^3 + \cdots + \lambda_5 L_5^3.$$

Algorithm 2.12 (Hilbert's ternary quintic reconstruction)**.** Let $F \in \mathbb{C}[x_0, x_1, x_2]_5$ be a general ternary quintic. Then $r_{\text{gen}} = 7$, and Hilbert's theorem gives a unique expression

$$F = \lambda_1 L_1^5 + \cdots + \lambda_7 L_7^5.$$

This case is not covered by the previous secant-intersection method. Laface–Massarenti–Rischter use instead a projection from the space of derivatives [LMR23, Proposition 4.1 and Algorithm 4.2].

Compute the span of second partial derivatives

$$H^2_{\partial F} \subset \mathbb{P}(\mathbb{C}[x_0, x_1, x_2]_3).$$

For a general quintic, $H^2_{\partial F} \simeq \mathbb{P}^5$, and it is disjoint from the cubic Veronese surface $V_3^2 \subset \mathbb{P}^9$. Hence the projection

$$\pi_{H^2_{\partial F}} : \mathbb{P}^9 \dashrightarrow \mathbb{P}^3$$

restricts to a morphism on V_3^2. Let

$$\overline{V} = \pi_{H^2_{\partial F}}(V_3^2) \subset \mathbb{P}^3.$$

If $F = \sum_{i=1}^7 \lambda_i L_i^5$, then the seven points $L_i^3 \in V_3^2$ span a $\mathbb{P}^6$, say

$$H_L = \langle L_1^3, \dots, L_7^3 \rangle.$$

Moreover $H^2_{\partial F}$ is a hyperplane in H_L. Therefore the projection contracts H_L to a point $p \in \mathbb{P}^3$. Since H_L meets V_3^2 at the seven points L_i^3, the point p is a point of multiplicity 7 of the projected surface $\overline{V}$.

The algorithm is therefore:

(1) compute $\overline{V} = \pi_{H^2_{\partial F}}(V_3^2) \subset \mathbb{P}^3$;
(2) compute the reduced subscheme $S_7 \subset \overline{V}$ of points of multiplicity 7;
(3) for a general quintic, $S_7 = \{p\}$ is a single point;
(4) form the linear span $H_p = \langle H^2_{\partial F}, p \rangle \subset \mathbb{P}^9$;
(5) compute
$$H_p \cap V_3^2 = \{L_1^3, \dots, L_7^3\};$$
(6) recover $L_1, \dots, L_7$ and solve
$$F = \lambda_1 L_1^5 + \cdots + \lambda_7 L_7^5.$$

The uniqueness of the multiplicity-7 point is the computational trace of Hilbert's canonical form.

Remark 2.13. Both algorithms recover the unknown linear forms from auxiliary geometry in a lower degree Veronese variety. In the Sylvester case, the key object is the finite set $H^1_{\partial F} \cap \mathbb{S}\text{ec}_2(V_2^3)$, which records the incidence geometry of the pentahedron. In the Hilbert case, the key object is the unique multiplicity-7 point of the projection of V_3^2 from $H^2_{\partial F}$. Thus both algorithms use spaces of derivatives, but the geometric signatures are different: incidence points for Sylvester, a high-multiplicity projected singularity for Hilbert.

3. Bronowski's conjecture

Assume that $X \subset \mathbb{P}^{hn+h-1}$ is irreducible, nondegenerate, and of dimension n. Then $\mathbb{S}\text{ec}_h(X)$ is expected to fill the ambient space. Thus the abstract h-secant map is expected to be generically finite, and the general $(h-1)$-tangential projection

$$\tau^X_{h-1} : X \dashrightarrow \mathbb{P}^n$$

has target $\mathbb{P}^n$.

Bronowski's philosophy is that identifiability should be detected by this tangential projection. In modern language, this can be stated as follows.

Conjecture 3.1 (Bronowski)**.** *Let $X \subset \mathbb{P}^{hn+h-1}$ be irreducible and nondegenerate of dimension n. Then X is h-identifiable if and only if the general $(h-1)$-tangential projection $\tau^X_{h-1} : X \dashrightarrow \mathbb{P}^n$ is birational.*

The implication from identifiability to birationality is known, and it is best seen by contraposition. Fix general points $x_1, \ldots, x_{h-1} \in X$ and set

$$\Lambda = \langle T_{x_1}X, \ldots, T_{x_{h-1}}X \rangle.$$

Assume that the tangential projection

$$\tau^X_{h-1} : X \dashrightarrow \mathbb{P}^n$$

is not birational. Then two general points $x, y \in X$ have the same image. This means that the line $\langle x, y \rangle$ meets the center Λ, or, equivalently,

$$\langle \Lambda, x \rangle = \langle \Lambda, y \rangle.$$

Now consider the abstract h-secant map near the point corresponding to the h-tuple

$$(x_1, \ldots, x_{h-1}, x).$$

Its differential is controlled by Terracini's lemma: the tangent directions obtained by moving the first $h-1$ points span

$$\Lambda = \langle T_{x_1}X, \ldots, T_{x_{h-1}}X \rangle,$$

while moving the last point x contributes T_xX. Thus the infinitesimal image is

$$\langle \Lambda, T_xX \rangle.$$

The equality $\langle \Lambda, x \rangle = \langle \Lambda, y \rangle$ says that, after projecting from Λ, the points x and y give the same point of the target. Hence, infinitesimally, the h-secant spaces obtained from

$$(x_1, \ldots, x_{h-1}, x) \qquad \text{and} \qquad (x_1, \ldots, x_{h-1}, y)$$

meet the same general fiber of the projection from Λ.

