

Towards Guiding Data Imputation for Scientific Data Analytics Via SHAP-like Scores

Omer Abramovich¹, Hadas Stiebel-Kalish² and Daniel Deutch¹

¹Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel

²Gray Faculty of Health Life Sciences, Tel Aviv University, Tel Aviv, Israel; Eye Laboratory, Felsenstein Medical Research Center, Tel Aviv, Israel; Department of Ophthalmology, Rabin Medical Center, Petah Tikva, Israel

Abstract

Missing data is prevalent in the context of scientific databases, arising from measurements failures, partially manually filled data, and other reasons. Missing data may adversely affect scientific analytics and needs to be dealt with, in a process referred to as *data imputation*. Standard techniques for data imputation either incur loss of data, introduce errors, or require significant manual labor to complete the values. In this paper we put forward a novel approach, that guides the data imputation process by adapting and extending a recently emerging technique for explainable computation, namely *attribution*. Attribution involves assigning importance scores to individual data items based on their contribution to the computation result, and concrete notions of such scores have been extensively studied in recent years based on measurements from game theory such as Shapley and Banzhaf values. Our observation is that these attribution scores are valuable not only for explaining the computation but also, in presence of missing values, for guiding data imputation in the sense of deciding what data to compute and what resources to allocate for it. We present the approach as well as a concrete use case in the domain of medical data, and highlight multiple directions for future research.

Declaration on Generative AI

Generative AI tools (GPT-5.2) were only used for proofreading and light editing of this paper. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

1. Introduction

This paper describes a novel approach for guiding data imputation in the context of scientific analysis. We start by recalling the problem of missing data in scientific analysis, highlighting the main challenges. We then briefly overview existing approaches for data imputation and their limitations. We follow with a description of notions of attribution and importance scores and then describe our high-level approach towards a solution.

Missing Data in Scientific Analytics Scientific studies typically involve the analysis of data that comes from various sources, and is often incomplete or difficult to obtain. Missing data (NULLs) can significantly affect the computation and the analytics results on one hand, but maybe difficult or even impossible to complete on the other hand.

Example 1.1. We will use as an example the case of *retrospective medical data analysis*, where data in retrospective studies often contains missing information. Leading causes of incomplete data include: patient failure to show up for appointments, physician work-overload leading to incomplete data notation, and lack of healthcare resource availability.

Published in the Proceedings of the Workshops of the EDBT/ICDT 2026 Joint Conference (March 24-27, 2026), Tampere, Finland

✉ omera1@mail.tau.ac.il (O. Abramovich); hadaskalish@tauex.tau.ac.il (H. Stiebel-Kalish); danielde@post.tau.ac.il (D. Deutch)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Data Imputation Techniques To overcome the challenges posed by missing data and to nevertheless perform analytics, multiple approaches have been developed which generally fall under the term *data imputation* [1, 2]. There is a large body of techniques on data imputation, including listwise deletion (namely, ignoring/deleting the entire tuple), manual curation of the missing values, algorithmic and ML-based completion mechanisms. Each of these techniques is inherently imperfect: listwise deletion comes with the cost of losing information (in particular with respect to values in other attributes of the same tuple), which may be prohibitively wasteful; automatic imputation (e.g. using regression, K-nearest Neighbors, etc.) is inherently imprecise; and manual labor is costly.

Example 1.2. A particular use case of interest is that of *rare disease analysis*. By definition, such analysis needs to deal with a scarcity of data even in the absence of NULLs, and as such listwise deletion of tuples in which a particular value is absent, is highly problematic. Since there is too few data, the analysis is also highly sensitive to each individual data value, and as such errors in automated data imputation may have significant effects on the analysis results. Manual curation and case-by-case data completion through contacting the medical staff may be possible, but is labor-intensive.

Ideally, we would like a way to identify which of the cells where data is missing are most important for the analysis purposes, and focus the imputation efforts on these cells.

Attribution for Explainable Analytics In the context of Machine Learning models, and recently also in that of database query evaluation, a prominent approach is that of *attribution*. Namely, one assigns a score to each feature (in Machine Learning) or input tuple (in databases) reflecting its contribution to prediction/query results. The measures used are typically ones that have been proposed in the context of game theory, namely *Shapley* [3] and *Banzhaf* [4, 5] values. In a nutshell, these measures aggregate, in different ways, the marginal contribution of each player to every subset of players (coalition). To apply them in different contexts such as database queries and ML models, one needs to define the *players* and the *game value function*. For instance, for database queries, the tuples are players and the query result is the game value function; for ML prediction, the players are features and the prediction result is the game value function. Similar applications of these measures have been successfully employed in other contexts.

