International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

# From Case Intake to Regulatory Reporting: an Integrated Review of Individual Case Safety Report (ICSR) Management in Pharmacovigilance

**Hari M[1], Thenmozhi R[2]**

[1]Student, Department of Pharmaceutical Science
[2]Student, Department of Pharmaceutical Technology

**Abstract**
Individual Case Safety Reports (ICSRs) form a central foundation of pharmacovigilance, converting initially reported adverse-event information into structured safety records for patient protection. Nevertheless, case processing requires a sequence of interconnected activities—including triage, data entry, medical coding, causality assessment, narrative writing, quality control, and regulatory reporting—that are frequently considered as separate activities.
This review considers these stages as one connected workflow of the ICSR lifecycle and the factors affecting the reliability of safety information. It highlights the challenges of applying standard tools like MedDRA, WHO-DD, the WHO-UMC system, and the Naranjo algorithm. Additionally, it covers established approaches to duplicate detection, quality control, and electronic submission via the ICH E2B(R3) framework. Finally, it explores how artificial intelligence and natural language processing are reshaping tasks from intake to coding, balancing automation with necessary human oversight.

**Keywords:** Individual Case Safety Report, Pharmacovigilance, MedDRA, WHO Drug Dictionary, Causality Assessment, Medical Narrative, Quality Control, Duplicate Detection, ICH E2B(R3), Artificial Intelligence

## Introduction

Pharmacovigilance (PV) ensures the safe use of medicines by detecting, assessing, understanding, and preventing adverse effects throughout a product's lifecycle. The Individual Case Safety Report (ICSR) is the source of safety information on which this process depends. An ICSR documents an individual patient's experience of a suspected adverse drug reaction and brings together key information about the patient, reporter, suspected medicinal product, and adverse event in a structured format. These individual reports form the foundation for broader safety evaluations, as they can be aggregated and analyzed to identify potential safety signals and support regulatory decisions. The completeness and accuracy of individual reports affect downstream activities, including signal detection, benefit–risk evaluation, product labeling, and regulatory decision-making [1].
Producing an ICSR that is both compliant and clinically useful requires considerably more than documenting an adverse event. Safety information may initially arrive in diverse forms—including

telephone calls, emails, published literature, clinical records, or digital and social media sources. Often incomplete or unstructured, this raw information must pass through a chain of processing stages, ranging from initial intake and medical coding to causality assessment, quality control, and final regulatory submission.

Every step in this sequence can influence the quality of the completed case. For example, decisions made during triage affect whether strict regulatory timelines are met, while inconsistent coding can skew adverse event identification during signal detection [1]. In the same way, the quality of causality assessments and medical narratives shapes how clearly a case is interpreted by medical reviewers, regulators, and inspectors. ICSR processing is not a series of isolated activities; a decision at one point carries consequences for every downstream stage.

Yet current literature tends to examine individual components of ICSR management in isolation. Publications focusing on medical coding using MedDRA or the WHO Drug Dictionary (WHO-DD) often overlook how coding decisions cascade into causality assessment. Discussions of causality tools like the WHO-UMC system or the Naranjo algorithm are often separated from practical challenges such as narrative writing, duplicate detection, and quality review [3]. The growing integration of artificial intelligence (AI) and natural language processing (NLP) is usually analyzed as a standalone technological development, rather than evaluated across the complete ICSR lifecycle.

Studying these stages separately hides how early-stage processing choices affect the overall reliability of safety data. An incomplete intake, a flawed coding decision, or an overlooked duplicate can compromise causality assessments, delay reporting, and degrade data consistency at a population level. To address this gap, this review presents a stage-by-stage overview of ICSR management. Beginning with case intake and continuing through triage, data entry, medical coding, causality assessment, narrative development, duplicate detection, quality control, and regulatory reporting under the ICH E2B(R3) framework, the review examines the practical application of key standards including MedDRA, WHO-DD, WHO-UMC, and Naranjo. It then addresses the emerging role of AI and NLP across the lifecycle, weighing automation against the need for expert human oversight [2]. Examining these processes as one workflow gives a fuller picture of what governs pharmacovigilance the reliability of safety information from the moment safety data is received.

