Armillary maps stellar spectra without labels, then uses a few labelled stars to calibrate the whole map.

Draft version October 7, 2026
Typeset using LaTeX **twocolumn** style in AASTeX631

arXiv:2610.07112v1 [astro-ph.IM] 5 Oct 2026

Armillary: A Label-Free Coordinate System for Stellar Spectra

Yuan-Sen Ting[1,2,3] and Serat Saad[1,2]

[1] *Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA*
[2] *Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA*
[3] *Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*

ABSTRACT

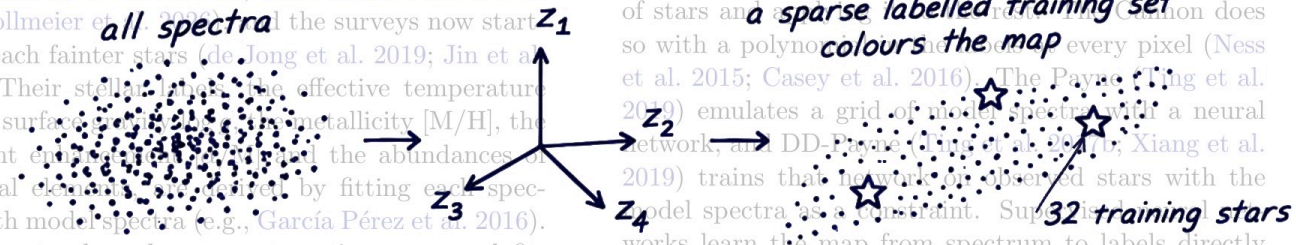
A stellar spectrum is determined by a few parameters, so a set of spectra lies on a low-dimensional manifold with those parameters as its coordinates. We introduce Armillary, which generalises the Sequencer of Baron and Ménard from an ordering to multidimensional coordinates, constructed from the spectra alone. Two stars are compared by the Wasserstein distance between their signed absorption profiles. Shortest paths through a nearest-neighbour graph measure distance along the manifold and give the coordinates, and the same graph transfers the labels of a few stars to all the others. The coordinates are more nearly linear in $T_{\rm eff}$ and $\log g$ than principal components, t-SNE, UMAP or an autoencoder. The label transfer, without deep learning, is more precise than the first three and comparable to the autoencoder. On 202 932 APOGEE DR17 stars, an illustrative draw of 32 training stars recovers the ASPCAP $T_{\rm eff}$, $\log g$ and [M/H] of held-out dwarfs and giants, with 1σ errors of 100 K, 0.18 and 0.083 dex. Resolving the two α sequences of the disc in [α/M] for giants takes 128 training stars, and with 2048 the four errors improve to 33 K, 0.06, 0.026 and 0.022 dex. On 104 799 DESI stars, 32 training labels from APOGEE recover the stellar parameters to 129 K, 0.25 and 0.203 dex, even for spectra degraded to a signal-to-noise of about 20. At a signal-to-noise of about 100, 2048 training labels resolve the two α sequences of the disc in DESI giants, with errors against APOGEE of 51 K, 0.07, 0.047 and 0.026 dex. The code is public at https://github.com/tingyuansen/armillary.

The practical prize is survey-to-survey calibration without refitting every spectrum from scratch.

Keywords: Stellar spectral lines (1630), Stellar abundances (1577), Astrostatistics techniques (1886), Dimensionality reduction (1943), Spectroscopy (1558)

1. INTRODUCTION

Spectroscopic surveys now contain tens of millions of stellar spectra (Cui et al. 2012; Majewski et al. 2017; Gaia Collaboration et al. 2023; DESI Collaboration et al. 2026; Kollmeier et al. 2026), and the surveys now start ing to reach fainter stars (de Jong et al. 2019; Jin et al. 2024). Their stellar labels, the effective temperature $T_{\rm eff}$, the surface gravity $\log g$, the metallicity [M/H], the α-element enhancement [α/M], and the abundances of individual elements, are derived by fitting each spectrum with model spectra (e.g., García Pérez et al. 2016). Model spectra have known systematic errors, and fitting them to every spectrum of a survey is expensive. The resulting labels disagree between pipelines (Jofré et al. 2014; Smiljanic et al. 2014) and between surveys (Hegedűs et al. 2023), at a level comparable to the astrophysical signals of interest (Shen et al. 2026). Placing one survey's spectra on another survey's label scale, without refitting every spectrum from first principles, is therefore a recurring need.



Unlabelled spectra are not dead weight; they determine the geometry before any labels arrive.

Data-driven models address this need by learning the relation between labels and spectra from a training set of stars and applying it to the rest. The Cannon does so with a polynomial in the labels for every pixel (Ness et al. 2015; Casey et al. 2016). The Payne (Ting et al. 2019) emulates a grid of model spectra with a neural network, and DD-Payne (Ting et al. 2019; Xiang et al. 2019) trains that network on observed stars with the model spectra as a prior constraint. Supervised neural works learn the map from spectrum to labels directly (Fabbro et al. 2018; Leung & Bovy 2019; Zhang et al. 2019). Authors use these models to transfer labels from one survey to another through the stars that the surveys share (e.g., Ho et al. 2017; Xiang et al. 2019; Das et al. 2025).

In all of these the labels enter with the spectra, and the model learns their relation only where labels exist. The training set must therefore span the label space (Ness

et al. 2015). The more flexible the model, the larger the set it needs, since a direct map from spectrum to labels can also learn correlations that hold in the training set but not in the physics. Such a model gets nothing from the unlabelled majority of a sample, and a star outside its label-space coverage is labelled by extrapolation. Label transfers between two surveys used large training sets of stars in common (Ho et al. 2017; Thomas et al. 2024; Das et al. 2025).

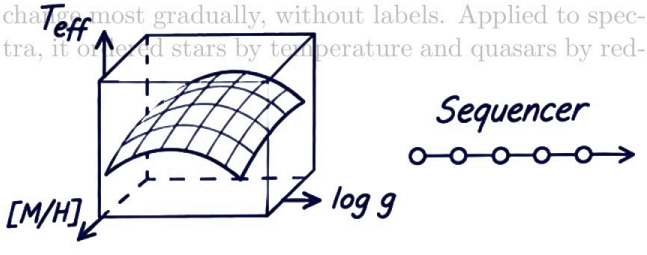
A supervised model learns only where labels exist; the manifold uses every spectrum first.

Representation learning exploits the structure of the data itself (Bengio et al. 2013). A spectrum is determined by a few parameters, so a set of spectra lies on a low-dimensional manifold in the space of pixel values. The parameters are coordinates on that manifold. Representation learning recovers the manifold from all the spectra, labelled or not. The labels are introduced afterwards, attached to a small training set from which every other star's labels are transferred, which in the language of deep learning is a few-shot use of labels (Snell et al. 2017). The unlabelled majority then contributes to the coordinates, and the training set need only cover the part of the manifold it labels. A star far from every training star on the manifold appears so in the coordinates.

Most current representation learning of spectra uses deep learning. Neural networks trained to compress spectra, or segments of them, through a bottleneck, known as autoencoders, return a latent representation of every spectrum (Yang & Li 2015; Sedaghat et al. 2021). Such representations can be expressive, but they are often entangled, with several physical parameters mixed in each latent variable (Bengio et al. 2013; Laroche & Speagle 2025). Disentangling them requires a supervised model of its own (Locatello et al. 2018; de Mijolla et al. 2021), and that model requires the labels the few-shot setting was intended to avoid. A latent space learned without labels can still hold a physical structure, as the two α sequences do for cool giants in one learned from Gaia BP/RP spectra (Laroche & Speagle 2025). Its latent variables, however, have no definition of their own.

The physics of the problem is simpler than that of most data to which deep neural networks are applied. The manifold has a few principal dimensions, the stellar parameters $T_{\rm eff}$, $\log g$ and [M/H] and the α enhancement, along which the spectra vary most. Recovering coordinates on a manifold of a few dimensions is a classical problem. We ask whether classical techniques recover them.

The Sequencer of Baron & Ménard (2021) orders a data set along the single dimension in which its objects change most gradually, without labels. Applied to spectra, it ordered stars by temperature and quasars by redshift. Among the distances it uses is one that measures where absorption lies along the wavelength axis, independently of the total amount of absorption. A graph of the pairwise distances is built, and the ordering follows the minimum spanning tree (Kruskal 1956) of that graph. An ordering, however, assigns a single coordinate to each object, whereas the manifold of stellar spectra has more than one dimension. A single ordering can follow only one of them.



We generalise the Sequencer from one dimension to as many as the spectra vary along. The distance and the graph are retained. The ordering is replaced by coordinates, obtained from the lengths of the shortest paths through the graph, and the number of coordinates is read from the data. We call the method Armillary, after the instrument whose rings are a coordinate system for the sky.

Section 2 presents Armillary and the reasoning behind each of its steps, and shows that a regular grid of labels is recovered from synthetic spectra. Sections 3 and 4 apply it to APOGEE and to DESI. Section 5 relates Armillary to earlier work and states its limitations, and Section 6 concludes.

The Sequencer found one direction; stars need several, so we replace an ordering with coordinates.

2. CONSTRUCTION OF THE COORDINATES

Armillary constructs the coordinates of the manifold in four steps. A distance between spectra is defined, and a graph of nearest neighbours is built from it. Coordinates are extracted from the shortest paths through the graph, and a local refinement reduces the noise the graph leaves in them. Figure 1 summarises the steps.

2.1. *The distance between two spectra*

The distances set the neighbourhoods and the initial coordinates. The refinement also uses the local variation of the normalised spectra. A given change in the parameters must therefore correspond to the same distance wherever it occurs on the manifold. Any unevenness in the distance is inherited by the graph and by the coordinates. The simplest distance, the Euclidean distance between the two spectra as pixel vectors, does not satisfy this requirement, because it is dominated by how much absorption each spectrum contains. A small step in temperature, for example, changes a cool metal-rich spectrum, which is rich in strong lines, by a large amount. It changes a hot metal-poor spectrum, with few lines, by a small one. The coordinates built on such a distance are compressed in the metal-poor corner of the label space and stretched in the metal-rich one.

The Sequencer compares its objects under several distances, each applied to contiguous segments of the data at several scales. Its default set is the Euclidean

The recipe is distance, neighbour graph, geodesics, coordinates, refinement, and finally label transfer.

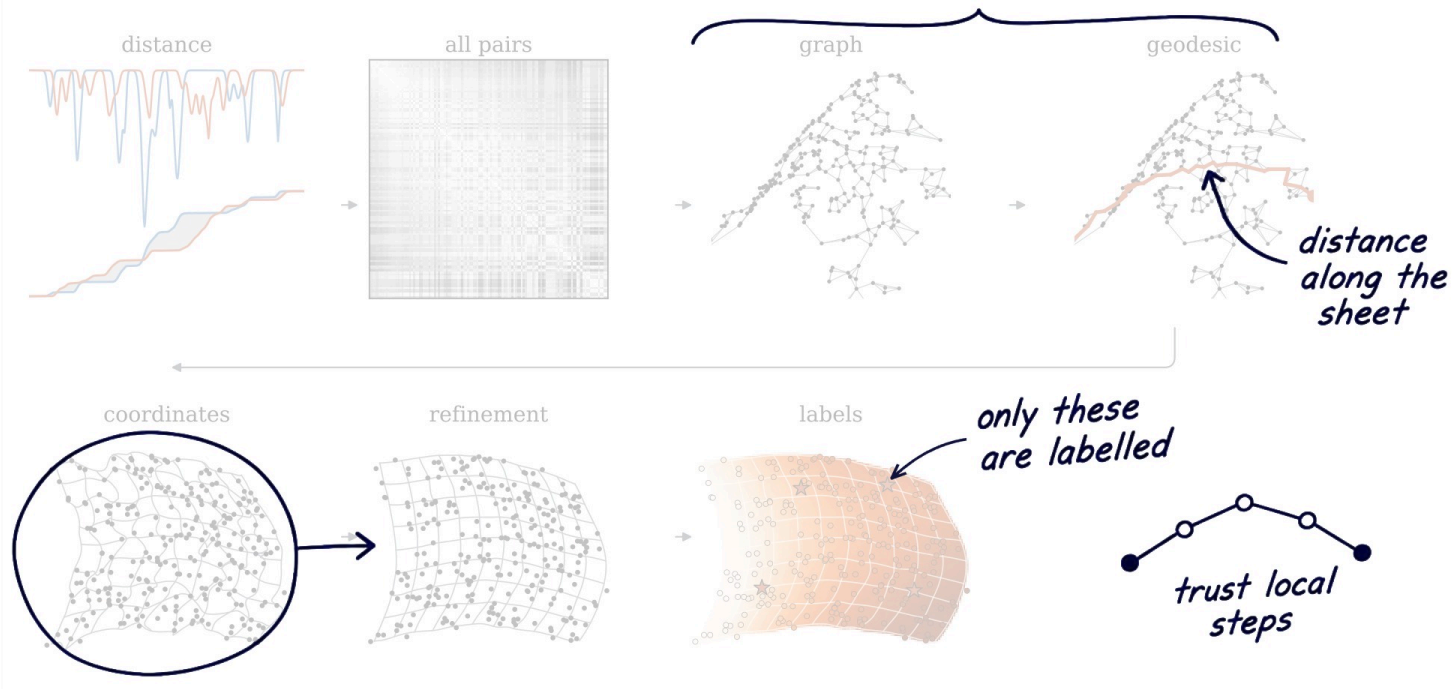


Figure 1. A schematic of the steps of Armillary, shown for points on a bent sheet. Top row, from left. Two spectra are compared by the Wasserstein distance (see Fig. 2) between their signed absorption profiles, the shaded area between their cumulative curves. In this schematic the distance is computed between every pair of stars. A graph joins each star to its nearest neighbours under that distance. The geodesic distance between two distant stars is the length of the shortest path along the graph (orange). Bottom row, the stars in the recovered coordinates, with a regular grid of the sheet drawn through them to show the shape of the manifold (see Fig. 3). The geodesic distances between all pairs are turned into coordinates, which follow the large-scale shape but scatter on small scales and crumple the grid. A local refinement reduces the scatter by penalising departures from the weighted neighbour relations. The labels of a few labelled stars (orange stars) then determine those of every other star through the same weights, here one label shaded from light to dark.

The graph turns many trustworthy short comparisons into one long-distance comparison.

distance, the Kullback–Leibler divergence (Kullback & Leibler 1951), the Earth Mover's Distance and the energy distance (Székely & Rizzo 2013). Baroll & Ménard (2021) note that the Earth Mover's Distance and the energy distance are sensitive to the magnitude of displacements along the pixel axis, whereas the Euclidean distance and the Kullback–Leibler divergence treat each pixel as a separate dimension. We adopt the Earth Mover's Distance, the Wasserstein distance of optimal transport (Rubner et al. 2000; Peyré & Cuturi 2018), in its cumulative form and with its segmentation into chunks, extending it to signed absorption depths so that noise above the continuum need not be clipped. The resulting distance compares the pattern of absorption after the net amount of absorption is normalised. The Euclidean distance serves below as the baseline. Another distance can be substituted where the data call for one.

We divide the wavelength range into N_{chunk} chunks, each a few to a few tens of nanometres wide. Within chunk c, star i has an absorption depth $1 - f_i(\lambda)$ at each wavelength λ. Here f_i is the flux divided by a local continuum, a straight line through the upper envelope of the flux over a window whose width is set for each survey. The depths define a signed absorption profile along the wavelength axis.

Dividing the depth by the net absorption in the chunk gives the normalised profile $\rho_i^{(c)}$. Its integral is one, whatever the strength of the lines, provided the net absorption is nonzero. A pixel where noise puts the flux above the continuum has a negative depth, and the depth is kept negative. Noise scatters the flux above and below the continuum by equal amounts, so on average it adds no net absorption.

The depth is also weighted where the pixel errors carry structure of their own, from the sky and the detector. Each pixel's depth is multiplied by a reliability weight $v_i(\lambda) = \sigma_i^{-2}(\lambda)/\langle\sigma_i^{-2}\rangle_c$, the pixel's inverse variance divided by its mean over the good pixels of the chunk. A noisy pixel then carries less of the chunk's absorption, and the rest of the chunk keeps its weight. Where the errors carry no such structure, $v_i = 1$. Integrating $\rho_i^{(c)}$ from the start of the chunk λ_c gives the cumulative curve

Normalising absorption prevents line-rich stars from dominating merely because they have more absorption.

The Wasserstein distance measures where absorption moves, not simply how much flux changes pixel by pixel.