Our Approach We propose a novel approach for leveraging attribution for targeted data imputation, which we next describe at a high-level. In query evaluation, as explained above, attribution helps us identify which tuples are most important for query results. The high-level idea is that these attribution scores are potentially highly useful for data imputation: as explained above, the “budget” for high-quality data imputation is typically limited, because it typically requires extensive manual labor. Via attribution, we can identify which of the tuples in which some of the values are missing, are most useful for the analysis, and allocate resources for imputing these values.

Example 1.3. As an extreme example, some tuples (e.g. particular examination results) may be filtered out early on in the analysis and have no effect over the query result. Asking experts to impute NULLs in these tuples would be wasteful. Even if a particular tuple t is not filtered out, there may be many *alternative* tuples to t that yield the same result, in which case the contribution of t is less significant than in the absence of such alternatives.

An analysis of the above flavor focuses on database *tuples*. By contrast, data imputation typically focuses on *cells* where NULLs currently stand instead of values. When we guide cell imputation based on importance of tuples, we lose in granularity.

Example 1.4. A tuple as a whole may be important, but the query may then project out its NULL attribute. As another example, if the query involves arithmetics, then it may e.g. perform a weighted average of attributes and then use it for decision making.

This requires a fine-grained analysis that is *cell-based*. To our knowledge, cell-based attribution has not been studied in the context of query evaluation. However, it is closely related to attribution methods commonly used in machine learning, in particular feature importance techniques such as SHAP [6]. The technical challenge in shapley- (or banzhaf-) based analysis where players are features/cells is that the game value function cannot be computed directly: a model may not be run on a subset of features and a query may not be executed on a subset of cells. SHAP addresses this by defining the game value function for a subset of features as the model’s prediction when missing features are filled in using a background distribution. We explore in this paper ways to adapt this idea to the setting of cell-based explanations to query results.

Contributions and scope of this paper This short paper puts forward the novel approach of using attribution to guide data imputation; it presents a cell-based attribution model and a sampling-based algorithm to attribution computation; and it outlines a potential concrete use-case in the domain of retrospective medical data analysis with missing data.

An implementation of our proposed approach, its deployment for the particular use-case we describe and others, and experimentation are left as future work and are beyond the scope of this paper.

2. Preliminaries

We next overview necessary preliminaries on databases, queries, data imputation and attribution.

Database and Queries A *relational database* D consists of a finite set of relations $\{\mathcal{R}_1, \dots, \mathcal{R}_m\}$. Each relation \mathcal{R} has a fixed schema $\text{schema}(\mathcal{R}) = \langle A_1, \dots, A_k \rangle$. A database instance is a mapping of each relation \mathcal{R} to a finite set of tuples over its schema. A query Q maps a database instance D to a Boolean or numeric value $Q(D)$. For queries that return a relation, we instead consider the residual query Q_t , which asks whether a tuple t appears in the output.

Example 2.1. Figure 1 presents a database of medical test information for different patients. The tests are stored in two relations, namely blood tests and eye tests. In addition, the database stores the weights of four linear models in the ClassificationCoefficients relation. The figure also presents a SQL query Q that encodes the linear classifier and returns the patients who “pass” it.

Missing Values Databases often contain missing values, for instance when patients fail to provide some information or when some test result is undocumented. These missing values are captured by NULLs. Analysts may treat NULLs by applying the standard three-value logic for SQL evaluation, essentially treating the NULLs as unknown values. Alternatively, one may *impute* NULLs. An imputation function \mathcal{I} maps a database instance D with NULLs to a complete instance $\mathcal{I}(D)$ without NULLs. Common approaches include listwise deletion (i.e. deleting the tuple), completing NULLs using the mean/median/most common value occurring in this attribute, performing regression, iterative imputation approaches, imputation based on some learned distribution, and others.

Example 2.2. Consider again the database in Figure 1. Consider an imputation function \mathcal{I} that fills each numerical (categorical) cell with the average (most common) value in the corresponding attribute. Consider the tuple in relation ‘EyeTest’ annotated by e_3 , with a missing ‘Vision’ attribute. \mathcal{I} completes this value to be the mean value for the attribute, namely 0.73. Consider now the tuple in relation ‘BloodTest’ annotated by b_4 with a missing ‘Type’ attribute; \mathcal{I} may fill in this value with the most common value, ‘A’. Eventually, $\mathcal{I}(D)$ returns a database without missing values \hat{D} .

