# Inducing Fourier Geometry: A Smoothness Prior Imposes Approximate Translation Equivariance in Multimodal Stellar Embeddings

**Ilay Kamai** *Technion - Israel Institute of Technology*

**Alex M. Bronstein** *ISTA - Institute of Science and Technology Austria, Technion - Israel Institute of Technology*

**Hagai B. Perets** *Technion - Israel Institute of Technology*

arXiv:2610.24292v1 [cs.LG] 16 Oct 2026

## Abstract

Learned representations often contain geometric structure induced by symmetries in the underlying data. We ask whether a known continuous parameter space can instead be used to *induce* such structure when the corresponding symmetry is only approximate. We introduce a graph-Laplacian smoothness prior that encourages representations to vary smoothly over a physical parameter space without specifying a representation geometry; in the continuum limit it selects Laplacian eigenfunctions and therefore predicts the Fourier-like geometry associated with translation symmetry. On multimodal stellar representations spanning six astronomical surveys, the prior induces Toeplitz Gram structure and sinusoidal modes along stellar parameters, including in modalities whose native representations lack this geometry, while modality-specific residuals preserve structure outside the imposed coordinates. Of the exact-symmetry predictions, stationarity in a recoverable warped coordinate survives; diagnostics assuming exact quantization or asymptotic rates do not. The prior also improves downstream linear-probe performance and label efficiency, suggesting a general strategy: where a meaningful continuous parameter space exists, its local geometry can serve as a data-informed prior.

## 1. Introduction

Recent work on representation geometry has shown that translation-symmetric structure in data statistics leaves a characteristic fingerprint in learned embeddings: when pairwise co-occurrence statistics are approximately invariant under shifts of an underlying coordinate, the representation Gram matrix is approximately Toeplitz, its eigenvectors are Fourier modes, and one-dimensional manifolds embed as Lissajous-like curves in successive principal components (Karkada et al., 2026). There the geometry *emerges*, discovered by generic objectives on data whose statistics happen to carry the symmetry. We ask the converse: can it be deliberately *induced* where the symmetry is only approximate, and is it useful to do so?

Scientific data provides a natural testbed. Stellar properties vary continuously with physical coordinates such as effective temperature ($T_{\rm eff}$) and surface gravity ($\log g$), making translation along them physically meaningful, yet stellar populations are not exactly stationary and different surveys carry distinct noise, resolution, and selection effects. We study the shared representation of a multimodal stellar embedding model in which six surveys are aligned into a common latent space, augmenting the alignment objective with a

© 2026 I. Kamai, A.M. Bronstein & H.B. Perets.

graph-Laplacian smoothness penalty in physical-label space. The penalty pulls nearby stars in $(T_{\rm eff}, \log g)$ together but specifies no Fourier geometry; in the continuum limit it becomes a Dirichlet energy whose fixed-variance minimizers are Laplacian eigenfunctions. Under approximate stationarity, the Fourier geometry of Karkada et al. (2026) is therefore a *prediction* of the prior rather than an empirical surprise.

The prior is not new in isolation; what differs is where we point it. A graph-Laplacian Dirichlet penalty is manifold regularization (Belkin et al., 2006), with the density-weighted Laplacian of diffusion maps as its continuum limit (Coifman and Lafon, 2006), and the formalism has been used for data fusion and alignment (Lafon et al., 2006; Ham et al., 2005; Wang and Mahadevan, 2009). Here the manifold is supplied externally, as a trusted physical coordinate space, and the regularized object is a shared representation in a deep network. That only part of a multimodal representation admits a common basis is long established (Tucker, 1958; Flury, 1984; Eynard et al., 2015; Klami et al., 2013; Tsai et al., 2018) and parallels our shared/residual split; recovering a withheld modality can also be seen as graph-regularized matrix completion (Kalofolias et al., 2014; Boyarski et al., 2019).

Together our results turn Fourier-like representation geometry from a diagnostic of symmetric data into a design target: a smoothness prior defined in physical-label space installs Toeplitz Gram structure and sinusoidal modes in all six surveys' shared representations, including three whose raw features lack them, and improves linear probing on every label we test.

## 2. Setting: data, model, and the smoothness prior

**Data.** We use six large-scale stellar surveys spanning spectroscopy (LAMOST, APOGEE, GALAH), spectrophotometry (*Gaia* XP), and photometric time series (*Kepler*, TESS), together with *Gaia* astrometry, cross-matched through the *Gaia* DR3 source identifier. We probe geometry along continuous stellar parameters, primarily $T_{\rm eff}$, $\log g$, and [Fe/H]; survey statistics and preprocessing are in Section 5.1.

**Model.** Each modality is mapped by a pretrained, *frozen* encoder to features $z_i$. A Perceiver-IO-style alignment module maps these into a $d = 64$ *shared* stream, trained for cross-survey agreement, and a $d = 192$ per-survey *residual* stream that can retain modality-specific structure. The shared stream is trained with VICReg (Bardes et al., 2021); astrometry enters as a conditioning vector.

**The smoothness prior.** We add a density-normalised graph-Laplacian penalty in physical-label space. Let $\bar{s}_a$ be the mean shared token of star $a$ and $c_a = (T_{\rm eff}, \log g)_a$ its normalized label coordinates. With a fixed-bandwidth Gaussian kernel $w_{ab} = \exp(-\|c_a - c_b\|^2/2\varepsilon^2)$ and degrees $q_a = \sum_b w_{ab}$,

$$\mathcal{L}_{\rm smooth} = \frac{1}{d_s} \frac{\sum_{(a,b)\in E} \tilde{w}_{ab} \, \|\bar{s}_a - \bar{s}_b\|^2}{\sum_{(a,b)\in E} \tilde{w}_{ab}}, \qquad \tilde{w}_{ab} = \frac{w_{ab}}{q_a \, q_b}, \tag{1}$$

where $E$ retains edges with $\|c_a - c_b\| < 3\varepsilon$. The $1/(q_a q_b)$ normalisation is the $\alpha = 1$ member of the Coifman–Lafon family (Coifman and Lafon, 2006), which makes the continuum limit the *unweighted* Dirichlet energy $\int \|\nabla_c s\|^2 \, dc$ rather than a density-weighted one and ensures the penalty is not dominated by the most densely populated regions of the HR diagram (Figure 3); the VICReg variance term, evaluated under a matching density-free measure,

fixes its scale (Section 5.4). Minimising then selects the low-lying eigenfunctions of the ordinary Laplacian, which for an approximately one-dimensional coordinate are sinusoids of increasing node count. The penalty encodes no Fourier structure explicitly: the geometry of Karkada et al. (2026) is a predicted minimiser of the loss rather than a feature we impose by hand.

**Measurements and scope.** Geometry is measured from bin-mean embeddings along individual coordinates (or coordinate pairs), on signature-consistent main-sequence samples to avoid survey-availability effects; we quantify Toeplitz structure by the fraction of off-diagonal Gram variance explained ($R^2_{\mathrm{Toep}}$), count sign changes in leading eigenmodes, and fit kernels to Gram-diagonal decay. Two scope conditions apply: stellar parameters enter parts of the training pipeline, so we claim induced consistency, not emergence from unlabeled data; and the physical symmetry is only approximate, so we test explicitly which exact-symmetry predictions survive. Throughout, "with $\mathcal{L}_{\mathrm{smooth}}$" denotes the full density-normalised prior—the $\alpha$=1 kernel together with the $1/q$-weighted variance and covariance statistics; the ablation removes both.

## 3. Results

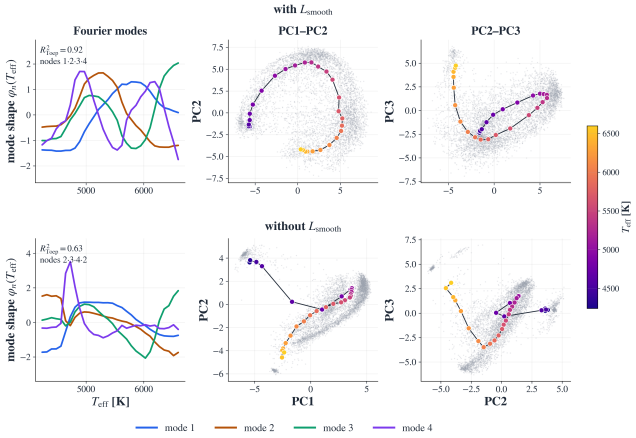



Figure 1: Translation symmetry in the shared latent space. **Upper row:** our full model, with $\mathcal{L}_{\mathrm{smooth}}$. Fourier modes in $T_{\mathrm{eff}}$ with frequency matching eigenmode degree, and principal-component manifolds matching the prediction of Karkada et al. (2026). **Lower row:** the same model without $\mathcal{L}_{\mathrm{smooth}}$, showing neither.

**The prior installs the geometry.** The direct test compares the shared stream with and without $\mathcal{L}_{\mathrm{smooth}}$, holding data, architecture, and seed fixed. With the term, the modality-averaged shared token's Gram matrix along $T_{\mathrm{eff}}$ is strongly Toeplitz ($R^2_{\mathrm{Toep}} = 0.92$) with leading eigenmodes of node count $1\cdot2\cdot3\cdot4$, tracing closed Lissajous curves in successive principal-component pairs (Figure 1). Without it the fit drops to 0.63 and the node ladder disorders; $\log g$, the second graph coordinate, behaves the same way (Figure 4).

**Where the geometry lives.** Figure 6 probes three layers of the with-prior model. In the raw features $z_i$ the symmetry is strong in the spectroscopic surveys (0.75–0.85) and weak in the photometric and spectrophotometric ones (0.35–0.60); the shared stream makes it stronger and uniform (all six in 0.85–0.93) while the residual stays uneven (0.45–0.85). The battery along $\log g$ reproduces the pattern (Figure 7). Along [Fe/H], a coordinate the prior does not name, the shared stream transfers the symmetry rather than taxing it: native spectroscopic structure is approximately preserved while the surveys lacking it gain it

outright (TESS 0.27 → 0.76; Figure 8). The transfer is carried by alignment, not the prior: the pooled [Fe/H] Gram is indistinguishable with and without $\mathcal{L}_{\text{smooth}}$ ($R^2_{\text{Toep}} = 0.76$ vs. 0.75; Figure 5). The induced geometry is targeted: installed along the named coordinates, inert off-graph.