Therefore the abstract secant map cannot be generically one-to-one: a general point of $\mathbb{S}\mathrm{ec}_h(X)$ is obtained from two distinct h-tuples, one with last point x and one with last point y. Thus, by contraposition, if X is h-identifiable, the general $(h-1)$-tangential projection must be birational. The converse is the delicate direction.

Ciliberto and Russo revisited Bronowski's assertion in the broader setting of varieties with minimal secant degree [CR06]. In their language one considers q-secant varieties and the number of apparent q-secant $(q-1)$-planes. A variety X is an MA_q-variety if its q-secant variety has minimal degree and if a general point of $\mathbb{S}\mathrm{ec}_q(X)$ lies on a unique q-secant $(q-1)$-plane, that is, if X is q-identifiable. The Ciliberto–Russo form of the generalized Bronowski principle predicts that, under the minimal-degree hypothesis, birationality of a general $(q-1)$-tangential projection should force this uniqueness. Thus the modern version of Bronowski's conjecture links three objects:

$$\text{identifiability} \quad \longleftrightarrow \quad \text{birational tangential projections} \quad \longleftrightarrow \quad \text{minimal secant degree.}$$

In this form the conjecture is true in several classical situations, for instance for curves and smooth surfaces in suitable ranges, but it is false in general. The obstruction is not visible from the birationality of τ^X_{h-1} alone: one must also control the Gauss map.

Theorem 3.2 (Massarenti–Mella counterexamples)**.** *Let $\Gamma \subset \mathbb{P}^N$ be a rational normal curve of degree $N \geq 7$, and let $r \geq 2$ be such that*

$$h = \frac{N+1}{2r}$$

is an integer. Then $X = \mathbb{S}\mathrm{ec}_r(\Gamma) \subset \mathbb{P}^N$ gives a counterexample to Bronowski's conjecture. Moreover, these examples also contradict the minimal-degree strengthening suggested by the Ciliberto–Russo point of view.

Idea of the construction. Set $X = \mathbb{S}\mathrm{ec}_r(\Gamma)$, where $\Gamma \subset \mathbb{P}^N$ is a rational normal curve and $N + 1 = 2rh$. Then

$$\dim X = 2r - 1 =: n, \qquad N = h(n+1) - 1,$$

so $X \subset \mathbb{P}^{hn+h-1}$ is in the Bronowski range. A general point $x \in X$ lies on a unique r-secant $(r-1)$-plane to Γ, say

$$x \in \langle a_1, \ldots, a_r \rangle, \qquad a_i \in \Gamma.$$

By Terracini's lemma,

$$T_xX = \langle T_{a_1}\Gamma, \ldots, T_{a_r}\Gamma \rangle.$$

In particular, T_xX depends only on the r points $a_1, \ldots, a_r$, and not on the point $x \in \langle a_1, \ldots, a_r \rangle$. Hence the Gauss map of X contracts the general r-secant $(r-1)$-plane of Γ. This is the basic tangential degeneracy of the construction.

Let $x_1, \dots, x_{h-1} \in X$ be general. Write

$$x_i \in \langle a_{i,1}, \dots, a_{i,r} \rangle, \qquad a_{i,j} \in \Gamma.$$

Then the center of the general $(h-1)$-tangential projection of X is

$$\Lambda = \langle T_{x_1}X, \dots, T_{x_{h-1}}X \rangle = \langle T_{a_{i,j}}\Gamma \mid i = 1, \dots, h-1,\ j = 1, \dots, r \rangle.$$

Thus the tangential projection of X is induced by the projection of Γ from the span of the tangent lines at the $r(h-1)$ points $a_{i,j}$. For a rational normal curve, projecting from these tangent lines lowers the degree by $2r(h-1)$, hence one obtains a rational normal curve

$$\Gamma' \subset \mathbb{P}^{N-2r(h-1)} = \mathbb{P}^{2r-1}.$$

Moreover, the image of $X = \mathbb{S}\mathrm{ec}_r(\Gamma)$ is $\mathbb{S}\mathrm{ec}_r(\Gamma')$. But $\mathbb{S}\mathrm{ec}_r(\Gamma') = \mathbb{P}^{2r-1}$, and the general point of $\mathbb{P}^{2r-1}$ has a unique expression as a point of an r-secant $(r-1)$-plane to Γ', by the classical Sylvester theorem for binary forms of degree $2r-1$. Therefore

$$\tau^X_{h-1} : X \dashrightarrow \mathbb{P}^{2r-1}$$

is birational. This is why the Bronowski test predicts h-identifiability.

