



Analyzing Soccer Players' Skill Ratings Over Time Using Tensor-Based Methods

Kenneth Verstraete¹ , Tom Decroos², Bruno Coussement²,
Nick Vannieuwenhoven² , and Jesse Davis² 

¹ Department of Chronic Diseases, Metabolism and Ageing, KU Leuven,
3000 Leuven, Belgium

`kenneth.verstraete@kuleuven.be`

² Department of Computer Science, KU Leuven, 3001 Leuven, Belgium
{`tom.decroos,nick.vannieuwenhoven,jesse.davis`}@cs.kuleuven.be

Abstract. Soccer players have a variety of skills such as passing, tackling, shooting and dribbling. However, their abilities are not fixed and evolve over time. Understanding this evolution could be interesting from many perspectives. We analyze player skill data from the FIFA video game series by EA Sports using tensor methods. This data can be organized as a tensor over three dimensions, namely players, skills, and age, which we explore in two different ways. First, we use a polyadic decomposition to uncover hidden structures among skills and see how these structures evolve over time. Second, we use a Tucker decomposition to predict how a specific player's skills will evolve over time.

1 Introduction

Playing soccer involves multiple technical skills such as passing, tackling, shooting and dribbling as well as physical skills like acceleration and endurance. Over time, a player's ability will change; with practice skills can improve, but as a player ages certain skills will inevitably decline. Understanding this evolution is interesting from several perspectives. First, gaining intuition about the relationship between different skills and how they change over time can yield insights into the game. Second, predicting how a player will evolve can help a club with roster decisions such as identifying potential transfer targets or deciding which current players to retain and for how long.

One potential source of skill data comes from the FIFA video game series, where EA Sports models all soccer players using ratings for in-game skills with the aim of reflecting their real-life soccer skills.¹ In this paper, we explore modeling temporal evolutions in this skill data set. These data can be naturally modeled using a tensor with three *modes* or axes: players, skills and age. We make three contributions. First, we highlight a variety of challenges that arose while analyzing the SoFIFA data. Second, we explored the SoFIFA data tensor using the *canonical polyadic decomposition* (CPD) in order to extract interpretable

¹ These skill ratings are accessible online at <https://www.sofifa.com>.

latent structures. Moreover, this decomposition can yield some insights into how these structures evolve over time. Third, we employ the *Tucker decomposition* of a tensor in order to project how a specific player’s skills will evolve over time. That is, we can predict a player’s ratings for each skill at a future age.

2 Tensors and the SoFIFA Data

We provide a brief background on tensors, describe our data set and highlight a number of challenges encountered when analyzing this data.

2.1 Tensors

Tensors are a generalization of vectors and matrices that allow us to model the interactions between different variables such as players, skills, and age. An order- d tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times \dots \times n_d}$ can be identified with a d -array. The different dimensions of the tensors are called the modes. This paper focuses on order-3 tensors, but all presented theory can be generalized to higher orders as well. See the survey by Kolda and Bader [4] for more details.

2.2 Data Description


We used all available data from the FIFA 07 game up to the FIFA 18 game to track players’ skill evolution. Human experts employed by EA Sports assign each soccer player ratings between 0 and 100 for 24 different skills (e.g., crossing, finishing, dribbling, etc; see Fig. 1). Which experts and how many assigned ratings to each player is not known. Updating the skill ratings is done in an open-source fashion; a community of 8,000 coaches, scouts and season ticket holders can submit inconsistencies which are then checked and fixed by a small team of 25 EA producers [10]. SoFIFA.com details each player’s biological data, wage, club, positions and skills in player cards (Fig. 1). These player cards are updated through time. From 2007 to 2012, the player cards were updated biannually, and thereafter EA Sports started updating them weekly.

For each player, we can capture the evolution of his skills as he ages in a matrix by scraping the skill values on his player card as they were on January 1st of each year from 2007 to 2018. We then stack these player matrices in a three-dimensional tensor of the form: *players* \times *skills* \times *ages*.