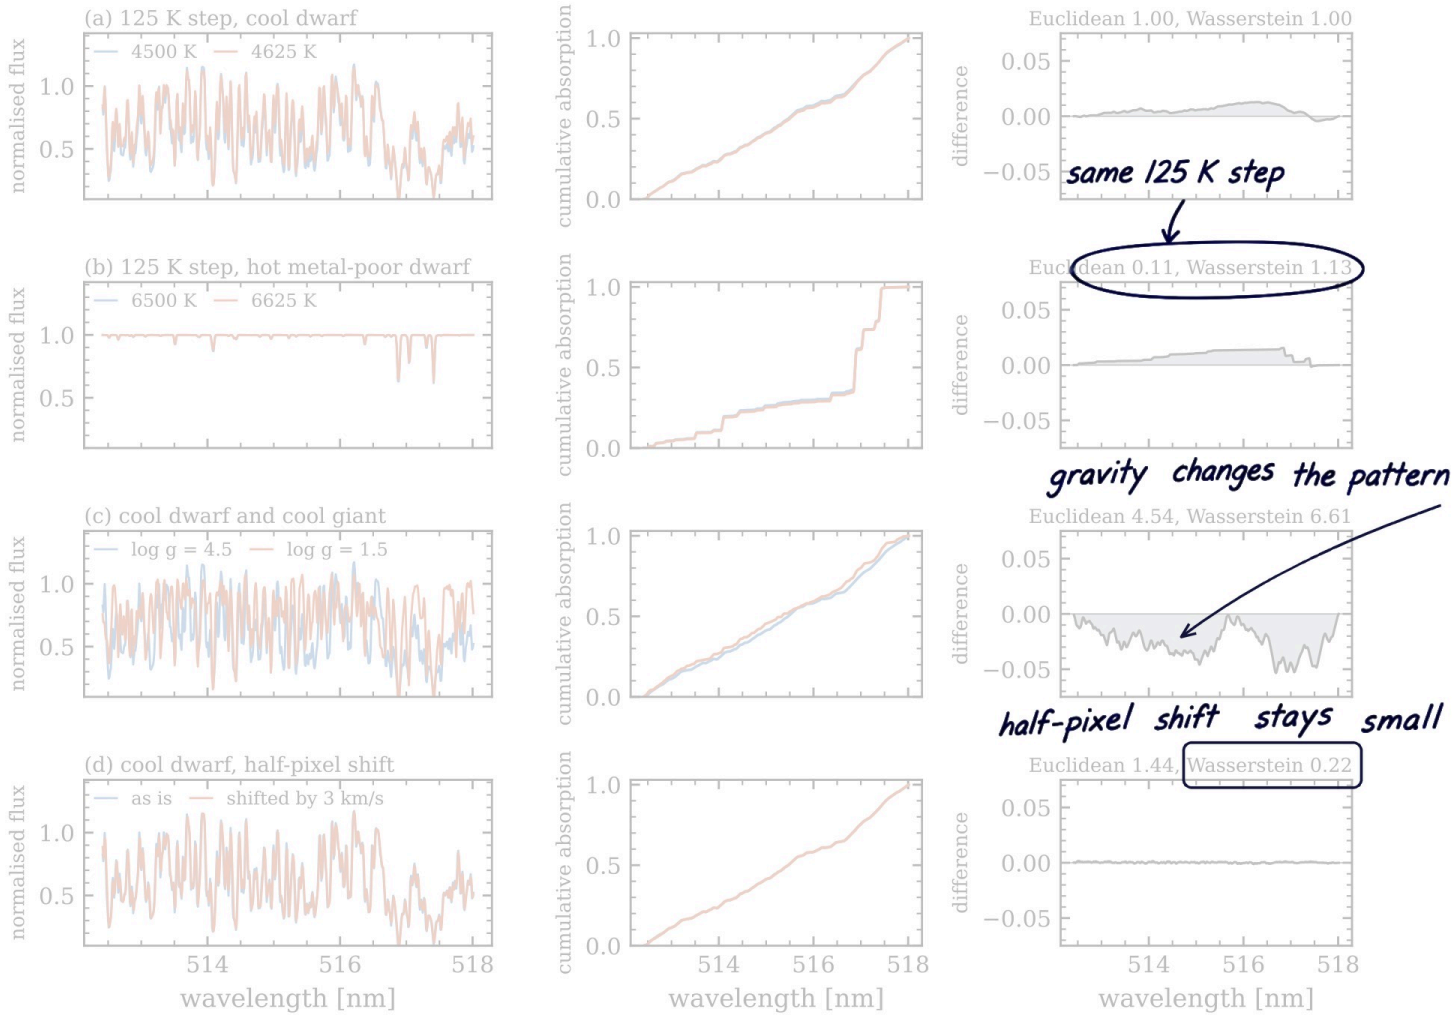


Figure 2. The Wasserstein distance measures a change of the spectrum more evenly than the Euclidean distance does, shown on four pairs of synthetic spectra, one pair per row, in the chunk containing the Mg b triplet. Left, the normalised flux. Middle, the cumulative absorption curves F. Right, their difference, whose integral (shaded) is the Wasserstein distance of the chunk. Above each right-hand panel are the Euclidean and Wasserstein distances between the two complete spectra, over all chunks, each in units of its own value in row (a). Every row then reads as a number of temperature steps of the cool dwarf under each distance. (a) A 125 K step for a cool dwarf (4500 K, $\log g = 4.5$, solar metallicity). (b) The same step for a hot metal-poor dwarf (6500 K, $\log g = 4.5$, [M/H] $= -2$). (c) The cool dwarf and a cool giant ($\log g = 1.5$) of the same temperature and metallicity. (d) The cool dwarf and itself shifted by half a pixel, 3 km s^{-1} at this resolution.

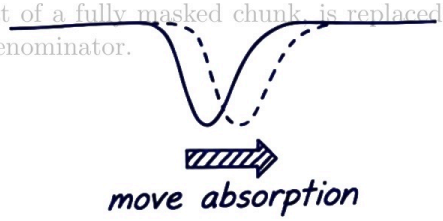
$F_i^{(c)}$,

$$\rho_i^{(c)}(\lambda) = \frac{v_i(\lambda)\,[1 - f_i(\lambda)]}{\int_c v_i(\lambda')\,[1 - f_i(\lambda')]\,\mathrm{d}\lambda'}, \tag{1}$$

$$F_i^{(c)}(\lambda) = \int_{\lambda_c}^{\lambda} \rho_i^{(c)}(\lambda')\,\mathrm{d}\lambda'. \tag{2}$$

The superscript names the chunk and the subscript the star. Each position's absorption is the cumulative curve rises through absorption lines and can fall where the flux lies above the continuum. In the pixel implementation, a net depth sum below 10^{-3} in absolute value, in practice only that of a fully masked chunk, is replaced by 10^{-3} in the denominator.

Euclidean distance badly shrinks the hot, metal-poor corner.



For nonnegative unit-normalised profiles, the one-dimensional Wasserstein distance, written W_1, is the minimum cost of moving absorption from one profile to the other, with cost equal to the amount of absorption moved times the wavelength distance it is moved. The subscript is the power to which that distance is raised in the cost, here one. The distance equals the area between the cumulative curves of the two profiles, and we keep that expression for the signed profiles,

$$W_1^{(c)}(i,j) = \int_c \left| F_i^{(c)}(\lambda) - F_j^{(c)}(\lambda) \right| \mathrm{d}\lambda, \tag{3}$$

shown in Figure 2. With negative depths the profiles are signed and F need not be a cumulative distribution function. The expression remains the L^1 distance between the two cumulative profiles. It extends the Wasserstein distance to signed profiles with the same integral, and we keep that name throughout.

Multiple chunk sizes let broad features and individual lines both vote.

The chunk distances are then combined rather than summed directly, they would be dominated by the widest or most line-rich chunks, because a wide chunk has more absorption over longer distances. We therefore divide each chunk's distance by its median over pairs of stars, the typical distance in that chunk, and sum over the chunks,

$$D_{ij} = \sum_{c=1}^{N_{\mathrm{chunk}}} \frac{W_1^{(c)}(i,j)}{\mathrm{median}_{(i',j')}\, W_1^{(c)}(i',j')}, \qquad (4)$$

so that every chunk contributes in the same units, the number of typical distances by which the two stars differ in that chunk. The wavelength range can be divided into chunks in more than one way, into a few wide chunks and into many narrow ones, and the sum then runs over the chunks of every division. A wide chunk compares the broad features of two spectra and a narrow one their individual lines.

Figure 2 compares the Euclidean and Wasserstein distances between complete spectra in three cases, with one chunk of each pair drawn. The spectra are synthetic, computed with Payne-Zero (Ting & Kim 2026), a Python implementation of Kurucz's ATLAS12 and SYNTHE (Kurucz 2005) with the synthesis on GPUs, at a resolving power of 20 000 over 480–680 nm without noise. They are divided into 32 chunks each about 6 nm wide.

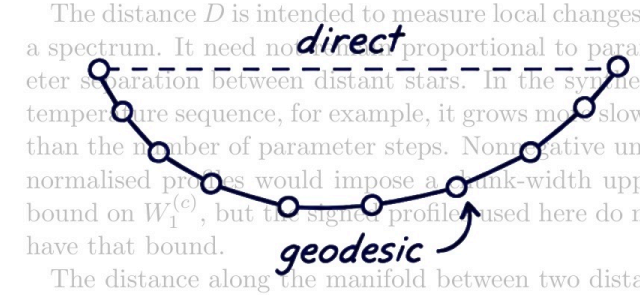
The first case, rows (a) and (b), is a small step in temperature for two stars. One is a cool dwarf of solar metallicity, whose spectrum is crowded with strong lines, and the other a hot metal-poor dwarf, whose spectrum has few and weak lines. Once the total is divided out, each has one unit of absorption in every chunk, and the Wasserstein distance does not see that the cool dwarf absorbs far more. It measures where along the chunk that unit lies, and the temperature step shifts that location by a similar amount in both stars. Under the Euclidean distance the step is nine times shorter for the hot metal-poor dwarf than for the cool dwarf. Under the Wasserstein distance the two are within fifteen per cent.

The second case, row (c), is the cool dwarf and a cool giant of the same temperature and metallicity. They have similar amounts of absorption but differ in its pattern, for example in the pressure-broadened wings of the strong lines. The Wasserstein distance separates them more widely than the Euclidean distance does, by six and a half steps against four and a half. It does so because the amount of absorption is divided out, and only the pattern remains.

The third case, row (d), is the cool dwarf against itself shifted by half a pixel, as by an uncorrected radial-velocity offset. The Euclidean distance compares the two spectra at matched wavelengths, so every shifted line appears as two residuals of opposite sign. That shift gives a larger distance than a temperature step. The Wasserstein distance responds to a displacement through changes in the cumulative absorption profile, so the same shift gives 0.22 of a temperature step. The Wasserstein distance is less sensitive to a small residual radial-velocity error, which changes none of the star's parameters.

2.2. *The neighbour graph and geodesic distances*

The distance D is intended to measure local changes of a spectrum. It need not be linear proportional to parameter separation between distant stars. In the synthetic temperature sequence, for example, it grows more slowly than the number of parameter steps. Nonnegative unit-normalised profiles would impose a chunk-width upper bound on $W_1^{(c)}$, but the signed profiles used here do not have that bound.



The distance along the manifold between two distant stars is therefore assembled from short distances. We build the neighbour graph $\mathcal{G}$ by finding the k nearest neighbours of every star under D. Stars a and b are joined by an edge if a is among the k nearest neighbours of b or b among those of a, and the edge has length D_{ab}. Such a graph can fall into a few separate pieces, since a group of unusual stars can have all its nearest neighbours within the group. A star in such a piece would have no path to the rest. Every smaller piece is therefore joined to the largest piece by a single edge, the shortest of all those that would connect it to the main graph. A bridge can span a spectroscopically sampled region where D no longer tracks the separation, so that piece's position relative to the rest is only approximate.

Direct distances saturate or curve; adding short edges keeps distance growing along the manifold.

The geodesic distance between stars i and j is the length of the shortest path through $\mathcal{G}$ between them, the sum of D_{ab} over the edges (a,b) of the path,

$$D^{\mathrm{g}}_{ij} = \min_{\text{paths } i\to j \text{ in } \mathcal{G}} \sum_{(a,b)\in\text{path}} D_{ab}, \qquad (5)$$

found by Dijkstra's algorithm (Dijkstra 1959), which grows a tree of shortest paths outward from i by always extending the closest star not yet reached, and so visits every star once. Apart from the bridges, the path uses neighbour distances, for which D is intended to be most

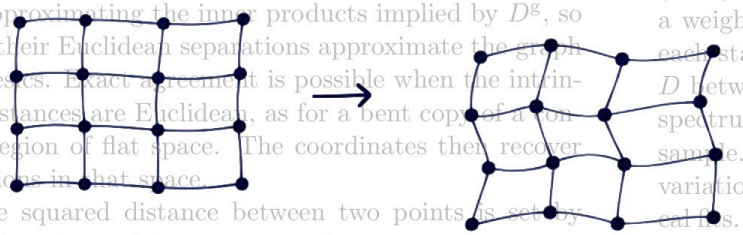
A bridge reconnects isolated components, but its absolute placement is necessarily approximate.

Classical multidimensional scaling converts graph geodesics into coordinates, then local linear weights smooth sampling noise.

informative. The geodesic distance can therefore keep growing with separation along the manifold where the direct distance D no longer tracks that separation. The Sequencer orders its objects by a walk through the minimum spanning tree of their combined distances, and it is that ordering we replace with coordinates.

2.3. *Coordinates from geodesic distances*

The geodesic distances form a matrix D^{g}, N^2 numbers between pairs of stars, and coordinates turn them into a position for each star. Classical multidimensional scaling (Torgerson 1952) gives points $\mathbf{c}_i$ in d dimensions by approximating the inner products implied by D^{g}, so that their Euclidean separations approximate the graph geodesics. Exact agreement is possible when the intrinsic distances are Euclidean, as for a bent copy of a convex region of flat space. The coordinates then recover positions in that space.



The squared distance between two points is set by their lengths and their inner product,

$$|\mathbf{c}_i - \mathbf{c}_j|^2 = |\mathbf{c}_i|^2 + |\mathbf{c}_j|^2 - 2\,\mathbf{c}_i \cdot \mathbf{c}_j, \tag{6}$$

so the matrix of squared distances is -2 times the matrix of inner products plus the squared lengths $|\mathbf{c}_i|^2$ and $|\mathbf{c}_j|^2$. The squared length of star i is the same in every entry of row i, and that of star j in every entry of column j. Subtracting the row and column means removes both and leaves -2 times the inner products, and the factor $-\frac{1}{2}$ then removes the -2. The centring matrix J performs the subtraction,

$$B = -\tfrac{1}{2}\, J\, (D^{\mathrm{g}})^2\, J, \qquad J = I - \tfrac{1}{N}\mathbf{1}\mathbf{1}^{\top}, \tag{7}$$

where N is the number of stars, I the identity, $\mathbf{1}$ the vector of ones, and $(D^{\mathrm{g}})^2$ is taken entry by entry. The inner products $B_{ij} = \mathbf{c}_i \cdot \mathbf{c}_j$ of coordinates whose mean lies at the origin.

Large positive eigenvalues tell us how many coordinate directions the data support.

In that case, writing the coordinates as a matrix C_{geo} with one row per star gives $B = C_{\mathrm{geo}} C_{\mathrm{geo}}^{\top}$. The eigenvectors of B, each scaled by the square root of its eigenvalue, are therefore the columns of C_{geo}. Keeping the eigenvectors of the d largest positive eigenvalues gives the d coordinates whose inner products reproduce B as closely as any d coordinates can. Each positive eigenvalue measures how much of B one coordinate accounts for. A nearly Euclidean manifold can give a small set of dominant eigenvalues, which guides how many coordinates to keep.

2.4. *A locally linear refinement*

Shortest paths preserve global shape but jitter locally because the sample never traces an ideal route.

The geodesic coordinates capture the large-scale arrangement but can scatter on smaller scales. A star's path through the graph depends on which stars sample the route. Detours through those stars can lengthen it, while edges across curved regions can shorten it, and the coordinates inherit these sampling errors. Even a small relative error in a long geodesic distance can displace a star's placement among its immediate neighbours.

The refinement reduces that scatter by using the local structure of the spectra, following Roweis & Saul (2000). On a smooth manifold every point is nearly a weighted average of its nearest neighbours. We take each star's nearest neighbours $\mathcal{N}(i)$ under the distance D between spectra. Let $\mathbf{x}_i$ be the projection of star i's spectrum onto the leading principal components of the sample. The projection retains the dominant spectral variation while reducing the pixel noise entering the local fits. The weights w_{ij}, summing to one, are those that best reconstruct $\mathbf{x}_i$ from the projections of its neighbours,

$$\mathbf{x}_i \approx \sum_{j \in \mathcal{N}(i)} w_{ij}\, \mathbf{x}_j, \qquad \sum_j w_{ij} = 1. \tag{8}$$

The weights are found in closed form, as one small linear solve per star. Let Z_i be the matrix whose rows are the differences $\mathbf{x}_j - \mathbf{x}_i$ from the star to its neighbours, and $H_i = Z_i Z_i^{\top}$ the matrix of their inner products. The weights solve $[H_i + \epsilon\, \mathrm{tr}(H_i) I]\, \mathbf{w}_i = \mathbf{1}$ with a small ridge ϵ and are then scaled to sum to one.

The weights describe how star i sits among its neighbours, so they can be imposed on the coordinates. They come from the spectra and D alone, so the noise of the geodesic coordinates does not enter them. We use them to regularise the geodesic coordinates, asking for coordinates that keep their large-scale shape while making every star the same weighted average of its neighbours.

Let W be the matrix with entries w_{ij}, one row per star like C_{geo}, and C the refined coordinates we seek, arranged the same way. The row i of $(I - W)\, C$ is star i's coordinate minus the weighted average of its neighbours', the amount by which the neighbour relation fails for that star. The coordinates that balance the two requirements are

The anchor prevents local smoothing from ironing away real large-scale structure.

$$C = \arg\min_{C} \; \big\|(I - W)\, C\big\|^2 + \mu_{\mathrm{c}} \big\|C - C_{\mathrm{geo}}\big\|^2, \tag{9}$$

in which the norms are sums of squares over every entry. The first term enforces the neighbour relations, the second holds the coordinates near C_{geo}, and the anchor strength μ_{c} sets the trade-off between them. Setting the gradient to zero gives one linear solve,

$$\big[(I - W)^{\top}(I - W) + \mu_{\mathrm{c}} I\big]\, C = \mu_{\mathrm{c}}\, C_{\mathrm{geo}}. \tag{10}$$

The same neighbour weights propagate sparse labels, and scalable approximations make the method practical for 200,000 stars.

In both equations μ_c stands for $\mu_c\,\bar{m}$, with $\bar{m}$ the mean diagonal entry of $(I-W)^\top(I-W)$, so that μ_c is dimensionless. As $\mu_c \to \infty$ the solution is $C_{\rm geo}$ itself. At smaller μ_c, coordinate variations that poorly satisfy the neighbour relations are suppressed more strongly, and this can include physical structure as well as noise. The geodesic anchor retains the global arrangement, and the recovered synthetic grid provides a check on the retained structure.

2.5. *Transferring labels through the graph*

For a fixed configuration the coordinates, graph and refinement weights are built without labels. Labelled stars used to choose that configuration are counted separately from the training set. The stars of the training set carry labels, and every other star receives its labels from them through the weights of the refinement.