**Which symmetry diagnostics survive approximation.** Karkada et al. (2026) predict three signatures of a distance-only latent coordinate: a Toeplitz Gram matrix, quantized eigenmode wavenumbers, and decoding error $\sim r^{-1/D}$. Only the first survives, and Toeplitz structure alone does not exclude a matched non-stationary kernel; the evidence is the *coordinate*: reparameterising by embedded arclength raises it to 0.97, and an independently fitted monotone warping agrees ($r = 0.92$). Mode counting is undecidable, its quantization ladder assuming an exponential kernel where our spectrum is Gaussian ($\ell \approx 440$ K), and decoding-error scaling deviates ($r^{-0.5}$ vs. $r^{-1}$ along $T_{\text{eff}}$; $r^{-0.8}$ vs. $r^{-1/2}$ on the plane), consistent with a finite lattice below the asymptotic regime.

**Does the induced geometry pay?** Linear probes on the frozen shared latent improve with the prior on every label at every training-set size (Figure 2), most where labels are scarce: age reaches $R^2 \approx 0.38 \pm 0.09$ from 50 labelled stars, comparable within error to the without-prior probe at 20,000. We attribute these gains to the prior as a whole; whether Fourier structure specifically, rather than smoothness generally, buys the improvement remains open.

**Physical coordinates, not generic smoothness.** We retrained the model changing only the graph the Laplacian is built on, with data, seed and recipe fixed (Table 3). A permutation null (identical kernel, bandwidth and neighbour count, with $(T_{\text{eff}}, \log g)$ shuffled across stars) falls below a no-Laplacian baseline on every label and destroys the geometry ($R^2_{\text{Toep}} = 0.39$). A label-free regulariser built on the model's own embedding also falls below it. Only the physical graph improves on no-Laplacian across labels, by +0.04 to +0.11 on five of six; a graph on the input embeddings, whose neighbours are only partly physical, helps only on $T_{\text{eff}}$ and $\log g$. Without the density correction ($\alpha$=0), the penalty's weight collapses onto the main sequence and red clump (Figure 3): the geometry on our main-sequence lattice is the cleanest of any arm ($R^2_{\text{Toep}} = 0.91$), but probes over the full population are weaker. A lattice score measured in one region therefore cannot rank models that weight regions differently. What the controls do show is that with the wrong neighbours, the geometry and the gain fail together.

**Implied distribution.** A Fourier geometry further implies that leading directions are arcsine-valued rather than Gaussian, the opposite of what distributional priors such as SIGReg (Balestriero and LeCun, 2025) target. This holds in all six surveys with the prior, in four without it, and in none of a vanilla VICReg latent (Section 5.6).

## 4. Discussion and conclusions

Emergence and induction are complementary: emergence reveals what data statistics already contain, induction supplies what they do not. Here the induction succeeds without any underlying physical invariance. What stellar physics guarantees is that observables vary smoothly with $T_{\text{eff}}$ and $\log g$, the same assumption behind spectral-grid interpolation, and approximate stationarity follows where information content varies slowly along the coordinate. The failures obey the same logic: age, whose governing physics changes regime across

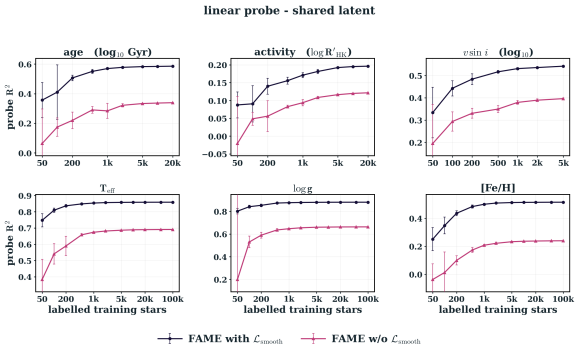



Figure 2: Effect of $\mathcal{L}_{\text{smooth}}$ on downstream tasks. Linear probes on the frozen shared latent, with $\mathcal{L}_{\text{smooth}}$ (black) and without (magenta), against the number of labelled training stars.

the Kraft break, shows no symmetry at any layer, and rotation only a partial one. The prior installs geometry where it is pointed and nowhere else. Importantly, the prior consistently improves performance over all labels, not only on graph coordinates. This highlights the importance of such priors in scientific domains.

Two questions follow. Our controls rule out generic smoothness, but not smoothness in physical coordinates: whether the Fourier structure itself drives the downstream gains, or comes along with them, remains open, since the two fail together when the graph's neighbours are wrong. And what replaces exact quantization and decoding-rate predictions under approximate symmetry is, to our knowledge, open. The physical sciences may be the natural place to settle both, since there the coordinates are real and the geometry can be checked against nature.

## References

Abdurro'uf, Katherine Accetta, Conny Aerts, Víctor Silva Aguirre, Romina Ahumada, Nikhil Ajgaonkar, N. Filiz Ak, Shadab Alam, Carlos Allende Prieto, Andrés Almeida, Friedrich Anders, Scott F. Anderson, Brett H. Andrews, Borja Anguiano, Erik Aquino-Ortíz, Alfonso Aragón-Salamanca, Maria Argudo-Fernández, Metin Ata, Marie Aubert, Vladimir Avila-Reese, Carles Badenes, Rodolfo H. Barbá, Kat Barger, Jorge K. Barrera-Ballesteros, Rachael L. Beaton, Timothy C. Beers, Francesco Belfiore, Chad F. Bender, Mariangela Bernardi, Matthew A. Bershady, Florian Beutler, Christian Moni Bidin, Jonathan C. Bird, Dmitry Bizyaev, Guillermo A. Blanc, Michael R. Blanton, Nicholas Fraser Boardman, Adam S. Bolton, Médéric Boquien, Jura Borissova, Jo Bovy, W. N. Brandt, Jordan Brown, Joel R. Brownstein, Marcella Brusa, Johannes Buchner, Kevin Bundy, Joseph N. Burchett, Martin Bureau, Adam Burgasser, Tuesday K. Cabang, Stephanie Campbell, Michele Cappellari, Joleen K. Carlberg, Fábio Carneiro Wanderley, Ricardo Carrera, Jennifer Cash, Yan-Ping Chen, Wei-Huai Chen, Brian Cherinka, Cristina Chiappini, Peter Doohyun Choi, S. Drew Chojnowski, Haeun Chung, Nicolas Clerc, Roger E. Cohen, Julia M. Comerford, Johan Comparat, Luiz da Costa, Kevin Covey, Jeffrey D. Crane, Irene Cruz-Gonzalez, Connor Culhane, Katia Cunha, Y. Sophia Dai, Guillermo Damke, Jeremy Darling, James W. Davidson, Jr., Roger Davies, Kyle Dawson, Nathan De Lee, Aleksandar M. Diamond-Stanic, Mariana Cano-Díaz, Helena Domínguez Sánchez, John Donor, Chris Duckworth, Tom Dwelly, Daniel J. Eisenstein, Yvonne P. Elsworth, Eric Emsellem, Mike Eracleous, Stephanie Escoffier, Xiaohui Fan, Emily Farr,

Shuai Feng, José G. Fernández-Trincado, Diane Feuillet, Andreas Filipp, Sean P. Fillingham, Peter M. Frinchaboy, Sebastien Fromenteau, Lluís Galbany, Rafael A. García, D. A. García-Hernández, Junqiang Ge, Doug Geisler, Joseph Gelfand, Tobias Géron, Benjamin J. Gibson, Julian Goddy, Diego Godoy-Rivera, Kathleen Grabowski, Paul J. Green, Michael Greener, Catherine J. Grier, Emily Griffith, Hong Guo, Julien Guy, Massinissa Hadjara, Paul Harding, Sten Hasselquist, Christian R. Hayes, Fred Hearty, Jesús Hernández, Lewis Hill, David W. Hogg, Jon A. Holtzman, Danny Horta, Bau-Ching Hsieh, Chin-Hao Hsu, Yun-Hsin Hsu, Daniel Huber, Marc Huertas-Company, Brian Hutchinson, Ho Seong Hwang, Héctor J. Ibarra-Medel, Jacob Ider Chitham, Gabriele S. Ilha, Julie Imig, Will Jaekle, Tharindu Jayasinghe, Xihan Ji, Jennifer A. Johnson, Amy Jones, Henrik Jönsson, Ivan Katkov, Arman Khalatyan, Dr., Karen Kinemuchi, Shobhit Kisku, Johan H. Knapen, Jean-Paul Kneib, Juna A. Kollmeier, Miranda Kong, Marina Kounkel, Kathryn Kreckel, Dhanesh Krishnarao, Ivan Lacerna, Richard R. Lane, Rachel Langgin, Ramon Lavender, David R. Law, Daniel Lazarz, Henry W. Leung, Ho-Hin Leung, Hannah M. Lewis, Cheng Li, Ran Li, Jianhui Lian, Fu-Heng Liang, Lihwai Lin, Yen-Ting Lin, Sicheng Lin, Chris Lintott, Dan Long, Penélope Longa-Peña, Carlos López-Cobá, Shengdong Lu, Britt F. Lundgren, Yuanze Luo, J. Ted Mackereth, Axel de la Macorra, Suvrath Mahadevan, Steven R. Majewski, Arturo Manchado, Travis Mandeville, Claudia Maraston, Berta Margalef-Bentabol, Thomas Masseron, Karen L. Masters, Savita Mathur, Richard M. McDermid, Myles Mckay, Andrea Merloni, Michael Merrifield, Szabolcs Meszaros, Andrea Miglio, Francesco Di Mille, Dante Minniti, Rebecca Minsley, and Antonela Monachesi. The Seventeenth Data Release of the Sloan Digital Sky Surveys: Complete Release of MaNGA, MaStar, and APOGEE-2 Data. ApJS, 259(2):35, April 2022. doi: 10.3847/1538-4365/ac4414.

