Modelling Pediatric Cardiometabolic Health Through Longitudinal Relationships and Personalized What-if Simulations

Nerma Begić
Jožef Stefan Institute
Department of Intelligent Systems
Ljubljana, Slovenia

Mitja Luštrek
Jožef Stefan Institute
Department of Intelligent Systems
Ljubljana, Slovenia
mitja.lustrek@ijs.si

Abstract

This paper examines how behavioural, anthropometric, fitness, metabolic, and blood-pressure variables are related over time and how these relationships can be used in individualized what-if simulations. Adjusted-lagged and change-based analyses, together with trajectory analysis, were used to examine temporal relationships and patterns of change and to identify the most consistently supported associations. The supported relationships were then used to build a health-state cascade, which describes how the measured variables are connected. The strongest relationships involved anthropometry, BMI, metabolic markers, and blood pressure.

The cascade was then used to simulate a healthier-lifestyle scenario in children and adolescents aged 7–18 years by moving physical activity, sleep, and diet to healthy target values using clinical and age-specific reference thresholds. The healthier behavioural values were then propagated through the cascade using separate XGBoost models trained for each predicted health variable. The models iteratively updated the remaining health variables until the predictions became stable. The largest model-predicted health-zone transitions were found for blood pressure. For SBP, 43.8% of baseline-yellow profiles moved to the green zone, while changes were also predicted for several metabolic markers. Overall, the results show how longitudinal relationships and what-if simulations can be combined to explore how a healthier lifestyle may be reflected across several connected health variables.

Keywords

pediatric health, cardiometabolic risk, longitudinal modelling, what-if simulation, XGBoost, health-state cascade

1 Introduction

Cardiometabolic health is shaped by several factors that develop together, including lifestyle, body composition, metabolic health, and blood pressure. These relationships are especially important during childhood and adolescence, when excess weight is common and persistent overweight is associated with later hypertension and metabolic disease [21, 9].

Many existing approaches focus on individual outcomes, which makes it harder to see how different parts of a child's health profile are connected. In this work, we therefore organize behavioural, anthropometric, metabolic, and blood-pressure variables into a structured health-state cascade.

The main contributions of this work are:

- We systematically evaluate longitudinal relationships among behavioural, anthropometric, metabolic, and blood-pressure variables.
- Based on the strongest and most consistently supported relationships, we construct an interpretable health-state cascade.
- We use the resulting cascade to generate individualized what-if scenarios based on changes in physical activity, sleep, and diet.

2 Related Work

Longitudinal studies have examined several links between lifestyle, body composition, metabolic health, and cardiovascular outcomes in children and adolescents. Physical activity, for example, has been associated with waist circumference and blood pressure [20], as well as with metabolic-syndrome markers including triglycerides, HDL cholesterol, glucose, and blood pressure [16]. Diet has also been studied in relation to later cardiometabolic health [13], while obesity and higher adiposity have been associated with adverse lipid, glycemic, and blood-pressure profiles [5]. Together, these studies show that different aspects of cardiometabolic health are closely connected. However, most focus on a small number of relationships rather than on how several of these links fit together over time.

Machine-learning studies have also combined a wider range of pediatric health information for prediction. Javidi et al. [6] used longitudinal clinical data to predict several pediatric metabolic diseases, while Calcaterra et al. [3] combined lifestyle, anthropometric, and clinical data to model several cardiometabolic outcomes within the same platform. These studies demonstrate the value of combining multiple health domains, but mainly focus on individual prediction targets rather than on how broader relationships can be organized into a connected health-state cascade.

Counterfactual and 'what-if' methods examine how a model's prediction would change under alternative input scenarios [19]. In healthcare, these methods have been used to explore personalized changes that could lead to better predicted outcomes. Lenatti et al. [11] used counterfactual explanations to identify changes associated with moving patients to a lower cardiovascular-risk group. These approaches usually focus on one outcome rather than on how changes may spread across several connected health variables.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
Information Society 2026, Ljubljana, Slovenia
© 2026 Copyright held by the owner/author(s).
https://doi.org/10.70314/is.2026.skui.126

3 Methodology

3.1 Overview of the Framework

The framework has two stages. We first examine longitudinal relationships among behavioural, anthropometric, fitness, metabolic, and blood-pressure variables. The complete cascade is not fixed in advance, instead, the most consistently supported relationships are used to define its structure. The resulting cascade is then used for individualized what-if simulations based on changes in physical activity, sleep, and diet. BMI is calculated from weight and height rather than predicted as a separate outcome.