The problem has the form of the refinement with two differences. The labels take the place of the coordinates, and the anchor acts only on the training set, pulling each of its stars toward its known labels, where the refinement pulled every star toward $C_{\rm geo}$. Let $\mathbf{y}_i$ be the known labels of star i and Y the labels of every star, one row per star. The labels are

$$Y = \arg\min_{Y} \; \big\|(I-W)\,Y\big\|^2 + \mu_{\rm y} \sum_{i\in\mathcal{T}} \big\|Y_i - \mathbf{y}_i\big\|^2, \quad (11)$$

All labels are solved together, so information travels through intermediate unlabelled stars.

where $\mathcal{T}$ is the training set, Y_i is row i of Y, and the anchor strength $\mu_{\rm y}$ is set by $\bar{m}$ as before, sets how strongly the training set is pulled toward its labels. As for the refinement, setting the gradient to zero gives one linear solve,

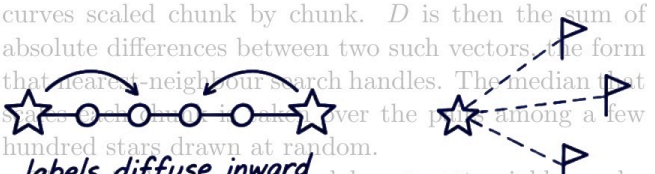
$$\big[(I-W)^\top(I-W) + \mu_{\rm y} P\big]\,Y = \mu_{\rm y}\,P\,\hat{Y}, \quad (12)$$

where P is the diagonal matrix with ones on the rows of the training set and zeros elsewhere, and $\hat{Y}$ holds the known labels on those rows and zeros elsewhere. The labels of all stars are determined together, so a star far from the training set receives its labels through the stars between it and the nearest training star. The label transfer adds one parameter, the anchor strength $\mu_{\rm y}$.

2.6. *Implementation*

A direct implementation would form several dense matrices over the stars. The matrix of D takes memory proportional to N^2, about 160 GB in single precision for $N = 2\times10^5$. Dijkstra's algorithm from all N stars on the sparse k-neighbour graph costs $O(kN^2\log N)$, or $O(N^2\log N)$ for fixed k, and storing all geodesic distances also takes $O(N^2)$ memory. The dense eigendecomposition of B and the dense refinement solve cost $O(N^3)$. We therefore replace these steps by computations that avoid the full dense matrices and give the same result, exactly or to a close approximation.

Only each star's nearest neighbours enter the graph, so the matrix of D is never needed, only a search for those neighbours. The integral that gives the chunk distance is a sum of absolute differences over the pixels, so each star can be written as one vector of its cumulative curves scaled chunk by chunk. D is then the sum of absolute differences between two such vectors, the form that nearest-neighbour search handles. The median that sets each chunk is taken over the pairs among a few hundred stars drawn at random.



The neighbours are found by nearest-neighbour descent (Dong et al. 2011), which starts from random neighbour lists and improves them by testing the neighbours of each star's current neighbours. Its computational cost depends on the sample geometry and the number of iterations, and it avoids evaluating every pair. For 256 stars drawn at random in each survey, the 100-neighbour lists overlap those from an exact search by more than 99.8 per cent on average.

The shortest paths and the eigendecomposition are run through a subset of stars, following de Silva & Tenenbaum (2004). A few hundred stars, M of them, are drawn at random as landmarks. Dijkstra's algorithm runs from those M stars alone, reducing the shortest-path cost to $O(MkN\log N)$, and the eigendecomposition then works on an $M\times M$ matrix of geodesic distances among them. Every other star is then placed from its geodesic distances to the landmarks, as a position is fixed from its distances to a few known points. Its squared distances, less the landmarks' mean squared distances among themselves, are the star's row of the centred matrix of squared distances. The inner products with the landmark coordinates that follow from it, divided by the eigenvalues, are its coordinates.

Landmarks replace an impossible all-pairs geodesic matrix with a few hundred reference distances.

The refinement and the label transfer are sparse linear systems, since a row of W has only as many entries as a star has neighbours. They are solved by conjugate gradients (Hestenes & Stiefel 1952), which does not form or invert the matrix. It needs only the matrix's product with a vector, two passes over the entries of W, repeated until the residual falls below a set tolerance.

Each replacement was checked against the dense original on a subset small enough to run both. The whole chain from spectra to coordinates runs in under an hour on one node for 2×10^5 stars, 60 per cent of it in the neighbour search.

The full coordinate pipeline finishes in under an hour on one node; neighbour search dominates.

On synthetic spectra, Armillary recovers the known three-dimensional label grid without folds or crossings.

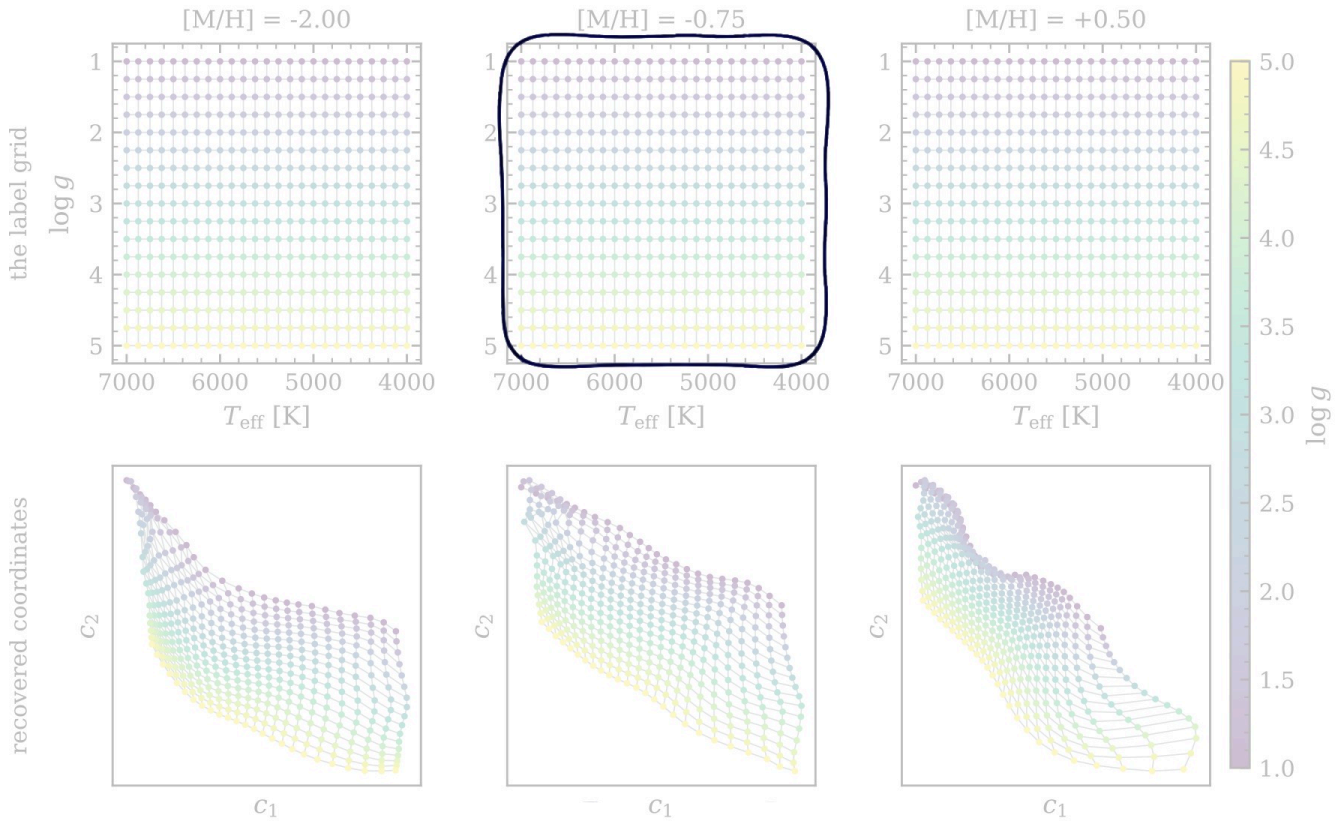


Figure 3. Armillary applied to the synthetic spectra of a $25 \times 17 \times 11$ grid in $T_{\rm eff}$, $\log g$ and [M/H] (spectra as in Fig. 2, steps as in Fig. 1). Each column is the grid's stars at one of its eleven metallicities. Top, those stars in $T_{\rm eff}$ and $\log g$. Bottom, the same stars in the first two recovered coordinates c_1 and c_2, with lines joining neighbouring grid points. The three metallicities are separated along c_3, which is not drawn. Colour is $\log g$ in both rows. The grid comes back with its lines in order and none crossing.

2.7. *Recovering a regular label grid from synthetic spectra*

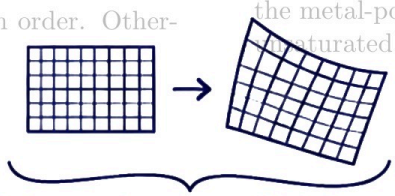
The aim of the construction is coordinates that are the stellar labels up to a smooth transformation. Real spectra cannot test this directly, because their labels are themselves estimates. Synthetic spectra can, since the labels that generated them are known exactly, and a regular grid of labels makes the test visual.

We build a grid of $25 \times 17 \times 11$ stars. $T_{\rm eff}$ runs from 4000 to 7000 K in steps of 125 K, $\log g$ from 1 to 5 in steps of 0.25, and [M/H] from -2 to $+0.5$ in steps of 0.25. Their spectra are synthesised with Payne-Zero as in Figure 2, and Armillary receives only the spectra. The settings are those adopted for APOGEE and DESI, adapted to the small grid, with the distance using the same 32 chunks as Figure 2 and all 4675 stars used as landmarks. If its coordinates are the labels up to a smooth transformation, the grid comes back as a smoothly deformed grid with its lines in order. Otherwise lines cross or fold.

Figure 3 shows the outcome at three of the eleven metallicities. The eigenvalues single out three coordinates, one per label, with the fourth eigenvalue 3.3 times smaller than the third. In every slice the grid comes back with temperature running one way, gravity the other, and no lines crossing.

The recovered grid is not a rectangle, since the coordinates record how much the spectra change from one grid point to the next and that change differs across the grid. A step of 125 K moves a cool spectrum further than a hot one. This is because many lines and the molecular bands respond to temperature at the cool end, while few lines remain at the hot end. The recovered grid is therefore stretched at its cool end and compressed at its hot end. A temperature step is 1.6 times longer at $T_{\rm eff} \leq 4500$ K than at $T_{\rm eff} \geq 6500$ K, and twice as long for the cool metal-rich dwarfs, whose crowded spectra respond most. A step in metallicity is 1.4 times longer at the metal-poor end than at the metal-rich end, because unsaturated lines grow in proportion to the abundance

This is the cleanest sanity check because the generating labels are exact, not pipeline estimates.



deformed is fine; folded is not

The recovered grid bends because equal label steps do not produce equal spectral changes everywhere.

and saturated ones barely grow (Gray 2005). Gravity steps are nearly even.

Under the Euclidean distance the same effect is larger, a factor of 41 between the largest and smallest median grid steps in the four combinations of cool or hot and metal-poor or metal-rich stars. That distance also [illegible] The Wasserstein distance pulls this proportion out and reduces the factor to 4. The grid is still not even, but the factor is ten times smaller than under the Euclidean distance.

Local continuum removal prevents calibration tilt from masquerading as shifted absorption.

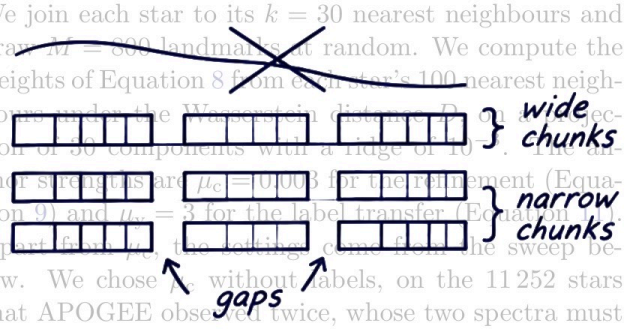
3. APPLICATION TO APOGEE SPECTRA

APOGEE (Majewski et al. 2017; Wilson et al. 2019) is a near-infrared survey at 1.51–1.70 μm and resolving power 22 500. Its pipeline ASPCAP (García Pérez et al. 2016; Jönsson et al. 2020) fits each spectrum with model spectra and reports $T_{\rm eff}$, $\log g$, the overall metallicity [M/H], [α/M], [C/M] and [N/M] from the fit of the whole spectrum. Individual abundances come from windows on their lines. The first four are the reference labels here. From DR17 (Abdurro'uf et al. 2022) we take every main-survey star with no ASPCAP bad bit, no star-level flag, finite labels and a signal-to-noise above 70. This gives 202 932 stars. Each enters with its combined spectrum from the reduction pipeline (Nidever et al. 2015), the one of highest signal-to-noise where a star has more than one.

We use the pipeline's pseudo-continuum-normalised spectra, which sit on three detectors with gaps between them. Chunks are built per detector so that none straddles a gap, in three chunkings of 11, 20 and 40 chunks, 16 to 4 nm wide. The local continuum is fitted in the chunks of the middle chunking, about eight nanometres wide, on top of the pipeline's normalisation. A slow tilt of that normalisation across a chunk would otherwise shift the absorption distribution as a real feature does. A pixel the pipeline masks has weight zero in the cumulative curve.

3.1. *The coordinates as physical axes*

Armillary is applied to the 202 932 spectra as follows. We join each star to its $k = 30$ nearest neighbours and draw $M = 800$ landmarks at random. We compute the weights of Equation 8 from each star's 100 nearest neighbours under the Wasserstein distance, $p = 0.5$, a projection of 50 components with a ridge of 10^{-3}. The anchor strengths are $\mu_{\rm c} = 0.003$ for the refinement (Equation 9) and $\mu_{\rm y} = 3$ for the label transfer (Equation 10). Apart from $\mu_{\rm c}$, the settings come from the sweep below. We chose $\mu_{\rm c}$, without labels, on the 11 252 stars that APOGEE observed twice, whose two spectra must receive the same coordinates. Their separation falls as $\mu_{\rm c}$ decreases and flattens below 0.003, which is adopted.



The other settings were fixed by one sweep on a random 60 000 of the stars. Each setting was varied in turn about the values above. Labels were transferred through the graph from training sets of 64, 256, 1024 and 16 384 stars, and the errors of the four labels were measured on held-out labelled giants. The number of neighbours of the weights was swept from 50 to 1600, $\mu_{\rm y}$ from 0.3 to 30, and the projection from 10 components to the spectrum itself. The ridge ϵ of the weights' solve was swept over [illegible] and the refinement [illegible] A [illegible] observed [illegible] ing sizes, 100 neighbours gives the minimum, [illegible] the error 15 per cent higher at 50 and 7 per cent higher at 400. Projections of 20 and 30 components are within 4 per cent of each other and 30 is adopted. The error decreases as $\mu_{\rm y}$ rises to 3 and changes by less than 4 per cent for larger values. The ridge of 10^{-3} is the best for training sets of 256 stars and fewer. Smaller ridges lower the [α/M] error by 15–16 per cent at 1024 training stars and 3–5 per cent at 16 384. The linear-probe R^2 of the coordinates changes by at most 0.05 in any label for graphs of 10 to 100 neighbours and 400 to 1600 landmarks. The paper's three scales of chunks give an error 1 to 3 per cent lower than one, two or four scales.

Repeated observations chose the refinement strength without consulting stellar labels.

Four coordinates are kept to represent the three stellar parameters and the α enhancement. The leading eigenvalues of B describe how the embedding is distributed across its dimensions. They fall as 1, 0.35, 0.14, 0.13, 0.09, 0.07, the fifth 30 per cent below the fourth. This is not a sharp separation of four physical axes, since weaker dimensions can contain noise and further spectral variation from individual abundances, and the geodesic distances need not be exactly Euclidean.

Figure [illegible] shows every pair of the four coordinates, coloured by [illegible] The first [illegible] runs along the Kiel diagram, [illegible] cool giants, the second across it, separating the two branches, and metallicity runs along the third. On the synthetic grid the coordinates were the labels up to a smooth transformation, and the question for the sample is how much of that holds. Coordinates built from distances are defined only up to a rotation, so no axis is privileged. The test is how closely a rotation and rescaling of the coordinates can follow the labels. That is a linear map $\mathbf{c} \to A\mathbf{c} + \mathbf{b}$ from the coordinates $\mathbf{c}$ to the labels, with A a matrix, which rotates and rescales, and $\mathbf{b}$ an offset. That map is the linear probe (Alain & Bengio 2016), and how much of the variance of a label it accounts for measures how linearly the coordinates follow that label.

Four dimensions are useful, but the eigenvalue drop is not a magical proof of exactly four physical parameters.

For APOGEE, four label-free coordinates emerge from 202,932 carefully normalised spectra.

The APOGEE coordinates visibly organise temperature, gravity, metallicity, and alpha enhancement before seeing any labels.