The prediction fails for a completely different reason. A general point $p \in \mathbb{P}^N$ has, again by Sylvester, a unique expression as a point of an rh-secant $(rh-1)$-plane to Γ:

$$p \in \langle b_1, \dots, b_{rh} \rangle, \qquad b_i \in \Gamma.$$

However, to express p as a point of an h-secant $(h-1)$-plane to $X = \mathbb{S}\mathrm{ec}_r(\Gamma)$, one may partition the set

$$\{b_1, \dots, b_{rh}\}$$

into h subsets of cardinality r. If

$$\{b_1, \dots, b_{rh}\} = B_1 \sqcup \dots \sqcup B_h, \qquad |B_i| = r,$$

then each $\langle B_i \rangle$ is an r-secant $(r-1)$-plane to Γ, hence it is contained in X. Grouping the summands of the binary decomposition of p according to this partition gives points

$$z_i \in \langle B_i \rangle \subset X, \qquad i = 1, \dots, h,$$

such that

$$p \in \langle z_1, \dots, z_h \rangle.$$

Different partitions give different h-secant decompositions of p with respect to X. Thus X is not h-identifiable.

The key point is that the tangential projection is blind to this phenomenon. It detects the birationality of the induced projection on the quotient rational normal curve, but it does not detect the fact that the Gauss map of X is degenerate. Along each general r-secant $(r-1)$-plane of Γ, the embedded tangent space to X is constant. Thus X is 1-tangentially weakly defective, and this tangential degeneracy is exactly what allows the same general point of $\mathbb{P}^N$ to be obtained from several different groupings of the unique binary decomposition.

This shows that the Gauss map hypothesis cannot be omitted. It also explains why the failure of Bronowski's conjecture is not a numerical accident: it is caused by a genuine tangential pathology.

The corrected statement proved in [MM24] requires nondegeneracy of the Gauss map and uses one higher secant variety.

Theorem 3.3. [MM24, Theorem 1.5] *Let $X \subset \mathbb{P}^N$ be irreducible and nondegenerate of dimension n. Let $h \geq 1$ and assume that:*

(i) $(h+1)n + h \leq N$;
(ii) *X has nondegenerate Gauss map;*
(iii) *X is not $(h+1)$-defective.*

Then X is h-identifiable.

Proof idea. We argue by contradiction. By assumption (i), the expected dimension of $\mathbb{S}\mathrm{ec}_{h+1}(X)$ is

$$(h+1)(n+1) - 1 = (h+1)n + h \leq N.$$

Since X is not $(h+1)$-defective, the abstract $(h+1)$-secant map is generically finite. Thus, to get a contradiction, it is enough to prove that π^X_{h+1} is of fiber type, or, by Terracini, that a general span of $h+1$ tangent spaces has dimension smaller than $(h+1)n + h$.

Assume that X is not h-identifiable. Let $x_1, \dots, x_h \in X$ be general and set

$$\Lambda = \langle T_{x_1}X, \dots, T_{x_h}X \rangle.$$

Non-identifiability forces a positive-dimensional h-tangential contact locus. Indeed, a general point $p \in \langle x_1, \ldots, x_h \rangle$ has another decomposition

$$p \in \langle y_1(p), \ldots, y_h(p) \rangle.$$

Terracini's lemma applied to the two decompositions gives

$$\Lambda = \langle T_{y_1(p)}X, \ldots, T_{y_h(p)}X \rangle.$$

As p moves in $\langle x_1, \ldots, x_h \rangle$, the points $y_j(p)$ cannot all remain fixed. Hence they sweep a positive-dimensional locus whose tangent spaces are contained in Λ. Thus the general h-tangential contact locus

$$\Gamma_h = \overline{\{z \in X_{\mathrm{reg}} \mid T_z X \subseteq \Lambda\}}$$

is nontrivial.

The hypothesis on the Gauss map is used at the opposite end: since the Gauss map is nondegenerate, X is not 1-tangentially weakly defective. Thus the contact phenomenon begins only at some higher order; it is not caused by tangent spaces being constant along positive-dimensional families.

There are two cases.

First assume that Γ_h is an irreducible curve. The structural results on contact loci say that the h-secant degree of X is then detected on Γ_h. Hence Γ_h is not h-identifiable. For curves, Bronowski's conjecture is true, so the corresponding tangential projection of Γ_h is not birational. Since the ambient h-tangential projection restricts to this projection on Γ_h, the map

$$\tau_h^X : X \dashrightarrow \mathbb{P}^{N-\dim \Lambda - 1}$$

has positive-dimensional general fibers. Lemma 3.2 of [MM24] compares these fibers with the fibers of the abstract $(h+1)$-secant map. Therefore π_{h+1}^X is of fiber type, contradicting the nondefectivity of $\mathbb{S}\mathrm{ec}_{h+1}(X)$.