Shapley Values Given a finite set of players A , A cooperative game is defined by a characteristic function $v : \mathcal{P}(A) \rightarrow \mathbb{R}$, such that $v(\emptyset) = 0$. The value $v(S)$ represents the total gain created by the subset of players S (a “coalition”). The *Shapley value* [3] then aims to measure the fair share of each individual player $a \in A$ in the total gain $v(A)$ for the cooperative game $v(A)$.

BloodTest					EyeTest				ClassificationCoefficients(exo)					
	Patient(exo)	Type	Hem	Chol		Patient(exo)	Vision	Field		Type	C:Hem	C:Chol	C:Vision	C:Field
b_1	P1	B	1.9	1.2	e_1	P1	1	0.8	c_1	A	0.25	0.4	0.2	0.15
b_2	P2	A	0.7	1.4	e_2	P2	0.8	NULL	c_2	B	0.4	0.3	0.1	0.2
b_3	P3	A	1.2	NULL	e_3	P3	NULL	0.6	c_3	AB	0.4	0.2	0.1	0.4
b_4	P4	NULL	0.8	0.6	e_4	P4	0.4	0.6	c_4	O	0.1	0.1	0.3	0.5

Query Q (Patient Classification)

```
SELECT DISTINCT Patient FROM BloodTest JOIN EyeTest USING(Patient) JOIN ClassificationCoefficients( USING(Type)
WHERE (Hem · C:Hem + Chol · C:Chol + Vision · C:Vision + Field · C:Field) ≥ 1
```

Figure 1: Example of a database of medical tests and a query classifying patients. Each tuple is annotated with a lineage variable. Missing values are marked by NULL. Exogenous tables or attributes are marked by (exo).

The Shapley value of a player a is defined by taking the expectation of the marginal contribution of a to coalitions form as we sample a random order of players.

$$\text{Shapley}(A, v, a) \stackrel{\text{def}}{=} \sum_{B \subseteq A \setminus \{a\}} \frac{|B|! \cdot (|A| - |B| - 1)!}{|A|!} (v(B \cup \{a\}) - v(B))$$

Shapley values capture the marginal contribution of a for each set, multiplied by $|B|! \cdot (|A| - |B| - 1)!$ which is the number of permutations of A such that a comes after B and before $(A \setminus \{a\}) \setminus B$.

3. Attribution Model for cells in Query Answering

We next present a simple attribution model for cells, based on SHAP-like scores.

3.1. SHAP-like scores for cells

To quantify the contribution of individual missing values to a query result using Shapley values, we model the database cells (i.e., individual attribute values) as players in a cooperative game. The game value is defined as the output of the query Q . The challenge is that, by contrast to the classic settings where players are tuples (and then a subset of tuples is simply a sub-database and the query semantics w.r.t. it is readily defined), here it is unclear what it means to evaluate the query on a subset of cells. To this end, we define the semantics as follows. Whenever the Shapley value formula requires the query result for a specific subset of cells, we treat all cells outside this subset as NULLs. We then apply a (black-box) imputation function to fill in these missing values, and finally evaluate the query on the completed database.

Concretely, let D be a database, let \mathcal{R} be a relation with attributes $\text{Schema}(\mathcal{R}) = \{\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_m\}$. Let $t \in D$ be a specific tuple in \mathcal{R} . We denote the value of the cell corresponding to \mathcal{A}_j in t by $t.j$. We partition the database D into endogenous cells and exogenous cells denoted by D_n , and respectively D_x such that $D = D_n \cup D_x$. Endogenous cells are those for which we want to measure their importance, where Exogenous cells are treated as constants of the database. This is analogous to the concept of endogenous and exogenous tuples presented in prior work [7, 8]. We then define:

Definition 3.1 (Cell Shapley Value). Let $D = D_n \cup D_x$ be a database, Q be a query, $t.c$ be an endogenous cell in a tuple t , and \mathcal{I} be an imputation function. The Shapley value of $t.c$ w.r.t. Q, D, \mathcal{I} is:

$$\text{Shapley}(Q, D_n, D_x, \mathcal{I}, t.c) \stackrel{\text{def}}{=} \sum_{P \subseteq D_n \setminus \{t.c\}} \mathcal{C}_{|P|} \cdot [Q(\mathcal{I}(P \cup \{t.c\} \cup D_x)) - Q(\mathcal{I}(P \cup D_x))] \quad (1)$$

Where $\mathcal{C}_{|P|}$ is the Shapley coefficient. For non-deterministic imputation functions, the Shapley value is defined as the expectation of Equation 1 over the imputation distribution.