Randall Balestriero and Yann LeCun. LeJEPA: Provable and Scalable Self-Supervised Learning Without the Heuristics. *arXiv e-prints*, art. arXiv:2511.08544, November 2025. doi: 10.48550/arXiv.2511.08544.

Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. *arXiv e-prints*, art. arXiv:2105.04906, May 2021. doi: 10.48550/arXiv.2105.04906.

Mikhail Belkin, Partha Niyogi, and Vikas Sindhwani. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *Journal of Machine Learning Research*, 7(85):2399–2434, 2006.

William J. Borucki, David Koch, Gibor Basri, Natalie Batalha, Timothy Brown, Douglas Caldwell, John Caldwell, Jørgen Christensen-Dalsgaard, William D. Cochran, Edna DeVore, Edward W. Dunham, Andrea K. Dupree, Thomas N. Gautier, John C. Geary, Ronald Gilliland, Alan Gould, Steve B. Howell, Jon M. Jenkins, Yoji Kondo, David W. Latham, Geoffrey W. Marcy, Søren Meibom, Hans Kjeldsen, Jack J. Lissauer, David G. Monet, David Morrison, Dimitar Sasselov, Jill Tarter, Alan Boss, Don Brownlee, Toby Owen, Derek Buzasi, David Charbonneau, Laurance Doyle, Jonathan Fortney, Eric B. Ford, Matthew J. Holman, Sara Seager, Jason H. Steffen, William F. Welsh, Jason Rowe, Howard Anderson, Lars Buchhave, David Ciardi, Lucianne Walkowicz, William

Sherry, Elliott Horch, Howard Isaacson, Mark E. Everett, Debra Fischer, Guillermo Torres, John Asher Johnson, Michael Endl, Phillip MacQueen, Stephen T. Bryson, Jessie Dotson, Michael Haas, Jeffrey Kolodziejczak, Jeffrey Van Cleve, Hema Chandrasekaran, Joseph D. Twicken, Elisa V. Quintana, Bruce D. Clarke, Christopher Allen, Jie Li, Haley Wu, Peter Tenenbaum, Ekaterina Verner, Frederick Bruhweiler, Jason Barnes, and Andrej Prsa. Kepler Planet-Detection Mission: Introduction and First Results. *Science*, 327 (5968):977, February 2010. doi: 10.1126/science.1185402.

Amit Boyarski, Sanketh Vedula, and Alexander M. Bronstein. Spectral geometric matrix completion. In *Mathematical and Scientific Machine Learning*, 2019. URL `https://api.semanticscholar.org/CorpusID:235363958`.

Sven Buder, Janez Kos, Xi Ella Wang, Madeleine McKenzie, Madeleine Howell, Sarah Martell, Michael R. Hayden, Daniel B. Zucker, Thomas Nordlander, Benjamin Montet, Gregor Traven, Joss Bland-Hawthorn, Gayandhi M. de Silva, Kenneth Freeman, Geraint Lewis, Karin Lind, Sanjib Sharma, Jeffrey D. Simpson, Dennis Stello, Tomaz Zwitter, Anish M. Amarsi, Joseph J. Armstrong, Kirsten Banks, Mark Beavis, Kevin-Luke Beeson, Boquan Chen, Ioana Ciucă, Gary S. da Costa, Richard de Grijs, Bailey Martin, David Moise Nataf, Melissa Ness, Adam D. Rains, Tim Scarr, Rok Vogrinčič, Zixian Purmortal Wang, Rob A. Wittenmyer, Yi Anne Xie, and The Galah Collaboration. The GALAH survey: Data release 4. PASA, 42:e051, May 2025. doi: 10.1017/pasa.2025.26.

Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. ISSN 1063-5203. doi: https://doi.org/10.1016/j.acha.2006.04.006. URL `https://www.sciencedirect.com/science/article/pii/S1063520306000546`. Special Issue: Diffusion Maps and Wavelets.

F. De Angeli, M. Weiler, P. Montegriffo, D. W. Evans, M. Riello, R. Andrae, J. M. Carrasco, G. Busso, P. W. Burgess, C. Cacciari, M. Davidson, D. L. Harrison, S. T. Hodgkin, C. Jordi, P. J. Osborne, E. Pancino, G. Altavilla, M. A. Barstow, C. A. L. Bailer-Jones, M. Bellazzini, A. G. A. Brown, M. Castellani, S. Cowell, L. Delchambre, F. De Luise, C. Diener, C. Fabricius, M. Fouesneau, Y. Frémat, G. Gilmore, G. Giuffrida, N. C. Hambly, S. Hidalgo, G. Holland, Z. Kostrzewa-Rutkowska, F. van Leeuwen, A. Lobel, S. Marinoni, N. Miller, C. Pagani, L. Palaversa, A. M. Piersimoni, L. Pulone, S. Ragaini, M. Rainer, P. J. Richards, G. T. Rixon, D. Ruz-Mieres, N. Sanna, L. M. Sarro, N. Rowell, R. Sordo, N. A. Walton, and A. Yoldas. Gaia Data Release 3. Processing and validation of BP/RP low-resolution spectral data. A&A, 674:A2, June 2023. doi: 10.1051/0004-6361/202243680.

G. M. De Silva, K. C. Freeman, J. Bland-Hawthorn, S. Martell, E. Wylie de Boer, M. Asplund, S. Keller, S. Sharma, D. B. Zucker, T. Zwitter, B. Anguiano, C. Bacigalupo, D. Bayliss, M. A. Beavis, M. Bergemann, S. Campbell, R. Cannon, D. Carollo, L. Casagrande, A. R. Casey, G. Da Costa, V. D'Orazi, A. Dotter, L. Duong, A. Heger, M. J. Ireland, P. R. Kafle, J. Kos, J. Lattanzio, G. F. Lewis, J. Lin, K. Lind, U. Munari, D. M. Nataf, S. O'Toole, Q. Parker, W. Reid, K. J. Schlesinger, A. Sheinis, J. D. Simpson, D. Stello, Y.-S. Ting, G. Traven, F. Watson, R. Wittenmyer, D. Yong, and M. Žerjal.

The GALAH survey: scientific motivation. MNRAS, 449(3):2604–2617, May 2015. doi: 10.1093/mnras/stv327.

Ming-Yi Ding, Jian-Rong Shi, Yue Wu, Hugh R. A. Jones, Hong-Liang Yan, Chun-Qian Li, Qi Gao, Tian-Yi Chen, Jing-Hua Zhang, Shuai Liu, Tai-Sheng Yan, and Xiao-Jin Xie. Stellar Atmospheric Parameters of M-type Stars from LAMOST DR8. ApJS, 260(2):45, June 2022. doi: 10.3847/1538-4365/ac6754.

Davide Eynard, Artiom Kovnatsky, Michael M. Bronstein, Klaus Glashoff, and Alexander M. Bronstein. Multimodal manifold analysis by simultaneous diagonalization of laplacians. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(12): 2505–2517, 2015. doi: 10.1109/TPAMI.2015.2408348.

Bernhard N. Flury. Common principal components in k groups. *Journal of the American Statistical Association*, 79(388):892–898, 1984. doi: 10.1080/01621459.1984.10477108. URL `https://doi.org/10.1080/01621459.1984.10477108`.

Gaia Collaboration, T. Prusti, J. H. J. de Bruijne, A. G. A. Brown, A. Vallenari, C. Babusiaux, C. A. L. Bailer-Jones, U. Bastian, M. Biermann, D. W. Evans, L. Eyer, F. Jansen, C. Jordi, S. A. Klioner, U. Lammers, L. Lindegren, X. Luri, F. Mignard, D. J. Milligan, C. Panem, V. Poinsignon, D. Pourbaix, S. Randich, G. Sarri, P. Sartoretti, H. I. Siddiqui, C. Soubiran, V. Valette, F. van Leeuwen, N. A. Walton, C. Aerts, F. Arenou, M. Cropper, R. Drimmel, E. Høg, D. Katz, M. G. Lattanzi, W. O'Mullane, E. K. Grebel, A. D. Holland, C. Huc, X. Passot, L. Bramante, C. Cacciari, J. Castañeda, L. Chaoul, N. Cheek, F. De Angeli, C. Fabricius, R. Guerra, J. Hernández, A. Jean-Antoine-Piccolo, E. Masana, R. Messineo, N. Mowlavi, K. Nienartowicz, D. Ordóñez-Blanco, P. Panuzzo, J. Portell, P. J. Richards, M. Riello, G. M. Seabroke, P. Tanga, F. Thévenin, J. Torra, S. G. Els, G. Gracia-Abril, G. Comoretto, M. Garcia-Reinaldos, T. Lock, E. Mercier, M. Altmann, R. Andrae, T. L. Astraatmadja, I. Bellas-Velidis, K. Benson, J. Berthier, R. Blomme, G. Busso, B. Carry, A. Cellino, G. Clementini, S. Cowell, O. Creevey, J. Cuypers, M. Davidson, J. De Ridder, A. de Torres, L. Delchambre, A. Dell'Oro, C. Ducourant, Y. Frémat, M. García-Torres, E. Gosset, J.-L. Halbwachs, N. C. Hambly, D. L. Harrison, M. Hauser, D. Hestroffer, S. T. Hodgkin, H. E. Huckle, A. Hutton, G. Jasniewicz, S. Jordan, M. Kontizas, A. J. Korn, A. C. Lanzafame, M. Manteiga, A. Moitinho, K. Muinonen, J. Osinde, E. Pancino, T. Pauwels, J.-M. Petit, A. Recio-Blanco, A. C. Robin, L. M. Sarro, C. Siopis, M. Smith, K. W. Smith, A. Sozzetti, W. Thuillot, W. van Reeven, Y. Viala, U. Abbas, A. Abreu Aramburu, S. Accart, J. J. Aguado, P. M. Allan, W. Allasia, G. Altavilla, M. A. Álvarez, J. Alves, R. I. Anderson, A. H. Andrei, E. Anglada Varela, E. Antiche, T. Antoja, S. Antón, B. Arcay, A. Atzei, L. Ayache, N. Bach, S. G. Baker, L. Balaguer-Núñez, C. Barache, C. Barata, A. Barbier, F. Barblan, M. Baroni, D. Barrado y Navascués, M. Barros, M. A. Barstow, U. Becciani, M. Bellazzini, G. Bellei, A. Bello García, V. Belokurov, P. Bendjoya, A. Berihuete, L. Bianchi, O. Bienaymé, F. Billebaud, N. Blagorodnova, S. Blanco-Cuaresma, T. Boch, A. Bombrun, R. Borrachero, S. Bouquillon, G. Bourda, H. Bouy, A. Bragaglia, M. A. Breddels, N. Brouillet, T. Brüsemeister, B. Bucciarelli, F. Budnik, P. Burgess, R. Burgon, A. Burlacu, D. Busonero, R. Buzzi, E. Caffau, J. Cambras, H. Campbell, R. Cancelliere, T. Cantat-Gaudin, T. Carlucci,