We may therefore assume that Γ_h is not an irreducible curve. If X were h-defective, then π_h^X would already be of fiber type, and so π_{h+1}^X would be of fiber type as well. This is impossible since X is not $(h+1)$-defective. Hence X is not h-defective, and

$$\dim \Lambda = h(n+1) - 1.$$

Now add one more general point $x \in X$ and consider the $(h+1)$-tangential contact locus associated with

$$x_1, \ldots, x_h, x.$$

Since Γ_h is not an irreducible curve, the structural dichotomy for contact loci gives enough room to move some of the marked points inside components of this larger contact locus. Along such curves, tangent spaces specialize to linear subspaces contained in the limiting tangent spans. In this way one produces, by semicontinuity, positive-dimensional subspaces of tangent spaces which remain constrained by the span

$$\langle T_{x_1}X, \ldots, T_{x_h}X, T_x X \rangle.$$

The key point is the following. If $T_x X$ were disjoint from Λ for a general point $x \in X$, then

$$\dim \langle \Lambda, T_x X \rangle = \dim \Lambda + n + 1 = (h+1)n + h,$$

that is, the expected dimension. The limiting tangent-space construction forces the opposite: for a general $x \in X$ one obtains

$$T_x X \cap \Lambda \neq \emptyset.$$

Therefore

$$\dim \langle \Lambda, T_x X \rangle < (h+1)n + h.$$

By Terracini's lemma this means that $\mathbb{S}\mathrm{ec}_{h+1}(X)$ is defective, contradicting assumption (iii). Both alternatives lead to a contradiction. Hence X must be h-identifiable.

Remark 3.4 (Sharpness)**.** Both extra hypotheses are essential. If the Gauss map is degenerate, the secant varieties of rational normal curves in Theorem 3.2 give counterexamples to Bronowski's conjecture and to its minimal-degree strengthening. The numerical inequality is also sharp: $V_6^2 \subset \mathbb{P}^{27}$ is not 10-defective, but a general point of $\mathbb{S}\mathrm{ec}_9(V_6^2)$ has exactly two Waring decompositions. Hence one cannot replace $(h+1)n + h \leq N$ by a weaker bound in Theorem 3.3.

4. From secants to neurovarieties

Let $\mathbf{n} = (n_0, \ldots, n_L)$ be a width vector, with $L \geq 2$, and let $\mathbf{d} = (d_1, \ldots, d_{L-1})$, with $d_i \geq 2$. A polynomial neural network has weight matrices

$$W_i = (\alpha^i_{r,s}) \in k^{n_i \times n_{i-1}}, \qquad i = 1, \ldots, L,$$

and activation maps

$$\sigma_i(u_1, \ldots, u_{n_i}) = (u_1^{d_i}, \ldots, u_{n_i}^{d_i}), \qquad i = 1, \ldots, L-1.$$

Its output is

$$F = W_L \circ \sigma_{L-1} \circ W_{L-1} \circ \cdots \circ W_2 \circ \sigma_1 \circ W_1 : k^{n_0} \longrightarrow k^{n_L}.$$

The coordinates of F are homogeneous forms of degree

$$D = \prod_{i=1}^{L-1} d_i.$$

The neurovariety $\mathcal{V}_{\mathbf{n},\mathbf{d}}$ is the Zariski closure of the image of the coefficient map sending the weights to the coefficients of these output forms [MM25].

The naive expected dimension is the number of effective parameters after the standard projective normalizations:

$$\operatorname{expdim}(\mathcal{V}_{\mathbf{n},\mathbf{d}}) = \min\left\{\sum_{i=1}^{L} n_i(n_{i-1} - 1),\ N\right\},$$

where

$$N = n_L \binom{n_0 - 1 + D}{n_0 - 1} - n_L.$$

The subtraction by n_L reflects the fact that each output is considered projectively. The variety $\mathcal{V}_{\mathbf{n},\mathbf{d}}$ is defective if its dimension is smaller than the expected one.

The recursive construction. We write the construction layer by layer. At step zero we start with $\mathbb{P}^{n_0-1}$ and homogeneous coordinates $[x_0 : \cdots : x_{n_0-1}]$. Applying W_1 gives n_1 linear forms

$$L_j = \sum_{i=0}^{n_0-1} \alpha^1_{j,i} x_i, \qquad j = 0, \ldots, n_1 - 1.$$

After the first activation and the second weight matrix one gets

$$F_{2,j} = \sum_{i=0}^{n_1-1} \alpha^2_{j,i} L_i^{d_1}, \qquad j = 0, \ldots, n_2 - 1.$$

At the next step

$$F_{3,j} = \sum_{i=0}^{n_2-1} \alpha^3_{j,i} F_{2,i}^{d_2}, \qquad j = 0, \ldots, n_3 - 1,$$

and, in general,

$$F_{k,j} = \sum_{i=0}^{n_{k-1}-1} \alpha^k_{j,i} F_{k-1,i}^{d_{k-1}}, \qquad k = 2, \ldots, L.$$

Thus $F_{L,j}$ is a homogeneous polynomial of degree D on $\mathbb{P}^{n_0-1}$.

There are two geometric interpretations. At the first step, raising the linear forms to the d_1-th power gives points on the Veronese variety $V_{d_1}^{n_0-1}$. At later steps, the same polynomial may be viewed either as a high degree form in the original variables or as a linear form after applying a new Veronese embedding. This is the reason why the geometry of the network is controlled by iterated Veronese embeddings followed by linear projections.

