Towards an Astronomical Foundation Model for Stars

Henry W. Leung[*] **Jo Bovy**
David A. Dunlap Department of Astronomy and Astrophysics
University of Toronto
50 St. George Street, Toronto
Ontario, M5S 3H4, Canada

Abstract

Rapid strides are currently being made in the field of artificial intelligence using Large Language Models (LLMs) with Transformers architecture. Aside from some use of the base technical components of Transformers—the attention mechanism—the real potential of Transformers in creating artificial intelligence in astronomy has not yet been explored. Here, we introduce a novel perspective on such model in data-driven astronomy by proposing a framework for astronomical data that use the same core techniques and architecture as used by natural-language LLMs. Using a variety of observations and labels of stars as an example, we build a prototype of a foundation model and we show that this model can be easily adapted and trained with cross-survey astronomical data sets. This single model has the ability to perform both discriminative and generative tasks even though the model was not trained to do any specific task that we test it on. This demonstrates that foundation models in astronomy are well within reach and will play a large role in the analysis of current and future large surveys. The full paper is available at `https://ui.adsabs.harvard.edu/abs/2024MNRAS.527.1494L/abstract`.

1 Introduction

Ever-expanding astronomical data set are being collected by large surveys like *Gaia* [1] and SDSS [2, 3] now, and LSST [4] and Euclid [5] in the future, across multiple areas such as spectroscopy, photometry, and time-domain observations. Data-driven analysis has become increasingly popular for these large data set. But so far, customized data-driven models are created for every separate task and data-driven models that focus on cross-survey and/or cross-domain analyses (like the work of Leung et al. 6) are only trained on the intersection instead of union of the relevant data, due to the lack of flexibility in model inputs and outputs. Data-driven science using deep neural networks in particular requires big data to be trained and it would be ideal if we can train such models on most of the available data to truly create a synergistic understanding of multiple surveys.

Currently, there is ongoing rapid development of Large Language Models (LLMs) like OpenAI's GPT [7] that have the ability to do some tasks thought to only be possible with a general intelligence model [8]. Science communities have been critical of LLMs due to the problem of hallucinations and because LLMs can easily fail at simple math problems. These issues mean that naively applying existing LLMs to science is difficult. Moreover, LLMs focus on natural language applications like chat-bots, which involve completely different kinds of data from the floating point astronomical data.

In this work, we present a novel perspective on the use of Transformers in astronomy by constructing a model that utilizes the core ideas and technology of natural-language LLMs without involving natural language. We train a proof-of-concept foundation model that is not trained on specific input/output pairs for specific supervised and unsupervised tasks, but rather is trained with a big data set to contain

[*]henrysky.leung@utoronto.ca

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

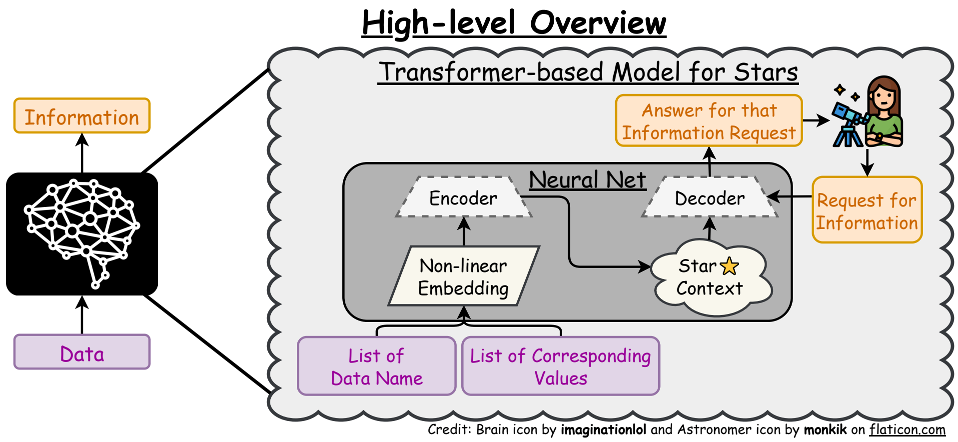

Figure 1: Structure of our foundation model for stars. The left part of the figure displays the overall goal of training a big neural network for astronomy that can turn a combination of observations to useful information. The right part of the figure shows the architecture of the specific model for stars in this work. The network architecture closely resembles that of a typical Transformer encoder-decoder. It consists of an encoder, which encodes information from the inputs to context of a star, and a decoder, which decodes the context of a star based on a request it is given by an agent and outputs the answer to the question. The inputs to the encoder come from a non-linear embedding of the astronomical observations that combines the type of observation and its magnitude (see Equation1).

general knowledge of the data set with self-supervised learning. From this perception of the data, one can later request information from the model. As a proof-of-concept, we specifically build a model for stars using cross-survey, cross-domain data from APOGEE, *Gaia* and 2MASS. Our approach allows us to think about building a big foundation model for astronomy, its potential role in artificial intelligence, and its application in astronomy [9].