Example 3.2. Consider the database D and the residual query of Q_{p_1} of the query Q from Figure 1 and consider computing the Shapley value of the cell $b_1.Hem$. To this end, consider an imputation function that replaces each attribute with its mean, and each categorical value with its most common value. We demonstrate the computation for the subset of endogenous cells $P = \{b_1.Chol, e_1.field\}$. The imputation function completes $b_1.Type$ with the most common value ('A') and $e_1.Vision$ with the mean (0.6). The resulting query value is $Q_{p_1}(\mathcal{I}(P \cup D_x)) = 0$. When adding $b_1.Hem$ to P , we get that $Q_{p_1}(\mathcal{I}(P \cup b_1.Hem \cup D_x)) = 1$. The marginal contribution ($1 - 0 = 1$) is then weighted by the Shapley coefficient. The final value is the weighted average of these contributions across all such subsets of endogenous cells.

3.2. Importance scores for NULL-valued cells

We next extend our definition of cell-based attribution to further quantify the importance of cells with missing (i.e. NULL) values. The definition is geared towards prioritizing cells to be imputed. Intuitively, a NULL should be prioritized if its possible underlying values lead to widely divergent results. When only one NULL affects the result, we could compute statistics about the query result when its value is drawn from the distribution, compared to its value under the imputation function. However, when many NULLS exist, comparing the effect of those NULLS becomes a more complex task. To analyze such cases, we assume the existence of a joint distribution Ω governing the missing values. Such distribution could be acquired using domain knowledge, or learned based on existing database values; if we have no sufficient such information, one may use the uniform distribution and then the attribution is still informative yet is based solely on the structure of the query.

We are now ready to define the ShapVar score of NULL-valued cells as the variance of their Shapley values over the distribution Ω .

Definition 3.3 (ShapVar score). Let D be a database, Q a query, \mathcal{I} an imputation function, $t.c$ a NULL-valued cell in D , and Ω a joint distribution over all NULL-valued cells in D . We define:

$$\text{ShapVar}(t.c, Q, D) = \text{var}_{\hat{D}, t.c \sim \Omega(D)}(\text{Shapley}(Q, \hat{D}, \mathcal{I}, t.c))$$

where var is the variance, \hat{D} is a complete database instance sampled from Ω , and $\text{Shapley}(\cdot)$ is the cell-level score defined in Equation 1 when only the NULL cells are endogenous.

Example 3.4. Consider again the database and query from Fig. 1, this time with the residual query Q_{p_2} for the tuple p_2 . Assume the NULL value in cell $e_2.Field$ is distributed uniformly over $[0, 1]$, and suppose the imputation function assigns its mean value. For every possible value of $e_2.Field$, patient p_2 passes the classification test. Thus, replacing the imputed value with any realization yields zero marginal contribution, and hence $\text{ShapVar} = 0$.

Now consider the classification of patient p_3 . There are two relevant NULL values, $b_3.Chol$ and $e_3.Vision$. Assume their distributions are independent and uniform in the ranges $[0.5, 2]$ and $[0.5, 1]$ respectively. We can numerically compute the ShapVar score by sampling from this distribution and computing the Shapley values. We get that the ShapVar score of $b_3.Chol$ and $e_3.Vision$ are 0.18 and 0.02, respectively. Intuitively, this is because $b_3.Chol$ ranges over broader values and has a higher coefficient in the classification test.