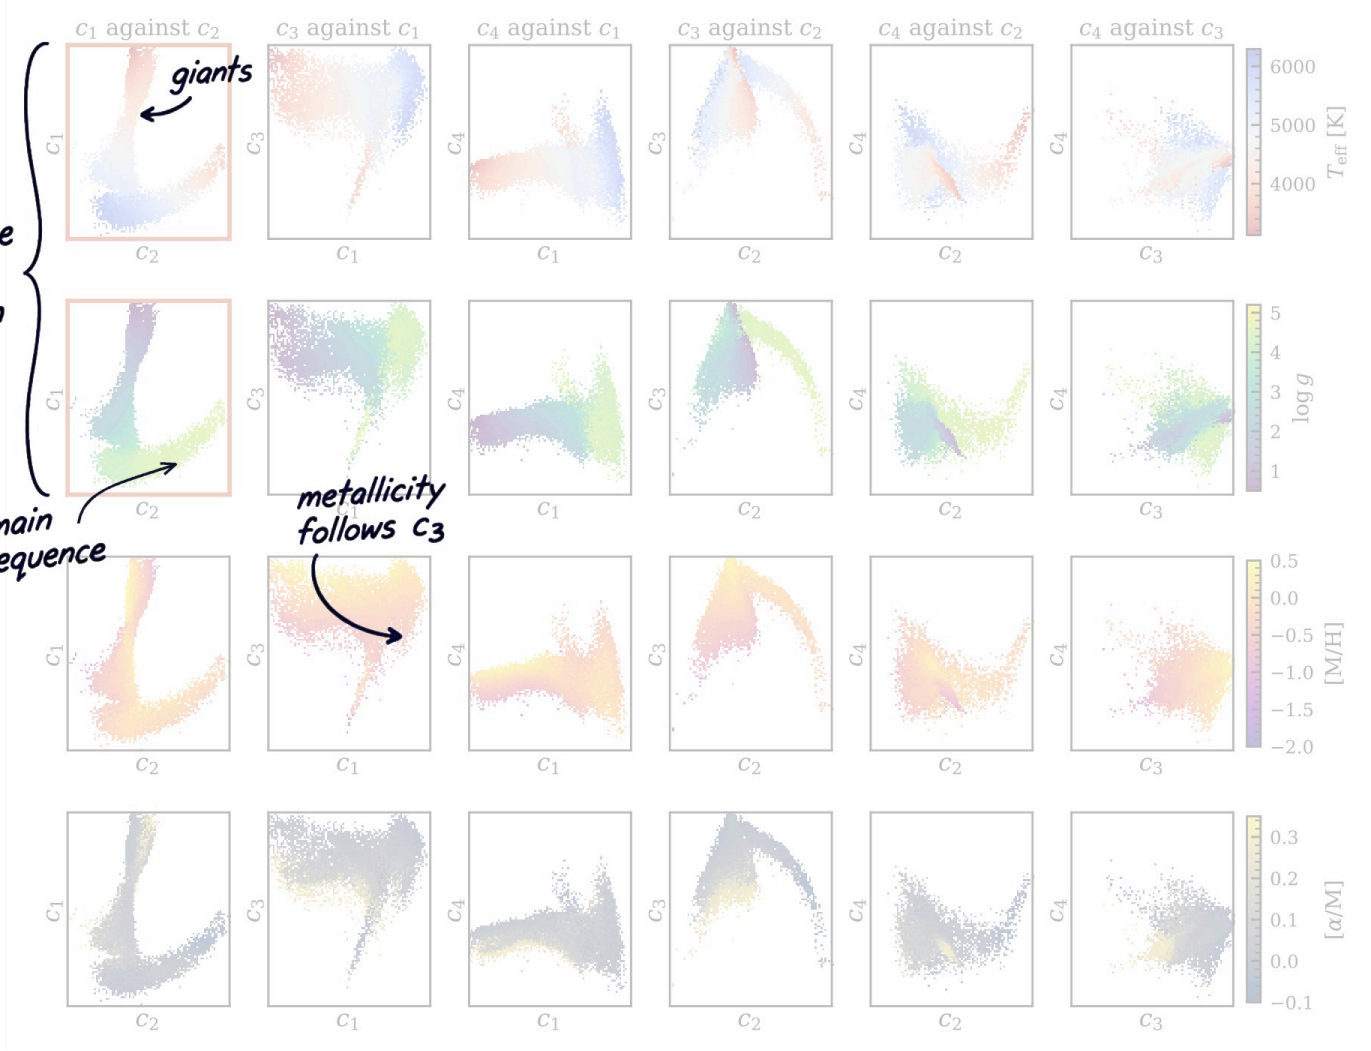


Figure 4. Armillary's four coordinates c_1 to c_4 of the 202 932 APOGEE stars. Each column plots one pair of coordinates, named above it as y against x in arbitrary units. Each row colours the same points by one ASPCAP label, $T_{\rm eff}$, $\log g$, [M/H] or [α/M], none of which entered the coordinates. Each cell shows the median label of the stars in it and is drawn where at least three fall. The boxed panels are the first two coordinates coloured by $T_{\rm eff}$ and by $\log g$. They already resemble the Kiel diagram, the giant branch rising to the upper right and the main sequence along the bottom, although no label was used in placing a star. Metallicity runs along c_3.

The same measure places Armillary beside other representations of the same spectra. Figure 5 puts the one-dimensional ordering of the Sequencer (Baron & Ménard 2021), principal component analysis (PCA, Pearson 1901; Hotelling 1933), t-SNE (van der Maaten & Hinton 2008), UMAP (McInnes et al. 2018) and an autoencoder (Hinton & Salakhutdinov 2006), each built on the same 202 932 stars, under the same linear map to $(T_{\rm eff}, \log g)$. Beside them is the Kiel diagram of the same stars drawn from ASPCAP's own labels. The map is fitted by least squares on all the stars and scored by R^2, separately for $T_{\rm eff}$ and $\log g$, which is 1 for an exact linear relation and 0 for none.

Instead, a linear probe asks whether a simple rotation and rescaling recover the physical axes.

Under the linear map Armillary's coordinates give the Kiel diagram of ASPCAP, and explain [illegible] cent of the variance of the two axes. [illegible] explains 76 and 91, PCA 75 and 91, UMAP 74 and 88 and t-SNE 69 and 76. The Sequencer ordering, one number per star, explains 15 and 56, and its map is a line through the Kiel diagram, since one number can hold one dimension of the variation. A linear map of Armillary's coordinates therefore recovers more of the sample's temperature and gravity variation than any of the other representations.

Coordinates are rotation-ambiguous; demanding one axis per label would be the wrong test.

The Sequencer ordering is built on Armillary's own graph. It ranks stars by their distance along the graph's minimum spanning tree from one end of its longest

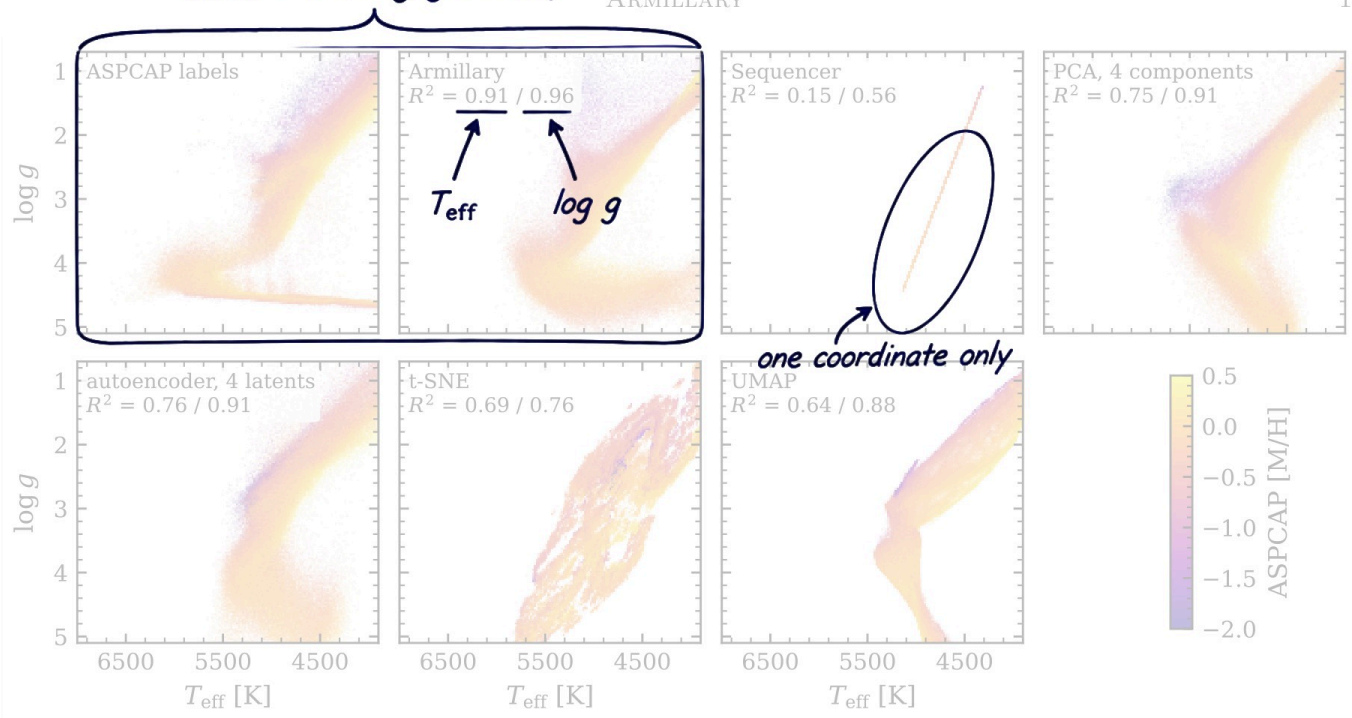


Figure 5. The Kiel diagram of the 202 932 APOGEE stars in ASPCAP's labels (top left), and the same stars placed by the linear probe, the best linear map from each representation's output to $(T_{\rm eff}, \log g)$. Each panel gives the R^2 of the map, the fraction of the variance of $T_{\rm eff}$ and of $\log g$ it accounts for. Colour is ASPCAP [M/H] throughout. Armillary, PCA, UMAP and the autoencoder are four-dimensional, t-SNE three, and the Sequencer ordering is one number per star, so its map is a line.

chain. PCA, t-SNE and UMAP are built on the same normalised spectra f_i that enter Armillary's distance, with masked pixels set to the pixel's mean over the sample, under Euclidean distances. PCA is kept to four components, the dimension Armillary uses, fitted on 30 000 of the stars and applied to all. UMAP is run in four dimensions, and t-SNE in three, the most its fast implementation allows at this sample size (van der Maaten 2014).

The hyperparameters of t-SNE and UMAP are chosen on a random draw of ten thousand stars. There each is run over a grid of them, t-SNE over perplexities from 5 to 100 with both of its initialisations, and UMAP over 5 to 100 neighbours and minimum distances from 0 to 0.5. The values with the highest linear-probe R^2 are kept for each, a choice made with the labels.

The autoencoder is a multilayer perceptron trained to reconstruct each spectrum through a bottleneck of four numbers, scoring the reconstruction only at unmasked pixels. Its encoder has two hidden layers, mirrored in the decoder. Its hyperparameters, the width and the regularisation, are chosen on the same draw. The 36 configurations are three widths, three weight decays, with and without dropout, and with and without noise added to the input at each pixel's own error. The choice is by the reconstruction error on a tenth of the spectra held out during training. The chosen configuration, 1024 and 256 hidden units with a weight decay of 10^{-5} and neither dropout nor input noise, is then trained on the whole sample. It scores [illegible] pixels that are unmasked for more than 95 per cent of the spectra.

We tuned the neural and nonlinear baselines rather than comparing against defaults.

The Sequencer becomes a line because one number cannot represent a two-dimensional stellar locus.

A quadratic map in place of the linear one can correct a smooth curvature of the coordinates but cannot undo a fold or a break in them. It separates the representations of Figure 5 in which the Kiel diagram is only curved from those in which parts of it are displaced. Fitted and evaluated the same way, it raises Armillary to 96 and 98 per cent, the autoencoder to 90 and 93, and PCA to 91 and 97, its broadened branches being curvature. The t-SNE and UMAP maps, which preserve each star's neighbours more closely than the arrangement of the neighbourhoods, reach 84 and 82, and 92 and 96. A quadratic map of Armillary's coordinates still recovers more of the temperature and gravity variation than any of the other representations.

3.2. *Transferring the labels from a few labelled stars*

Beyond placing the stellar parameters without a label, a representation is useful if a few labelled stars suffice to give every other star its labels, the few-shot setting of deep learning. Here the labels of n stars, the training set, are transferred through Armillary's graph, and the result is evaluated against ASPCAP on every other star. Throughout, the error of a label is a 1σ error, half the

A few dozen labelled APOGEE stars recover the broad stellar diagram, while chemical sequences need broader training coverage.

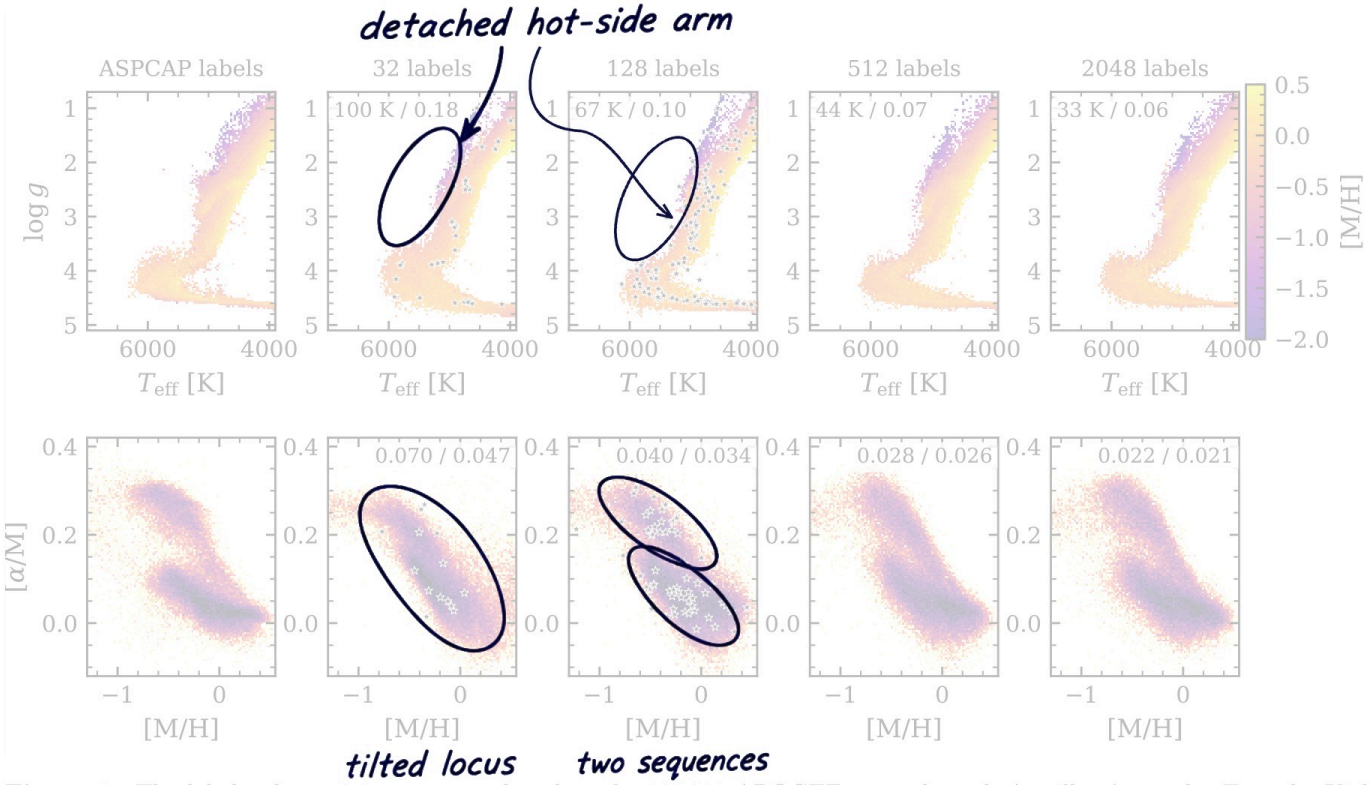


Figure 6. The labels of a training set transferred to the 202 932 APOGEE stars through Armillary's graph. Top, the Kiel diagram coloured by [M/H], and bottom, the [α/M]–[M/H] plane of the giants ($\log g < 3.5$) as a density on a logarithmic scale. The first column is ASPCAP's labels. The next four are the labels transferred from training sets of 32, 128, 512 and 2048 stars, every axis and the colour being the transferred labels. The star symbols on the second and third columns mark the 32 and 128 training stars at their ASPCAP labels, the giants among them below. The numbers are the 1σ errors against ASPCAP, $T_{\rm eff}$ and $\log g$ above over every star not in the training set, and [M/H] and [α/M] below over the giants shown.

width of the central 68 per cent of its differences from ASPCAP. It is the width of a Gaussian of the same core, which the few stars the graph places wrongly do not inflate.

A training set drawn at random follows the sample's own density, so the metal-poor giants and the high-α sequence, which are few, are the last regions it reaches. Until it does, the label transfer has no training star among the metal-poor giants and places them off the giant branch, on its hot side.

The coordinates exist before any label, so in practice they can decide which stars to follow up to obtain labels. The training set is drawn with probability proportional to the local density of the sample over the coordinates to the power $-1/4$, the density taken from the distance to the twentieth nearest star. This moves training stars toward the sparse regions. This matters most for the smallest training sets, in the extreme case of 32 stars. The exponent was set by a sweep on both surveys, and larger exponents place the training stars at the edges of the sample and reduce the precision at every training-set size.

Figure 6 shows the sample with labels transferred from training sets of 32, 128, 512 and 2048 stars, as the Kiel diagram and as the [α/M]–[M/H] plane of the giants ($\log g < 3.5$). The two sequences of the Galactic disc are seen in the solar neighbourhood (Bensby et al. 2003; Adibekyan et al. 2012; Anders et al. 2018) and across the Galaxy (Nidever et al. 2014; Hayden et al. 2015). From every training set the Kiel diagram shows the main sequence, the subgiant branch and the giant branch with its metallicity gradient. The two disc sequences are distinct from 128 stars, and from 32 they merge into a single tilted locus.

The training stars, marked by star symbols on the panels, explain the merging. At 32 every training giant below a metallicity of -0.45 lies on the high-α sequence, so the low-α stars there take their [α/M] from high-α training stars and the band tilts. At 128 the training giants cover both sequences down to a metallicity of -0.6 and the two separate. The label transfer takes labels to vary smoothly along the graph. The two sequences are two sheets of the manifold with few stars between them and at their metal-poor ends. Until a training star lies

Sparse coordinate regions are oversampled; random training sets reach rare metal-poor giants last.

At 32 labels, the α sequences merge because low-metallicity training giants cover only the high-α sequence.

Agreement with ASPCAP improves as the training set grows from 32 to 2,048 stars.