2.3 Data Challenges

Analyzing real-world data is challenging. In this paper, most of the challenges are related to the quality of the skill ratings. We discuss five challenges.

Cold start. Players have played few matches at the start of their career. Hence it is hard for experts to construct a player card with accurate skill ratings for new players based on a small amount of data.

Lionel Messi (ID: 158023) 

Age 31 (Jun 24, 1987) 170cm 72kg

CF RW ST

Value €110.5M Wage €565K

94 Overall Rating

94 Potential

Attacking

86 Crossing

95 Finishing

70 Heading Accuracy

92 Short Passing

Skill

97 Dribbling

94 FK Accuracy

89 Long Passing

96 Ball Control

Movement

91 Acceleration

86 Sprint Speed

95 Reactions

Goalkeeping

6 GK Diving

11 GK Handling

15 GK Kicking

14 GK Positioning

8 GK Reflexes

Power

85 Shot Power

72 Stamina

66 Strength

94 Long Shots

Mentality

48 Aggression

75 Penalties

Defending

33 Marking

28 Standing Tackle

Fig. 1. Lionel Messi's player card with his ratings for the 24 different skills [2]. Each skill rating is a number between 0 and 100 that quantifies how proficient the player is at that skill estimated by human experts. The overall rating is a weighted average of the skill ratings. The exact weighting scheme employed by FIFA is unknown, as is the way they predict the potential of a player, which is the highest overall rating a player can attain in the future.

Human subjectivity. Rating players' skills is inherently a subjective task. A difference in experience between human experts can lead to different opinions about the same player.

Ratings for irrelevant skills. Each player receives a rating for every skill available, including those that are not relevant for their position (e.g., field players receive ratings for goalkeeping skills).

Artificial boosts. EA Sports is known to manually boost the ratings for well-known players if they disagree with the experts' opinion [5], which disrupts the quality and consistency of the ratings in the data set.

Disruptive skill corrections. Three types of corrections can occur that are unrelated to actual changes in player skill. First, due to the cold start problem, early-career players can receive wildly inaccurate initial skill ratings, which are later (when the player has played more matches) corrected by a large value, causing a disruptive jump in the skill evolution. Second, a collective correction can happen where rating for a specific skill is dramatically altered for a large group of players. For example, all field players had their rating for the *GK Kicking* skill significantly lowered in the FIFA 11 edition of the game. Finally, the newer FIFA games receive updates on a weekly basis and hence, rating adjustments have become more event-related now compared to older FIFA games, when updates only happened once or twice a year.

For example, Messi’s penalty skill was lowered from 78 to 75 in the update from 21 December 2017 after he missed a penalty kick on 17 December 2017 [6, 7].

3 Exposing Hidden Structure with CPD

Since we are dealing with temporal data that describes the evolution of players over time, tensor decompositions provide a straightforward and natural way to discover patterns and predict evolution. A decomposition that allows us to discover intelligible hidden, underlying patterns in the player data is the canonical polyadic decomposition [3]. The CPD approximates a tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ as a sum of r pure tensors \mathcal{A}_i , where each \mathcal{A}_i can be written as the tensor product of the vectors \mathbf{p}_i , \mathbf{s}_i , and \mathbf{a}_i , living in respectively \mathbb{R}^{n_1} , \mathbb{R}^{n_2} and \mathbb{R}^{n_3} :

$$\mathcal{A} \approx \sum_{i=1}^r \mathcal{A}_i = \sum_{i=1}^r \mathbf{p}_i \otimes \mathbf{s}_i \otimes \mathbf{a}_i. \quad (1)$$

Recall that for $\mathbf{x} \in \mathbb{R}^{n_1}$, $\mathbf{y} \in \mathbb{R}^{n_2}$ and $\mathbf{z} \in \mathbb{R}^{n_3}$ we have

$$\mathbf{x} \otimes \mathbf{y} \otimes \mathbf{z} := [x_i y_j z_k]_{(i,j,k)=(1,1,1)}^{(n_1, n_2, n_3)} \in \mathbb{R}^{n_1 \times n_2 \times n_3}.$$