2 Transformer-based Model as Foundation Model

Generally, foundation model refers to a kind of large model trained with vast quantity of data that can be fine-tuned for downstream tasks, where Transformers give huge flexibility in input and output node. Such large models significantly outperform smaller models, because in general neural network capabilities scale well with the number of parameters [10]. The number of parameters is limited by compute power and the size of the training data. Models like LLMs gain knowledge about natural languages through a process of pre-training, where the LLM learns to predict, e.g., the next word when provided with the starts of sentences. Even without a concrete training target, these models are able to learn structure (e.g., grammar), the meaning of individual words, as well as ideas about our world.

Building a model that learns general knowledge about scientific areas, such as observations of stars, would obviously be useful in many scientific applications and we propose here that such models can be built by adopting the same core technologies and ideas of LLMs, but applying them to tasks that do not involve natural language. Such model trained with large amount of data can acquire embodied knowledge of stars in an implicit way. Once trained, we can give a context of an astronomical object like a list of known parameters or observations and later request for information like unmeasured parameters about that object.

2.1 Model Implementation

The model has ~ 8.8 millions trainable parameters with 64 context window length. A high level overview of the model is given in Figure 1; this mimics a typical Transformer encoder-decoder. Our model has the basic ability to interact with a user in the decoder which may seem redundant but is critical to train such model to give flexibility to the output node. The input astronomical data passes through a similar embedding processing as occurs in natural language processing. Embedding refers to learning the vectorization of the input data through training of the model. We have implemented a custom embedding process that we name "non-linear embedding" and that embeds data of kind x through

$$y_x = f(w_x \cdot M_x) + w_{b,x} \tag{1}$$

where y_x is the final vectorized data for a particular kind of data x that will be fed into the encoder. The function f is a typical activation function used in neural networks, w_x are the embedding weights, M_x is the magnitude of the data, and $w_{b,x}$ is a bias weight; all of these are particular to the kind of data x. Bias weights are necessary, because the neural network will get a vector of zeros for all data with zero magnitude, so the neural network has no way of knowing which observation that is while zero magnitude has different meaning for different kinds of data.

The "unit vector" w_x in Equation 1 for the non-linear embedding is also used as the request "token" given to the decoder (shown in Figure 1). We train the decodes such that its output is a scalar value corresponding to the requested information vector along with predictive uncertainty estimation.