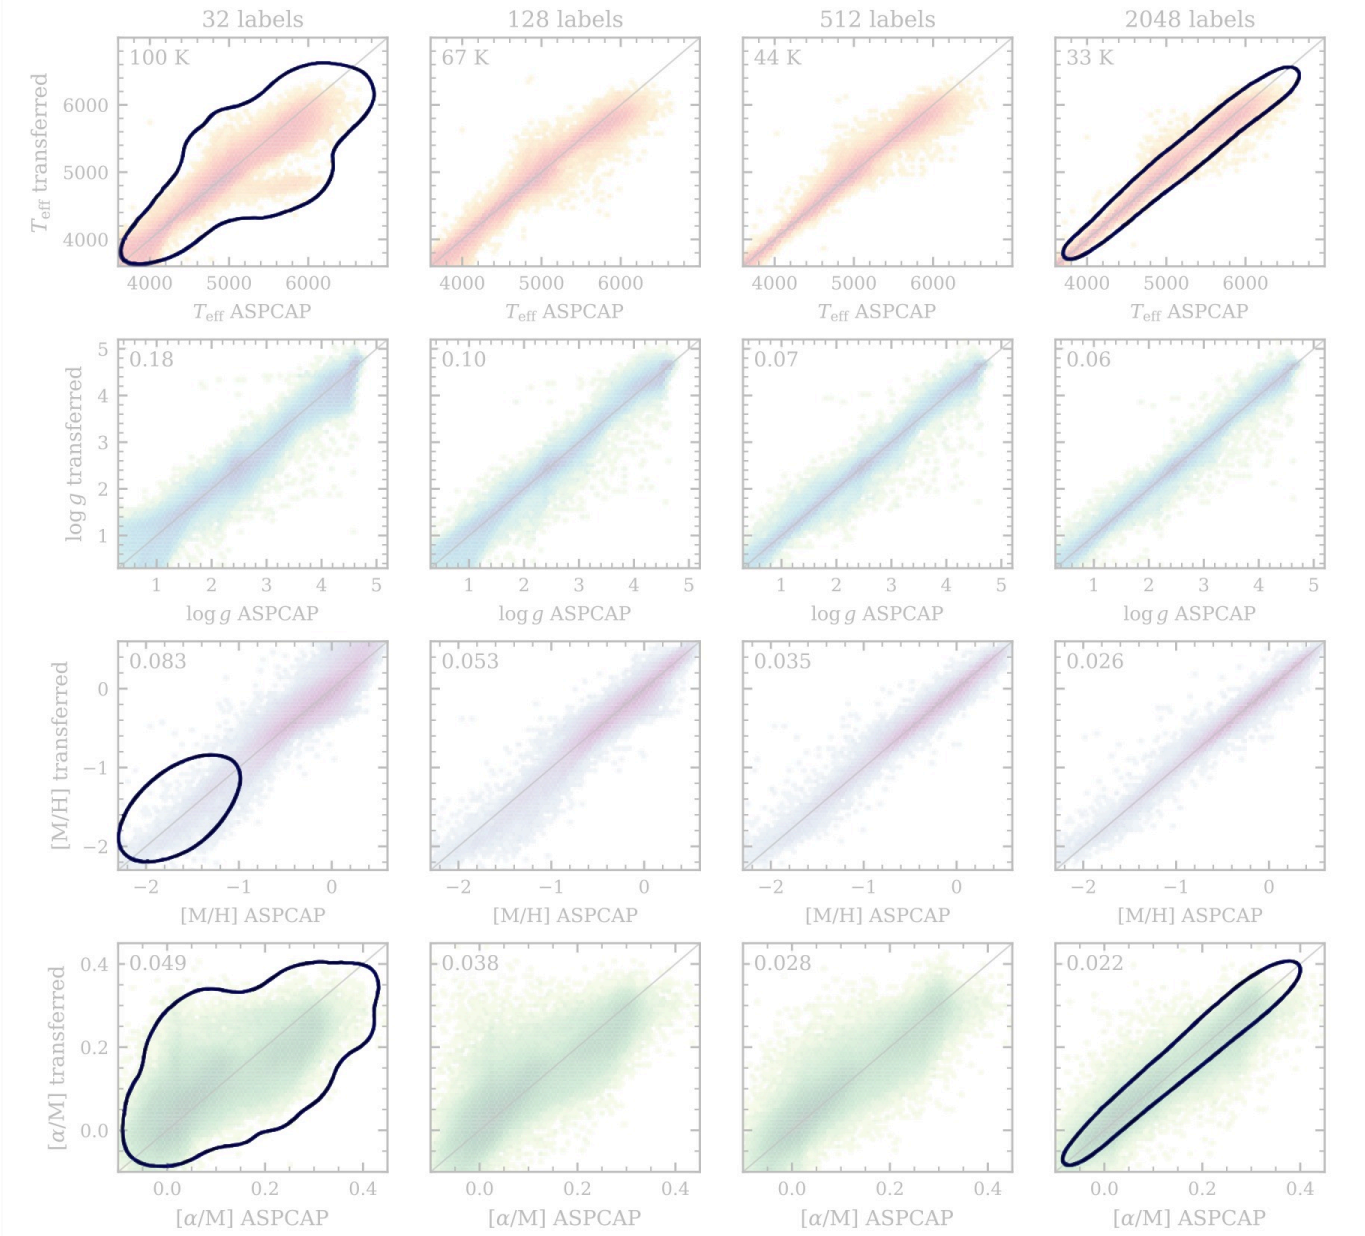


Figure 7. The labels transferred to the APOGEE stars against ASPCAP's, one row per label and one column per training set of 32, 128, 512 and 2048 stars, on every star not in the training set. Each panel is the density of the stars on a logarithmic scale, with the one-to-one line, and its number is the 1σ error of the difference. All four rows use the same held-out stars, including dwarfs and giants.

on each sheet in a region, the label transfer interpolates across the gap there as it would across any sparse region.

The Kiel diagram shows the same sparse coverage. At 32 and 128 training stars the metal-poor giants split off from the giant branch as a separate arm on its hot side, where ASPCAP's diagram is continuous. Only 2 of the 32 training stars and 5 of the 128 are giants below [M/H] $= -1$, so the labels of the whole strand rest on a few stars. By 512 training stars the arm rejoins the branch. The hot end of the main sequence is short at 32 and 128 training stars for the same reason. The stars that ASPCAP places above 6000 K receive 5850 K at the median from 32 training stars. A star at the edge of the sample has its neighbours, and so its training stars, on one side only.

Figure 7 shows each transferred label against ASPCAP's on every star not in the training set. Against ASPCAP, from a training set of 32 stars the 1σ errors are 100 K, 0.18, 0.083 and 0.049 dex. From 128 they are 67 K, 0.10, 0.053 and 0.038 dex. From 512 they are

The hot edge is biased inward because its neighbours lie on only one side.

These widths include ASPCAP's own errors, so they are not pure Armillary uncertainties.

Armillary matches the autoencoder with few training stars and is most precise from about 2,000.

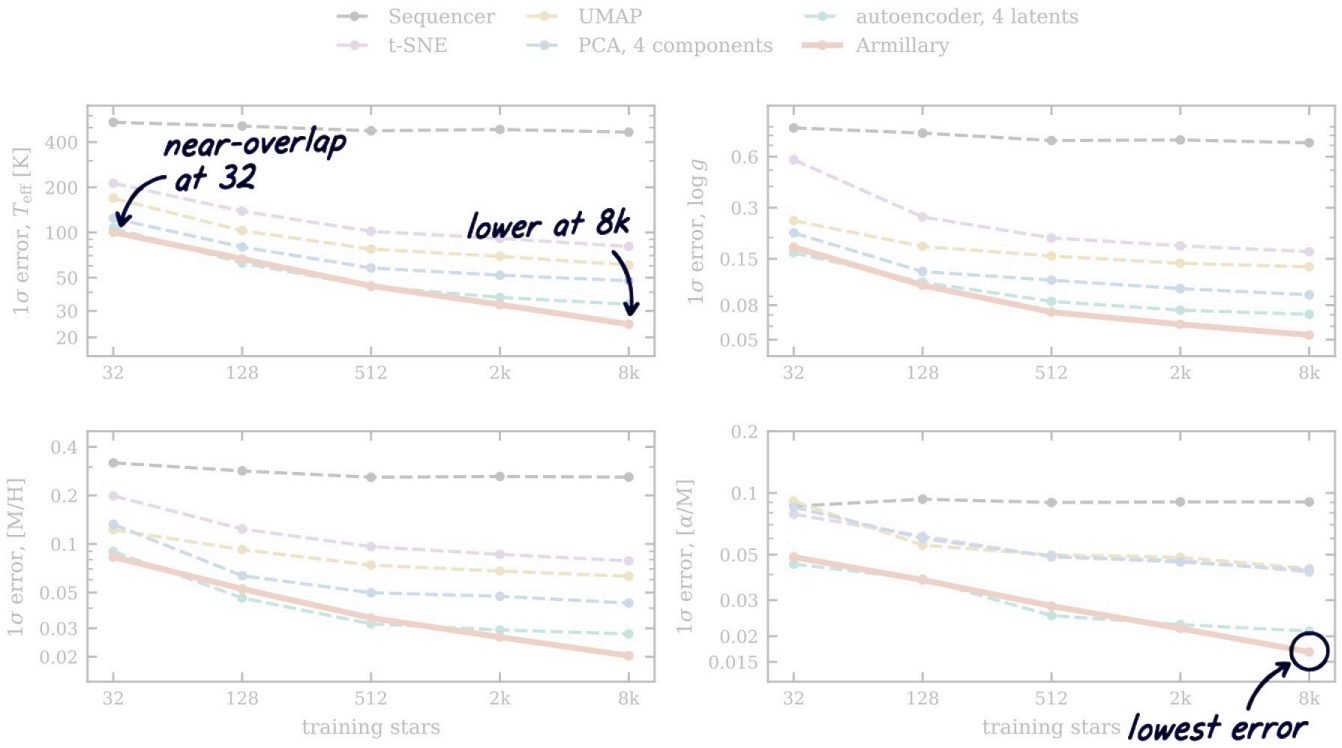


Figure 8. The 1σ error of each label against the size of the training set, on the 202 932 APOGEE stars, evaluated on every star outside each training set. The same training set is given to every representation. Armillary transfers its labels through the graph (orange, solid), and every other representation (dashed) by a multilayer perceptron from its output to the labels. Each perceptron is fitted on the training set at the weight decay of lowest error, chosen separately for each of the four labels.

44 K, 0.07, 0.035 and 0.028 dex. From 2048 they are 33 K, 0.06, 0.026 and 0.022 dex. The illustrated draw was selected for its recovered diagrams. Across three fixed draw seeds, the 32-star temperature error ranges from 100 to 147 K, narrowing to 33.0–33.8 K at 2048 stars.

The α enhancement lags the other labels because it changes the lines of a few elements, where the stellar parameters change the whole spectrum. Because the distance compares the absorption of a chunk as a whole, it separates two stars that differ in [α/M] more reluctantly than it separates stars that differ in [M/H] or $T_{\rm eff}$, and so it needs more training stars to resolve that difference.

Alpha enhancement changes fewer lines than temperature or gravity, so it needs more labelled stars.

Figure 8 compares Armillary with the other representations on all four labels, over training sets from 32 to 8192 stars, a factor of four apart. The same training set is given to each representation. Armillary transfers its labels through the graph. Each other representation is read by a multilayer perceptron from its output to the labels, with two hidden layers of 128 units, fitted on the training set to convergence. Its weight decay is swept from 10^{-4} to 10^{-1}, and the value of lowest error on the evaluated stars is kept, separately for each of the four labels.

From a training set of two thousand stars Armillary is the most precise representation, and its lead over the autoencoder trained on the sample grows with the training set. PCA, t-SNE and UMAP are less precise than both, and the Sequencer ordering is the least precise. Armillary therefore matches the neural network from a few dozen training stars and leads it from two thousand.

Every error above is measured against ASPCAP, whose labels carry errors of their own, and so includes them. How repeatable Armillary's labels are on their own is measured on the 11 252 stars observed twice. Their second spectra were reduced independently and are added to the sample for this test. Each of these stars then enters twice and must receive the same coordinates and the same labels from either spectrum. The median separation of the two coordinate vectors, divided by $\sqrt{2}$ and by the median distance to the tenth nearest neighbour, falls from 2.99 to 1.81 after refinement.

For the labels, a training set of 1024 stars is drawn as above from the stars observed once, so that neither spectrum of a twice-observed star is in it. Every twice-observed star then receives a predicted label from each spectrum. Dividing the 1σ scatter of the pair differences by $\sqrt{2}$ gives an external repeatability of 27 K, 0.03, 0.013 and 0.011 dex. The same training

Repeated observations test internal coordinate repeatability, independently of the ASPCAP labels.

set measured against ASPCAP gives 38 K, 0.07, 0.029 and 0.023 dex. The corresponding pair scatter divided by $\sqrt{2}$ for ASPCAP's own two determinations, one from each spectrum, is 18 K, 0.03, 0.012 and 0.009 dex. The two transferred labels share a graph and training set, so their errors can be correlated. This repeatability does not measure independent absolute accuracy.

4. APPLICATION TO DESI SPECTRA

The DESI Milky Way Survey (Cooper et al. 2023) observes stars at a resolving power of 2000 to 5000 over 360 to 980 nm with cameras called b, r and z (DESI Collaboration et al. 2023). Stellar parameters, and abundances above all, are hard to extract at this resolution, although the information is there (Ting et al. 2017a). Data-driven models recover individual abundances from APOGEE H band DESI spectra (Ting et al. 2017b; Xiang et al. 2019; Zhang et al. 2024). They train on labels from a high-resolution survey, for DESI mostly through LAMOST labels tied to APOGEE. A prior from model spectra keeps the learned effect of each element on the pixels where model spectra say it has lines. A model trained on the labels alone can otherwise read an abundance from whatever correlates with it in the training set. Armillary applies to DESI what APOGEE showed. The question is whether a label free map of DESI spectra, with APOGEE labels transferred through it, gives the Kiel diagram and the two α sequences with no model of the spectra. It does, for the bright majority of stars.

From Data Release 1 (DESI Collaboration et al. 2026) we take the 4799 stars of the APOGEE sample that DESI observed with a successful flagged fit in the survey's stellar pipeline (Koposov et al. 2024, 2026) and a signal-to-noise above 100. This selection uses the catalogue's red-camera signal-to-noise, $\mathrm{SN_R}$. To them we add 100 000 DESI stars without an ASPCAP label. These stars are above the same signal-to-noise, and by the survey pipeline's own $T_{\rm eff}$, $\log g$ and [M/H] they lie within the region that the cross-matched stars cover. The graph then has many stars where the label transfer is evaluated. Only the 4799 carry ASPCAP labels.

The spectra are shifted to rest frame with the survey's own velocity and resampled uniformly in $\ln\lambda$. On that grid a pixel spans the same velocity interval at every wavelength. Each camera's flux and native inverse-variance profile are interpolated linearly, with the latter supplying the weights when cameras overlap. DESI's spectra come in calibrated flux (Guy et al. 2023), so the normalisation that APOGEE's pipeline supplied is done here by a running continuum, the 90th percentile of the flux over a window of 100 pixels in each camera.

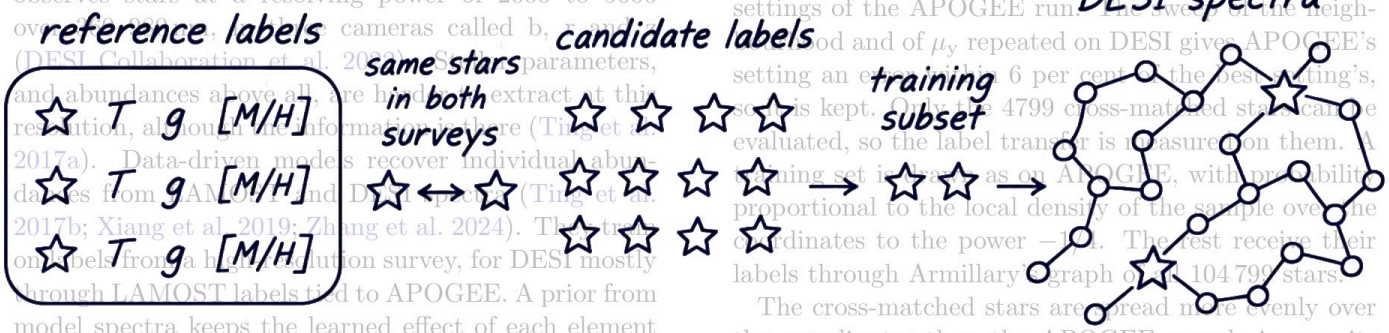
The interpolated inverse-variance profile, combined across cameras, also supplies the relative distance weights. We mask the telluric bands and then fit the local continuum as for APOGEE, over windows of 20 to 60 nm. We build the chunks per camera in four chunkings of 10, 20, 40 and 80 chunks, of equal width in $\ln\lambda$, so that a chunk spans a similar velocity interval wherever it lies. A chunk is 47 nm wide at 500 nm in the coarsest chunking and 6 nm in the finest.

The same chain places the 104 799 DESI stars with the settings of the APOGEE run. The sweep of the neighbourhood and of μ_y repeated on DESI gives APOGEE's setting an error inside 6 per cent of the best setting's, so this is kept. Only the 4799 cross-matched stars can be evaluated, so the label transfer is measured on them. A training set is chosen as on APOGEE, with probability proportional to the local density of the sample over the coordinates to the power $-1/2$. The rest receive their labels through Armillary's graph of the 104 799 stars.