3.2 Longitudinal Relationship Analysis

The adjusted-lagged analysis was treated as the primary longitudinal analysis, while change-based and trajectory analyses provided complementary evidence. The final cascade was defined from the combined pattern of relationships across all three analyses rather than from any single method alone.

The primary adjusted-lagged model was [12]

$$Y_{i,T1} = \beta_0 + \beta_1 X_{i,T0} + \beta_2 Y_{i,T0} + \boldsymbol{\gamma}^T \mathbf{Z}_i + \varepsilon_i, \tag{1}$$

where X is the source variable, Y the target variable, and $\mathbf{Z}$ contains age, gender, country, and the time between measurements. Here, $T0$ denotes the earlier observation and $T1$ the following observation.

To further examine the consistency of the temporal relationships, we also investigated whether changes in the source and target occurred together between measurements. We defined these changes as

$$\Delta X_i = X_{i,T1} - X_{i,T0}, \qquad \Delta Y_i = Y_{i,T1} - Y_{i,T0}.$$

The change scores were analysed with baseline source and target values, age, gender, country, and time between measurements as covariates [12, 4].

To examine relationships across more than two measurements, we fitted participant-specific linear trends over age for source and target variables in participants with at least three measurements. Associations between the slopes were adjusted for baseline values, age, number of measurements, follow-up duration, gender, and country, and repeated for participants with at least four measurements as a sensitivity check.

For each tested relationship, we calculated a standardized regression coefficient (β), which describes the strength and direction of the association. Larger absolute values of β indicate stronger associations, while the sign indicates whether the relationship is positive or negative. For example, in the adjusted-lagged analysis, $\beta = 0.20$ means that a one-standard-deviation higher source value is associated with a 0.20-standard-deviation higher target value at the following measurement, after considering the target value at baseline and the other covariates.

Variables were grouped into behaviour, anthropometry, metabolic markers, blood pressure, fitness, resting heart rate, and BMI. For each pathway, absolute standardized coefficients of its individual relationships were averaged to obtain mean $|\beta|$.

Relationships were considered supported when the 95% confidence interval excluded zero and the FDR-adjusted $q < 0.05$, using the Benjamini–Hochberg correction [1].

The adjusted-lagged analysis provided the primary evidence for direction, with change-based and trajectory analyses providing complementary support. Pathways were qualitatively summarized as strong, moderate, or weaker based on coefficient magnitude and consistency across analyses.

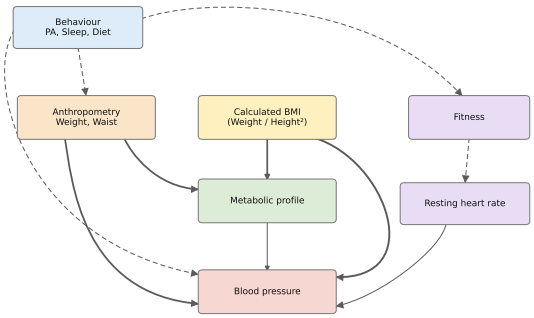


Figure 1: Final health-state cascade. Thick, thin, and dashed arrows represent strongly, moderately, and less consistently supported relationships, respectively. Arrow direction indicates the modelled relationship direction. BMI is calculated from weight and height.

The strongest and most consistently supported relationships across the analyses were used to define the final health-state cascade shown in Figure 1.

3.3 Personalized What-if Simulation

The cascade was used to simulate a standardized healthy-behaviour scenario. Blood pressure, waist circumference, and metabolic markers were mapped to green, yellow, and red health zones using clinical and age-specific thresholds; waist thresholds also depended on gender and blood-pressure thresholds on height percentile in younger children.

Physical activity, sleep, and diet were treated as modifiable inputs and moved to healthy targets when outside their healthy range. The healthy sleep and diet targets remained within the observed training support. The physical activity target of 7 h/week was also within the training range but above the 99th percentile (3.95 h/week), making the combined target profile idealized rather than typical of the training data.

