
Knowledge-based Graphical Method for Safety Signal Detection in Clinical Trials

Francois Vandenhende¹ Anna Georgiou¹ Michalis Georgiou¹
Theodoros Psaras¹ Ellie Karekla¹ Elena Hadjicosta¹

¹ClinBAY Limited, Limassol, Cyprus
Correspondence: francois@clinbay.com
<https://app.clinbay.com/safeterm>

Abstract

We present a graphical, knowledge-based method for reviewing treatment-emergent adverse events (AEs) in clinical trials. The approach enhances MedDRA by adding a hidden medical knowledge layer (Safeterm) that captures semantic relationships between terms in a 2-D map. Using this layer, AE Preferred Terms can be re-grouped automatically into similarity clusters, and their association to the trial disease may be quantified. The Safeterm map is available online and connected to aggregated AE incidence tables from ClinicalTrials.gov. For signal detection, we compute treatment-specific disproportionality metrics using shrinkage incidence ratios. Cluster-level EBGM values are then derived through precision-weighted aggregation. Two visual outputs support interpretation: a semantic map showing AE incidence and an expectedness-versus-disproportionality plot for rapid signal detection. Applied to three legacy trials, the automated method clearly recovers all expected safety signals. Overall, augmenting MedDRA with a medical knowledge layer improves clarity, efficiency, and accuracy in AE interpretation for clinical trials.

Keywords: Safety Review, Semantic Similarity, Graphical Methods, Signal Detection, Clinical Trials, MedDRA

1 Introduction

The evaluation of treatment-emergent adverse events (AEs) is a central component of safety analysis in clinical trials. Traditional approaches rely heavily on tabular summaries of MedDRA-coded events, supplemented by selected graphical displays such as volcano plots or forest plots [1, 2]. While these methods are widely adopted [3, 4], they often provide limited support for uncovering latent structure in AE data or for identifying nuanced safety patterns across related medical concepts [5]. As a result, safety reviewers may miss important semantic relationships between Preferred Terms (PTs), particularly when large and heterogeneous AE profiles are involved.

MedDRA itself provides a rich and hierarchical vocabulary for AE coding, but it does not explicitly encode clinical knowledge beyond its predefined parent-child taxonomy [6, 7]. Common aggregation strategies, such as using Standardised MedDRA Queries (SMQs) [8, 9], FDA Office of Custom Medical Queries [10], SOC-level or hybrid [11] summaries, offer broad grouping, but do not fully reflect latent semantic or mechanistic relationships specific to a drug, indication, or clinical trial. Consequently, analyses that treat MedDRA strictly as a hierarchical dictionary or using predefined groupings of terms may fail to capture deeper medical connections that could enhance AE interpretation and signal identification.

In response to these limitations, the intersection of machine learning, natural language processing, and Bayesian inference has introduced powerful methods to enrich AE evaluation. Embedding-based models and semantic similarity measures are gaining traction in pharmacovigilance for their ability to represent conceptual relationships beyond rigid hierarchies [12], combine terms in relevant groups [13] or improve disproportionality assessment in PV using information sharing [14].

These emerging approaches are promising but have seen limited application in clinical trial settings, where AE data are often sparse and more structured than post-marketing spontaneous reports. Disproportionality methods such as Empirical Bayes Geometric Mean (EBGM) or Bayesian hierarchical models remain standard in signal detection [12, 15], yet typically ignore semantic relationships among PTs. Similarly, statistical tools for safety review still largely rely on hierarchical groupings or co-occurrence statistics [16].

In this work, we introduce a knowledge-based graphical framework that augments MedDRA with an additional, hidden medical knowledge layer. This layer, encoded in the *Safeterm* system [17], provides a set of semantic descriptors and inter-term relationships that reflect clinical reasoning and knowledge not captured in MedDRA’s hierarchy. Our method combines high-dimensional semantic embeddings of MedDRA PTs, with a set of downstream applications tailored for clinical trial safety review: semantic clustering, Bayesian shrinkage disproportionality metrics at both the PT and cluster levels, and indication-specific expectedness scores by comparing PT embeddings to the semantic representation of the patient population.

To ease trial safety review, we present two intuitive visualizations: a two-dimensional semantic map of AEs and an expectedness-versus-disproportionality plot.