Example 4.1. Consider the architecture

$$\mathbf{n} = (2, 2, 2, 1), \qquad \mathbf{d} = (2, 2).$$

Then

$$F = c_{0,0}\left(b_{0,0}(a_{0,0}x_0 + a_{0,1}x_1)^2 + b_{0,1}(a_{1,0}x_0 + a_{1,1}x_1)^2\right)^2 + c_{0,1}\left(b_{1,0}(a_{0,0}x_0 + a_{0,1}x_1)^2 + b_{1,1}(a_{1,0}x_0 + a_{1,1}x_1)^2\right)^2.$$

As a polynomial in x_0, x_1, this is a binary quartic, hence gives a rational map $s : \mathbb{A}^5 \dashrightarrow \mathbb{P}^4$ after a normalization of parameters. If instead one sets

$$z_0 = x_0^2, \qquad z_1 = x_0 x_1, \qquad z_2 = x_1^2,$$

then F becomes a quadratic polynomial $G(z_0, z_1, z_2)$, hence gives a rational map $f : \mathbb{A}^5 \dashrightarrow \mathbb{P}^5$. The two maps fit in the diagram

$$\begin{array}{ccc} & & \mathbb{P}^5 \\ & \overset{f}{\nearrow} & \Big\downarrow{\scriptstyle \pi} \\ \mathbb{A}^5 & \underset{s}{\dashrightarrow} & \mathbb{P}^4 \end{array}$$

where π is the linear projection from the point of $\mathbb{P}^5$ representing the quadratic relation

$$z_0 z_2 - z_1^2.$$

Thus the passage from the intermediate polynomial in the variables z_i to the original binary quartic is a linear projection. This is the basic mechanism repeated at all deeper levels.

The single-output case as a constrained secant construction. Assume now that $n_L = 1$. The last layer takes a linear combination of n_{L-1} points produced by the previous layers. Thus the last operation is secant-like.

When $L = 2$, one recovers the usual Veronese secant variety:

$$\mathcal{V}_{(n_0,n_1,1),(d_1)} = \mathbb{S}\mathrm{ec}_{n_1}\big(V_{d_1}^{n_0-1}\big).$$

For $L \geq 3$, the construction is not the whole secant variety of a Veronese: the n_{L-1} last points are constrained by all previous layers.

The geometry may be organized as follows. Let

$$V_1 = V_{d_1}^{n_0-1} \subset \mathbb{P}^{N_1}.$$

Inductively, apply the next Veronese embedding and then project linearly:

$$\begin{array}{ccccccc} \mathbb{P}^{n_0-1} \xrightarrow{\nu_{d_1}} & V_1 \subset \mathbb{P}^{N_1} & \xrightarrow{\nu_{d_2}} & V_2 \subset \mathbb{P}^{M_2} & \xrightarrow{\nu_{d_3}} \cdots \xrightarrow{\nu_{d_{L-1}}} & V_{L-1} \subset \mathbb{P}^{M_{L-1}} \\ & \Big\downarrow{\scriptstyle \pi_2} & & \Big\downarrow{\scriptstyle \pi_3} & & \Big\downarrow{\scriptstyle \pi_{L-1}} \\ & \widetilde{V}_2 \subset \mathbb{P}^{N_2} & \xrightarrow[\widetilde{\nu}_{d_3}]{} & \widetilde{V}_3 \subset \mathbb{P}^{N_3} & \longrightarrow \cdots \xrightarrow[\widetilde{\nu}_{d_{L-1}}]{} & \widetilde{V}_{L-1} \subset \mathbb{P}^{N_{L-1}}. \end{array}$$

Here the vertical arrows are the linear projections coming from the relations imposed by the previous Veronese stages. The variety $\widetilde{V}_{L-1} \subset \mathbb{P}^{N_{L-1}}$ is the final projected iterated Veronese seen by the last layer.

A point of the neurovariety is obtained by choosing, recursively, points

$$p_i^1 \in V_1, \qquad p_i^2 \in \widetilde{\nu}_{d_2}\big(\langle p_1^1, \ldots, p_{n_1}^1 \rangle\big), \quad \ldots, \quad p_i^{L-1} \in \widetilde{\nu}_{d_{L-1}}\big(\langle p_1^{L-2}, \ldots, p_{n_{L-2}}^{L-2} \rangle\big),$$

and then taking a point in the final span

$$\langle p_1^{L-1}, \ldots, p_{n_{L-1}}^{L-1} \rangle \subset \mathbb{P}^{N_{L-1}}.$$

Thus the single-output neurovariety is the image of a constrained secant construction:

$$\mathcal{V}_{\mathbf{n},\mathbf{d}} = \overline{\bigcup \langle p_1^{L-1}, \ldots, p_{n_{L-1}}^{L-1} \rangle} \subset \mathbb{P}^{N_{L-1}},$$

where the points p_i^{L-1} are not arbitrary points of $\widetilde{V}_{L-1}$, but must arise from the previous layers.