J. M. Carrasco, M. Castellani, P. Charlot, J. Charnas, P. Charvet, F. Chassat, A. Chiavassa, M. Clotet, G. Cocozza, R. S. Collins, P. Collins, and G. Costigan. The Gaia mission. A&A, 595:A1, November 2016. doi: 10.1051/0004-6361/201629272.

Gaia Collaboration, A. Vallenari, A. G. A. Brown, T. Prusti, J. H. J. de Bruijne, F. Arenou, C. Babusiaux, M. Biermann, O. L. Creevey, C. Ducourant, D. W. Evans, L. Eyer, R. Guerra, A. Hutton, C. Jordi, S. A. Klioner, U. L. Lammers, L. Lindegren, X. Luri, F. Mignard, C. Panem, D. Pourbaix, S. Randich, P. Sartoretti, C. Soubiran, P. Tanga, N. A. Walton, C. A. L. Bailer-Jones, U. Bastian, R. Drimmel, F. Jansen, D. Katz, M. G. Lattanzi, F. van Leeuwen, J. Bakker, C. Cacciari, J. Castañeda, F. De Angeli, C. Fabricius, M. Fouesneau, Y. Frémat, L. Galluccio, A. Guerrier, U. Heiter, E. Masana, R. Messineo, N. Mowlavi, C. Nicolas, K. Nienartowicz, F. Pailler, P. Panuzzo, F. Riclet, W. Roux, G. M. Seabroke, R. Sordo, F. Thévenin, G. Gracia-Abril, J. Portell, D. Teyssier, M. Altmann, R. Andrae, M. Audard, I. Bellas-Velidis, K. Benson, J. Berthier, R. Blomme, P. W. Burgess, D. Busonero, G. Busso, H. Cánovas, B. Carry, A. Cellino, N. Cheek, G. Clementini, Y. Damerdji, M. Davidson, P. de Teodoro, M. Nuñez Campos, L. Delchambre, A. Dell'Oro, P. Esquej, J. Fernández-Hernández, E. Fraile, D. Garabato, P. García-Lario, E. Gosset, R. Haigron, J.-L. Halbwachs, N. C. Hambly, D. L. Harrison, J. Hernández, D. Hestroffer, S. T. Hodgkin, B. Holl, K. Janßen, G. Jevardat de Fombelle, S. Jordan, A. Krone-Martins, A. C. Lanzafame, W. Löffler, O. Marchal, P. M. Marrese, A. Moitinho, K. Muinonen, P. Osborne, E. Pancino, T. Pauwels, A. Recio-Blanco, C. Reylé, M. Riello, L. Rimoldini, T. Roegiers, J. Rybizki, L. M. Sarro, C. Siopis, M. Smith, A. Sozzetti, E. Utrilla, M. van Leeuwen, U. Abbas, P. Ábrahám, A. Abreu Aramburu, C. Aerts, J. J. Aguado, M. Ajaj, F. Aldea-Montero, G. Altavilla, M. A. Álvarez, J. Alves, F. Anders, R. I. Anderson, E. Anglada Varela, T. Antoja, D. Baines, S. G. Baker, L. Balaguer-Núñez, E. Balbinot, Z. Balog, C. Barache, D. Barbato, M. Barros, M. A. Barstow, S. Bartolomé, J.-L. Bassilana, N. Bauchet, U. Becciani, M. Bellazzini, A. Berihuete, M. Bernet, S. Bertone, L. Bianchi, A. Binnenfeld, S. Blanco-Cuaresma, A. Blazere, T. Boch, A. Bombrun, D. Bossini, S. Bouquillon, A. Bragaglia, L. Bramante, E. Breedt, A. Bressan, N. Brouillet, E. Brugaletta, B. Bucciarelli, A. Burlacu, A. G. Butkevich, R. Buzzi, E. Caffau, R. Cancelliere, T. Cantat-Gaudin, R. Carballo, T. Carlucci, M. I. Carnerero, J. M. Carrasco, L. Casamiquela, M. Castellani, A. Castro-Ginard, L. Chaoul, P. Charlot, L. Chemin, V. Chiaramida, A. Chiavassa, N. Chornay, G. Comoretto, G. Contursi, W. J. Cooper, T. Cornez, S. Cowell, F. Crifo, M. Cropper, M. Crosta, C. Crowley, C. Dafonte, A. Dapergolas, M. David, P. David, P. de Laverny, F. De Luise, and R. De March. Gaia Data Release 3. Summary of the content and survey properties. A&A, 674:A1, June 2023. doi: 10.1051/0004-6361/202243940.

Jihun Ham, Daniel D. Lee, and Lawrence K. Saul. Semisupervised alignment of manifolds. In *International Conference on Artificial Intelligence and Statistics*, 2005. URL `https://api.semanticscholar.org/CorpusID:2913820`.

Chelsea X. Huang, Andrew Vanderburg, Andras Pál, Lizhou Sha, Liang Yu, Willie Fong, Michael Fausnaugh, Avi Shporer, Natalia Guerrero, Roland Vanderspek, and George Ricker. Photometry of 10 Million Stars from the First Two Years of TESS Full Frame Im-

ages: Part II. *Research Notes of the American Astronomical Society*, 4(11):206, November 2020. doi: 10.3847/2515-5172/abca2d.

Dean R. H. Johnson and David R. Soderblom. Calculating Galactic Space Velocities and Their Uncertainties, with an Application to the Ursa Major Group. AJ, 93:864, April 1987. doi: 10.1086/114370.

Vassilis Kalofolias, Xavier Bresson, Michael M. Bronstein, and Pierre Vandergheynst. Matrix completion on graphs. *ArXiv*, abs/1408.1717, 2014. URL `https://api.semanticscholar.org/CorpusID:14746688`.

Dhruva Karkada, Daniel J. Korchinski, Andres Nava, Matthieu Wyart, and Yasaman Bahri. Symmetry in language statistics shapes the geometry of model representations. *arXiv e-prints*, art. arXiv:2602.15029, February 2026. doi: 10.48550/arXiv.2602.15029.

Arto Klami, Seppo Virtanen, and Samuel Kaski. Bayesian canonical correlation analysis. *J. Mach. Learn. Res.*, 14(1):965–1003, April 2013. ISSN 1532-4435.

S. Lafon, Y. Keller, and R.R. Coifman. Data fusion and multicue data matching by diffusion maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1784–1797, 2006. doi: 10.1109/TPAMI.2006.223.

N. R. Lomb. Least-Squares Frequency Analysis of Unequally Spaced Data. Ap&SS, 39(2): 447–462, February 1976. doi: 10.1007/BF00648343.

Steven R. Majewski, Ricardo P. Schiavon, Peter M. Frinchaboy, Carlos Allende Prieto, Robert Barkhouser, Dmitry Bizyaev, Basil Blank, Sophia Brunner, Adam Burton, Ricardo Carrera, S. Drew Chojnowski, Kátia Cunha, Courtney Epstein, Greg Fitzgerald, Ana E. García Pérez, Fred R. Hearty, Chuck Henderson, Jon A. Holtzman, Jennifer A. Johnson, Charles R. Lam, James E. Lawler, Paul Maseman, Szabolcs Mészáros, Matthew Nelson, Duy Coung Nguyen, David L. Nidever, Marc Pinsonneault, Matthew Shetrone, Stephen Smee, Verne V. Smith, Todd Stolberg, Michael F. Skrutskie, Eric Walker, John C. Wilson, Gail Zasowski, Friedrich Anders, Sarbani Basu, Stephane Beland, Michael R. Blanton, Jo Bovy, Joel R. Brownstein, Joleen Carlberg, William Chaplin, Cristina Chiappini, Daniel J. Eisenstein, Yvonne Elsworth, Diane Feuillet, Scott W. Fleming, Jessica Galbraith-Frew, Rafael A. García, D. Aníbal García-Hernández, Bruce A. Gillespie, Léo Girardi, James E. Gunn, Sten Hasselquist, Michael R. Hayden, Saskia Hekker, Inese Ivans, Karen Kinemuchi, Mark Klaene, Suvrath Mahadevan, Savita Mathur, Benoît Mosser, Demitri Muna, Jeffrey A. Munn, Robert C. Nichol, Robert W. O'Connell, John K. Parejko, A. C. Robin, Helio Rocha-Pinto, Matthias Schultheis, Aldo M. Serenelli, Neville Shane, Victor Silva Aguirre, Jennifer S. Sobeck, Benjamin Thompson, Nicholas W. Troup, David H. Weinberg, and Olga Zamora. The Apache Point Observatory Galactic Evolution Experiment (APOGEE). AJ, 154(3):94, September 2017. doi: 10.3847/1538-3881/aa784d.