The addition of a hidden knowledge layer to MedDRA offers a context-aware, interpretable, and scalable enhancement to conventional safety signal detection. It provides an effective way to identify clinically meaningful safety patterns and reduce the manual burden of AE review in trials.

2 Methods

2.1 Hidden Medical Knowledge Layer

The hidden medical knowledge layer, implemented in the *Safeterm* system, encodes MedDRA terminology into high-dimensional numerical vectors, or embeddings. These embeddings capture semantic, clinical, and contextual relationships among Preferred Terms (PTs). PTs with similar meanings or shared mechanistic characteristics occupy nearby positions in this space, enabling multivariate analyses to explore term dependencies and latent structures that are not explicitly encoded in MedDRA’s hierarchy.

While this high-dimensional representation is useful for computational analyses, it is not directly interpretable. Therefore, a set of decoding tools was developed to project signals back onto the MedDRA domain, supporting human interpretability.

2.2 Decoding Toolbox for PT Understanding

The decoding toolbox currently consists of three key components: PT clustering, two-dimensional interpretability projection, and expectedness scoring. These tools are applied to the hidden layer to facilitate understanding, visualization, and prioritization of PTs. Additional tools will be added in future releases of the *Safeterm* app.

2.2.1 Semantic Clustering of PTs

The embeddings of observed PTs in each trial were first projected into a reduced-dimensional space using Principal Component Analysis (PCA) [18]. This step preserves the most informative variance in the embedding space while improving cluster separability, computational efficiency, and identifiability of downstream similarity structures. Next, we applied a clustering algorithm to the PCA-reduced embeddings in order to identify groups of semantically related PTs and extract them from the random noise.

Cluster naming was performed using an AI-based decoding method that assigns interpretable labels to clusters based on their content centroid in the hidden medical knowledge space. Terms that did not group into any cluster were flagged in our reports for transparency.

2.2.2 Safeterm Map: Two-Dimensional Interpretability Projection

To enable interactive exploration, the high-dimensional embeddings were projected into a two-dimensional space for visualization. Each PT is represented as a point, and proximity between points reflects semantic similarity. Incidence data from the clinical trial were overlaid on the map by treatment arm, highlighting clustering patterns and enabling the identification of treatment-specific mapping differences. The Safeterm map facilitates intuitive exploration of AE relationships and provides a context-aware view of trial safety data.

2.2.3 Indication-Specific Expectedness Score

Patient population descriptors were extracted from the clinical protocol and mapped into the hidden space using the Safeterm encoder. Cosine similarity scores between the patient population descriptors and PT embeddings were computed, providing an expectedness score for each PT. Cosine similarity is widely used in natural language processing and biomedical embeddings [19, 13]. Higher scores indicate that a PT is more likely to be observed given the trial population. This measure allows safety reviewers to contextualize AE occurrences relative to the expected background profile of the trial population.

2.3 Statistical methods

The method uses aggregated incidence tables of subjects with adverse events, by Preferred Term (PT) and treatment group, consistent with standard clinical trial AE reporting practice. To identify potential safety signals beyond simple incidence counts, we computed the **Empirical Bayes Geometric Mean (EBGM)** statistics, which provide a robust disproportionality measure adjusted for sparse event counts [20].

For each PT, let n_i denote the observed count of subjects with events and N_i the total number of subjects at risk under treatment i ($i = 1, \dots, T$). The *expected count of events* E_i for treatment i is calculated as the product of the treatment sample size by the overall incidence of the PT across all treatment arms:

$$E_i = N_i \cdot \frac{\sum_i n_i}{\sum_i N_i}.$$

This choice of denominator differs from classical pharmacovigilance methods, which often use an external null or historical incidence [15]. Using the trial-specific overall incidence allows for an *internally consistent background* for relative comparisons between treatment arms, simplifies automated processing, and introduces only minor loss for exploratory signal detection.

The EBGM was then computed with prior stabilization as:

$$EBGM_i = \frac{n_i + \alpha}{E_i + \beta},$$

where α and β are prior parameters that reduce variance for low-count PTs. The posterior distribution of EBGM can then be obtained using conjugated gamma priors [20].