Separate XGBoost models predicted weight, waist circumference, fitness, resting heart rate, metabolic markers, SBP, and DBP; BMI was recalculated from weight and height. Behavioural targets were kept fixed while predicted health variables were iteratively updated through the cascade. Updates stopped when all changes fell below variable-specific numerical tolerances (0.00025–0.005) for two consecutive iterations. All test profiles met this criterion within five iterations, and the cascade contained no feedback loops.

Model performance was evaluated using R^2, MAE, and RMSE. What-if results were summarized by marker changes and health-zone transitions. Transition percentages were calculated per observation, with 95% confidence intervals from 1,000 participant-clustered bootstrap resamples. Sensitivity analysis was repeated using only models with validation $R^2 \geq 0.50$ (weight, waist circumference, triglycerides and SBP), while other outcomes were kept at baseline.

4 Experiments

4.1 Dataset

We used the SmartCHANGE dataset [15, 10], containing 265,852 records from 8,310 participants and 26 variables. Data came from

Table 1: Adjusted-lagged results.

Path	Mean $\|\beta\|$	Strongest link
B ⟶ A	0.0261	Sleep ⟶ Weight
A ⟶ M	0.0991	Waist ⟶ HDL
A ⟶ BP	0.2009	Waist ⟶ DBP
M ⟶ BP	0.0381	TG ⟶ DBP
BP ⟶ B	0.0307	SBP ⟶ Diet
BMI ⟶ M	0.0791	BMI ⟶ HDL
BMI ⟶ BP	0.1616	BMI ⟶ DBP
B ⟶ F	0.0409	Sleep ⟶ Fitness
RHR ⟶ F	0.0448	RHR ⟶ Fitness
RHR ⟶ BP	0.0441	RHR ⟶ DBP

B: behaviour; A: anthropometry; M: metabolic profile; BP: blood pressure; F: fitness (final speed reached in the shuttle-run test); RHR: resting heart rate; TG: triglycerides; SBP/DBP: systolic/diastolic blood pressure.

SLOfit [8], the Leuven Growth Study [2, 17], ABCD [18], and the Young Finns Study [14].

Before merging, corresponding variables were harmonized and comparable measures standardized where needed, following the SmartCHANGE data-integration framework [10, 7]. Missing values were addressed using model-based prediction during dataset construction [7]. The present analyses used complete values for the required variables. Country was included as a contextual variable/covariate rather than a separate study identifier, so residual study-specific differences may remain.

4.2 Experimental Setup

4.2.1 Longitudinal Analysis. Participants with at least two eligible measurements were included, resulting in 16,791 consecutive measurement pairs from 5,146 participants. The predefined training, validation, and test splits were used, but we ensured that all measurements from the same participant were kept in the same split to avoid participant overlap between the training, validation, and test sets. The training set was used to identify candidate relationships, the validation set to check their stability, and the test set for final evaluation.

4.2.2 Personalized What-if Simulation. After restricting the data to ages 7–18 years, 75,991 observations from 5,592 participants remained across the training, validation, and test sets. The training set contained 49,512 observations, the validation set 11,564, and the test set 14,915. The training set was used for predictive modelling, the validation set to evaluate the models, and the test set to generate the final what-if results.

5 Results

5.1 Longitudinal Analysis

The adjusted-lagged analysis showed the strongest relationships for anthropometry to blood pressure (mean $|\beta| = 0.2009$), followed by BMI to blood pressure (mean $|\beta| = 0.1616$) (Table 1). Relationships involving behaviour, fitness, and resting heart rate were generally weaker.

The change-based analysis showed a similar overall pattern (Table 2). The largest mean coefficients were found for anthropometry to blood pressure (0.2088), BMI to blood pressure (0.2077), BMI to metabolic markers (0.2063), and anthropometry to metabolic markers (0.1904).

The trajectory analysis provided an additional view of how variables changed across repeated measurements (Table 3). The largest trajectory associations were observed between blood pressure and BMI (mean $|\beta| = 0.4825$) and between anthropometry and blood pressure (mean $|\beta| = 0.4614$). Anthropometry and metabolic markers also showed a notable relationship (mean $|\beta| = 0.2880$). In contrast, relationships involving behaviour, fitness, and resting heart rate were generally weaker.