This can be encoded by an incidence diagram. Let $\mathcal{S}_{\mathbf{n},\mathbf{d}}$ be the parameter space of all admissible tuples

$$\Big((p_1^1, \ldots, p_{n_1}^1), \ldots, (p_1^{L-1}, \ldots, p_{n_{L-1}}^{L-1})\Big)$$

satisfying the recursive conditions above, and let

$$\mathcal{I}_{\mathbf{n},\mathbf{d}} = \Big\{ ((p_i^j), p) \mid p \in \langle p_1^{L-1}, \ldots, p_{n_{L-1}}^{L-1} \rangle \Big\}.$$

Then

$$\begin{array}{ccc} & \mathcal{I}_{\mathbf{n},\mathbf{d}} & \\ {\scriptstyle \varphi}\swarrow & & \searrow{\scriptstyle \psi} \\ \mathcal{S}_{\mathbf{n},\mathbf{d}} & & \mathbb{P}^{N_{L-1}} \end{array}$$

and $\mathcal{V}_{\mathbf{n},\mathbf{d}} = \overline{\psi(\mathcal{I}_{\mathbf{n},\mathbf{d}})}$. The map ψ is the analogue of the abstract secant map. Therefore the defectiveness problem for single-output neurovarieties becomes a constrained secant problem.

Since

$$\dim \mathcal{S}_{\mathbf{n},\mathbf{d}} = \sum_{i=1}^{L-1} n_i(n_{i-1}-1),$$

one has

$$\begin{aligned}\dim \mathcal{I}_{\mathbf{n},\mathbf{d}} &= \sum_{i=1}^{L-1} n_i(n_{i-1}-1) + (n_{L-1}-1)\\ &= \sum_{i=1}^{L} n_i(n_{i-1}-1),\end{aligned}$$

where $n_L = 1$. This gives the refined expected dimension

$$\text{(N1)} \qquad \operatorname{expdim}(\mathcal{V}_{\mathbf{n},\mathbf{d}}) = \min\left\{\sum_{i=1}^{L} n_i(n_{i-1}-1), \sum_{i=1}^{L-2} n_i(n_{i-1}-1) + \binom{n_{L-2}-1+d_{L-1}}{n_{L-2}-1}, N\right\}.$$

Here

$$N = \binom{n_0-1+D}{n_0-1} - 1.$$

The second term in (N1) appears since the last secant construction may fill the projective space spanned by the last Veronese stage, while the whole neurovariety still does not fill the ambient coefficient space.

Remark 4.2 (Identifiability in geometric language)**.** Permutations of hidden neurons and element-wise rescalings are the standard symmetries of polynomial neural networks. In the projective secant construction these symmetries are already quotiented out. Thus, for a general $F \in \mathcal{V}_{\mathbf{n},\mathbf{d}}$, global identifiability means that there is a unique admissible set

$$\{p_1^{L-1}, \dots, p_{n_{L-1}}^{L-1}\} \subset \widetilde{V}_{L-1}$$

such that

$$F \in \langle p_1^{L-1}, \dots, p_{n_{L-1}}^{L-1}\rangle.$$

This is the usual identifiability notion for secant varieties, adapted to the constrained family of secant spaces produced by the network.

A Terracini lemma for neurovarieties. The tangent theory mirrors Terracini's lemma, but with extra summands coming from the intermediate layers. Work in the single-output case $n_L = 1$. Let

$$F = \sum_{s=1}^{n_{L-1}} \alpha_{1,s}^{L} \, (p_s)^{d_{L-1}}$$

be a general output, where the p_s are the points produced at level $L-1$. The last layer contributes the usual secant tangent directions. The earlier layers contribute additional spaces, called level-wise normal spaces.

For $j = 1, \dots, L-2$, fix a point p_i^j at level j. Vary only the parameters in the i-th row of W_j, and push this variation through all later layers. The span of the resulting derivative columns is the level-j normal space

$$\mathcal{N}^j_{p_i^j} \subset \mathbb{P}^{N_{L-1}}.$$

The total level-j normal space is

$$\mathcal{N}^j = \langle \mathcal{N}^j_{p_1^j}, \dots, \mathcal{N}^j_{p_{n_j}^j}\rangle.$$

These spaces are called "normal" since, after pushing forward through the later Veronese and projection maps, they are generally not tangent to the last projected Veronese. They record the transversal motion induced by changing earlier layers.

Lemma 4.3 (Terracini lemma for neurovarieties)**.** *Let $n_L = 1$, and let $F \in \mathcal{V}_{\mathbf{n},\mathbf{d}}$ be a general network output. Let*

$$p_1^{L-1}, \dots, p_{n_{L-1}}^{L-1} \in \widetilde{V}_{L-1}$$

be the last-level points associated to F. Then

$$\text{(N2)} \qquad T_F \mathcal{V}_{\mathbf{n},\mathbf{d}} = \left\langle \mathcal{N}^1, \dots, \mathcal{N}^{L-2}, T_{p_1^{L-1}} \widetilde{V}_{L-1}, \dots, T_{p_{n_{L-1}}^{L-1}} \widetilde{V}_{L-1} \right\rangle.$$

The last summands in (N2) are the classical Terracini contribution for the last secant construction. The spaces $\mathcal{N}^1, \dots, \mathcal{N}^{L-2}$ are the new feature of deep networks. Thus nondefectivity is reduced to proving that these level-wise normal blocks and the final secant block contribute the expected number of independent directions.