2.3.1 Cluster-level EBGM

To integrate semantic information from the hidden Safeterm space, EBGM values were aggregated at the cluster level. Because patient-level information was not available, a meta-analytic weighting strategy was applied.

For each Preferred Term (PT) j within a cluster, the sampling variance was approximated as:

$$\text{Var}_{\text{EBGM},j} = \frac{\text{EBGM}_j^2}{n_j + \alpha}, \quad (1)$$

where n_j is the number of subjects reporting the PT and α is the prior parameter used in the EBGM shrinkage model. The precision was then defined as:

$$w_j = \frac{1}{\text{Var}_{\text{EBGM},j}}. \quad (2)$$

The cluster-level EBGM for each treatment group was computed as a precision-weighted average:

$$\text{EBGM}_{\text{cluster}} = \frac{\sum_j w_j \cdot \text{EBGM}_j}{\sum_j w_j}. \quad (3)$$

This approach stabilizes estimates for rare events while preserving semantic coherence among PTs within a cluster [13]. Although future extensions could quantify uncertainty at the cluster level (e.g., via Bayesian hierarchical modeling), this study focuses on point estimates only, reflecting its exploratory safety-review objective. Overall, the method provides a practical and interpretable summary measure for cluster-level disproportionality.

2.3.2 Expectedness versus Disproportionality Plot

To contextualize signals relative to trial-specific expectations, we developed the Expectedness versus Disproportionality (EVD) plot. In this plot:

- The x-axis represents the *expectedness score* of each PT, determining how the PT relates to the studied indication, or patient descriptors.
- The y-axis represents the *EBGM value*, reflecting how the AE incidence compares in a treatment arm versus in the entire study population.

PTs are grouped according to Safeterm clusters. This visualization provides a synthetic, interpretable overview of major safety patterns, highlighting PTs and clusters that are both unexpected and disproportionately observed in a treatment arm.

2.4 Performance Evaluation

The performance of the proposed method was explored using three clinical trials with publicly available AE data from ClinicalTrials.gov. For each study, the incidence of subjects with AEs was computed as the sum of counts for serious and other (non-serious) events. Only AEs coded as valid MedDRA PTs were included.

The selected trials were:

1. **NCT05096221** – Phase 3, randomized, double-blind, placebo-controlled study evaluating SRP-9001 in Duchenne Muscular Dystrophy (DMD) [21].
2. **NCT02348593** – Twelve-week, double-blind, placebo-controlled, randomized study of JZP-110 in narcolepsy [22].
3. **NCT05008224** – Phase 2 sequential oncology study of pembrolizumab and chemotherapy in classical Hodgkin lymphoma (KEYNOTE-C11) [23].

These trials were chosen to illustrate the method across different therapeutic areas, patient populations, and AE profiles. NCT05096221 involves a gene therapy in DMD with known transient liver AEs; NCT02348593 is a dose-dependent stimulant study in narcolepsy; and NCT05008224 is a sequential oncology trial with non-randomized chemotherapy groups, decided based on PET scan outcome. It illustrates treatment-safety confounding by disease severity.




Figure 1: Safeterm map of incidence of subjects with AEs by PT, treatment and cluster for the DMD trial.

3 Results

3.1 Duchenne Muscular Dystrophy Trial

This phase III trial reported a total of 72 distinct Preferred Terms (PTs), of which 44 were clustered into 7 Safeterm groups. Three main groups are detailed in Table 1, combining PTs from multiple SOC.

A notable cluster labeled *Liver damage* was identified, encompassing PTs such as *Elevated liver function tests*, *Hepatic enzyme increased*, and *Liver injury*. This cluster captured the spectrum of liver-related adverse events (AEs) observed in the trial, facilitating a more comprehensive safety analysis.

Figure 1 presents the Safeterm 2-D map for the reported PTs. Each dot represents a PT, with its size proportional to the proportion of subjects affected. PTs are colored by Safeterm cluster, with ungrouped terms displayed in brown. Clusters correspond to PTs that occupy similar coordinates in the projected space. For example, the *Liver damage* cluster is located in the central-left quadrant of the map. The map shows more and larger dots for liver-related AEs in the active treatment groups compared to placebo. This visual pattern suggests a treatment-related increase in liver AEs.