Taken together, the three longitudinal analyses were used to define the final cascade shown in Figure 1. Based on the obtained β coefficients and the consistency of the relationships across the three analyses, the cascade pathways were classified as strong, moderate, or weaker. Consistent support was found for behaviour-to-anthropometry, anthropometry-to-metabolic, anthropometry-to-blood-pressure, and BMI-related pathways.

Table 2: Change-based results.

Path	Mean $\|\beta\|$	Strongest link
B ⟶ A	0.0832	Diet ⟶ Weight
A ⟶ M	0.1904	Waist ⟶ TG
A ⟶ BP	0.2088	Weight ⟶ SBP
M ⟶ BP	0.0626	TG ⟶ DBP
B ⟶ BP	0.0744	Diet ⟶ SBP
BMI ⟶ M	0.2063	BMI ⟶ TG
BMI ⟶ BP	0.2077	BMI ⟶ SBP
B ⟶ F	0.0513	PA ⟶ Fitness
F ⟶ RHR	0.0030	F ⟶ RHR
RHR ⟶ BP	0.0646	RHR ⟶ SBP

PA: physical activity; other abbreviations as in Table 1.

Table 3: Trajectory results for participants with at least three visits.

Path	Mean $\|\beta\|$	Strongest link
A ⟶ B	0.1370	Weight ⟶ Diet
A ⟶ M	0.2880	Weight ⟶ HDL
A ⟶ BP	0.4614	Weight ⟶ SBP
BP ⟶ M	0.1750	SBP ⟶ HDL
BP ⟶ B	0.1224	SBP ⟶ Sleep
M ⟶ BMI	0.2210	HDL ⟶ BMI
BP ⟶ BMI	0.4825	SBP ⟶ BMI
B ⟶ F	0.0928	PA ⟶ Fitness
F ⟶ RHR	0.1032	Fitness ⟶ RHR
BP ⟶ RHR	0.0258	SBP ⟶ RHR

Abbreviations as in Table 2.

Table 4: Performance of the XGBoost models for each marker.

Target	R^2	MAE	RMSE
Weight [kg]	0.919	2.76	4.10
Waist circumference [cm]	0.914	2.20	3.10
Fitness [km/h]	0.453	0.52	0.77
Resting heart rate [bpm]	0.282	6.06	9.78
Glucose [mmol/L]	0.040	0.07	0.41
HDL [mmol/L]	0.245	0.10	0.18
LDL [mmol/L]	0.122	0.26	0.50
Triglycerides [mmol/L]	0.610	0.10	0.20
Total cholesterol [mmol/L]	0.147	0.29	0.53
DBP [mmHg]	0.240	3.18	5.79
SBP [mmHg]	0.514	3.36	6.11

5.2 Personalized What-if Simulation

Before generating the what-if profiles, the predictive models were evaluated on the validation set (Table 4). The strongest predictive performance was obtained for weight ($R^2 = 0.919$) and waist circumference ($R^2 = 0.914$). Triglycerides ($R^2 = 0.610$) and SBP ($R^2 = 0.514$) also showed relatively good performance. Fitness showed moderate performance, while the remaining metabolic, resting-heart-rate, and DBP models were less accurate.

Table 5: Model-generated health-zone transitions under the what-if scenario, reported per observation with 95% participant-clustered bootstrap CIs. R/Y/G denote red/yellow/green; Δ is the mean predicted change.