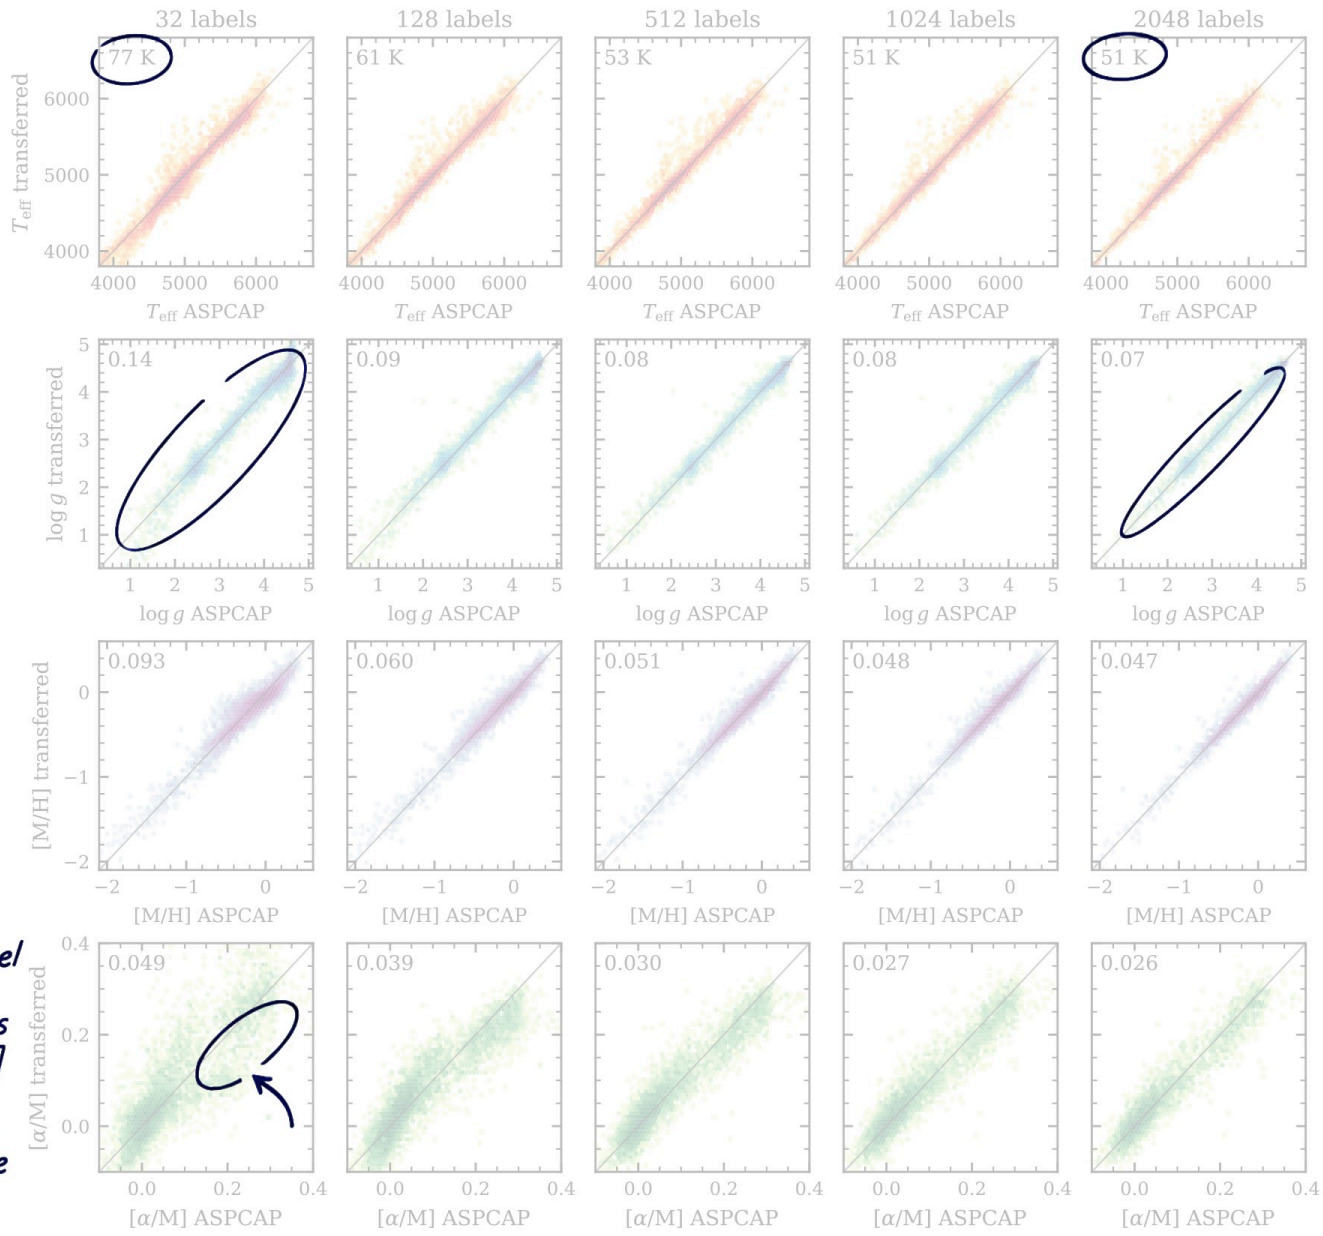
The cross-matched stars are spread more evenly over the coordinates than the APOGEE sample is over its own, so the draw changes less here. The draw still places a training star near the metal-poor dwarfs of the halo, which are few among the cross-matched stars and many in the survey. A random training set of fewer than a thousand stars leaves them to be placed off the main sequence.

Most of the survey is fainter, with a median signal-to-noise near 23 in catalogue $\mathrm{SN_R}$, and the distant halo stars of most interest are fainter still. Fainter DESI stars have no match to supply labels, so the sample is degraded instead. Independent Gaussian noise, scaled to each pixel's error, is added to lower the red-camera signal-to-noise to one of five levels. The graph is rebuilt and the labels transferred at each level, with the same spectra and training stars throughout.

Here signal-to-noise is the median flux divided by its error over usable red-camera pixels, after the noise is added. Its median over the stars is 148 as observed, and 113, 87, 63, 41 and 20 at the five degraded levels. Figure 8 shows the label transfer by Armillary against ASPCAP at four of these levels, with 1000 training stars. The training set is the same at every level, so the errors isolate the effect of the spectra alone. They contain roughly fifty times fewer held-out stars than Figure 7, so sparse tails are less visible, and the quoted errors provide the quantitative comparison. $T_{\rm eff}$, $\log g$ and [M/H] follow the one-to-one line to 1σ errors of 77 K, 0.14 and 0.093 dex from 32 training stars and 53 K, 0.08 and 0.051 dex from 512. The three stellar parameters, at least, are transferred from a few



At DESI signal-to-noise near 100, even 32 APOGEE-labelled stars recover the three main stellar parameters.



The 32-label panel compresses high-[α/M] values below the one-to-one line.

Figure 9. The labels transferred to the cross-matched DESI stars against ASPCAP's at signal-to-noise about 100. Each row is one label, and each column is one training set of 32, 128, 512, 1024 and 2048 stars. The panels show the cross-matched stars not in the training set. Each panel is the density of the stars on a logarithmic scale, with the one-to-one line, and its number is the 1σ error of the difference. All four rows use the same held-out stars, including dwarfs and giants.

dozen training stars. From merely 8 training stars the errors are 117 K, 0.21 and 0.21 dex, the Kiel diagram coarse but recognisable, although not shown.

Figure 10 shows the Kiel diagram of the whole graph at each signal-to-noise and training set. The giant branch is in place from the smallest training set at every signal-to-noise tested. A broad band at first, the main sequence narrows with the size of the training set and broadens with the noise. The errors from 1024 training stars rise with the noise from 51 K, 0.08 and 0.048 dex at signal-to-noise about 100 to 84 K, 0.15 and 0.089 dex at about 20. From 32 training stars at signal-to-noise about 20 the errors are 129 K, 0.25 and 0.203 dex and a streak of stars runs from the turn-off to the giant branch. The recovery of the Kiel diagram from small training sets down to a signal-to-noise of about 20 indicates that the graph already captures much of the stellar-parameter structure before the labels are attached.

For comparison, the survey's own pipeline labels, as published and without the metallicity calibration that

The DESI Kiel diagram persists at low S/N, while sparse training labels and noise broaden the stellar locus.

ASPCAP comparison

The graph captures much stellar-parameter structure before labels are attached; labels then map that structure to parameters.

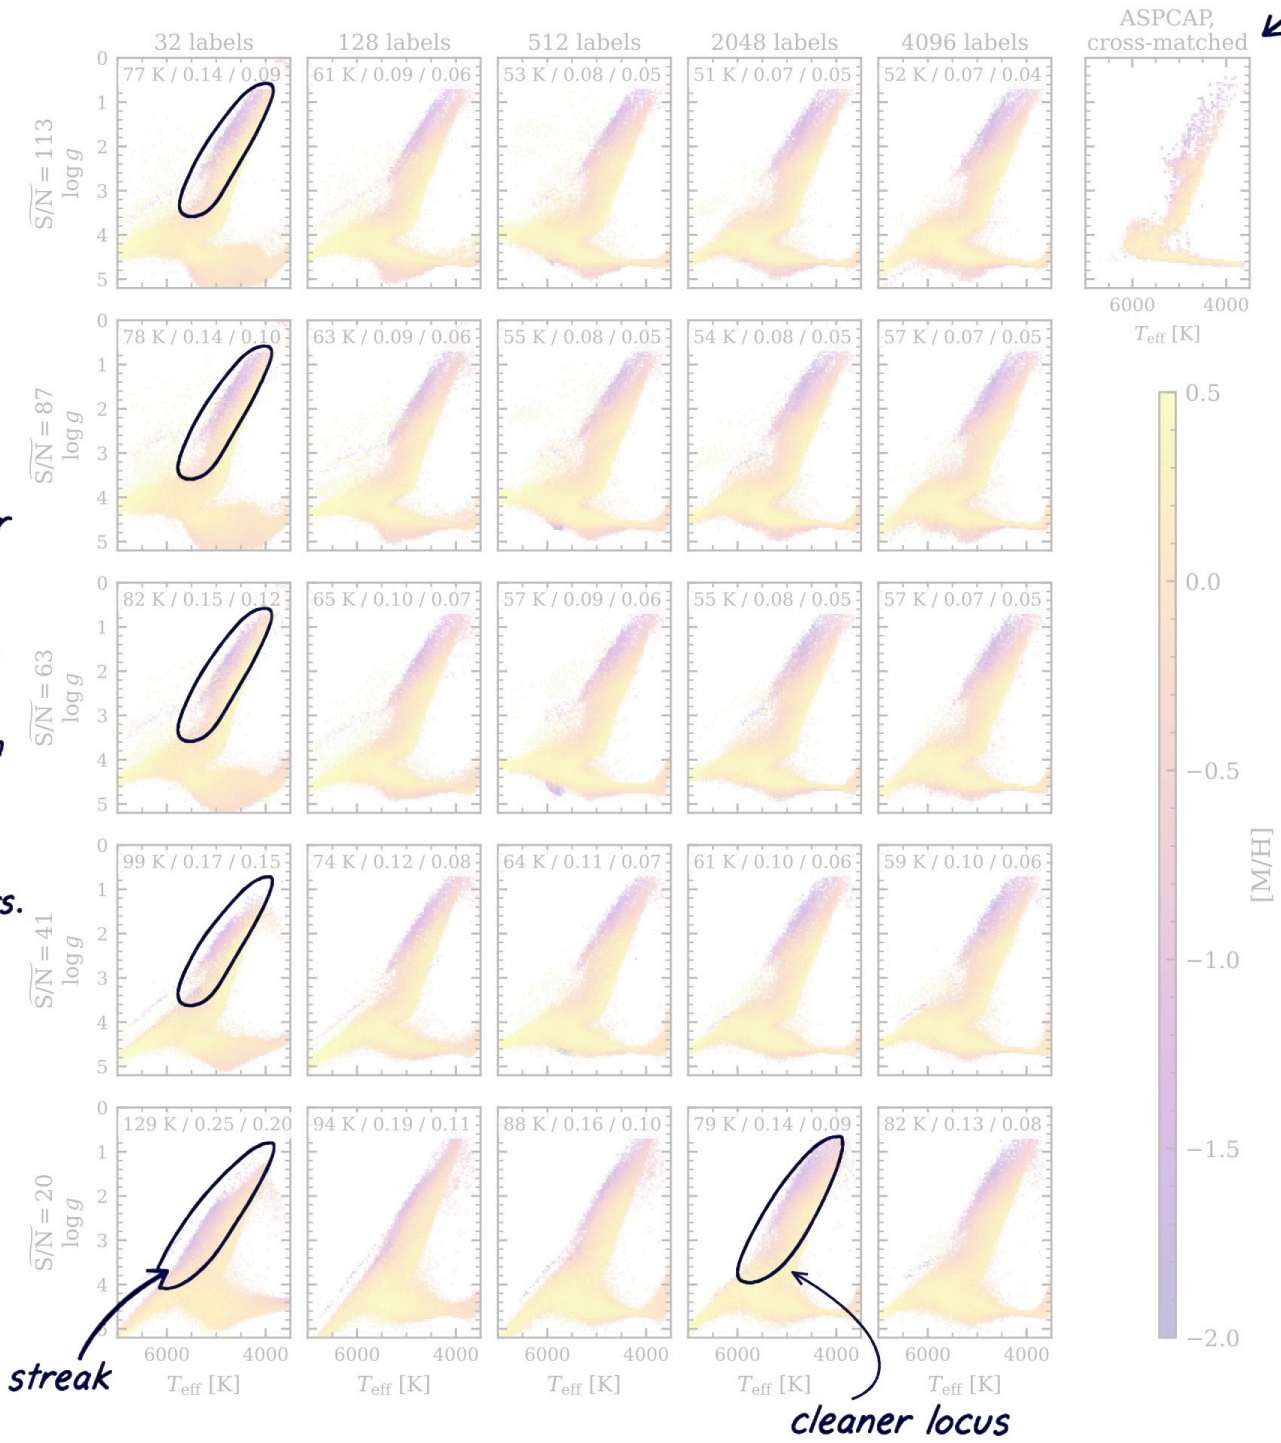


Figure 10. APOGEE's labels transferred to the DESI spectra through Armillary's graph, against the median red-grid signal-to-noise (rows) and the size of the training set (columns). Each panel is the Kiel diagram of the 104 799 DESI stars, coloured by [M/H], with the labels transferred from a training set of the cross-matched stars. The panel above the colour bar is the Kiel diagram of the 4799 cross-matched stars in ASPCAP's labels. The numbers are the 1σ errors against ASPCAP in $T_{\rm eff}$, $\log g$ and [M/H] over the cross-matched stars not in the training set.

At S/N = 20, more labels sharpen the locus, though errors remain larger than at higher S/N.

DESI's two α sequences become clearer with higher signal-to-noise and larger training sets.