Figure 2 shows the Expectedness versus Disproportionality (EVD) plot for the clustered PTs. The x-axis represents the expectedness score for Duchenne muscular dystrophy, while the y-axis shows the Empirical Bayes Geometric Mean (EBGM) reflecting disproportionality relative to the overall trial population. PTs associated with musculoskeletal trauma have the highest expectedness, whereas terms from acute respiratory illness (e.g., *Nasopharyngitis*, *Epistaxis*) have the lowest. Each dot is a PT, with size proportional to incidence.

Focusing on the *Liver damage* cluster, all terms show a consistent EBGM greater than 1 for active treatment versus placebo. PTs such as *GGT increased* and *Blood bilirubin increased* demonstrate the highest disproportionality, with EBGM point estimates above 2, indicating a strong association with treatment. This observation is corroborated by the cluster-level EBGM analysis (see Table 2), which provides a comparative treatment overview across all PTs forming relevant clusters.

3.2 Dose-Response Narcolepsy Trial

The second study illustrates the utility of our method in detecting dose-dependent pharmacological effects.

A total of 18 distinct Preferred Terms (PTs) were reported in this trial, of which 15 were clustered into 2 categories: *Stress response* and *Respiratory infection*.




Figure 2: EVD plot of incidence of subjects with AEs by treatment for clustered PTs (DMD trial).




Figure 3: EVD plot of incidence of subjects with AEs by PT, treatment and cluster for the narcolepsy trial.

The *Stress response* cluster included 12 PTs such as *Increased heart rate*, *Anxiety*, *Nervousness*, and *Insomnia*, reflecting heightened sympathetic activity consistent with the pharmacological action of JZP-110 in narcolepsy.

This cluster shows the strongest connection to the disease and treatment mechanism (see Figure 3). A clear monotonic dose-response relationship is observed, with the incidence of stress-related AEs increasing with higher JZP-110 doses (see Table 3). Most stress-related events were non-serious, with the exception of a single episode of *Anxiety* at the 150 mg dose.

These results demonstrate that our augmented MedDRA analyses can identify PT patterns beyond standard safety, aligned with known pharmacological effects of drugs, supporting both safety and mechanistic interpretations in clinical trials.

3.3 Hodgkin's Lymphoma Trial

In this sequential oncology trial, a total of 104 distinct Preferred Terms (PTs) were reported, which clustered into 12 identifiable semantic groups (see Figure 4).




Figure 4: Safeterm map of incidence of subjects with AEs by PT, treatment and cluster for the lymphoma trial




Figure 5: EVD plot of incidence of subjects with AEs by PT, treatment and cluster for the lymphoma trial

As shown in Figure 5, the *Bone marrow failure* cluster was most closely associated with the disease and its treatment. Patients receiving escBEACOPP, selected based on unfavorable PET scan outcomes, exhibited a higher incidence of bone marrow-related adverse events compared to those receiving AVD chemotherapy or Pembrolizumab monotherapy.

Despite the trial's non-randomized design and unbalanced treatment group sizes, our method provided a rapid and informative overview of safety signals contextualized to the disease. The tri-variate association among patient characteristics, treatment, and adverse events enabled clearer interpretation of safety signals in this complex study setting.

4 Discussion

This study demonstrates the potential of augmenting MedDRA terminology with a hidden medical knowledge layer to enhance safety and pharmacology signal detection in clinical trials. We applied the approach to three distinct trials, each illustrating different challenges in safety review: a gene therapy trial in Duchenne Muscular Dystrophy (DMD) with known liver-related adverse events (AEs); a dose-response study in narcolepsy where a mechanistic effect was detected; and a sequential oncology trial where treatment and disease population are confounded.

In all cases, the combination of automated term grouping, AE-disease expectedness evaluation, and two graphical displays (Safeterm Map and EVD plot) proved highly effective in identifying relevant safety signals. The method provided a rapid and intuitive overview of AE patterns, demonstrating clear advantages over traditional inspection of AE tables.

While these results are promising, further work is needed to systematically quantify performance improvements across a broader range of clinical trials. Future research could incorporate additional AE attributes, such as severity, duration, and onset time. Although the current method is exploratory, further development is needed to enable causal inference, either at the study level or across trials.