A. McQuillan, S. Aigrain, and T. Mazeh. Measuring the rotation period distribution of field M dwarfs with Kepler. MNRAS, 432(2):1203–1216, June 2013. doi: 10.1093/mnras/stt536.

George R. Ricker, Joshua N. Winn, Roland Vanderspek, David W. Latham, Gáspár Á. Bakos, Jacob L. Bean, Zachory K. Berta-Thompson, Timothy M. Brown, Lars Buchhave, Nathaniel R. Butler, R. Paul Butler, William J. Chaplin, David Charbonneau, Jørgen Christensen-Dalsgaard, Mark Clampin, Drake Deming, John Doty, Nathan De Lee, Courtney Dressing, Edward W. Dunham, Michael Endl, Francois Fressin, Jian Ge, Thomas Henning, Matthew J. Holman, Andrew W. Howard, Shigeru Ida, Jon M. Jenkins, Garrett Jernigan, John Asher Johnson, Lisa Kaltenegger, Nobuyuki Kawai, Hans Kjeldsen, Gregory Laughlin, Alan M. Levine, Douglas Lin, Jack J. Lissauer, Phillip MacQueen, Geoffrey Marcy, Peter R. McCullough, Timothy D. Morton, Norio Narita, Martin Paegert, Enric Palle, Francesco Pepe, Joshua Pepper, Andreas Quirrenbach, Stephen A. Rinehart, Dimitar Sasselov, Bun'ei Sato, Sara Seager, Alessandro Sozzetti, Keivan G. Stassun, Peter Sullivan, Andrew Szentgyorgyi, Guillermo Torres, Stephane Udry, and Joel Villasenor. Transiting Exoplanet Survey Satellite (TESS). *Journal of Astronomical Telescopes, Instruments, and Systems*, 1:014003, January 2015. doi: 10.1117/1.JATIS.1.1.014003.

J. D. Scargle. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. ApJ, 263:835–853, December 1982. doi: 10.1086/160554.

STScI. Kepler lc, q0-q17, 2016. URL `http://archive.stsci.edu/doi/resolve/resolve.html?doi=10.17909/T9488N`.

Yao-Hung Hubert Tsai, Paul Pu Liang, Amir Zadeh, Louis philippe Morency, and Ruslan Salakhutdinov. Learning factorized multimodal representations. *ArXiv*, abs/1806.06176, 2018. URL `https://api.semanticscholar.org/CorpusID:49303347`.

Ledyard R. Tucker. An inter-battery method of factor analysis. *Psychometrika*, 23:111–136, 1958. URL `https://api.semanticscholar.org/CorpusID:123711893`.

Jeffrey E. Van Cleve, Steve B. Howell, Jeffrey C. Smith, Bruce D. Clarke, Susan E. Thompson, Stephen T. Bryson, Mikkel N. Lund, Rasmus Handberg, and William J. Chaplin. That's How We Roll: The NASA K2 Mission Science Products and Their Performance Metrics. PASP, 128(965):075002, July 2016. doi: 10.1088/1538-3873/128/965/075002.

Chang Wang and Sridhar Mahadevan. Manifold alignment without correspondence. IJCAI'09, page 1273–1278, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc.

Gang Zhao, Yongheng Zhao, Y. Chu, Y. Jing, and L. Deng. Lamost spectral survey — an overview. *Research in Astronomy and Astrophysics*, 2012. doi: 10.1088/1674-4527/12/7/002. URL `https://www.semanticscholar.org/paper/6dd33c0f7d83c6cbc8e2ee544cee259518d6f474`.

## 5. appendix

### 5.1. data and pre processing

***Kepler.*** The *Kepler* mission (Borucki et al., 2010) monitored a single $\sim$115 deg$^2$ field in Cygnus–Lyra for four years at 30-minute (long) cadence, delivering the highest-precision,

longest-baseline optical photometry available for main-sequence stars. We use the long-cadence PDCSAP light curves (STScI (2016)) on their native, uniform 48 samples day$^{-1}$ grid, with gaps zero-filled and no rebinning, cut into 450-day windows (21,600 points). Each window is lightly smoothed (Savitzky–Golay, $\sim$0.5 d) and expanded into six channels: standardized flux, median-normalized flux, two activity-amplitude envelopes (a slow envelope on a $6\,P_{\rm rot}$ window and a high-passed one on a $1\,P_{\rm rot}$ window), the autocorrelation function (ACF, McQuillan et al. (2013)), and the Lomb-Scargle periodogram (LS, Lomb (1976), Scargle (1982)). The resulting sample is 193,951 stars.

**TESS.** TESS (Ricker et al., 2015) trades *Kepler*'s depth for sky coverage, observing nearly the full sky in $\sim$27-day sectors at multiple cadences. We use light curves from the Quick-Look Pipeline (Huang et al., 2020), restricted to $T_{\rm mag} < 16$ and to targets observed in more than one sector. Each light curve is resampled to a 30-minute grid and truncated to 1000 points in four channels: standardized flux, median-filtered flux, ACF, and LS. These are the same channels as in *Kepler*, but without the activity envelopes, as the sectors are too short for it. Compared to *Kepler*, TESS has a wider coverage (1,503,260 encoded stars) but has a shorter per-sector baseline, bigger pixel size ($21''$ vs. $4''$), and is generally noisier (Van Cleve et al. (2016).

**LAMOST.** LAMOST (Zhao et al., 2012) is a 4-m telescope with 4000 fibers, conducting the largest optical spectroscopic survey of the northern sky. We use low-resolution ($R\sim1800$) spectra from Ding et al. (2022), continuum-normalized and resampled onto a common 4096-point wavelength grid. Its combination of scale (5.1M spectra) and northern footprint makes LAMOST the dominant spectroscopic partner for both *Kepler* and TESS.

**APOGEE.** APOGEE (Majewski et al., 2017) is a high-resolution ($R\sim22{,}500$) $H$-band spectrograph. Working in the near-infrared lets it observe through dust into the Galactic plane and bulge, and it targets predominantly giants—a population complementary to the dwarf-dominated *Kepler* and GALAH samples. We use the continuum-normalized combined spectra from Abdurro'uf et al. (2022), resampled to 8576 points.

**GALAH.** GALAH (De Silva et al., 2015) is a high-resolution ($R\sim28{,}000$) southern optical survey using the HERMES spectrograph on the Anglo-Australian Telescope, designed for Galactic archaeology and delivering up to $\sim$30 elemental abundances per star. We use the continuum-normalized spectra of Buder et al. (2025) over 4700–7900 Å, concatenated across the four HERMES arms onto a 14,420-point grid, and learning, in addition to basic stellar parameters and $v \sin i$, also 11 abundance elements ( [Mg/Fe], [Si/Fe], [Ca/Fe], [Ti/Fe], [Na/Fe], [Al/Fe], [Mn/Fe], [Ni/Fe], [Cr/Fe], [Ba/Fe], [Y/Fe]). GALAH's southern footprint is disjoint from *Kepler*'s.

***Gaia.*** *Gaia* (Gaia Collaboration et al., 2016) provides all-sky astrometry and low-resolution spectrophotometry, and serves here as the universal hub. We use two products from DR3 (Gaia Collaboration et al., 2023): (i) The BP/RP (XP) *sampled* spectra (De Angeli et al., 2023): 340 flux points spanning 330–1050 nm, scaled per-spectrum with a robust (median/IQR) scaler and deliberately *not* continuum-normalized, so that the absolute spectral-energy-distribution shape, which carries temperature and extinction information, is preserved. (ii) The five-parameter astrometric solution plus radial velocities, from which we

Table 1: Survey census and pairwise overlaps in the cross-match master (unique stars; *Gaia* is the anchor and covers essentially all 1,386,210 stars). The diagonal gives each survey's total; off-diagonal entries give the number of stars observed by both surveys.

| | *Kepler* | TESS | LAMOST | APOGEE | GALAH |
|---|---|---|---|---|---|
| *Kepler* | 91,213 | 69,701 | 66,615 | 20,712 | 0 |
| TESS | | 1,247,706 | 935,262 | 226,347 | 188,513 |
| LAMOST | | | 1,053,269 | 135,429 | 60,388 |
| APOGEE | | | | 310,347 | 55,623 |
| GALAH | | | | | 247,444 |

| Modality | Architecture | Params | Input Length | $d_i$ |
|---|---|---|---|---|
| *Gaia* XP | Spectral ViT | 38.5 M | 340 | 512 |
| LAMOST | Spectral ViT | 615.3 M | 4096 | 2048 |
| APOGEE | Spectral ViT | 154.2 M | 8576 | 1024 |
| GALAH | Spectral ViT | 154.4 M | 14,420 | 1024 |
| *Kepler* | Dual-view LC | 129.6 M | $6 \times 21{,}600$ | 768 |
| TESS | Dual-branch LC | 139.5 M | $4 \times 1000$ | 512 |

Table 2: The six frozen unimodal encoders. Input Length is the number of flux bins (spectra) or channels and time samples (light curves). $d_i$ is the feature dimensionality.

derive $UVW$ space velocities via the equatorial-to-Galactic transformation of Johnson and Soderblom (1987).

### 5.2. Cross-match

All surveys are cross-matched onto a common star identity keyed by the *Gaia* DR3 `source_id`. The master cross-match retains every star observed by at least two of the five non-*Gaia* surveys, yielding **1,613,142** observations of **1,386,210** unique stars. Table 1 lists the per-survey census and the pairwise overlaps. The pair frequencies span roughly two orders of magnitude—from TESS×LAMOST (935,262 stars) down to *Kepler*×APOGEE (20,712). Notably, *Kepler* and GALAH have *zero* common stars (northern vs. southern footprints): these two modalities are never observed together and can only be related transitively, through the shared latent space. Counting the number of non-*Gaia* surveys per star, 1,225,913 stars have exactly two, 143,035 have three, and 17,262 have four; no star is covered by all five.