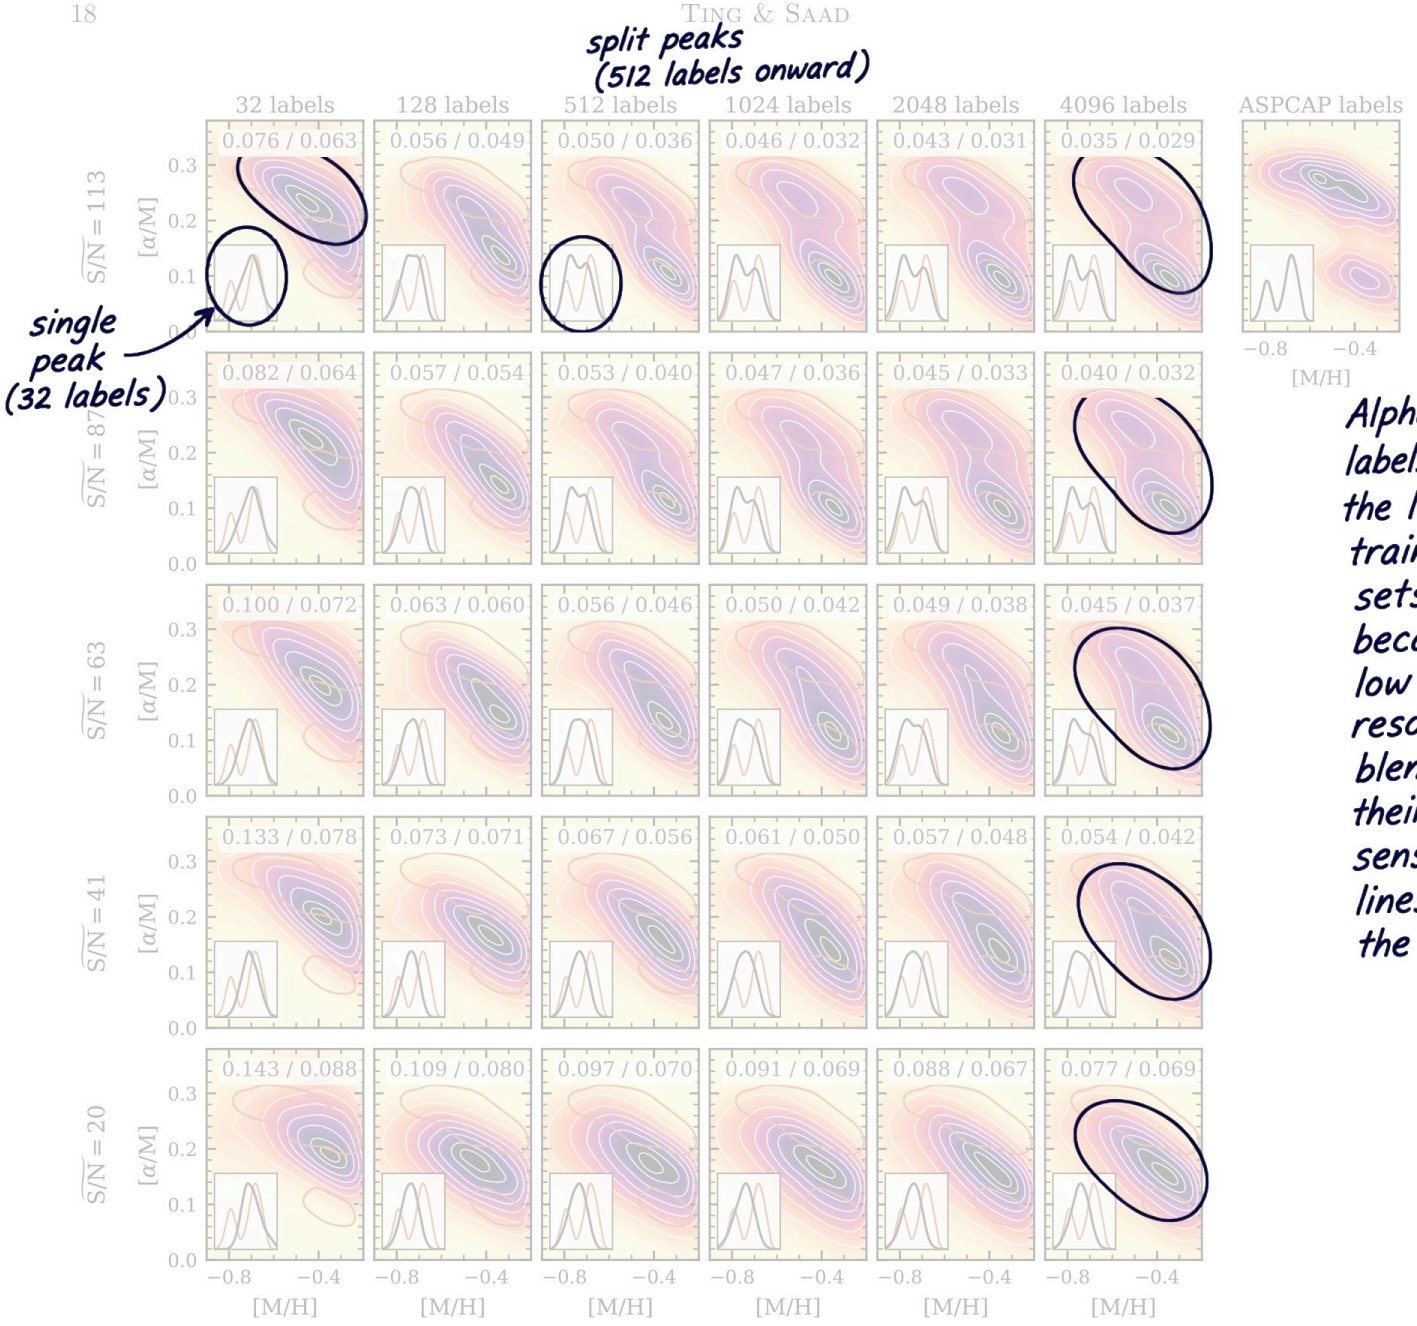


Figure 11. The emergence of α structure in DESI against the signal-to-noise of the spectra and the size of the training set, in the [α/M]–[M/H] plane of every giant of the graph by its transferred labels. Down the figure the noise increases to the median signal-to-noise named at the left, and the labels are transferred again through Armillary's graph. Across it the number of training stars increases through 32, 128, 512, 1024, 2048 and 4096. The panel at the top right is ASPCAP's labels for the cross-matched giants. Each panel is a smoothed density with white contours at the same fractions of its peak throughout, and the orange contour repeats one contour of the ASPCAP panel. The inset of each panel is the distribution of [α/M] of the giants with $-0.55 <$ [M/H] < -0.3, ASPCAP's in orange. The numbers at the top are the 1σ errors of the transferred [M/H] and [α/M] against ASPCAP's over the cross-matched giants not in the training set.

Koposov et al. (2026) provide, differ from ASPCAP. On the same held-out stars used for the 1024-star label transfer, the differences are 108 K, 0.26, 0.15 and 0.06 dex. Here we compare the pipeline's [Fe/H] and [α/Fe] to ASPCAP's [M/H] and [α/M]. The label transfer from 1024 training stars at signal-to-noise about 100 differs by 51 K, 0.08, 0.048 and 0.027 dex. The transferred labels lie closer to ASPCAP than the survey's own pipeline does, as expected of a label transfer that

Even with 4096 labels, increasing noise weakens the two branches and raises the errors.

Armillary transfers stellar parameters from few labels, but α structure needs larger training sets.

takes ASPCAP's labels as its training set. The two surveys then share one label scale.

[α/M] needs a larger training set. As on APOGEE, the α enhancement changes the lines of a few elements where the stellar parameters change the whole spectrum. Here at low resolution those lines are blended with the rest, so the label transfer needs the most training stars for this label.

Transferred labels share ASPCAP's scale because ASPCAP supplies the training labels.

Figure 12 shows [α/M] plane of every giant [illegible] training set. Each panel has an inset with the distribution of [α/M] of the giants with $-0.55 < [M/H] < -0.3$, where both of ASPCAP's sequences are present. From the smallest training sets the range of [α/M] is compressed, the high-α stars pulled down and the low-α stars up, and the compression eases with the size of the training set. At signal-to-noise about 100, two branches emerge from about 512 training stars and become more apparent with larger sets.

The noise weakens the branches and raises the errors. At 1024 training stars the errors of [M/H] and [α/M] in giants rise from 0.046 and 0.032 dex at about 100 to 0.091 and 0.069 dex at about 20. At about 80 the branches emerge from the same training set size of 512 but are shallower. At about 60 there are hints at 512 and definite shoulders from 2048. At about 40 a tentative hint of bimodality remains at 4096 training stars. The label-free map of DESI therefore gives the stellar parameters from a few dozen training stars and the two α sequences from about five hundred, at signal-to-noise above about 60.

5. DISCUSSION

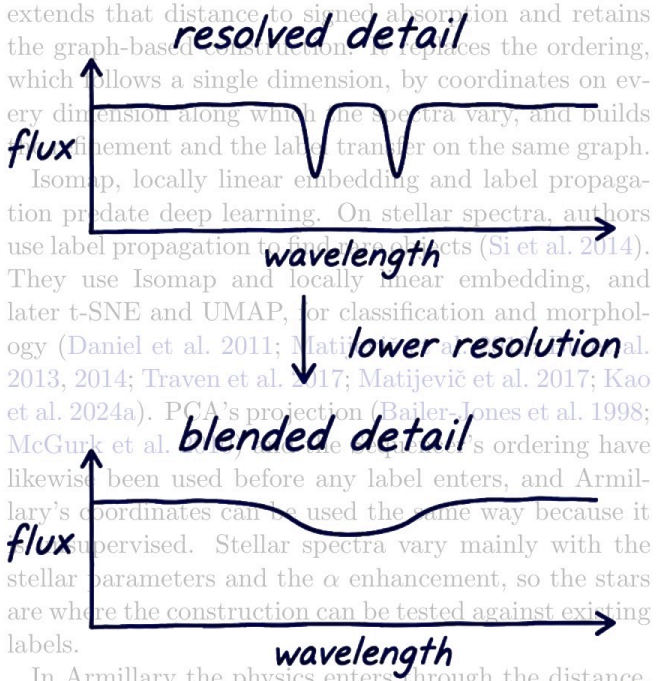
For DESI, roughly five hundred training stars reveal alpha structure when signal-to-noise exceeds about sixty.

[illegible] DESI [illegible] has a structure [illegible] spectra, and they lie on a manifold of a few dimensions. A representation that recovers that manifold from all the spectra lets a few labelled stars supply the labels of every other one. Deep learning makes that aim practical. Self-supervised neural networks now learn from unlabelled spectra first and take labels afterwards (Sedaghat et al. 2021; Zhao et al. 2026), and foundation models trained on spectra and labels together are adapted to new tasks (Leung & Bovy 2024; Koblischke & Bovy 2024). A deep neural network can learn flexible nonlinear relations. It returns, however, a latent representation in which the physical parameters are often entangled, and which usually needs labels for a physical interpretation.

The question asked here is how much of that structure classical machine learning tools recover when they are given the right ingredients. Armillary defines a distance between spectra that respects the physics of absorption, builds the graph of nearest neighbours under it, takes coordinates from the shortest paths through the graph, and transfers labels through it. On APOGEE the four coordinates follow the three stellar parameters, temperature, gravity and metallicity, closely under a linear probe. On both surveys the graph transfers the stellar parameters from a small training set.

5.2. *Classical manifold methods*

Armillary follows the key ideas of classical manifold learning. The coordinate step is Isomap (Tenenbaum et al. 2000), the refinement's weights are those of locally linear embedding (Roweis & Saul 2000), and the label transfer is the label propagation of semi-supervised learning (Zhu & Ghahramani 2002; Zhou et al. 2004) in its form through locally linear weights (Wang & Zhang 2008). The Sequencer (Baron & Ménard 2021) applied the Wasserstein distance of optimal transport, among other distances, to spectra and showed that a single ordering built on them follows temperature. Armillary extends that distance to signed absorption and retains the graph-based construction. It replaces the ordering, which follows a single dimension, by coordinates on every dimension along which the spectra vary, and builds refinement and the label transfer on the same graph.



Isomap, locally linear embedding and label propagation predate deep learning. On stellar spectra, authors use label propagation to find [illegible] objects (Si et al. 2014). They use Isomap and locally linear embedding, and later t-SNE and UMAP, for classification and morphology (Daniel et al. 2011; Matijevič et al. 2012; Schiavon et al. 2013, 2014; Traven et al. 2017; Matijevič et al. 2017; Kao et al. 2024a). PCA's projection (Bailer-Jones et al. 1998; McGurk et al. 2015; [illegible]) and the Sequencer's ordering have likewise been used before any label enters, and Armillary's coordinates can be used the same way because it is unsupervised. Stellar spectra vary mainly with the stellar parameters and the α enhancement, so the stars are where the construction can be tested against existing labels.

In Armillary the physics enters through the distance. The Wasserstein distance compares where the absorption lies along the wavelength axis once the total amount of absorption is divided out. A step in temperature or gravity gives a similar distance for a cool star crowded with lines and for a hot star with few. A small shift in wavelength changes the distance only slightly. The Euclidean distance is dominated by how much absorption a spectrum contains, so the same step is large for the [illegible] built on it is stretched in one corner of the label space

Blending hides the signature of a few elements, so abundance transfer needs broader label coverage.

The method is interpretable and fast, but it remains a sample-level graph rather than a per-spectrum generative model.

and compressed in another, and the coordinates inherit that unevenness. The Wasserstein distance makes the geometry of the sampled manifold more even in these tests, improving the local distances on which the graph and coordinates depend.

Beyond the samples analysed here, the coordinates of quasar or galaxy spectra, shaped by more than one process, can show along which dimensions a sample varies at all. The same holds for white dwarf spectra or stellar light curves, each with a distance that respects its own physics. Authors apply manifold methods to the spectra of galaxies, quasars and white dwarfs (Vanderplas & Connolly 2009; Richards et al. 2009; Kao et al. 2024b,a). A spectrum far from every neighbour, from a reduction defect or a peculiar object, shows as such on the same manifold (Matijevič et al. 2012).

5.3. *Combining Armillary with spectral model fitting*

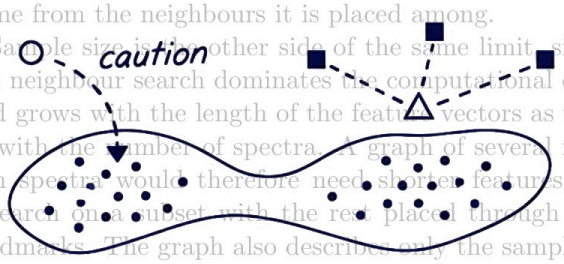
The supervised data-driven models complement Armillary. The Cannon (Ness et al. 2015; Casey et al. 2016) and DD-Payne (Ting et al. 2017b; Xiang et al. 2019) learn a function from labels to spectra. Trained on thousands of labelled stars, as in their survey-scale applications (Casey et al. 2016; Xiang et al. 2019), they fit the whole spectrum of each star and recover the abundances of individual elements. Armillary needs a few dozen to a few hundred labelled stars. Where labelled stars are few, as for a new instrument or a wavelength range with few reference stars, the graph transfers their labels from the handful that exist. The individual elements are out of its reach as it stands, because of the distance. It compares the shape of the absorption over chunks of a few nanometres. It therefore resolves what many lines carry together, the stellar parameters and the α enhancement, before an element with one or two weak lines.

The α enhancement already shows where that limit lies. The label transfer of the stellar parameters needs a few dozen training stars. That of the α enhancement needs a hundred on APOGEE and five hundred on DESI at signal-to-noise above about 60. The difference reflects the weaker imprint of the α lines on the distance, which weighs them with every other line of the chunk. Weighting the pixels that a line list singles out can improve sensitivity to one element. A distance made sensitive to every element would, however, reduce to fitting a physical model of the spectra. That would defeat the aim of this study, which is to find how far a purely data-driven, unsupervised construction can go.

5.4. *Limitations*

Armillary's limits follow from its construction. The graph is built once, from the distances among all the spectra of a sample, and the labels are transferred through it. A new spectrum therefore has no place in it until the construction is rerun, and there is no fitting of new spectra to it. Each star is analysed as a whole this matters little. For a stream of new spectra one route is the landmark placement of the implementation, which positions a star from its distances to a few hundred landmarks without rebuilding the graph (de Silva & Tenenbaum 2004). Its labels would then come from the neighbours it is placed among.

New spectra currently require placement through landmarks or rebuilding the graph.



Sample size is the other side of the same limit, since the neighbour search dominates the computational cost and grows with the length of the feature vectors as well as with the number of spectra. A graph of several million spectra would therefore need shorter features, or a search on a subset with the rest placed through the landmarks. The graph also describes only the sample it was built from, a limit it shares with other data-driven models. A star of a kind the sample does not otherwise contain, a binary, a chemically peculiar star or a spectrum with a reduction defect, is placed wherever its spectrum is closest. It receives the labels of that neighbourhood, and only the geometry on the manifold warns of it.

A genuinely novel star is assigned to its nearest familiar neighbourhood; geometry warns us, but labels may be wrong.

The noise is handled less completely than in a fit. A model fitted to a spectrum, as the Cannon or the Payne do, weighs each pixel by its known uncertainty. Our distance averages the pixel noise within a chunk, and the reliability profile discounts the noisiest pixels, but neither weighs one pixel against another as a fit does.

Mixing stars of widely different signal-to-noise can also distort the neighbourhoods used for label transfer. Noise perturbs the distances, and wavelength-dependent noise can introduce structure that differs from star to star. The experiments therefore apply one noise prescription to the same sample at each level, although the achieved signal-to-noise still varies across stars. The label transfer of the stellar parameters and the α enhancement degrades gradually with the signal-to-noise, as Figures 10 and 11 show.

6. CONCLUSION

Pixel uncertainties enter less completely than in a full forward fit.

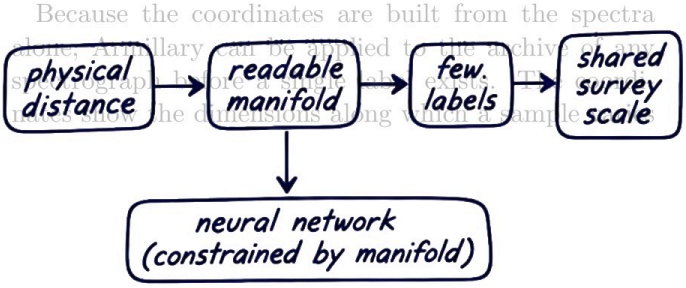
Most representation learning on spectra now uses deep neural networks, which learn a latent space from the unlabelled spectra and capture flexible nonlinear relations, although interpreting their latent representation usually requires labels. We ask, alongside that approach, how much of the structure of a stellar spectral sample the classical tools of manifold learning recover when they

A physics-aware distance lets classical manifold learning recover useful stellar coordinates without a black box.

are given the right ingredients. The construction that answers it and its main results are summarised below.

- The distance decides the geometry of the manifold of spectra. The Wasserstein distance between normalised absorption profiles compares where absorption lies once the total amount of absorption is divided out. It makes the manifold more even than the Euclidean distance does. The Euclidean distance stretches one corner of the manifold and compresses another.
- Under that distance the Sequencer's ordering generalises to the whole manifold. Shortest paths through the neighbour graph give coordinates in as many dimensions as the eigenvalues single out, in place of the one dimension an ordering can follow. A label-free refinement and the transfer of labels are built on the same graph.
- The coordinates come closest to temperature and gravity under a linear map. A linear map of the coordinates gives the Kiel diagram, recovering 91 per cent of the variance of $T_{\rm eff}$ and 96 per cent of that of $\log g$. This is more than PCA, t-SNE, UMAP or a tuned autoencoder give under the same map, so the geometry of the manifold carries the parameters before any of them is fitted.
- Given the same training sets, Armillary's label transfer is more precise than PCA, t-SNE, UMAP and the Sequencer ordering, and comparable to a tuned autoencoder. On the stellar parameters and the α enhancement the classical chain rivals representation learning with a neural network, and it is interpretable.
- A few dozen training stars recover the stellar parameters on APOGEE and on DESI. From 32 training stars the Kiel diagram of either survey is recognisable. The 1σ errors against ASPCAP are 100 K, 0.18 and 0.083 dex on APOGEE, and 77 K, 0.14 and 0.093 dex on DESI at signal-to-noise about 100. They improve to 33 K, 0.06 and 0.026 dex from 2048 stars on APOGEE and 51 K, 0.07 and 0.047 dex from 2048 on DESI. The two α sequences of the disc are distinct in giants from a hundred training stars on APOGEE, and on DESI they emerge from about five hundred at signal-to-noise above about 60.

Because the coordinates are built from the spectra alone, Armillary can be applied to the archive of any survey as soon as a few labelled stars exist. The [illegible] tells how the dimensions along which a sample of stars and the objects that lie apart from the rest. A few labelled stars, for example those shared with another survey, are then enough to place every other star on that survey's label scale, and so to homogenise the labels across surveys.



Armillary can suit the broad, low-resolution spectra of Gaia (Gaia Collaboration et al. 2023), Euclid (Euclid Collaboration et al. 2025), the Roman Space Telescope (Akeson et al. 2019) and SPHEREx (Bock et al. 2026), which are arriving, or about to arrive, in vast numbers. Its distance compares the broad shape of the absorption, which is what such spectra carry. Fast, label-free classical methods of this kind can be a practical way to extract stellar parameters at that scale.