Individual patient-level data could also be analyzed, including medical history, concomitant medications, and AEs. These could be mapped onto the medical knowledge space to better understand their relationships and reveal drug-disease-event interactions.

5 Conclusion

Augmenting MedDRA with a hidden medical knowledge layer enhances the analysis of adverse events in clinical trials. Semantic clustering, 2-D interpretability maps, and Expectedness versus Disproportionality plots allow reviewers to identify clinically meaningful AE patterns efficiently and in context.

Applied to diverse trials, including gene therapy, dose-response, and oncology studies, the method facilitated rapid detection of treatment-specific safety signals and highlighted relevant pharmacological effects. The Safeterm online app, connected to clinicaltrials.gov, provides an accessible, interactive platform for exploring trial AE data, supporting more informed and streamlined safety review.

Statements and Declarations

Funding: No funding was provided for this study.

Conflicts of interest: The authors are affiliated with ClinBAY Ltd; they declare no financial or commercial conflict beyond this.

Author contributions: All authors contributed equally to concept, methodology, data analysis, and writing.

Ethics approval: Not applicable; retrospective analysis of existing data.

Data availability:

<https://clinicaltrials.gov/study/NCT05096221>
<https://clinicaltrials.gov/study/NCT02348593>
<https://clinicaltrials.gov/study/NCT05008224>

References

- [1] Richard C. Zink, Russell D. Wolfinger, and Geoffrey Mann. Summarizing the incidence of adverse events using volcano plots and time intervals. *Clinical Trials*, 10(3):398–406, 2013. doi: 10.1177/1740774513485311.
- [2] Victoria Cornelius, Suzie Cro, and Rachel Phillips. Advantages of visualisations to evaluate and communicate adverse event information in randomised controlled trials. *Trials*, 21(1):1028, 2020. doi: 10.1186/s13063-020-04903-0.
- [3] International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH). ICH E2A: Clinical Safety Data Management: Definitions and Standards for Expedited Reporting. Guideline, ICH, 1994. URL <https://www.ich.org/page/efficacy-guidelines>. Accessed 2025-11-19.
- [4] International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use (ICH). ICH E9: Statistical Principles for Clinical Trials. Guideline, ICH, 1998. URL <https://www.ich.org/page/efficacy-guidelines>. Accessed 2025-11-19.
- [5] Riaz Qureshi, Xiwei Chen, Carsten Goerg, Evan Mayo-Wilson, Stephanie Dickinson, Lilian Golzarri-Arroyo, Hwanhee Hong, Rachel Phillips, Victoria Cornelius, Mara McAdams De-Marco, Eliseo Guallar, and Tianjing Li. Comparing the value of data visualization methods for communicating harms in clinical trials. *Epidemiologic Reviews*, 44(1):55–66, 2022. doi: 10.1093/epirev/mxac005.
- [6] E. G. Brown, L. Wood, and S. Wood. The medical dictionary for regulatory activities (meddra). *Drug Safety*, 20(2):109–117, 1999. doi: 10.2165/00002018-199920020-00002.
- [7] Patricia Mozzicato. MedDRA: An Overview of the Medical Dictionary for Regulatory Activities. *Pharmaceutical Medicine*, 23(2):65–75, 2009. doi: 10.1007/BF03256752. URL <https://doi.org/10.1007/BF03256752>.
- [8] Patricia Mozzicato. Standardised meddra queries: Their role in signal detection. *Drug Safety*, 30(7):617–619, 2007. doi: 10.2165/00002018-200730070-00009. URL <https://www.ovid.com/journals/drusa/fulltext/00002018-200730070-00009>.
- [9] Lin Chau Chang, Riaz Mahmood, Samina Qureshi, and Christopher D. Breder. Patterns of use and impact of standardised meddra query analyses on the safety evaluation and review of new drug and biologics license applications. *PLOS ONE*, 12(6):e0178104, 2017. doi: 10.1371/journal.pone.0178104. URL <https://doi.org/10.1371/journal.pone.0178104>.
- [10] Scott Proestel, Vaishali Popat, Ellis F. Unger, and Linda J. B. Jeng. The development and use of office of new drugs custom medical queries for safety analyses of clinical trial data. *Drug Safety*, 48(9):1331–1337, 2025. doi: 10.1007/s40264-025-01582-1. URL <https://doi.org/10.1007/s40264-025-01582-1>.
- [11] Marie Dupuch, Laëtitia Dupuch, Amandine Perinet, Thierry Hamon, and Natalia Grabar. Grouping the pharmacovigilance terms with a hybrid approach. *Studies in Health Technology and Informatics*, 180:235–239, 2012. URL <https://pubmed.ncbi.nlm.nih.gov/22874187>. Combines semantic similarity and terminology structuring; references SMQs explicitly.
- [12] William DuMouchel. Bayesian data mining in large frequency tables, with an application to the fda spontaneous reporting system. *The American Statistician*, 53(3):177–190, 1999. doi: <https://doi.org/10.1080/00031305.1999.10474456>.
- [13] Jeffery L. Painter, François Haguinet, Gregory E. Powell, and Andrew Bate. Ontology-based semantic similarity measures for clustering medical concepts in drug safety. 2025. Preprint from GlaxoSmithKline.
- [14] François Haguinet, Jeffery L. Painter, Gregory E. Powell, Andrea Callegaro, and Andrew Bate. Bayesian dynamic borrowing considering semantic similarity between outcomes for disproportionality analysis. 2025. Preprint from GlaxoSmithKline (IC SSM approach).