We then define the player factor matrix as $P := [\mathbf{p}_1 \ \mathbf{p}_2 \ \dots \ \mathbf{p}_r]$, the skill factor matrix as $S := [\mathbf{s}_1 \ \mathbf{s}_2 \ \dots \ \mathbf{s}_r]$ and the age factor matrix as $A := [\mathbf{a}_1 \ \mathbf{a}_2 \ \dots \ \mathbf{a}_r]$. If the CPD is of rank r and all the elements in the rank-1 terms are non-negative, then the rank-1 terms can be seen as the r dominant building blocks of our tensor, which is conceptually similar to non-negative matrix factorization. A crucial feature of CPD that distinguishes it from many matrix-based approaches is that the factorization into the pure tensors \mathcal{A}_i is unique up to the order of the summands; see [1] for details.

To illustrate a CPD-based analysis, we take the rank-7 CPD of a tensor containing skill ratings of 17,859 players for all 24 skills over the age range [23..29]. Each term can be interpreted as a collection of correlated skills that make up a common characteristic or trait observed in different soccer players. Figure 2a illustrates each term’s skill vector containing the weights of all 24 skills. Each weight in such a vector represents how important the skill is in the term. Each term’s evolution over player age is visualized in Fig. 2b. For most of the terms the evolution is near-constant. We now briefly discuss the terms.

1. **Mean.** The first term of the decomposition represents the mean of the data as it is the most dominant part of the data.
2. **Evolution term.** The second term describes the overall positive evolution of skills up to the age of 28–29.
3. **Defender trait.** The third term represents the trait of a defender, showing high weights for the *Marking* and *Standing Tackle* skills.
4. **Heading trait.** The fourth term represents the trait of players that are good at heading, showing high weights for *Strength* and *Aggression*. Interestingly, this trait is also correlated with *Acceleration* and *Sprint Speed* according to our decomposition.

5. **Goalkeeper trait.** The fourth term represents the trait of goalkeepers, with high weights for all goalkeeper skills (e.g., *GK Diving*, *GK Handling*, *GK Positioning*) and more general skills also associated with goalkeepers (e.g., *Strength*, *Reactions*).
6. **Correction term.** Prior to FIFA 11, all players had a high rating for *GK Kicking*. Starting from the FIFA 11 edition, all field players' ratings for this skill were significantly lowered. This term accounts for this correction and illustrates how CPD can expose hidden structure in the data as we were initially unaware of this correction in the data.
7. **Striker trait.** The seventh term is the trait of players who are good at skills typically associated with strikers such as *Finishing*, *Dribbling*, *Shot Power*, *Long Shots*. and *Penalties*.

Player Examples. We illustrate how each player is essentially a linear combination of these traits. The coefficients for each player can be found in the player factor matrix. We give an example for a famous player in each position. Their coefficients can be found in Table 1.


Table 1. Coefficients for Messi, Modrić, Chiellini and Buffon.

Player	Position	Mean	Evolution	Defender	Heading	Goalkeeper	Correction	Striker
Messi	ATT	811.82	27.42	0.0	0.0	0.0	6.03	46.30
Modrić	MID	766.61	34.63	26.55	0.0	0.26	9.70	14.88
Chiellini	DEF	572.59	14.92	56.53	21.05	5.55	17.24	7.81
Buffon	GK	98.32	12.17	7.62	4.52	75.19	2.06	29.51


Lionel Messi. His coefficients for the mean term (811.82) and the striker trait (46.30) are extremely high. This matches his reputation as one of the best offensive soccer players ever. His coefficients for the defender, heading and goalkeeper traits are equal to zero, which also matches his playing style and small stature.

Luka Modrić. With midfielders being a combination of attackers and defenders, there is no real midfielder trait in the CPD. The Croatian player does not possess particularly high coefficients for any of the traits compared to other players, but is reasonably strong in all of them, except for heading.

Giorgio Chiellini. A typical classic defender. With his tall physical presence, he is good at heading, as illustrated by his high coefficient for the heading trait (21.05). Compared to Messi, whose coefficient for the heading trait is zero, Chiellini shows significantly higher ratings than Messi for heading skills such as *Heading Accuracy* (83 vs 70) and *Strength* (89 vs 59).