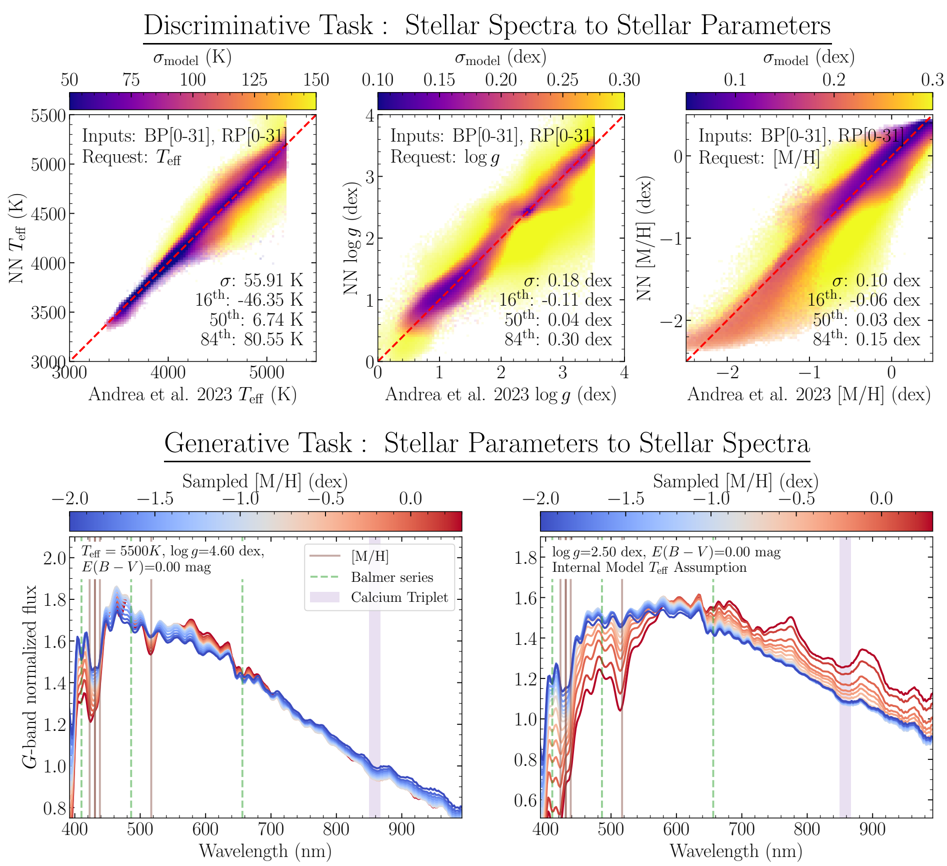

Figure 2: Our single Transformer-based model can do both "generative" and "discriminative" tasks in the traditional sense. The top three panels show a scenario where we give *Gaia* XP spectra for each star, and we request the prediction for T_{eff} (left), $\log g$ (middle) and $[\text{M/H}]$ (right) from the model. The prediction from our model is compared to an external catalog [11] color-coded by our model uncertainty. The bottom two panels (left: main-sequence; right: red-giants) show the opposite scenario where we give stellar parameters with varying metallicity to the model and request the prediction for how should G-band normalized *Gaia* XP spectra look like. We emphasize that the model has never been specifically trained to do either of these task (i.e., these specific combinations of inputs and outputs).

3 Datasets and Training

For our model, we construct a small, low-dimensional heterogeneous data set that is perfect for fast proof-of-concept prototyping. We choose *Gaia* DR3 photometry and XP spectra [12, 13], 2MASS photometry [14], interstellar extinction map [15], APOGEE DR17 [16] stellar properties for $396,718$ stars with majority of them being sub-giants and red giants.

Gaia XP spectra are low-resolution optical to near-infrared ($R \sim 30$ to 100, 330 to $1050\,\mathrm{nm}$; Carrasco et al. 17) spectra obtained from Blue and Red Photometers (BP/RP) aboard the *Gaia* spacecraft. Unlike usual stellar spectra, *Gaia* XP spectra were released as 110 coefficients of an orthogonal basis function expansion, where lower-order coefficients explain large-scale features of the spectra (hence more information like T_{eff}) and higher-order coefficients explain small-scale features including the noise. We simply treat each coefficient as a kind of data. We normalize the XP coefficients by the *Gaia* G-band apparent flux.

In the data set, each star has a row of available observations in which missing data (represented by `NaN`) are replaced by a special padding token (the same idea used to mask empty spaces in sentences), which is masked automatically in the model. In order to train the model without a clear learning objective, we pick a random set of data for each star as input and a random one as output which may or may not be included in the input already even for stars in the same batch during training. This is enabled by the ability to interact with the decoder, because we do not have a fixed output target during training. This way of training is similar to LLM pre-training, where the goal is to predict the next words given the starting part of sentences. But here we do not care about the relative ordering of the input data, but simply with learning the general relationship.

4 Results

To illustrate the general capabilities that our trained model with *all* neurons always participating, we present the result of both a discriminative task and a generative task. Both of these tasks have been investigated by others: e.g., discriminative learning using tree-based machine learning methods [18] and generative modelling using a feed-forward neural network [19]. Our model performs similarly well compared to these previous methods but with a *single* model.

Our model has never been specifically trained to predict the requested labels from the specific set of inputs that we test. The result is shown in Figure 2. By comparing to external catalog like [11], our model predict stellar labels at a similar accuracy compared to works like [20, 19] while having reasonable uncertainty estimation. Our model also have a generative capability that can show what spectral features vary with metallicity. To the best of our knowledge, this is the first model in astronomy that can accomplish both discrimiative and generative task in a single model (i.e., all weights of the model are used for predictions in both tasks). Models like conditional autoencoders can do both tasks, but only the encoder participates in the discriminative task and only the decoder in the generative task.