- [15] G. Niklas Norén, Rolf Sundberg, Andrew Bate, and I. Ralph Edwards. A statistical methodology for drug–drug interaction surveillance. *Statistics in Medicine*, 27(16):3057–3070, 2008. doi: 10.1002/sim.3247. Erratum published in *Statistics in Medicine*, 2008; 27(29):6271–6272.
- [16] Laetitia de Abreu Nunes, Richard Hooper, Patricia McGettigan, and Rachel Phillips. Statistical methods leveraging the hierarchical structure of adverse events for signal detection in clinical trials: a scoping review of the methodological literature. *BMC Medical Research Methodology*, 24(1):253, 2024. doi: 10.1186/s12874-024-02369-1.
- [17] ClinBay Ltd. Safeterm: Interactive Safety Signal Review Tool. <https://app.clinbay.com/safeterm>, 2025. Accessed 2025-11-20.
- [18] I. T. Jolliffe and J. Cadima. *Principal Component Analysis*, volume 374, 2016. doi: 10.1098/rsta.2015.0202. Comprehensive review of PCA, including applications for dimensionality reduction of high-dimensional embeddings.
- [19] Daniel Sanchez and Montse Batet. Semantic similarity estimation in the biomedical domain: an ontology-based information-theoretic perspective. *J Biomed Inform*, 44(5):749–759, 2011. doi: 10.1016/j.jbi.2011.03.013. URL <https://www.ncbi.nlm.nih.gov/pubmed/21463704>.
- [20] Travis Canida and John Ihrie. openebgm: An r implementation of the gamma-poisson shrinker data mining model. *The R Journal*, 9:499–519, 2017. ISSN 2073-4859. doi: 10.32614/RJ-2017-063. <https://doi.org/10.32614/RJ-2017-063>.
- [21] Jerry R. Mendell, Francesco Muntoni, Craig M. McDonald, Eugenio M. Mercuri, Emma Ciafaloni, Hirofumi Komaki, Carmen Leon-Astudillo, Andrés Nascimento, Crystal Proud, Ulrike Schara-Schmidt, Aravindhan Veerapandiyar, Craig M. Zaidman, Maitea Guridi, Alexander P. Murphy, Carol Reid, Christoph Wandel, Damon R. Asher, Eddie Darton, Stefanie Mason, Rachael A. Potter, Teji Singh, Wenfei Zhang, Paulo Fontoura, Jacob S. Elkins, and Louise R. Rodino-Klapac. AAV gene therapy for Duchenne muscular dystrophy: the EMBARK phase 3 randomized trial. *Nature Medicine*, 31(1):332–341, 2025. doi: 10.1038/s41591-024-03304-z. ClinicalTrials.gov Identifier: NCT05096221.
- [22] Russell Rosenberg, Michael J. Thorpy, Yves Dauvilliers, Paula K. Schweitzer, Gary Zammit, Mark Gotfried, Shay Bujanover, Brian Scheckner, and Atul Malhotra. Incidence and duration of common early-onset adverse events in randomized controlled trials of solriamfetol for treatment of excessive daytime sleepiness in obstructive sleep apnea and narcolepsy. *Journal of Clinical Sleep Medicine*, 18(1):235–244, 2022. doi: 10.5664/jcsm.9550.
- [23] Ranjana Advani, Abraham Avigdor, Anna Balari, David Lavie, Stefan Hohaus, Jan Zaucha, Vu Hua, Vittorio Zilioli, Raimundo Gazitua, Muhit Ozcan, Amos Odeleye-Ajakaye, Nishitha Reddy, Patricia Marinello, and Jane Winter. Efficacy and safety of pembrolizumab and chemotherapy in newly-diagnosed, early unfavorable or advanced classic hodgkin lymphoma: The phase 2 keynote-c11 study. *Blood*, 142:3068–3068, 11 2023. doi: 10.1182/blood-2023-181171.