(a) The skill vectors of the seven terms in the CPD of our data tensor.



(b) The age vectors of the seven terms in the CPD of our data tensor. The orange evolution term describes the positive evolution of skills up to the age of 28-29. The brown correction term accounts for the reduction in *GK Kicking* all players received at the introduction of FIFA 11. The other five terms are relatively stable over all ages.

Fig. 2. The vectors of the rank-7 CPD visualized. (Color figure online)

Gianluigi Buffon. Buffon has a low coefficient for the mean (98.32) because goalkeepers have low ratings for most skills. Naturally, his coefficient for the goalkeeper trait (75.19) is extremely high compared to the field players. Having a low value for the coefficient of the mean can cause other coefficients to have illogical values; a small peak in any non-goalkeeper skill has to be fitted by a non-goalkeeper trait with a large coefficient. As an example, Buffon has a strong value for the striker trait (29.51) but that does not make him a good striker. Interpreting the values of non-goalkeeper traits for goalkeepers is often not useful.

4 Predicting Skill Ratings Using Tucker Decomposition

Discovering underlying patterns in player data is certainly interesting from a research perspective. However, a more interesting direction from an application perspective is predicting a player's evolution, as this can directly influence decision making on player acquisition and player retainment. Our task is then the following:

Given: Skill ratings of players for ages $[h_1..h_k]$

Predict: Skill ratings of the same players for ages $[h_{k+1}..h_N]$

We address this task by learning latent structures in a full tensor with training data for ages $[h_1..h_N]$ using the Tucker decomposition. Unlike the CPD, the factor matrices of this decomposition are less intuitive as they are not unique.

4.1 Tucker Decomposition Theory

The multilinear multiplication is a tensor multiplication in which a tensor is multiplied with a matrix in each mode [4]. The multilinear multiplication in which a tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is multiplied in each mode i by matrix $M_i \in \mathbb{R}^{m_i \times n_i}$, is denoted by

$$(M_1, M_2, M_3) \cdot \mathcal{A} \in \mathbb{R}^{m_1 \times m_2 \times m_3}.$$

The Tucker decomposition can be seen as a higher-order analogue of the principal component analysis [8]. In this decomposition, a tensor $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ is factorized in a core tensor $\mathcal{S} \in \mathbb{R}^{r_1 \times r_2 \times r_3}$ multiplied with factor matrices along each mode. In the three-dimensional case, the factorization is

$$\mathcal{A} = (U_1, U_2, U_3) \cdot \mathcal{S},$$

with $\mathcal{A} \in \mathbb{R}^{n_1 \times n_2 \times n_3}$ the original tensor, $\mathcal{S} \in \mathbb{R}^{r_1 \times r_2 \times r_3}$ the core tensor and $U_1 \in \mathbb{R}^{n_1 \times r_1}$, $U_2 \in \mathbb{R}^{n_2 \times r_2}$, $U_3 \in \mathbb{R}^{n_3 \times r_3}$ the factor matrices.

If $n_i > r_i$, $i = 1, 2, 3$, the core tensor \mathcal{S} can be seen as a compressed version of the tensor \mathcal{A} . It describes how and to which extent the elements of the tensor interact with each other using the factor matrices.

4.2 Predicting Skill Ratings

Given the full data tensor $\mathcal{A} \in \mathbb{R}^{p_{\text{total}} \times M \times N}$ with p_{total} players, M skills and the data from age h_1 until h_N , the tensor is split into a training tensor $\mathcal{A}_{\text{train}} \in \mathbb{R}^{p_{\text{train}} \times M \times N}$ with p_{train} players and a test tensor $\mathcal{A}_{\text{test}} \in \mathbb{R}^{p_{\text{test}} \times M \times N}$ with p_{test} players, such that $p_{\text{total}} = p_{\text{train}} + p_{\text{test}}$. From the test set, only the first h_k ages are used. The predictions are made for the ages $[h_{k+1}..h_N]$.