5 Conclusion

We have introduced a novel framework of utilizing a Transformer-based model towards building a foundation model in astronomy which we hope will accelerate the development such model in the future (a review of neural networks and the role of foundation model in astronomy can be found in Smith and Geach 9). Our model adopts and adapts technology from natural-language LLMs to inherit their advantages (and disadvantages).

We have demonstrated this idea and how it can allow training general astronomy models on big cross-survey, multi-domain astronomical data sets with missing data due to different footprint of survey, etc. This opens up new avenues for training Transformer-based foundation models in the future. Transformer-based models not only provide a way to utilize the advantages of LLMs, but also provide an opportunity for astronomers to build, train, and incorporate our domain expertise in the model. For example, our past experiences of machine learning on a wide range of astronomical data will be crucial in the construction of foundation model(s).

References

[1] Gaia Collaboration et al. The Gaia mission. A&A, 595:A1, November 2016. doi: 10.1051/0004-6361/201629272.

[2] Michael R. Blanton, Matthew A. Bershady, Bela Abolfathi, Franco D. Albareti, Carlos Allende Prieto, Andres Almeida, Javier Alonso-García, Friedrich Anders, Scott F. Anderson, Brett Andrews, Erik Aquino-Ortíz, Alfonso Aragón-Salamanca, Maria Argudo-Fernández, Eric Armengaud, Eric Aubourg, Vladimir Avila-Reese, Carles Badenes, Stephen Bailey, Kathleen A. Barger, Jorge Barrera-Ballesteros, Curtis Bartosz, Dominic Bates, Falk Baumgarten, Julian Bautista, Rachael Beaton, Timothy C. Beers, Francesco Belfiore, Chad F. Bender, Andreas A. Berlind, Mariangela Bernardi, Florian Beutler, Jonathan C. Bird, Dmitry Bizyaev, Guillermo A. Blanc, Michael Blomqvist, Adam S. Bolton, Médéric Boquien, Jura Borissova, Remco van den Bosch, Jo Bovy, William N. Brandt, Jonathan Brinkmann, Joel R. Brownstein, Kevin Bundy, Adam J. Burgasser, Etienne Burtin, Nicolás G. Busca, Michele Cappellari, Maria Leticia Delgado Carigi, Joleen K. Carlberg, Aurelio Carnero Rosell, Ricardo Carrera, Nancy J. Chanover, Brian Cherinka, Edmond Cheung, Yilen Gómez Maqueo Chew, Cristina Chiappini, Peter Doohyun Choi, Drew Chojnowski, Chia-Hsun Chuang, Haeun Chung, Rafael Fernando Cirolini, Nicolas Clerc, Roger E. Cohen, Johan Comparat, Luiz da Costa, Marie-Claude Cousinou, Kevin Covey, Jeffrey D. Crane, Rupert A. C. Croft, Irene Cruz-Gonzalez, Daniel Garrido Cuadra, Katia Cunha, Guillermo J. Damke, Jeremy Darling, Roger Davies, Kyle Dawson, Axel de la Macorra, Flavia Dell'Agli, Nathan De Lee, Timothée Delubac, Francesco Di Mille, Aleks Diamond-Stanic, Mariana Cano-Díaz, John Donor, Juan José Downes, Niv Drory, Hélion du Mas des Bourboux, Christopher J. Duckworth, Tom Dwelly, Jamie Dyer, Garrett Ebelke, Arthur D. Eigenbrot, Daniel J. Eisenstein, Eric Emsellem, Mike Eracleous, Stephanie Escoffier, Michael L. Evans, Xiaohui Fan, Emma