**Sources of Safety Information**

Safety information reaches pharmacovigilance systems through several distinct channels, each with its own reporting characteristics and data quality profile. Spontaneous reports come unprompted from patients, caregivers, or healthcare professionals who voluntarily notify a manufacturer, regulator, or poison control center about a suspected adverse event [1]. These reports form the traditional backbone of post-marketing surveillance, though they vary widely in completeness since reporters are not obligated to follow a structured format.

Solicited reports arise from organized data collection activities, including patient support programs, market research surveys, disease registries, and named-patient or compassionate-use programs [1]. Because these programs actively prompt patients or healthcare providers for information, solicited reports tend to be more complete than spontaneous ones, but they also require careful classification to distinguish genuine adverse events from routine program feedback.

Scientific and medical literature represents another major source, with adverse event information appearing in published case reports, clinical trial results, and post-marketing studies [1]. Literature

monitoring requires systematic searching across databases to identify relevant publications, followed by extraction of case-level detail sufficient to meet ICSR minimum criteria—a task complicated by variation in how journals report adverse event data.

Clinical trials generate a further stream of safety information, distinct from post-marketing sources in that reporting obligations, causality assessment by investigators, and blinding status introduce additional procedural requirements [2]. Regulatory authorities themselves also forward cases to manufacturers when reports are received directly by an agency, requiring reconciliation to avoid duplicate processing.

Digital and social media sources have become an increasingly relevant channel for potential safety information, with patients increasingly describing their treatment experiences on public platforms [1, 21]. These sources present particular challenges: identifying which posts meet the four minimum criteria for a valid ICSR, verifying reporter and patient identifiability, and managing the volume of content requiring screening. Regulatory guidance in this area continues to evolve, and organizations vary in how systematically they monitor these channels [1, 22].

Across all source types, the common challenge is translating heterogeneous, often informal information into a structured case record without losing clinically relevant detail—a task that sets the foundation for every processing stage that follows.

**Case Validity & Triage**

Once safety information reaches a pharmacovigilance system, the first structured task is to determine whether it qualifies as a valid ICSR. For an ICSR to meet the minimum validity requirements, four elements must be present: an identifiable reporter, an identifiable patient, at least one suspect drug, and at least one adverse event [1]. Reports missing any of these elements are not processed as formal cases, though follow-up contact with the original reporter may be attempted to recover the missing detail before a final determination is made.

Beyond meeting these baseline criteria, each valid case must be assessed for seriousness. A case is classified as serious if the adverse event resulted in death, was life-threatening, required or prolonged hospitalization, caused persistent or significant disability, or resulted in a congenital anomaly [1]. This classification directly determines which regulatory reporting timeline the case must follow, with serious and unexpected events typically requiring expedited submission within a fixed window under ICH E2A guidelines, while other cases proceed through standard periodic reporting [2].

Triage builds on the validity and seriousness assessment by ranking cases for processing priority. Cases nearing an expedited reporting deadline are moved ahead of routine cases, and staff must weigh seriousness against expectedness—whether the event is already listed in the product's reference safety information—since unexpected serious events carry the greatest regulatory urgency [2]. A triage error at this stage has consequences beyond the individual case: a serious event misclassified as non-serious can cause an organization to miss its expedited reporting window, while over-classification can divert limited processing resources away from genuinely urgent reports.

Triage teams must also distinguish new cases from duplicate reports of an already-known event, since the same adverse event is sometimes reported through more than one channel or by more than one party [1]. Left unresolved, this can inflate case counts and distort later signal detection. Where information is incomplete, triage staff issue follow-up queries to reporters, requesting details such as dosage, event onset and duration, concomitant medications, or outcome—information that later stages of processing, particularly coding and causality assessment, depend on to proceed accurately.

**Data Entry and Case Documentation**

Following triage, the available case information is transferred into the safety database. Information from the original report is converted into structured fields that capture patient details, medical history, suspect and concomitant medicines, adverse events, and other clinical information [1]. Case entry is more than a direct transcription of the source material. Case processors need to interpret the information provided and enter it into the appropriate fields without adding details that are not present in the source report.