Table 1: List of MedDRA PTs in Selected Safeterm Groupings for Duchene Muscular Dystrophy Trial

SafeTerm Group	SOC	Term
Abdominal pathology	Infections and infestations	Appendicitis Anal abscess
	Gastrointestinal disorders	Abdominal pain upper Abdominal pain
Gastrointestinal disturbance	Infections and infestations	Rotavirus infection Gastroenteritis viral Gastroenteritis
		Vomiting Nausea Diarrhoea Constipation
	Gastrointestinal disorders	Transaminases increased Hepatic enzyme increased Glutamate dehydrogenase increased Gamma-glutamyltransferase increased Blood bilirubin increased Liver injury Hepatotoxicity
Liver damage	Investigations	

Table 2: EBGM at the PT Group Level for Duchene Muscular Dystrophy Trial

Safeterm PT Group	Part 1: Delandistrogene Moxeparvovec	Part 1: Placebo	Part 2: Delandistrogene Moxeparvovec	Part 2: Placebo
Liver damage	1.30	0.29	1.59	0.34
Gastrointestinal disturbance	1.15	0.60	1.41	0.33
Abdominal pathology	1.16	1.18	1.29	0.21
Acute respiratory illness	1.03	1.15	0.85	0.71
Behavioral disturbance	0.91	0.87	0.86	0.68
Trauma to musculoskeletal system	0.89	0.99	1.07	0.73
Cardiac muscle damage	0.78	0.92	0.69	0.98

Table 3: Incidence of Patients with AEs by PT and Treatment For Stress Response Cluster in Narcolepsy Trial

Term	150 mg JZP-110	300 mg JZP-110	75 mg JZP-110	Placebo
Insomnia	0/59 (0.0%)	3/59 (5.08%)	2/59 (3.39%)	0/59 (0.0%)
Anxiety	4/59 (6.78%)	5/59 (8.47%)	1/59 (1.69%)	1/59 (1.69%)
Headache	14/59 (23.73%)	18/59 (30.51%)	6/59 (10.17%)	3/59 (5.08%)
Dizziness	1/59 (1.69%)	3/59 (5.08%)	2/59 (3.39%)	2/59 (3.39%)
Decreased appetite	5/59 (8.47%)	9/59 (15.25%)	5/59 (8.47%)	1/59 (1.69%)
Weight increased	0/59 (0.0%)	1/59 (1.69%)	2/59 (3.39%)	3/59 (5.08%)
Weight decreased	1/59 (1.69%)	3/59 (5.08%)	1/59 (1.69%)	0/59 (0.0%)
Heart rate increased	0/59 (0.0%)	4/59 (6.78%)	0/59 (0.0%)	0/59 (0.0%)
Fatigue	2/59 (3.39%)	3/59 (5.08%)	0/59 (0.0%)	0/59 (0.0%)
Nausea	6/59 (10.17%)	6/59 (10.17%)	3/59 (5.08%)	1/59 (1.69%)
Dyspepsia	2/59 (3.39%)	3/59 (5.08%)	1/59 (1.69%)	0/59 (0.0%)
Constipation	1/59 (1.69%)	0/59 (0.0%)	3/59 (5.08%)	1/59 (1.69%)