Marker	R→Y (% [95% CI])	Δ R→Y	Y→G (% [95% CI])	Δ Y→G	R→G (% [95% CI])	Δ R→G
Waist circumference	4.0 [1.7–6.9]	−4.108 cm	8.2 [3.5–13.1]	−3.486 cm	0.5 [0.0–2.0]	−8.345 cm
HDL	13.0 [0.5–35.6]	+0.048 mmol/L	11.4 [2.8–21.1]	+0.040 mmol/L	0.0 [0.0–0.0]	–
LDL	6.6 [3.1–11.0]	−0.213 mmol/L	10.1 [6.1–14.8]	−0.158 mmol/L	0.0 [0.0–0.0]	–
Triglycerides	13.8 [9.1–19.1]	−0.359 mmol/L	16.1 [7.8–25.3]	−0.156 mmol/L	16.0 [10.9–21.3]	−0.595 mmol/L
Total cholesterol	7.5 [4.1–11.6]	−0.241 mmol/L	7.8 [6.0–9.7]	−0.176 mmol/L	0.0 [0.0–0.0]	–
SBP	16.3 [7.4–25.2]	−3.16 mmHg	43.8 [31.3–57.4]	−3.69 mmHg	8.0 [1.7–17.0]	−5.86 mmHg
DBP	33.9 [13.5–55.0]	−5.06 mmHg	63.7 [43.4–81.8]	−5.31 mmHg	17.1 [3.8–34.2]	−12.96 mmHg

After moving physical activity, sleep, and diet to their healthy target values, the simulation predicted health-zone transitions for several outcomes (Table 5). For DBP, 33.9% moved from red to yellow, 63.7% from yellow to green, and 17.1% directly from red to green; however, these results remain exploratory because of the model's lower predictive performance. The corresponding SBP transitions were 16.3%, 43.8%, and 8.0%.

As a sensitivity analysis, we repeated the simulation while updating only outcomes with validation $R^2 \geq 0.50$; lower-performing outcomes were kept at their baseline values. The main predicted patterns for weight, waist circumference, triglycerides, and SBP remained similar.

6 Conclusion

We developed an interpretable framework combining longitudinal analysis with individualized what-if simulation. The most consistent relationships involved anthropometry, metabolic markers, BMI, and blood pressure, while behavioural relationships were weaker. Evidence from all three analyses was used to construct the final cascade.

Using the resulting cascade, we simulated healthier physical activity, sleep, and diet values. The what-if simulations predicted changes in blood pressure and several metabolic markers. Results from lower-performing models should be interpreted more cautiously. These simulations are hypothesis-generating and require prospective validation before any causal interpretation.

Acknowledgments

This work was supported by the SmartCHANGE project, funded by the European Union's Horizon Europe research and innovation programme under grant agreement No. 101080965. The SLOfit dataset was provided by the University of Ljubljana (courtesy of Gregor Jurak et al.), the Leuven Growth Study dataset by KU Leuven (courtesy of Martine Thomis), the Amsterdam Born Children and their Development dataset by Amsterdam University Medical Centers, and the Young Finns Study dataset by the University of Turku.

References

[1] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B*, 57, 1, 289–300.

[2] Gaston P. Beunen, Robert M. Malina, Marc A. Van't Hof, Johan Simons, Michel Ostyn, Roland Renson, and Dirk Van Gerven. 1988. *Adolescent Growth and Motor Performance: A Longitudinal Study of Belgian Boys. HKP Sport Science Monograph Series*. Vol. 1. Human Kinetics, Champaign, IL.

[3] Valeria Calcaterra, Umberto Ciriello, Samuele Medici, Valter Pagani, Cristina Campoy, Lucia Labati, et al. 2026. PODiaCarD: a prototype of a digital twin platform for the management of pediatric obesity and related cardiometabolic complications. *European Journal of Pediatrics*, 185, 2, 77.

[4] Oliver Dukes, Zach Shahn, and Audrey Renson. 2025. Change scores and baseline adjustment: splitting the difference (in differences). *International Journal of Epidemiology*, 54, 3, dyaf042.

[5] Yun Huang, Karolina Sulek, Sara E. Stinson, et al. 2025. Lipid profiling identifies modifiable signatures of cardiometabolic risk in children and adolescents with obesity. *Nature Medicine*, 31, 1, 294–305.

[6] Hamed Javidi, Arshiya Mariam, Lina Alkhaled, Kevin M. Pantalone, and Daniel M. Rotroff. 2024. An interpretable predictive deep learning platform for pediatric metabolic diseases. *Journal of the American Medical Informatics Association*, 31, 6, 1227–1238.

[7] Marko Jordan, Nina Reščič, Sebastjan Kramar, Marcel Založnik, and Mitja Luštrek. 2024. SmartCHANGE risk prediction tool: demonstrating risk assessment for children and youth. In *Proceedings of the Slovenian Conference on Artificial Intelligence 2024*. doi:10.70314/is.2024.scai.8844.