### 5.3. Training

Pre-training was done on a mixture of H100 and B200 GPUs. Each encoder was trained until convergence, using early stopping criteria. Training times vary between several hundred and several thousand GPU-hours. Training the alignment model is fast: of the order of 10 B200 GPU-hours. The alignment model is trained on a fixed stratified subset that deliberately reserves half its budget for single-modality stars, so the shared space is shaped by the long unimodal tail and not only by the well-covered multimodal core. Multimodal stars are

weighted by their rarest constituent modality pair. The invariance and variance terms in VICReg were given a weight of 25. All other loss terms had a weight of 1.

### 5.4. Why a continuity penalty produces a Fourier geometry

We sketch why minimising the smoothness penalty organises the shared coordinates into sinusoidal modes. The argument motivates the observed geometry rather than proving a theorem about the full objective; idealisations are stated as we go.

**The unnormalised penalty carries the density squared.** Write $s(c)$ for the shared embedding at parameter $c$. An unnormalised graph penalty $\frac{1}{n^2}\sum_{a\neq b} w_{ab}\|s_a - s_b\|^2$ is a sum over *pairs* of stars drawn i.i.d. from the sampling density $p$, so its expectation integrates against $p(x)p(y)$. Expanding at small bandwidth,

$$\frac{1}{n^2}\sum_{a\neq b} w_{ab}\|s_a - s_b\|^2 \;\longrightarrow\; \varepsilon^2 \int \|s'(c)\|^2\, p(c)^2\, dc, \tag{2}$$

one factor of $p$ per summation index. More generally, dividing the kernel by $q_a^\alpha q_b^\alpha$ (with $q_a \propto p(c_a)$ the degree) gives a continuum limit weighted by $p^{2-2\alpha}$ (Coifman and Lafon, 2006). A global normaliser such as $\sum w$ cannot change this exponent; only the per-node division can.

**Normalizing the density** Taking $\alpha = 1$, i.e. $\tilde{w}_{ab} = w_{ab}/(q_a q_b)$, cancels the density to leading order and leaves the unweighted Dirichlet energy $E[s] = \int \|s'(c)\|^2\, dc$. Its Euler–Lagrange operator, under Neumann boundary conditions, is the ordinary Laplacian $-d^2/dc^2$.

**The scale constraint carries the density too.** The energy alone is minimised by a constant; a scale must be fixed, and this is done by the VICReg variance and covariance terms. But these are batch averages, hence estimates of integrals against $p$: the variance constraint is "unit variance under $p$", which reintroduces the density through the measure and warps the minimisers by a factor $\propto \sqrt{p}$. We therefore apply the same $\alpha = 1$ correction a second time, importance-weighting every variance and covariance moment by $1/q_a \propto 1/p(c_a)$ before the (unchanged) hinge and off-diagonal penalties. Under a quadratic-constraint idealisation the effective constraint becomes "unit variance under Lebesgue measure", and the Rayleigh–Ritz problem is the density-free

$$-s_n'' = \lambda_n\, s_n. \tag{3}$$

Normalising the kernel alone leaves the $\sqrt{p}$ warp; normalising both is what makes the geometry density-free. The invariance term is left unweighted: it compares two modalities of the same star, a per-pair quantity with no density to cancel.

**Fourier modes and their robustness.** On $[0, L]$ with Neumann boundaries the eigenfunctions are $s_n(c) = \cos(n\pi c/L)$, $\lambda_n = (n\pi/L)^2$: the observed ladder of increasingly oscillatory coordinates. Even where the correction is imperfect, the operator remains of Sturm–Liouville type, whose $n$-th eigenfunction has exactly $n$ nodes, so the node ladder is preserved regardless of residual density weighting. This robustness is why the ladder survives approximate stationarity.

**Why two coordinates degrade.** The one-dimensional argument does not extend directly to two dimensions. The distribution in $(T_{\text{eff}}, \log g)$ is neither translation invariant nor defined on a rectangular domain, so higher eigenfunctions distort and mix through near-degenerate eigenvalues. The lowest-frequency mode is far more robust, explaining why only the fundamental remains clean in the two-dimensional analysis.

**Bandwidth.** $\varepsilon$ is fixed (chosen for $\sim$10 graph neighbours per batch); edges are retained within a $3\varepsilon$ ball. Optimising the bandwidth, including per-coordinate, is left to future work.

### 5.5. Additional Results

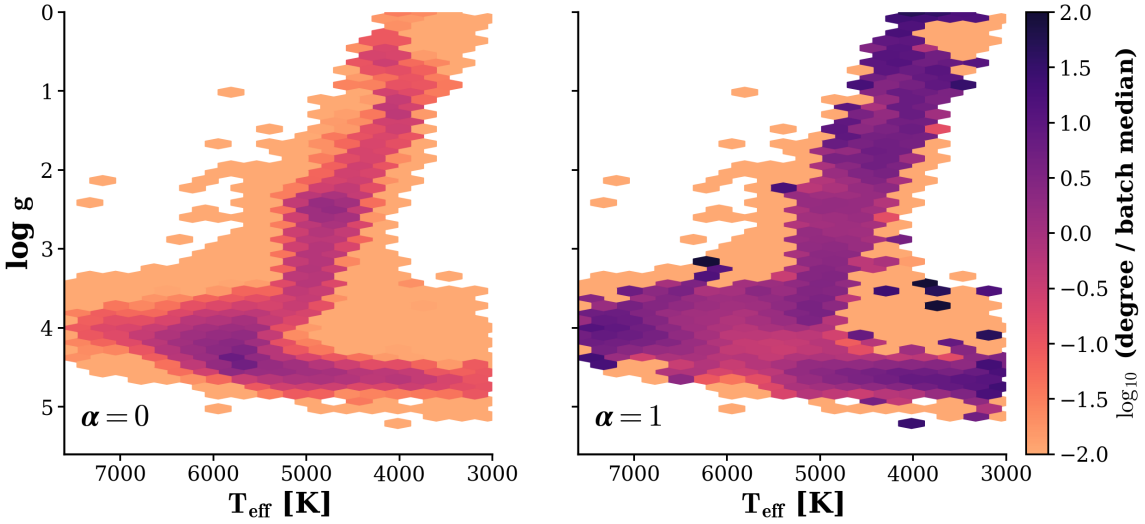



Figure 3: Where the smoothness penalty acts. Each star's graph degree, relative to the batch median ($\log_{10}$), binned in $(T_{\rm eff}, \log g)$. **Left:** without the density correction ($\alpha$=0), the weight concentrates on the densest regions, the main sequence and red clump, and giants, hot stars and subgiants are nearly unweighted. **Right:** with it ($\alpha$=1, our model), the weight spreads across a wider part of the HR diagram; the few very high-degree bins at the sparse edges are stars the correction up-weights where the density estimate is thin.

### 5.6. Marginal shape

A Fourier geometry constrains the feature distribution: a direction close to $\cos(k\pi c)$ should be bimodal across stars (the arcsine law). We test this per principal direction on frozen embeddings, calibrated against a synthetic Fourier manifold (Table 5); per-block coverage is in Table 4. Past the third direction every model is heavy-tailed rather than Gaussian, so "arcsine head over a Gaussian bulk" is not the geometry these encoders produce.

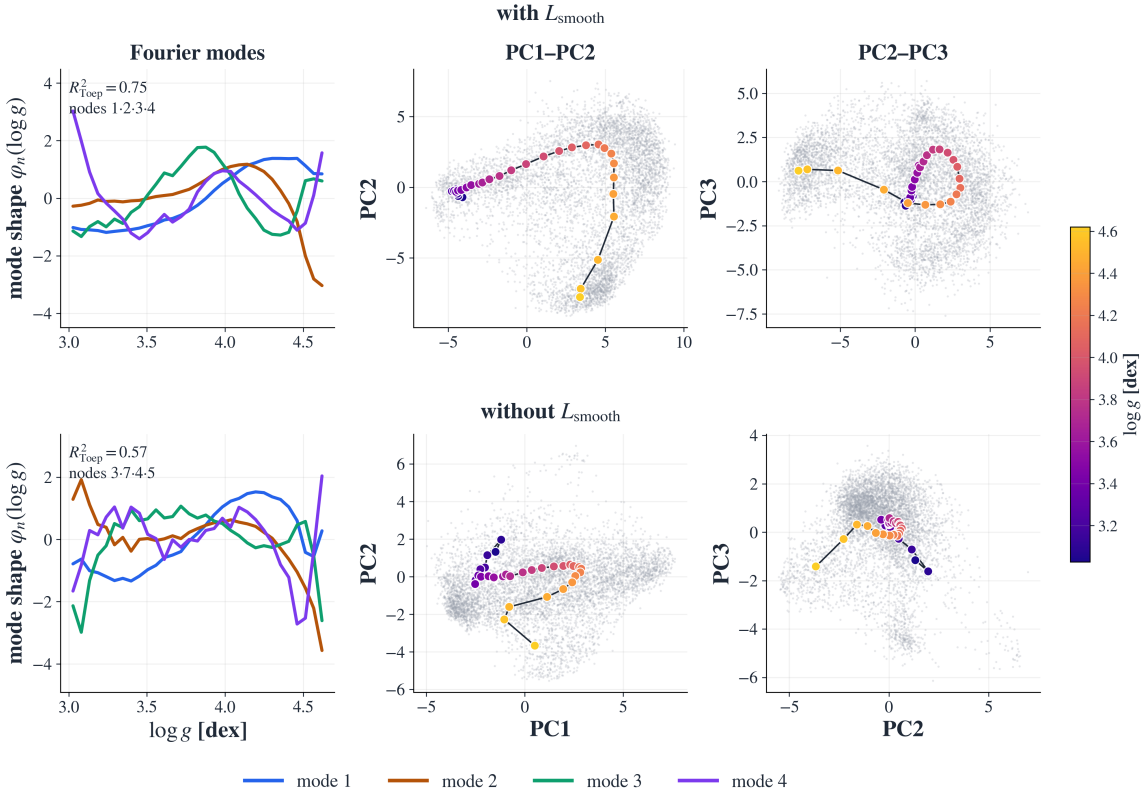



Figure 4: Same as Figure 1, but for $\log g$.