Example 4.4 (A tangent-space computation)**.** Consider

$$\mathbf{n} = (2, 2, 2, 1), \qquad \mathbf{d} = (3, 3).$$

Set

$$L_1 = a_{1,1}x + y, \qquad L_2 = x + a_{2,2}y.$$

At level two take

$$R = b_{1,1}L_1^3 + L_2^3, \qquad S = b_{2,1}L_1^3 + L_2^3,$$

and at the final layer

$$F = c_{1,1}R^3 + S^3 \in \operatorname{Sym}^9 k^2.$$

On the affine chart $a_{1,2} = a_{2,1} = 1$, fix the two level-one points

$$a_{1,1} = 0, \qquad a_{2,2} = 0.$$

Then

$$L_1 = y, \qquad L_2 = x,$$

and

$$R = b_{1,1}y^3 + x^3, \qquad S = b_{2,1}y^3 + x^3.$$

Writing $F = [y_0 : \dots : y_9]$ in the basis

$$[x^9, x^8y, \dots, xy^8, y^9],$$

the only nonzero coordinates at this base point are

$$y_0 = c_{1,1} + 1, \qquad y_3 = 3(c_{1,1}b_{1,1} + b_{2,1}),$$

$$y_6 = 3(c_{1,1}b_{1,1}^2 + b_{2,1}^2), \qquad y_9 = c_{1,1}b_{1,1}^3 + b_{2,1}^3.$$

Thus F lies in the 3-space

$$\Lambda = (y_1 = y_2 = y_4 = y_5 = y_7 = y_8 = 0).$$

The Jacobian columns split as

$$\left(\underbrace{\partial_{a_{2,2}}, \partial_{a_{1,1}}}_{\text{level-1 normal block}} \;\middle|\; \underbrace{\partial_{b_{1,1}}, \partial_{b_{2,1}}, \partial_{c_{1,1}}}_{\text{last-level secant block}} \right).$$

For general parameters, the first two columns span $\mathcal{N}^1$, the last three columns span the tangent space to the last secant block, and the five columns are independent. Therefore

$$\dim T_F \mathcal{V}_{\mathbf{n},\mathbf{d}} = 5,$$

which is the expected dimension for the neurovariety associated to $(2, 2, 2, 1)$ and $(3, 3)$.

The Alexander–Hirschowitz theorem for neurovarieties. Write $\mathbf{n}' = (n_0, \dots, n_{L-1}, 1)$. Thus $\mathcal{V}_{\mathbf{n}',\mathbf{d}}$ is the single-output sibling of the multi-output architecture.

Theorem 1. [MM25, Theorem A] *Let* $\mathbf{n} = (n_0, \dots, n_L)$ *with* $L \geq 2$*, and let* $\mathbf{d} = (d_1, \dots, d_{L-1})$ *over a subfield of* $\mathbb{C}$*. Assume that:*

(i) $n_{i-1} + n_i - 1 < \binom{n_{i-1}-1+d_i}{n_{i-1}-1}$ *for every* $i = 1, \dots, L-1$*;*
(ii) $V_{d_{L-1}}^{n_{L-2}-1} \subset \mathbb{P}^{N_{L-1}}$ *is not* n_{L-1}*-defective.*

Then:

(a) *if* $n_L = 1$*, the neurovariety* $\mathcal{V}_{\mathbf{n},\mathbf{d}} \subset \mathbb{P}^{N_{L-1}}$ *is nondefective;*
(b) *if* $n_L \geq 2$ *and*

$$\operatorname{expdim}(\mathcal{V}_{\mathbf{n}',\mathbf{d}}) = \sum_{i=1}^{L} n'_i(n'_{i-1} - 1),$$

then $\mathcal{V}_{\mathbf{n},\mathbf{d}} \subset (\mathbb{P}^{N_{L-1}})^{\times n_L}$ *is nondefective and the associated neural network is globally identifiable.*

Idea of proof. For $n_L = 1$, apply the neuro-Terracini formula (N2). The room condition in (i) says that each intermediate layer has enough ambient space, so the level-wise normal block $\mathcal{N}^j$ contributes the expected number $n_j(n_{j-1} - 1)$ of independent directions. Assumption (ii) says that the final Veronese secant block is nondefective. Hence the tangent space has the expected dimension (N1).

For $n_L \geq 2$, use the single-output sibling $\mathcal{V}_{\mathbf{n}',\mathbf{d}}$. A single output gives one general point in an admissible n_{L-1}-secant space Λ. Nondefectivity and the non-filling hypothesis imply that through this point there are only finitely many admissible secant spaces. A multi-output network gives instead a general n_L-tuple of points in the same Λ. Since $n_L \geq 2$, this tuple avoids the proper closed subsets obtained by intersecting Λ with the other finitely many admissible secant spaces. Thus Λ is uniquely recovered. Once Λ is recovered, the last hidden layer is recovered from

$$\Lambda \cap \widetilde{V}_{L-1}.$$

This gives global identifiability, up to the standard permutations and rescalings.

Theorem 2. [MM25, Theorem B] *Over $\mathbb{R}$, the polynomial neural network with architecture $(n_0, \ldots, n_L)$ and activation exponents $(d_1, \ldots, d_{L-1})$ is globally identifiable if*

(i) $2n_i < n_{i-1}(n_{i-1} - 1)$ *for* $i < L$;
(ii) $n_L \geq 2$;
(iii) $d_{L-1} > 4$;
(iv) $\operatorname{expdim}(\mathcal{V}_{\mathbf{n}',\mathbf{d}}) = \sum_{i=1}^{L} n'_i(n'_{i-1} - 1)$.