[8] Gregor Jurak et al. 2020. SLOfit surveillance system of somatic and motor development of children and adolescents: upgrading the slovenian sports educational chart. *Acta Universitatis Carolinae Kinanthropologica*, 56, 1, 28–40.

[9] Kelemu Tilahun Kibret, Claudia Strugnell, Kathryn Backholer, Anna Peeters, Teketo Kassaw Tegegne, and Melanie Nichols. 2024. Life-course trajectories of body mass index and cardiovascular disease risks and health outcomes in adulthood: systematic review and meta-analysis. *Obesity Reviews*, 25, 4, e13695.

[10] Sebastjan Kramar, Marko Jordan, Marcel Založnik, Nina Reščič, and Mitja Luštrek. 2024. Health risk prediction for children built on a combination of heterogeneous datasets. In *ETAI 2024 Conference Proceedings* number 1. Vol. 3. Paper ID 39.

[11] Marta Lenatti, Alberto Carlevaro, Aziz Guergachi, Karim Keshavjee, Maurizio Mongelli, and Alessia Paglialonga. 2025. Estimation and conformity evaluation of multi-class counterfactual explanations for chronic disease prevention. *IEEE Journal of Biomedical and Health Informatics*, 29, 9, 6132–6142.

[12] Oliver Lüdtke and Alexander Robitzsch. 2025. ANCOVA versus change score for the analysis of two-wave data. *The Journal of Experimental Education*, 93, 2, 363–395.

[13] Oonagh Markey, Rebecca Pradeilles, Sophie Goudet, Paula L. Griffiths, Benjamin Boxer, Christopher Carroll, and Emily K. Rousham. 2023. Unhealthy food and beverage consumption during childhood and risk of cardiometabolic disease: a systematic review of prospective cohort studies. *The Journal of Nutrition*, 153, 1, 176–189.

[14] Olli T. Raitakari et al. 2008. Cohort profile: the cardiovascular risk in young finns study. *International Journal of Epidemiology*, 37, 6, 1220–1226.

[15] Nina Reščič et al. 2023. Smartchange: ai-based long-term health risk evaluation for driving behaviour change strategies in children and youth. In *2023 International Conference on Applied Mathematics Computer Science (ICAMCS)*, 81–89.

[16] Alejandro Sánchez-Delgado, Sandra Sánchez-Parente, David Martínez-Gómez, Sonia Gómez-Martínez, Oscar L. Veiga, Ascensión Marcos, José Castro-Piñero, and Alejandro Pérez-Bey. 2023. Objective measured physical activity and metabolic syndrome score in children and adolescents: the up&down longitudinal study. *Scandinavian Journal of Medicine & Science in Sports*, 33, 11, 2299–2312.

[17] Johan Simons, Gaston Beunen, Roland Renson, Albrecht L. M. Claessens, Bart Vanreusel, and Johan A. V. Lefevre. 1990. *Growth and Fitness of Flemish Girls: The Leuven Growth Study. HKP Sport Science Monograph Series*. Vol. 3. Human Kinetics, Champaign, IL.

[18] Manon van Eijsden, Tanja G. M. Vrijkotte, Reinoud J. B. J. Gemke, and Marcel F. van der Wal. 2011. Cohort profile: the amsterdam born children and their development (ABCD) study. *International Journal of Epidemiology*, 40, 5, 1176–1186.

[19] Sahil Verma, Varich Boonsanong, Minh Hoang, Keegan E. Hines, John P. Dickerson, and Chirag Shah. 2024. Counterfactual explanations and algorithmic recourses for machine learning: a review. *ACM Computing Surveys*, 56, 12, 1–42.

[20] Sally P. Waterworth, Catherine J. Kerr, Christopher J. McManus, Henry C. Chung, Brandon S. Shaw, Ina Shaw, and Gavin R. Sandercock. 2024. Four-year longitudinal associations of physical activity, waist circumference, and blood pressure in UK adolescents. *Pediatric Research*, 95, 3, 736–743.

[21] Xinyue Zhang et al. 2024. Global prevalence of overweight and obesity in children and adolescents: a systematic review and meta-analysis. *JAMA Pediatrics*, 178, 8.