Table 3: Graph controls: linear-probe $R^2$ on the shared token ($n_{\rm train} = 1000$, 30 paired draws) and $R^2_{\rm Toep}$ along $T_{\rm eff}$ on 11,347 clean main-sequence singles with LAMOST, APOGEE and *Gaia*, fixed so that survey coverage cannot drive the comparison. All arms share data, seed and recipe; only the Laplacian's graph differs. The two FAME rows are the same recipe trained twice and set the noise level ($\lesssim 0.03$ in probe $R^2$, 0.04 in $R^2_{\rm Toep}$). One training seed per arm; best checkpoints fall at different epochs (3 to 79).

| Graph | $T_{\rm eff}$ | $\log g$ | [Fe/H] | Age | Act. | $v \sin i$ | $R^2_{\rm Toep}$ |
|---|---|---|---|---|---|---|---|
| Physical (rerun) | 0.872 | 0.873 | 0.530 | 0.557 | 0.165 | 0.481 | 0.852 |
| Physical (paper) | 0.860 | 0.875 | 0.502 | 0.562 | 0.167 | 0.509 | 0.889 |
| No Laplacian | 0.781 | 0.762 | 0.435 | 0.490 | 0.123 | 0.474 | 0.759 |
| Permuted | 0.548 | 0.449 | 0.180 | 0.302 | 0.005 | 0.399 | 0.394 |
| Latent self-graph | 0.705 | 0.504 | 0.204 | 0.448 | 0.109 | 0.442 | 0.775 |
| Input $z_i$ | 0.826 | 0.869 | 0.326 | 0.478 | 0.124 | 0.486 | 0.846 |
| Physical, $\alpha$=0 | 0.812 | 0.813 | 0.147 | 0.508 | 0.103 | 0.468 | 0.909 |
| Physical, warped | 0.879 | 0.889 | 0.442 | 0.555 | 0.158 | 0.427 | 0.861 |

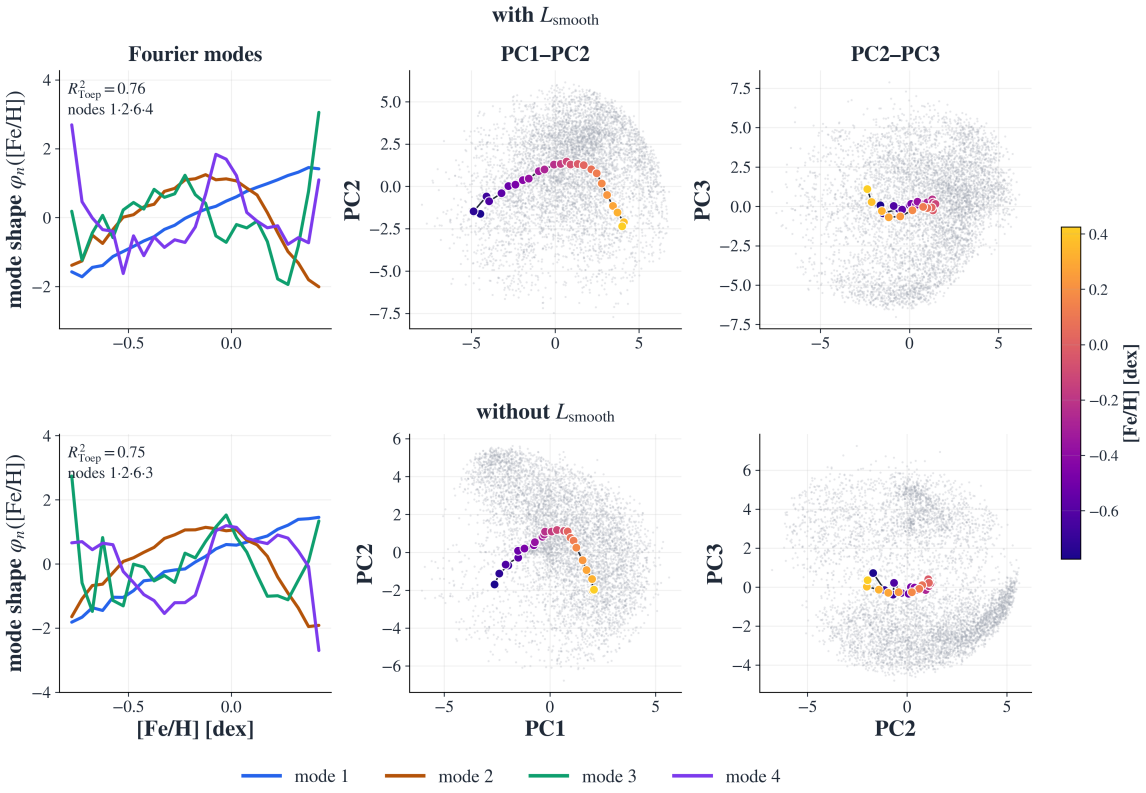



Figure 5: Same as Figure 1, but for [Fe/H].

Table 4: Marginal shape by survey block: excess kurtosis of PC1 ($\gamma_2$; pure 1-D arcsine limit $-1.5$, Gaussian 0), with *strong depth* as a subscript: the number of leading contiguous PCs with $\gamma_2 < -1$. Bold marks $\gamma_2$(PC1) $< -1$. Star sets differ per block ($n$ in the header row), so this table reports *coverage*, not cross-arm magnitude; for the latter see Table 5. *Kepler* is the minimum stratum ($n$ = 1,819) and its verdict is the least secure. $\mathcal{L}_{\rm smooth}$ is the density-normalised ($\alpha$=1) form throughout: a fixed-$\epsilon$ kernel divided by $q_a q_b$, with the VICReg variance and covariance statistics importance-weighted by $1/q$. The prior takes the signature from three blocks to six and tightens it onto the arcsine value; without the prior the same architecture retains it in four blocks, marginally, and loses TESS and GALAH.

| | *Kepler* | LAMOST | APOGEE | TESS | *Gaia* XP | GALAH |
|---|---|---|---|---|---|---|
| $n$ | 1,819 | 14,161 | 20,000 | 20,000 | 20,000 | 20,000 |
| Raw $z_i$ | $+449.38_0$ | $\mathbf{-1.16_2}$ | $\mathbf{-1.56_1}$ | $+22.44_0$ | $+4.21_0$ | $\mathbf{-1.69_1}$ |
| Vanilla VICReg | $-0.95_0$ | $-0.19_0$ | $-0.24_0$ | $-0.05_0$ | $-0.31_0$ | $+0.34_0$ |
| Shared, no $\mathcal{L}_{\rm smooth}$ | $\mathbf{-1.26_1}$ | $\mathbf{-1.13_2}$ | $\mathbf{-1.00_2}$ | $+0.04_0$ | $\mathbf{-1.07_2}$ | $+0.82_0$ |
| Shared, with $\mathcal{L}_{\rm smooth}$ | $\mathbf{-1.40_2}$ | $\mathbf{-1.37_2}$ | $\mathbf{-1.50_2}$ | $\mathbf{-1.33_2}$ | $\mathbf{-1.38_2}$ | $\mathbf{-1.54_1}$ |
| Residual, with $\mathcal{L}_{\rm smooth}$ | $+446.25_0$ | $\mathbf{-1.10_2}$ | $-0.89_0$ | $+5.73_0$ | $+3.35_0$ | $\mathbf{-1.64_1}$ |

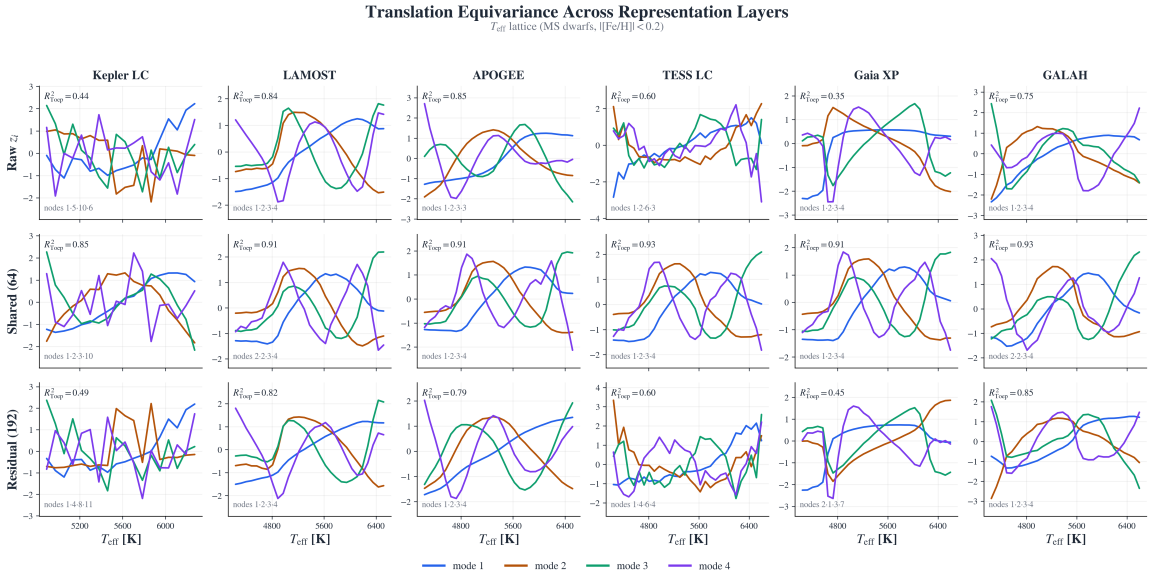



Figure 6: Fourier geometry per modality and layers in the model for $T_{\rm eff}$. Each column represents a specific modality embedding. Each row represents different parts in the model - raw pretrained features, $z_i$, (upper), shared space (middle), and residual space (lower). We see that the shared space improves Fourier geometry compared to $z_i$ across all modalities. The residual stream preserves the geometry of $z_i$.