3.3. Computational Aspects

Computing the cell Shapley value poses significant computational challenges. Naively, its definition requires considering exponentially many subsets of cells, and evaluating the imputation function and query on each such subset. Even when query evaluation and imputation are efficient, this combinatorial explosion renders exact computation infeasible in general. Nevertheless, prior work on Shapley value computation in databases and ML suggests that computation may be tractable in many practical settings. For database tuples, techniques based on lineage, factorization, and knowledge compilation have been

used to avoid explicit enumeration of subsets [9, 10, 11]. For SHAP scores, sampling and model based approximations were shown to work well in practice [12, 6, 13]. These complementary lines of work indicate that similar strategies could be adapted to the cell setting. In particular:

Proposition 3.5. Let Q be a query and \mathcal{I} an imputation function, both computable in polynomial time with respect to the size of the database D . Then, the cell Shapley value $\text{Shapley}(Q, D, \mathcal{I}, t.c)$ admits a Fully Polynomial-time Randomized Approximation Scheme.

We leave for future work the development of approximation schemes for ShapVar and the identification of tractable fragments for exact computation.

4. Use Case: Rare Disease Treatment Planning

More than ten thousand conditions fit the definition of a rare disease, affecting approximately 300-400 million people worldwide. Due to small sample sizes and logistic constraints, perfect data or randomized controlled trials are unlikely to be performed for most of these conditions; thus, actionable information may be most likely to be obtained from existing retrospective data and meticulous analysis of treatments across patients [14]. However, data collected from retrospective studies often contains missing information, due to e.g. patient's failure to show up for appointments, physician work-overload leading to incomplete data notation, and lack of resource availability. These imperfections, compounded with the fact that the disease is rare, lead to a paucity of quality data. In turn, missing data poses significant challenges that complicate data analysis and the derivation of conclusions such as treatment guidelines.

Existing Data Imputation techniques do not adequately address the challenges posed by missing data in this context. Automated data imputation is based on statistical estimation of the missing values and is inherently imprecise. Coupled with the small amount of data and the high stakes associated with the analysis, statistical errors are especially harmful. Manual curation, on the other hand, is highly laborious, especially in the context of retrospective studies, where reaching out to the patients for questions regarding the missing data may be costly or even impossible.

A particular medical question of interest is that of deciding the optimal treatment for optic neuritis (ON), a rare, yet potentially blinding, inflammatory disorder of the optic nerve. ON most commonly presents in the context of multiple sclerosis (MS), Neuromyelitis Optica Spectrum Disorder (NMOSD), and Myelin Oligodendrocyte Glycoprotein Antibody-Associated Disease (MOGAD). Most MS-ON cases respond well to intravenous methylprednisolone (IVMP) with excellent visual outcomes, whereas NMOSD-ON and MOGAD-ON are frequently associated with severe, permanent visual loss [15, 16, 17]. The common treatment algorithm for ON is based on retrospective case series, often containing missing data. Such missing data lead to treatment-algorithm distortions by preventing the correct prediction of which MS-ON will not have a good visual outcome, which NMOSD-ON patient is rapidly responding to steroids and may not require immunoadsorption, etc.

This particular use case is expected to benefit significantly from our approach of importance-guided data imputation, which allocates costly imputation efforts only to data cells that materially affect the analysis. As a simple example, if a particular patient is excluded from the analysis, then naturally there is no need to contact the clinic to complete missing information regarding their treatment. Importantly, different patients included in the analysis differ in their level of influence on the analysis result; for example, a patient with MS-ON with missing data will affect the analysis less than a patient with NMOSD with missing data, because NMOSD-Optic neuritis is a very rare condition, while MS-ON is more common with a wealth of data for this disorder.

We intend to deploy our solutions using a retrospective clinical database of 70 patients treated for acute demyelinating ON with documented 3–6 month visual outcomes. To assess factors affecting final visual outcome, we will collect: (1) visual acuity at nadir; (2) time from vision loss to IVMP initiation; (3) visual acuity at 5 ± 2 days; (4) visual acuity at 30 ± 7 days; (5) visual acuity at 3–6 months (main outcome); (6) antibody status at 2 weeks, when escalation therapy decisions are made. Additional

covariates include age, sex, and escalation therapy use. The database contains missing values, such as visual acuity at critical time points and delayed antibody test results. We expect the study to (1) provide insight into the quality of our measures and algorithms and (2) support improved approaches to ON treatment. Implementation and experimentation are ongoing.

5. Conclusions

We have presented a novel approach for guiding data imputation using attribution in query answering. While attribution has been used for explanations, we demonstrate a novel application of it, for imputation. In this context, we have identified a need to extend attribution solutions to account for the granularity of cells rather than tuples, and have introduced a novel attribution model to this effect. We have then presented a concrete use case for data imputation in the context of medical data analysis where some data is missing. This is ongoing work and we plan to implement algorithms for cell-based attribution as well as a solution that leverages these algorithms for guiding data imputation. We further intend to deploy these solutions in the context of medical data analysis in line with the outlined use-case.