Fernández-Alvar, J. G. Fernandez-Trincado, Diane K. Feuillet, Alexis Finoguenov, Scott W. Fleming, Andreu Font-Ribera, Alexander Fredrickson, Gordon Freischlad, Peter M. Frinchaboy, Carla E. Fuentes, Lluís Galbany, R. Garcia-Dias, D. A. García-Hernández, Patrick Gaulme, Doug Geisler, Joseph D. Gelfand, Héctor Gil-Marín, Bruce A. Gillespie, Daniel Goddard, Violeta Gonzalez-Perez, Kathleen Grabowski, Paul J. Green, Catherine J. Grier, James E. Gunn, Hong Guo, Julien Guy, Alex Hagen, ChangHoon Hahn, Matthew Hall, Paul Harding, Sten Hasselquist, Suzanne L. Hawley, Fred Hearty, Jonay I. Gonzalez Hernández, Shirley Ho, David W. Hogg, Kelly Holley-Bockelmann, Jon A. Holtzman, Parker H. Holzer, Joseph Huehnerhoff, Timothy A. Hutchinson, Ho Seong Hwang, Héctor J. Ibarra-Medel, Gabriele da Silva Ilha, Inese I. Ivans, KeShawn Ivory, Kelly Jackson, Trey W. Jensen, Jennifer A. Johnson, Amy Jones, Henrik Jönsson, Eric Jullo, Vikrant Kamble, Karen Kinemuchi, David Kirkby, Francisco-Shu Kitaura, Mark Klaene, Gillian R. Knapp, Jean-Paul Kneib, Juna A. Kollmeier, Ivan Lacerna, Richard R. Lane, Dustin Lang, David R. Law, Daniel Lazarz, Youngbae Lee, Jean-Marc Le Goff, Fu-Heng Liang, Cheng Li, Hongyu Li, Jianhui Lian, Marcos Lima, Lihwai Lin, Yen-Ting Lin, Sara Bertran de Lis, Chao Liu, Miguel Angel C. de Icaza Lizaola, Dan Long, Sara Lucatello, Britt Lundgren, Nicholas K. MacDonald, Alice Deconto Machado, Chelsea L. MacLeod, Suvrath Mahadevan, Marcio Antonio Geimba Maia, Roberto Maiolino, Steven R. Majewski, Elena Malanushenko, Viktor Malanushenko, Arturo Manchado, Shude Mao, Claudia Maraston, Rui Marques-Chaves, Thomas Masseron, Karen L. Masters, Cameron K. McBride, Richard M. McDermid, Brianne McGrath, Ian D. McGreer, Nicolás Medina Peña, Matthew Melendez, Andrea Merloni, Michael R. Merrifield, Szabolcs Meszaros, Andres Meza, Ivan Minchev, Dante Minniti, Takamitsu Miyaji, Surhud More, John Mulchaey, Francisco Müller-Sánchez, Demitri Muna, Ricardo R. Munoz, Adam D. Myers, Preethi Nair, Kirpal Nandra, Janaina Correa do Nascimento, Alenka Negrete, Melissa Ness, Jeffrey A. Newman, Robert C. Nichol, David L. Nidever, Christian Nitschelm, Pierros Ntelis, Julia E. O'Connell, Ryan J. Oelkers, Audrey Oravetz, Daniel Oravetz, Zach Pace, Nelson Padilla, Nathalie Palanque-Delabrouille, Pedro Alonso Palicio, Kaike Pan, John K. Parejko, Taniya Parikh, Isabelle Pâris, Changbom Park, Alim Y. Patten, Sebastien Peirani, Marcos Pellejero-Ibanez, Samantha Penny, Will J. Percival, Ismael Perez-Fournon, Patrick Petitjean, Matthew M. Pieri, Marc Pinsonneault, Alice Pisani, Radosław Poleski, Francisco Prada, Abhishek Prakash, Anna Bárbara de Andrade Queiroz, M. Jordan Raddick, Anand Raichoor, Sandro Barboza Rembold, Hannah Richstein, Rogemar A. Riffel, Rogério Riffel, Hans-Walter Rix, Annie C. Robin, Constance M. Rockosi, Sergio Rodríguez-Torres, A. Roman-Lopes, Carlos Román-Zúñiga, Margarita Rosado, Ashley J. Ross, Graziano Rossi, John Ruan, Rossana Ruggeri, Eli S. Rykoff, Salvador Salazar-Albornoz, Mara Salvato, Ariel G. Sánchez, D. S. Aguado, José R.