Idea of proof. The inequality $2n_i < n_{i-1}(n_{i-1} - 1)$ implies the room condition in Theorem A since $d_i \geq 2$. The condition $d_{L-1} > 4$ excludes the Alexander–Hirschowitz defective cases for the last Veronese, hence assumption (ii) of Theorem A holds after base change from $\mathbb{R}$ to $\mathbb{C}$. The last hypothesis is exactly the non-filling condition for the single-output sibling. Theorem A then gives global identifiability.

Remark 4.5 (What the theorem says geometrically)**.** The single-output network is a constrained secant construction. Its dimension is computed by a Terracini-type formula with two kinds of tangent directions: the usual final secant directions and the level-wise normal directions created by earlier layers. If these directions are independent and the final Veronese secant is nondefective, then the neurovariety has the expected dimension.

The multi-output hypothesis $n_L \geq 2$ turns finite ambiguity into uniqueness. One output may lie on finitely many admissible secant spaces. A general tuple of at least two outputs in the same secant space determines that secant space uniquely. This is the geometric reason why Theorem A gives global identifiability in the multi-output range.

References

[AH95] James Alexander and André Hirschowitz. Polynomial interpolation in several variables. *Journal of Algebraic Geometry*, 4(2):201–222, 1995.

[BBC18] Edoardo Ballico, Alessandra Bernardi, and Luca Chiantini. On the dimension of contact loci and the identifiability of tensors. *Arkiv för Matematik*, 56(2):265–283, 2018.

[BCO14] Cristiano Bocci, Luca Chiantini, and Giorgio Ottaviani. Refined methods for the identifiability of tensors. *Annali di Matematica Pura ed Applicata*, 193(6):1691–1702, 2014.

[BV18] Alessandra Bernardi and Davide Vanzo. A new class of non-identifiable skew-symmetric tensors. *Annali di Matematica Pura ed Applicata*, 197(5):1499–1518, 2018.

[CC02] Luca Chiantini and Ciro Ciliberto. Weakly defective varieties. *Transactions of the American Mathematical Society*, 354(1):151–178, 2002.

[CC06] Luca Chiantini and Ciro Ciliberto. On the concept of k-secant order of a variety. *Journal of the London Mathematical Society*, 73(2):436–454, 2006.

[CO21] Luca Chiantini and Giorgio Ottaviani. A footnote to a footnote to a paper of B. Segre, 2021.

[Cob22] Arthur B. Coble. Associated sets of points. *Transactions of the American Mathematical Society*, 24(1):1–20, 1922.

[COV17] Luca Chiantini, Giorgio Ottaviani, and Nick Vannieuwenhoven. On generic identifiability of symmetric tensors of subgeneric rank. *Transactions of the American Mathematical Society*, 369(6):4021–4042, 2017.

[CR06] Ciro Ciliberto and Francesco Russo. Varieties with minimal secant degree and linear systems of maximal dimension on surfaces. *Advances in Mathematics*, 200(1):1–50, 2006.

[Dol04] Igor V. Dolgachev. On certain families of elliptic curves in projective space. *Annali di Matematica Pura ed Applicata*, 183(3):317–331, 2004.

[GM19] Francesco Galuppi and Massimiliano Mella. Identifiability of homogeneous polynomials and Cremona transformations. *Journal für die reine und angewandte Mathematik*, 2019(757):279–308, 2019.

[Hil88] David Hilbert. Lettre adressée à M. Hermite. *Journal de Mathématiques Pures et Appliquées*, pages 249–256, 1888.

[IK99] Anthony Iarrobino and Vassil Kanev. *Power Sums, Gorenstein Algebras, and Determinantal Loci*, volume 1721 of *Lecture Notes in Mathematics*. Springer, Berlin, 1999.

[LMR23] Antonio Laface, Alex Massarenti, and Rick Rischter. Decomposition algorithms for tensors and polynomials. *SIAM Journal on Applied Algebra and Geometry*, 7(1):264–290, 2023.

[Mel06] Massimiliano Mella. Singularities of linear systems and the Waring problem. *Transactions of the American Mathematical Society*, 358(12):5523–5538, 2006.

[MM24] Alex Massarenti and Massimiliano Mella. Bronowski's conjecture and the identifiability of projective varieties. *Duke Mathematical Journal*, 173(17):3293–3316, 2024.

[MM25] Alex Massarenti and Massimiliano Mella. The Alexander–Hirschowitz theorem for neurovarieties, 2025. Preprint.

[Syl04] James Joseph Sylvester. *The Collected Mathematical Papers of James Joseph Sylvester*, volume 1. Cambridge University Press, Cambridge, 1904.

[Ter11] Alessandro Terracini. Sulle V_k per cui la variet'a degli S_h secanti ha dimensione minore dell'ordinario. *Rendiconti del Circolo Matematico di Palermo*, 31:392–396, 1911.

Alex Massarenti, Dipartimento di Matematica e Informatica, Università di Ferrara, Via Machiavelli 30, 44121 Ferrara, Italy

Email address: `msslxa@unife.it`