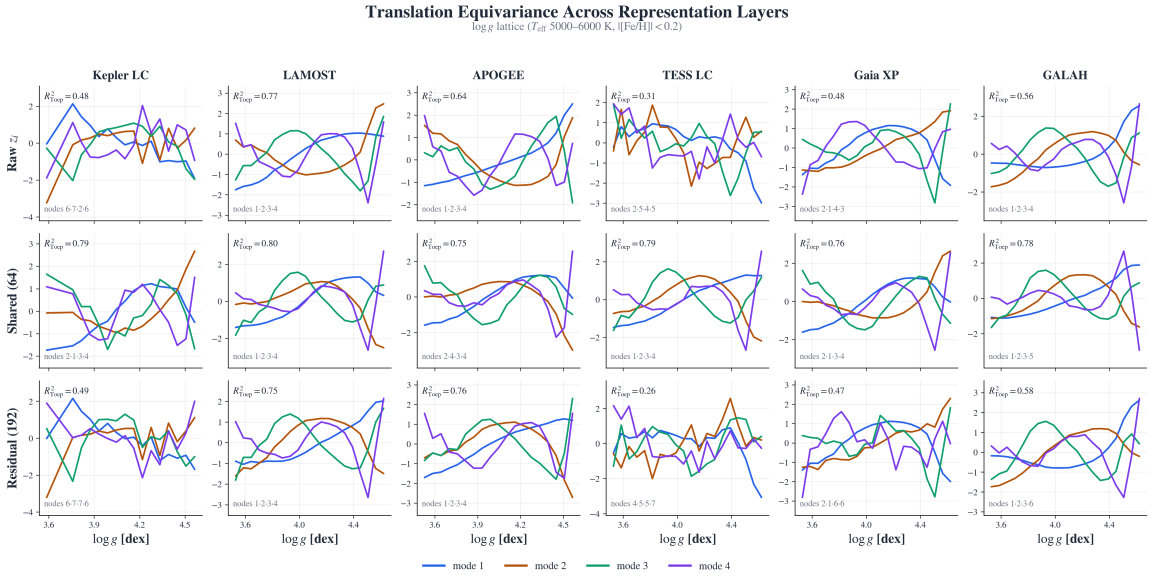



Figure 7: same as Figure 6, for $\log g$

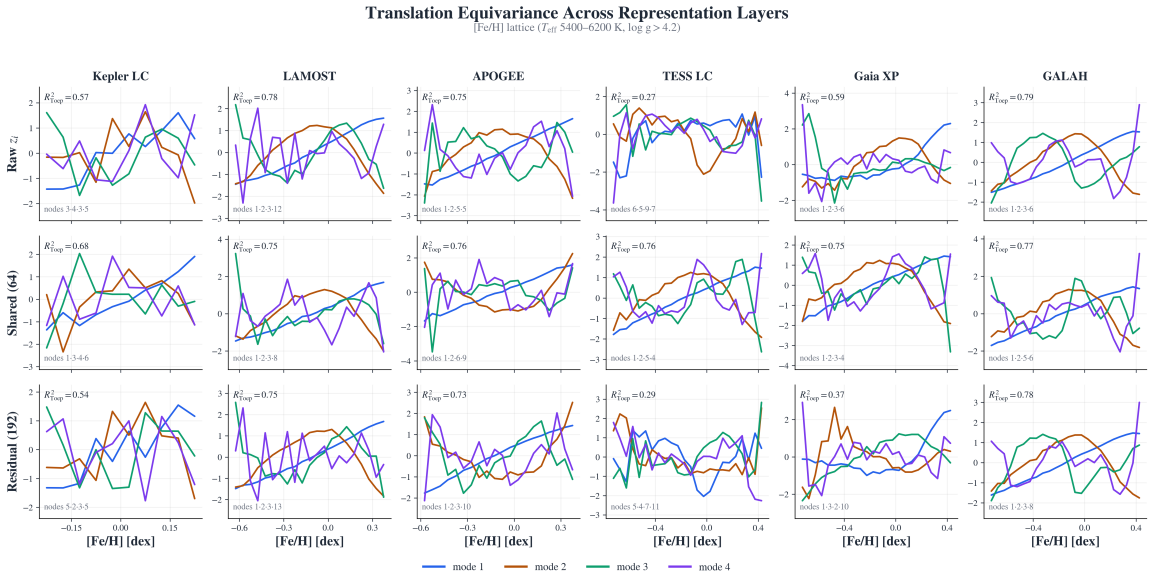



Figure 8: same as Figure 6, for [Fe/H]. Note that this is not a graph coordinate.

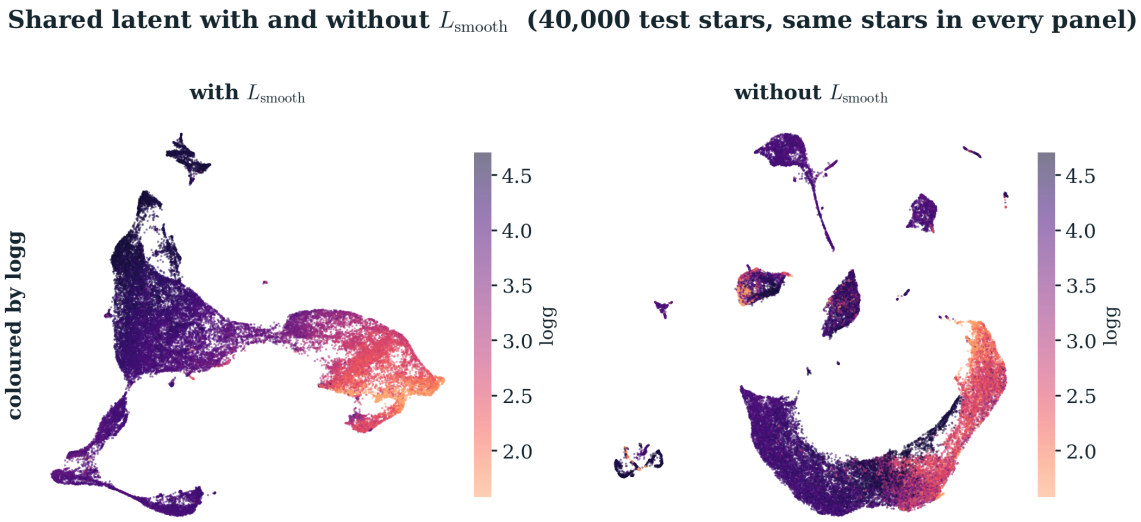



Figure 9: Effect of $\mathcal{L}_{\rm smooth}$ on the latent space. **Left:** UMAP of the shared latent space with $\mathcal{L}_{\rm smooth}$. **Right:** the same model without it. The prior yields a more continuous latent space with a clean $\log g$ gradient.

Table 5: Common star set and calibration. *Upper block:* tier `cross` (clean main-sequence stars with both LAMOST and APOGEE present, $n$ = 10,214; *pooled* is the modality-averaged token on the largest exact availability cell, $n$ = 4,584). Identical stars in every row, so this is the only cross-arm comparison we make. *Lower block:* synthetic manifolds pushed through the identical pipeline at the measured signal fraction (0.899). The density-matched control is the honest reference: a *perfect* Fourier manifold at this sample's own $T_{\text{eff}}$ density reaches strong depth 2, so the achievable depth here is two to three directions. $\mathcal{L}_{\text{smooth}}$ is the density-normalised ($\alpha$=1) form throughout. Single seed, $\varepsilon = 0.0139$.

| Arm | Panel | $\gamma_2$(PC1) | $\gamma_2$(PC2) | Strong depth | Med. PC1–8 | Bulk |
|---|---|---|---|---|---|---|
| Raw $z_i$ | LAMOST | −1.38 | −1.05 | 2 | +0.10 | 0.00 |
| | APOGEE | −1.53 | −0.45 | 1 | +1.04 | 0.00 |
| Vanilla VICReg | LAMOST | −0.29 | −0.32 | 0 | −0.02 | 0.00 |
| | APOGEE | −0.42 | −0.52 | 0 | −0.08 | 0.00 |
| | pooled | −0.48 | −0.60 | 0 | −0.08 | 0.00 |
| Shared, no $\mathcal{L}_{\text{smooth}}$ | LAMOST | −0.86 | −1.10 | 0 | +1.52 | 0.00 |
| | APOGEE | −0.84 | −1.11 | 0 | +1.52 | 0.00 |
| | pooled | −0.22 | −1.21 | 0 | +1.30 | 0.00 |
| Shared, with $\mathcal{L}_{\text{smooth}}$ | LAMOST | **−1.54** | −1.13 | 2 | +2.73 | 0.00 |
| | APOGEE | **−1.54** | −1.13 | 2 | +2.74 | 0.00 |
| | pooled | **−1.58** | −1.09 | 2 | +2.05 | 0.00 |
| Residual, with $\mathcal{L}_{\text{smooth}}$ | LAMOST | **−1.37** | −0.78 | 1 | +1.73 | 0.00 |
| | APOGEE | −0.88 | −0.65 | 0 | +0.52 | 0.00 |
| Control: Fourier, uniform coordinate | — | −1.49 | −1.43 | 4 | −0.73 | 0.90 |
| Control: Fourier, measured density | — | −1.51 | −1.32 | 2 | −0.08 | 0.88 |
| Control: isotropic Gaussian | — | −0.06 | −0.01 | 0 | +0.02 | 0.81 |