In order to make predictions, a model has to be trained first. Unlike most learning algorithms, there are no iterative steps in which a model is trained. The latent structures are extracted by computing a single decomposition. Our approach consists of the following three steps (see Fig. 3).

1. **Extracting latent structures.** The latent structures of the data are extracted using the Tucker decomposition on the training tensor $\mathcal{A}_{\text{train}}$. These structures are represented by the factor matrices of the decomposition; \hat{U}_1 for the player factor matrix, \hat{U}_2 for the skill factor matrix and \hat{U}_3 for the age factor matrix (Fig. 3a).
2. **Finding the player factor matrix.** We find a player factor matrix \hat{U}_1^{test} such that the test tensor $\mathcal{A}_{\text{test}}$ is best approximated by a Tucker decomposition using the core \hat{S} , the skill factor matrix \hat{U}_2 and the truncated age factor matrix \hat{U}_3^{trunc} using only the rows of ages $[h_1..h_k]$ from the train tensor $\mathcal{A}_{\text{train}}$ (Fig. 3b). Finding this player factor matrix can be reduced to solving a least-squares problem.
3. **Completing the test tensor.** We can now complete the missing data in $\mathcal{A}_{\text{test}}$ by multiplying the player factor matrix \hat{U}_1^{test} with the core \hat{S} , the skill factor matrix \hat{U}_2 , and the other part of the age factor matrix \hat{U}_3^{other} corresponding to the ages $[h_{k+1}..h_N]$ (Fig. 3c).

4.3 Experiments


To evaluate how well our tensor-based method can predict future skill ratings, we address three different instances of our prediction task. $h_k : [h_{k+1}..h_N]$ denotes the task of using data of the players at age h_k to predict their skill ratings at the ages $[h_{k+1}..h_N]$. We predict the evolution of young players (18 : [19..26]), mid-career players (23 : [24..31]), and older players (26 : [27..34]). The goalkeeper skills *GK Diving*, *GK Handling*, *GK Kicking*, *GK Positioning*, and *GK Reflexes* are disregarded, as these skills are irrelevant for most players.

For each task, we report the mean absolute error (MAE) over all players over all skills over ages $[h_{k+1}..h_N]$ and compare our tensor-based method against two baseline models. The first baseline model predicts players to have no evolution at all, i.e., we use a player’s skill ratings at age h_k as the predictions for his skill ratings at ages $[h_{k+1}..h_N]$. The second baseline model uses the well-known k -nearest neighbors algorithm (KNN) in which the predictions are made using the k closest neighbors in the data set. The results can be found in Table 2.


Table 2. The mean absolute error of the models for the three prediction tasks.

Prediction method	18 : [19..26]	23 : [24..31]	26 : [27..34]
No evolution baseline	10.45	8.23	6.90
KNN baseline ($k = 10$)	8.31	7.82	7.67
Tensor-based method	8.57	7.71	6.74


If we focus on short-term predictions (1–2 years into the future), then the no evolution model shows the best performance. This is expected, as soccer players tend to improve gradually and thus show little evolution over 1–2 years.



(a) The Tucker decomposition of the training tensor $\mathcal{A}_{\text{train}}$ is computed. The i^{th} row of the player factor matrix, \hat{U}_1 , corresponds to the data of the i^{th} player in the training tensor.



(b) Given $\mathcal{A}_{\text{test}}$ containing the data for ages $[h_1..h_k]$ of the test players, core tensor \hat{S} , skill factor matrix \hat{U}_2 and the truncated age factor matrix \hat{U}_3^{trunc} , the coefficients for all players in the test set, \hat{U}_1^{test} , are computed.



(c) With \hat{U}_1^{test} computed, the predictions \mathcal{P} are now calculated by using the multilinear multiplication with \hat{U}_3^{other} : $\mathcal{P} = (\hat{U}_1^{\text{test}}, \hat{U}_2, \hat{U}_3^{\text{other}}) \cdot \hat{S}$.