Sánchez-Gallego, Felipe A. Santana, Basílio Xavier Santiago, Conor Sayres, Ricardo P. Schiavon, Jaderson da Silva Schimoia, Edward F. Schlafly, David J. Schlegel, Donald P. Schneider, Mathias Schultheis, William J. Schuster, Axel Schwope, Hee-Jong Seo, Zhengyi Shao, Shiyin Shen, Matthew Shetrone, Michael Shull, Joshua D. Simon, Danielle Skinner, M. F. Skrutskie, Anže Slosar, Verne V. Smith, Jennifer S. Sobeck, Flavia Sobreira, Garrett Somers, Diogo Souto, David V. Stark, Keivan Stassun, Fritz Stauffer, Matthias Steinmetz, Thaisa Storchi-Bergmann, Alina Streblyanska, Guy S. Stringfellow, Genaro Suárez, Jing Sun, Nao Suzuki, Laszlo Szigeti, Manuchehr Taghizadeh-Popp, Baitian Tang, Charling Tao, Jamie Tayar, Mita Tembe, Johanna Teske, Aniruddha R. Thakar, Daniel Thomas, Benjamin A. Thompson, Jeremy L. Tinker, Patricia Tissera, Rita Tojeiro, Hector Hernandez Toledo, Sylvain de la Torre, Christy Tremonti, Nicholas W. Troup, Octavio Valenzuela, Inma Martinez Valpuesta, Jaime Vargas-González, Mariana Vargas-Magaña, Jose Alberto Vazquez, Sandro Villanova, M. Vivek, Nicole Vogt, David Wake, Rene Walterbos, Yuting Wang, Benjamin Alan Weaver, Anne-Marie Weijmans, David H. Weinberg, Kyle B. Westfall, David G. Whelan, Vivienne Wild, John Wilson, W. M. Wood-Vasey, Dominika Wylezalek, Ting Xiao, Renbin Yan, Meng Yang, Jason E. Ybarra, Christophe Yèche, Nadia Zakamska, Olga Zamora, Pauline Zarrouk, Gail Zasowski, Kai Zhang, Gong-Bo Zhao, Zheng Zheng, Zheng Zheng, Xu Zhou, Zhi-Min Zhou, Guangtun B. Zhu, Manuela Zoccali, and Hu Zou. Sloan Digital Sky Survey IV: Mapping the Milky Way, Nearby Galaxies, and the Distant Universe. AJ, 154(1):28, July 2017. doi: 10.3847/1538-3881/aa7567.

[3] Juna A. Kollmeier, Gail Zasowski, Hans-Walter Rix, Matt Johns, Scott F. Anderson, Niv Drory, Jennifer A. Johnson, Richard W. Pogge, Jonathan C. Bird, Guillermo A. Blanc, Joel R. Brownstein, Jeffrey D. Crane, Nathan M. De Lee, Mark A. Klaene, Kathryn Kreckel, Nick MacDonald, Andrea Merloni, Melissa K. Ness, Thomas O'Brien, Jose R. Sanchez-Gallego, Conor C. Sayres, Yue Shen, Ani R. Thakar, Andrew Tkachenko, Conny Aerts, Michael R. Blanton, Daniel J. Eisenstein, Jon A. Holtzman, Dan Maoz, Kirpal Nandra, Constance Rockosi, David H. Weinberg, Jo Bovy, Andrew R. Casey, Julio Chaname, Nicolas Clerc, Charlie Conroy, Michael Eracleous, Boris T. Gänsicke, Saskia Hekker, Keith Horne, Jens Kauffmann, Kristen B. W. McQuinn, Eric W. Pellegrini, Eva Schinnerer, Edward F. Schlafly, Axel D. Schwope, Mark Seibert, Johanna K. Teske, and Jennifer L. van Saders. SDSS-V: Pioneering Panoptic Spectroscopy. *arXiv e-prints*, art. arXiv:1711.03234, November 2017. doi: 10.48550/arXiv.1711.03234.

[4] Željko Ivezić et al. LSST: From Science Drivers to Reference Design and Anticipated Data Products. ApJ, 873(2):111, March 2019. doi: 10.3847/1538-4357/ab042c.

[5] R. Laureijs et al. Euclid Definition Study Report. *arXiv e-prints*, art. arXiv:1110.3193, October 2011. doi: 10.48550/arXiv.1110.3193.

[6] Henry W. Leung, Jo Bovy, J. Ted Mackereth, and Andrea Miglio. A variational encoder-decoder approach to precise spectroscopic age estimation for large Galactic surveys. MNRAS, 522(3): 4577–4597, July 2023. doi: 10.1093/mnras/stad1272.

[7] OpenAI. GPT-4 Technical Report. *arXiv e-prints*, art. arXiv:2303.08774, March 2023. doi: 10.48550/arXiv.2303.08774.

[8] Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of Artificial General Intelligence: Early experiments with GPT-4. *arXiv e-prints*, art. arXiv:2303.12712, March 2023. doi: 10.48550/arXiv.2303.12712.

[9] Michael J. Smith and James E. Geach. Astronomia ex machina: a history, primer, and outlook on neural networks in astronomy. *arXiv e-prints*, art. arXiv:2211.03796, November 2022. doi: 10.48550/arXiv.2211.03796.