Beyond these applications, a distance that respects the physics gives a manifold that classical machine learning tools recover, so for tasks of this kind a black box can be unnecessary. Neural networks, foundation models among them, can be trained under such a distance or held to the geometry it defines, and so return a latent representation that can be read. Understanding what such simple ingredients recover is a step toward understanding what large deep learning models learn from spectra.

This is most attractive for vast low-resolution archives where labels are scarce and spectra are abundant.

ACKNOWLEDGMENTS

We thank Todd A. Thompson, Christopher S. Kochanek, Krzysztof Z. Stanek and Xiaosheng Zhao for useful discussions, and Tomasz Różański and Jiadong Li for their work with YST on the earlier attempts described below. YST acknowledges a Humboldt Research Fellowship from the Alexander von Humboldt Foundation, held at the Max Planck Institute for Astronomy.

This project is the accumulation of several years of YST's failed attempts to make the Sequencer and autoencoders say something useful about stellar spectra, for the reasons laid out in this paper. It came back to life during work in preparation with Xiaosheng Zhao on foundation models built on joint-embedding predictive architectures. Their success at learning cleaner representations suggested that a clean representation of stellar spectra was possible, and left the question of whether it could also be understood. The way in was inspired by a discussion with Todd A. Thompson, Christopher S. Kochanek and Krzysztof Z. Stanek on distances in the spirit of the Wasserstein distance and on which spectra should count as adjacent. The manifold ideas are older. Isomap and locally linear embedding date back to before graduate school, when YST was trying to measure the dimensionality of chemical abundance space with principal component analysis (Ting et al. 2012). Some of

The broader lesson is not "never use neural networks", but "give geometry physical meaning first".

them were explored then, and came back more than a decade later.

The code and much of the brainstorming were done with GPT-5.6, GPT-6, Claude Opus 5, Claude Opus 5.1 and Claude Opus 5.5. The paper began as an extensive outline by the authors. The authors wrote the first draft of the introduction, motivation, discussion and conclusion. Claude drafted the method and the results, where holding every ablation detail made it more accurate, and did the copy editing throughout, although left to itself it writes Claudish rather than English, verbose and rarely the more helpful for it. The language models also canvassed the literature, which helped to map what had already been done. The direction of the project, and every result and sentence of the paper, were set or validated by the authors.

Software: NumPy (Harris et al. 2020), SciPy (Virtanen et al. 2020), scikit-learn (Pedregosa et al. 2011), Numba (Lam et al. 2015), PyNNDescent, umap-learn (McInnes et al. 2018), openTSNE (Poličar et al. 2024), PyTorch (Paszke et al. 2019), Astropy (Astropy Collaboration et al. 2013, 2018, 2022), Matplotlib (Hunter 2007)

REFERENCES

Abdurro'uf, Accetta, K., Aerts, C., et al. 2022, ApJS, 259, 35, doi: 10.3847/1538-4365/ac4414

Adibekyan, V. Z., Sousa, S. G., Santos, N. C., et al. 2012, A&A, 545, A32, doi: 10.1051/0004-6361/201219401

Akeson, R., Armus, L., Bachelet, E., et al. 2019, arXiv e-prints, arXiv:1902.05569, doi: 10.48550/arXiv.1902.05569

Alain, G., & Bengio, Y. 2016, arXiv e-prints, arXiv:1610.01644, doi: 10.48550/arXiv.1610.01644

Anders, F., Chiappini, C., Santiago, B. X., et al. 2018, A&A, 619, A125, doi: 10.1051/0004-6361/201833099

Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: 10.1051/0004-6361/201322068

Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, AJ, 156, 123, doi: 10.3847/1538-3881/aabc4f

Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, ApJ, 935, 167, doi: 10.3847/1538-4357/ac7c74

Bailer-Jones, C. A. L., Irwin, M., & von Hippel, T. 1998, MNRAS, 298, 361, doi: 10.1046/j.1365-8711.1998.01596.x

Baron, D., & Ménard, B. 2021, ApJ, 916, 91, doi: 10.3847/1538-4357/abfc4d

Bengio, Y., Courville, A., & Vincent, P. 2013, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35, 1798, doi: 10.1109/TPAMI.2013.50

Bensby, T., Feltzing, S., & Lundström, I. 2003, A&A, 410, 527, doi: 10.1051/0004-6361:20031213

Bock, J. J., Aboobaker, A. M., Adamo, J., et al. 2026, ApJ, 999, 139, doi: 10.3847/1538-4357/ae2be2

Bu, Y., Chen, F., & Pan, J. 2014, NewA, 28, 35, doi: 10.1016/j.newast.2013.09.007

Bu, Y., Pan, J., Jiang, B., & Wei, P. 2013, PASJ, 65, 81, doi: 10.1093/pasj/65.4.81

Casey, A. R., Hogg, D. W., Ness, M., et al. 2016, arXiv e-prints, arXiv:1603.03040, doi: 10.48550/arXiv.1603.03040

Cooper, A. P., Koposov, S. E., Allende Prieto, C., et al. 2023, ApJ, 947, 37, doi: 10.3847/1538-4357/acb3c0

Cui, X.-Q., Zhao, Y.-H., Chu, Y.-Q., et al. 2012, Research in Astronomy and Astrophysics, 12, 1197, doi: 10.1088/1674-4527/12/9/003

Daniel, S. F., Connolly, A., Schneider, J., Vanderplas, J., & Xiong, L. 2011, AJ, 142, 203, doi: 10.1088/0004-6256/142/6/203

Das, P. B., Zucker, D. B., De Silva, G. M., et al. 2025, MNRAS, 538, 605, doi: 10.1093/mnras/staf169

de Jong, R. S., Agertz, O., Berbel, A. A., et al. 2019, The Messenger, 175, 3, doi: 10.18727/0722-6691/5117

de Mijolla, D., Ness, M. K., Viti, S., & Wheeler, A. J. 2021, ApJ, 913, 12, doi: 10.3847/1538-4357/abece1

de Silva, V., & Tenenbaum, J. B. 2004, Sparse multidimensional scaling using landmark points, Tech. rep., Stanford University. https://graphics.stanford.edu/courses/cs468-05-winter/Papers/Landmarks/Silva_landmarks5.pdf

DESI Collaboration, Abareshi, B., Aguilar, J., et al. 2022, AJ, 164, 207, doi: 10.3847/1538-3881/ac882b

DESI Collaboration, Abdul Karim, M., Adame, A. G., et al. 2026, AJ, 171, 285, doi: 10.3847/1538-3881/ae4c43

Dijkstra, E. W. 1959, Numerische Mathematik, 1, 269, doi: 10.1007/BF01386390

Dong, W., Charikar, M., & Li, K. 2011, in Proceedings of the 20th International Conference on World Wide Web, 577–586, doi: 10.1145/1963405.1963487

Euclid Collaboration, Mellier, Y., Abdurro'uf, et al. 2025, A&A, 697, A1, doi: 10.1051/0004-6361/202450810

Fabbro, S., Venn, K. A., O'Briain, T., et al. 2018, MNRAS, 475, 2978, doi: 10.1093/mnras/stx3298

Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, A&A, 674, A1, doi: 10.1051/0004-6361/202243940

García Pérez, A. E., Allende Prieto, C., Holtzman, J. A., et al. 2016, AJ, 151, 144, doi: 10.3847/0004-6256/151/6/144
Gray, D. F. 2005, The Observation and Analysis of Stellar Photospheres, 3rd edn. (Cambridge University Press), doi: 10.1017/CBO9781316036570
Guy, J., Bailey, S., Kremin, A., et al. 2023, AJ, 165, 144, doi: 10.3847/1538-3881/acb212
Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: 10.1038/s41586-020-2649-2
Hayden, M. R., Bovy, J., Holtzman, J. A., et al. 2015, ApJ, 808, 132, doi: 10.1088/0004-637X/808/2/132
Hegedűs, V., Mészáros, S., Jofré, P., et al. 2023, A&A, 670, A107, doi: 10.1051/0004-6361/202244813
Hestenes, M. R., & Stiefel, E. 1952, Journal of Research of the National Bureau of Standards, 49, 409, doi: 10.6028/jres.049.044
Hinton, G. E., & Salakhutdinov, R. R. 2006, Science, 313, 504, doi: 10.1126/science.1127647
Ho, A. Y. Q., Ness, M. K., Hogg, D. W., et al. 2017, ApJ, 836, 5, doi: 10.3847/1538-4357/836/1/5
Hotelling, H. 1933, Journal of Educational Psychology, 24, 417, doi: 10.1037/h0071325
Hunter, J. D. 2007, Computing in Science and Engineering, 9, 90, doi: 10.1109/MCSE.2007.55
Jin, S., Trager, S. C., Dalton, G. B., et al. 2024, MNRAS, 530, 2688, doi: 10.1093/mnras/stad557
Jofré, P., Heiter, U., Soubiran, C., et al. 2014, A&A, 564, A133, doi: 10.1051/0004-6361/201322440
Jönsson, H., Holtzman, J. A., Allende Prieto, C., et al. 2020, AJ, 160, 120, doi: 10.3847/1538-3881/aba592
Kao, M. L., Hawkins, K., Rogers, L. K., et al. 2024a, ApJ, 970, 181, doi: 10.3847/1538-4357/ad5d6e
Kao, W.-B., Zhang, Y., & Wu, X.-B. 2024b, PASJ, 76, 653, doi: 10.1093/pasj/psae037
Koblischke, N., & Bovy, J. 2024, arXiv e-prints, arXiv:2411.04750, doi: 10.48550/arXiv.2411.04750
Kollmeier, J. A., Rix, H.-W., Aerts, C., et al. 2026, AJ, 171, 52, doi: 10.3847/1538-3881/ae0576
Koposov, S., Li, T. S., Allende Prieto, C., et al. 2026, The Open Journal of Astrophysics, 9, 55260, doi: 10.33232/001c.155260
Koposov, S. E., Allende Prieto, C., Cooper, A. P., et al. 2024, MNRAS, 533, 1012, doi: 10.1093/mnras/stae1842
Kruskal, J. B. 1956, Proceedings of the American Mathematical Society, 7, 48, doi: 10.1090/S0002-9939-1956-0078686-7
Kullback, S., & Leibler, R. A. 1951, The Annals of Mathematical Statistics, 22, 79, doi: 10.1214/aoms/1177729694
Kurucz, R. L. 2005, Memorie della Societa Astronomica Italiana Supplementi, 8, 14
Lam, S. K., Pitrou, A., & Seibert, S. 2015, in Proc. Second Workshop on the LLVM Compiler Infrastructure in HPC, 1–6, doi: 10.1145/2833157.2833162
Laroche, A., & Speagle, J. S. 2025, ApJ, 979, 5, doi: 10.3847/1538-4357/ad9607
Leung, H. W., & Bovy, J. 2019, MNRAS, 483, 3255, doi: 10.1093/mnras/sty3217
—. 2024, MNRAS, 527, 1494, doi: 10.1093/mnras/stad3015
Locatello, F., Bauer, S., Lucic, M., et al. 2018, arXiv e-prints, arXiv:1811.12359, doi: 10.48550/arXiv.1811.12359
Majewski, S. R., Schiavon, R. P., Frinchaboy, P. M., et al. 2017, AJ, 154, 94, doi: 10.3847/1538-3881/aa784d
Matijevič, G., Zwitter, T., Bienaymé, O., et al. 2012, ApJS, 200, 14, doi: 10.1088/0067-0049/200/2/14
Matijevič, G., Chiappini, C., Grebel, E. K., et al. 2017, A&A, 603, A19, doi: 10.1051/0004-6361/201730417
McGurk, R. C., Kimball, A. E., & Ivezić, Ž. 2010, AJ, 139, 1261, doi: 10.1088/0004-6256/139/3/1261
McInnes, L., Healy, J., & Melville, J. 2018, arXiv e-prints, arXiv:1802.03426, doi: 10.48550/arXiv.1802.03426
Ness, M., Hogg, D. W., Rix, H.-W., Ho, A. Y. Q., & Zasowski, G. 2015, ApJ, 808, 16, doi: 10.1088/0004-637X/808/1/16
Nidever, D. L., Bovy, J., Bird, J. C., et al. 2014, ApJ, 796, 38, doi: 10.1088/0004-637X/796/1/38
Nidever, D. L., Holtzman, J. A., Allende Prieto, C., et al. 2015, AJ, 150, 173, doi: 10.1088/0004-6256/150/6/173
Paszke, A., Gross, S., Massa, F., et al. 2019, arXiv e-prints, arXiv:1912.01703, doi: 10.48550/arXiv.1912.01703
Pearson, K. 1901, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 2, 559, doi: 10.1080/14786440109462720
Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825, doi: 10.48550/arXiv.1201.0490
Peyré, G., & Cuturi, M. 2018, arXiv e-prints, arXiv:1803.00567, doi: 10.48550/arXiv.1803.00567
Poličar, P. G., Stražar, M., & Zupan, B. 2024, Journal of Statistical Software, 109, 1, doi: 10.18637/jss.v109.i03
Richards, J. W., Freeman, P. E., Lee, A. B., & Schafer, C. M. 2009, ApJ, 691, 32, doi: 10.1088/0004-637X/691/1/32
Roweis, S. T., & Saul, L. K. 2000, Science, 290, 2323, doi: 10.1126/science.290.5500.2323
Rubner, Y., Tomasi, C., & Guibas, L. J. 2000, International Journal of Computer Vision, 40, 99, doi: 10.1023/A:1026543900054

Sedaghat, N., Romaniello, M., Carrick, J. E., & Pineau, F.-X. 2021, MNRAS, 501, 6026, doi: 10.1093/mnras/staa3540

Shen, J., Speagle, J. S., & Ho, S. 2026, ApJ, 1007, 69, doi: 10.3847/1538-4357/ae81b3

Si, J., Luo, A., Li, Y., et al. 2014, Science China Physics, Mechanics, and Astronomy, 57, 176, doi: 10.1007/s11433-013-5374-0

Smiljanic, R., Korn, A. J., Bergemann, M., et al. 2014, A&A, 570, A122, doi: 10.1051/0004-6361/201423937

Snell, J., Swersky, K., & Zemel, R. S. 2017, arXiv e-prints, arXiv:1703.05175, doi: 10.48550/arXiv.1703.05175

Székely, G. J., & Rizzo, M. L. 2013, Journal of Statistical Planning and Inference, 143, 1249, doi: 10.1016/j.jspi.2013.03.018

Tenenbaum, J. B., de Silva, V., & Langford, J. C. 2000, Science, 290, 2319, doi: 10.1126/science.290.5500.2319

Thomas, G. F., Battaglia, G., Gran, F., et al. 2024, A&A, 690, A54, doi: 10.1051/0004-6361/202450198

Ting, Y.-S., Conroy, C., Rix, H.-W., & Cargile, P. 2017a, ApJ, 843, 32, doi: 10.3847/1538-4357/aa7688

—. 2019, ApJ, 879, 69, doi: 10.3847/1538-4357/ab2331

Ting, Y.-S., Freeman, K. C., Kobayashi, C., De Silva, G. M., & Bland-Hawthorn, J. 2012, MNRAS, 421, 1231, doi: 10.1111/j.1365-2966.2011.20387.x

Ting, Y.-S., & Kim, E. M. 2026, The Open Journal of Astrophysics, 9, 71333, doi: 10.33232/001c.171333

Ting, Y.-S., Rix, H.-W., Conroy, C., Ho, A. Y. Q., & Lin, J. 2017b, ApJL, 849, L9, doi: 10.3847/2041-8213/aa921c

Torgerson, W. S. 1952, Psychometrika, 17, 401, doi: 10.1007/BF02288916

Traven, G., Matijevič, G., Zwitter, T., et al. 2017, ApJS, 228, 24, doi: 10.3847/1538-4365/228/2/24

van der Maaten, L. 2014, Journal of Machine Learning Research, 15, 3221. https://www.jmlr.org/papers/v15/vandermaaten14a.html

van der Maaten, L., & Hinton, G. 2008, Journal of Machine Learning Research, 9, 2579. https://www.jmlr.org/papers/v9/vandermaaten08a.html

Vanderplas, J., & Connolly, A. 2009, AJ, 138, 1365, doi: 10.1088/0004-6256/138/5/1365

Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nature Methods, 17, 261, doi: 10.1038/s41592-019-0686-2

Wang, F., & Zhang, C. 2008, IEEE Transactions on Knowledge and Data Engineering, 20, 55, doi: 10.1109/TKDE.2007.190672

Wilson, J. C., Hearty, F. R., Skrutskie, M. F., et al. 2019, PASP, 131, 055001, doi: 10.1088/1538-3873/ab0075

Xiang, M., Ting, Y.-S., Rix, H.-W., et al. 2019, ApJS, 245, 34, doi: 10.3847/1538-4365/ab5364

Yang, T., & Li, X. 2015, MNRAS, 452, 158, doi: 10.1093/mnras/stv1210

Zhang, M., Xiang, M., Ting, Y.-S., et al. 2024, ApJS, 273, 19, doi: 10.3847/1538-4365/ad51dd

Zhang, X., Zhao, G., Yang, C. Q., Wang, Q. X., & Zuo, W. B. 2019, PASP, 131, 094202, doi: 10.1088/1538-3873/ab2687

Zhao, X., Huang, Y., Xue, G., et al. 2026, ApJ, 998, 189, doi: 10.3847/1538-4357/ae2c7e

Zhou, D., Bousquet, O., Lal, T. N., Weston, J., & Schölkopf, B. 2004, in Advances in Neural Information Processing Systems 16, 321–328. https://papers.nips.cc/paper/2506-learning-with-local-and-global-consistency

Zhu, X., & Ghahramani, Z. 2002, Learning from labeled and unlabeled data with label propagation, Tech. Rep. CMU-CALD-02-107, Carnegie Mellon University. https://mlg.eng.cam.ac.uk/zoubin/papers/CMU-CALD-02-107.pdf