Acknowledgments

This study was supported by the Clair and Amedee Maratier Institute for the study of Blindness and Visual Disorders, Faculty of Medical & Health Sciences, Tel-Aviv University; the Israeli Science Foundation (Grant 1476/24); the Len Blavatnik and the Blavatnik Family foundation; and the Deutsch foundation.

References

- [1] D. B. Rubin, Inference and missing data, *Biometrika* 63 (1976) 581–592.
- [2] J. Yoon, J. Jordon, M. Schaar, Gain: Missing data imputation using generative adversarial nets, in: *International conference on machine learning*, PMLR, 2018, pp. 5689–5698.
- [3] L. S. Shapley, A value for n -person games, *Contributions to the Theory of Games* 2 (1953) 307–317. URL: <http://www.library.fu.ru/files/Roth2.pdf#page=39>.
- [4] L. S. Penrose, The elementary statistics of majority voting, *J. Royal Stats. Soc.* 109 (1946) 53–57. URL: <http://www.jstor.org/stable/2981392>.
- [5] J. F. Banzhaf III, Weighted voting doesn't work: A mathematical analysis, *Rutgers Law Review* 19 (1965) 317–343. URL: <https://heinonline.org/HOL/LandingPage?handle=hein.journals/rutlr19&div=19&id=&page=>.
- [6] S. M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, *Advances in neural information processing systems* 30 (2017).
- [7] E. Livshits, L. Bertossi, B. Kimelfeld, M. Sebag, The shapley value of tuples in query answering, *Logical Methods in Computer Science* 17 (2021).
- [8] O. Abramovich, D. Deutch, N. Frost, A. Kara, D. Olteanu, Banzhaf values for facts in query answering, *Proc. ACM Manag. Data* 2 (2024). URL: <https://doi.org/10.1145/3654926>. doi:10.1145/3654926.
- [9] D. Deutch, N. Frost, B. Kimelfeld, M. Monet, Computing the shapley value of facts in query answering, in: *Proceedings of the 2022 International Conference on Management of Data*, 2022, pp. 1570–1583.
- [10] O. Abramovich, D. Deutch, N. Frost, A. Kara, D. Olteanu, Banzhaf values for facts in query answering, *Proceedings of the ACM on Management of Data* 2 (2024) 1–26.
- [11] O. Abramovich, D. Deutch, N. Frost, A. Kara, D. Olteanu, Advancing fact attribution for query answering: Aggregate queries and novel algorithms, *Proc. VLDB Endow.* 18 (2025) 3996–4008. URL: <https://doi.org/10.14778/3749646.3749670>. doi:10.14778/3749646.3749670.

- [12] E. Štrumbelj, I. Kononenko, Explaining prediction models and individual predictions with feature contributions, *Knowledge and information systems* 41 (2014) 647–665.
- [13] R. Okhrati, A. Lipani, A multilinear sampling algorithm to estimate shapley values, in: *2020 25th International Conference on Pattern Recognition (ICPR)*, IEEE, 2021, pp. 7992–7999.
- [14] T. R. Frieden, Evidence for health decision making—beyond randomized, controlled trials, *New England Journal of Medicine* 377 (2017) 465–475.
- [15] N. Moheb, J. J. Chen, The neuro-ophthalmological manifestations of nmosd and mogad—a comprehensive review, *Eye* 37 (2023) 2391–2398.
- [16] J. J. Chen, E. P. Flanagan, S. J. Pittock, N. C. Stern, N. Tisavipat, M. T. Bhatti, K. D. Chodnicki, D. A. Tajfirouz, S. Jamali, A. Kunchok, et al., Visual outcomes following plasma exchange for optic neuritis: an international multicenter retrospective analysis of 395 optic neuritis attacks, *American journal of ophthalmology* 252 (2023) 213–224.
- [17] J. S. Graves, F. C. Oertel, A. Van der Walt, S. Collorone, E. S. Sotirchos, G. Pihl-Jensen, P. Albrecht, E. A. Yeh, S. Saidha, J. Frederiksen, et al., Leveraging visual outcome measures to advance therapy development in neuroimmunologic disorders, *Neurology: Neuroimmunology & Neuroinflammation* 9 (2021) e1126.