Fig. 3. The three steps of our tensor-based method to predict skill ratings.

If we focus on long-term predictions (5–7 years into the future), then KNN shows the best performance. This matches the observations of Vroonen et al. who showed the superiority of KNN over the no evolution baseline for long-term predictions [9]. Our tensor-based method seems to hit a sweet spot between these two prediction methods, as evidenced by its lowest overall MAE for two of the three prediction tasks in Table 2.

Note that the predictions in our experiments are based on only one year of data. The performance of the tensor-based method degrades if we use more than one year of data. We can think of two reasons for this. First, the algorithm to decompose a tensor aims to optimally compress the entire data set and not necessarily to optimally predict missing data for certain ages. There is thus a technical misalignment between our prediction task and the task that the tensor decomposition is solving. Fixing this technical misalignment could be interesting future work. Second, skill ratings barely change in one or two years. Hence, taking three years instead of one year, for example, does not add much

more information to the prediction model. Vroonen et al. [9] also came to the conclusion that increasing the number of years of given data to make predictions, has a limited effect on performance.

5 Conclusions

In this paper, we made three contributions to the field of soccer analytics. First, we highlighted a variety of challenges that arose while processing the SoFIFA data. Second, we explored the SoFIFA data tensor using the canonical polyadic decomposition (CPD) in order to extract interpretable latent structures. We showed how these latent structures group together related skills, how they evolve over players' ages, and how each player can be summarized as a linear combination of these latent structures. Third, we employed the Tucker decomposition of a tensor in order to project how a specific player's skills will evolve as he ages.

Acknowledgments. Tom Decroos is supported by the Research Foundation Flanders (FWO). Nick Vannieuwenhoven is supported by a Postdoctoral Fellowship of the Research Foundation Flanders with project 12E8119N. Jesse Davis is partially supported by KU Leuven Research Fund (C14/17/070 and C32/17/036), Interreg V A project NANO4Sports and the Flemish Government under the "Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen" programme.

References

1. Chiantini, L., Ottaviani, G., Vannieuwenhoven, N.: Effective criteria for specific identifiability of tensors and forms. *SIAM J. Matrix Anal. Appl.* **38**, 656–681 (2017). <https://doi.org/10.1137/16M1090132>
2. EA Vancouver: FIFA 19, video game (2018)
3. Hitchcock, F.L.: The expression of a tensor or a polyadic as a sum of products. *J. Math. Phys.* **6**(1–4), 164–189 (1927). <https://doi.org/10.1002/sapm192761164>
4. Kolda, T., Bader, B.: Tensor decompositions and applications. *SIAM Rev.* **51**(3), 455–500 (2009). <https://doi.org/10.1137/07070111X>
5. Mulani, D.: The embarrassing story behind Thomas Muller's FIFA rating (2017). <https://www.sportskeeda.com/esports/the-embarrassing-story-behind-thomas-mullers-fifa-rating>. Accessed 16 Dec 2018
6. SoFIFA.com. <https://sofifa.com/player/158023/lionel-messi/changeLog>. Accessed 27 Dec 2018
7. transfermarkt.com. <https://www.transfermarkt.com/lionel-messi/elfmetertore/spieler/28003>. Accessed 27 Dec 2018
8. Tucker, L.R.: Some mathematical notes on three-mode factor analysis. *Psychometrika* **31**(3), 279–311 (1966). <https://doi.org/10.1007/bf02289464>
9. Vroonen, R., Decroos, T., Haaren, J.V., Davis, J.: Predicting the potential of professional soccer players. In: Davis, J., Kaytoue, M., Zimmermann, A. (eds.) 4th Workshop on Machine Learning and Data Mining for Sports Analytics (MLSA), CEUR Workshop Proceedings, Aachen, vol. 1971, pp. 1–10 (2017). <http://ceur-ws.org/Vol-1971/paper-02.pdf>
10. WIRED: Meet the data master behind EA Sports' popular FIFA franchise. <https://datamakespossible.westerndigital.com/meet-data-master-ea-sports-fifa/>. Accessed 15 Dec 2018