[10] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling Laws for Neural Language Models. *arXiv e-prints*, art. arXiv:2001.08361, January 2020. doi: 10.48550/arXiv.2001.08361.

[11] René Andrae, Hans-Walter Rix, and Vedant Chandra. Robust Data-driven Metallicities for 175 Million Stars from Gaia XP Spectra. ApJS, 267(1):8, July 2023. doi: 10.3847/1538-4365/acd53e.

[12] 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. Fremat, 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. *arXiv e-prints*, art. arXiv:2206.06143, June 2022. doi: 10.48550/arXiv.2206.06143.

[13] P. Montegriffo, F. De Angeli, R. Andrae, M. Riello, E. Pancino, N. Sanna, M. Bellazzini, D. W. Evans, J. M. Carrasco, R. Sordo, G. Busso, C. Cacciari, C. Jordi, F. van Leeuwen, A. Vallenari, G. Altavilla, M. A. Barstow, A. G. A. Brown, P. W. Burgess, M. Castellani, S. Cowell, M. Davidson, F. De Luise, L. Delchambre, C. Diener, C. Fabricius, Y. Fremat, M. Fouesneau, G. Gilmore, G. Giuffrida, N. C. Hambly, D. L. Harrison, S. Hidalgo, S. T. Hodgkin, G. Holland, S. Marinoni, P. J. Osborne, C. Pagani, L. Palaversa, A. M. Piersimoni, L. Pulone, S. Ragaini, M. Rainer, P. J. Richards, N. Rowell, D. Ruz-Mieres, L. M. Sarro, N. A. Walton, and A. Yoldas. Gaia Data Release 3: External calibration of BP/RP low-resolution spectroscopic data. *arXiv e-prints*, art. arXiv:2206.06205, June 2022. doi: 10.48550/arXiv.2206.06205.

[14] M. F. Skrutskie, R. M. Cutri, R. Stiening, M. D. Weinberg, S. Schneider, J. M. Carpenter, C. Beichman, R. Capps, T. Chester, J. Elias, J. Huchra, J. Liebert, C. Lonsdale, D. G. Monet, S. Price, P. Seitzer, T. Jarrett, J. D. Kirkpatrick, J. E. Gizis, E. Howard, T. Evans, J. Fowler, L. Fullmer, R. Hurt, R. Light, E. L. Kopan, K. A. Marsh, H. L. McCallon, R. Tam, S. Van Dyk, and S. Wheelock. The Two Micron All Sky Survey (2MASS). AJ, 131(2):1163–1183, February 2006. doi: 10.1086/498708.

[15] Jo Bovy, Hans-Walter Rix, Gregory M. Green, Edward F. Schlafly, and Douglas P. Finkbeiner. On Galactic Density Modeling in the Presence of Dust Extinction. ApJ, 818(2):130, February 2016. doi: 10.3847/0004-637X/818/2/130.

[16] Abdurro'uf and et al. 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.

[17] J. M. Carrasco, M. Weiler, C. Jordi, C. Fabricius, F. De Angeli, D. W. Evans, F. van Leeuwen, M. Riello, and P. Montegriffo. Internal calibration of Gaia BP/RP low-resolution spectra. A&A, 652:A86, August 2021. doi: 10.1051/0004-6361/202141249.

[18] Rene Andrae, Hans-Walter Rix, and Vedant Chandra. Robust Data-driven Metallicities for 120 Million Stars from Gaia XP Spectra. *arXiv e-prints*, art. arXiv:2302.02611, February 2023. doi: 10.48550/arXiv.2302.02611.

[19] Xiangyu Zhang, Gregory M. Green, and Hans-Walter Rix. Parameters of 220 million stars from Gaia BP/RP spectra. *arXiv e-prints*, art. arXiv:2303.03420, March 2023. doi: 10.48550/arXiv.2303.03420.

[20] Hans-Walter Rix, Vedant Chandra, René Andrae, Adrian M. Price-Whelan, David H. Weinberg, Charlie Conroy, Morgan Fouesneau, David W. Hogg, Francesca De Angeli, Rohan P. Naidu, Maosheng Xiang, and Daniela Ruz-Mieres. The Poor Old Heart of the Milky Way. ApJ, 941 (1):45, December 2022. doi: 10.3847/1538-4357/ac9e01.