Source documents may also be reviewed to confirm the information entered into the safety database. This is particularly relevant for clinical trial cases and other cases where supporting records are available. Documents such as clinical records, laboratory reports, and correspondence can be used to verify details such as treatment dates, dose, route of administration, and event onset [1]. Identifying differences between the source information and the database record at this stage can prevent errors from being carried into later stages of case processing.

An ICSR may also change as new information is received during follow-up. Additional information may clarify the adverse event, treatment history, laboratory findings, or patient outcome. The safety database should therefore record the history of case updates, including the date information was received and the changes made during processing [1]. This provides traceability and allows reviewers to follow how the case developed over time.

Information that is not available should not be assumed or created during data entry. If an important detail remains unavailable after follow-up, this should be documented in the case. Accurate data entry is important because the information captured at this stage is used for medical coding, causality assessment, narrative development, duplicate detection, quality control, and regulatory reporting [3, 25]. An incorrect event date or an omitted concomitant medicine, for example, may affect the assessment of the time relationship between drug exposure and an event or the identification of possible drug interactions. Data entry is therefore an important part of ICSR processing and can influence the quality of the activities that follow.

**Medical Coding with MedDRA**

Medical coding transforms free-text descriptions of adverse events, medical history, and diagnostic information recorded in an ICSR into standardized terminology using the Medical Dictionary for Regulatory Activities (MedDRA) [4]. A common terminology permits safety information originating from different countries, languages, and reporting systems to be represented consistently. This supports the retrieval, comparison, and analysis of cases and is important for signal detection within safety databases [23].

MedDRA uses a five-level hierarchy consisting of Lowest Level Terms (LLTs), Preferred Terms (PTs), High Level Terms (HLTs), High Level Group Terms (HLGTs), and System Organ Classes (SOCs) [5]. LLTs represent the level closest to the wording used in the original report, while PTs represent standardized medical concepts. HLTs, HLGTs, and SOCs provide broader levels of classification. During coding, the processor selects the MedDRA term that best represents the information reported, while preserving the meaning of the original description.

Choosing the most suitable term may require consideration of the clinical context. For example, the term "swelling" may refer to localized edema, swelling of a specific body part, or another clinical finding depending on the context in which it was reported. The coder must therefore consider the information available in the case before selecting the appropriate term. Coding should reflect the reported information

and should not introduce a diagnosis or clinical interpretation that is not supported by the source.

Inconsistent coding of similar events distributes them across different terms, which can affect case retrieval, data review, and signal detection. Consistent application of coding conventions can therefore improve the comparability of cases within a safety database.

MedDRA is updated twice each year, and new versions may introduce, modify, or retire terminology [5]. Organizations therefore need to manage MedDRA versions when processing cases and conducting safety analyses. Changes between dictionary versions can affect the representation of previously coded information and should be considered during activities such as data review, trend analysis, and periodic safety reporting. Coding practices for situations not fully addressed by the terminology may also be supported by organizational coding conventions and documented procedures [5].

Quality control is used to assess whether coding decisions are consistent with the information provided in the source report. Depending on the organization's procedures, coded terms may be reviewed against the original case information to confirm that the selected LLT and resulting PT appropriately represent the reported event. Such review can identify coding differences and provide feedback to improve consistency across case processors.

**Drug Coding with WHO-DD**

Drug coding assigns standardized terminology to medicinal products recorded in an ICSR, including suspect drugs, concomitant medicines, and treatments for adverse events, into standardized terminology using the WHO Drug Dictionary (WHO-DD) [6]. Products may be reported by trade name, active ingredient, or other product details. The reported information is matched with the appropriate WHO-DD entry so that medicinal products can be identified and analyzed consistently within a safety database.

WHO-DD links medicinal products with information such as their active ingredients and product characteristics [6]. The same trade name may refer to different products or formulations in different countries. Coders therefore need to consider details such as the country of report, product name, active ingredient, dosage form, strength, and route of administration when selecting the appropriate entry. When information is incomplete, the coding decision should remain consistent with the details available in the source report.

Combination products can create additional coding challenges because a single product may contain more than one active ingredient. Their representation should follow the applicable coding conventions and database requirements so that relevant product information is retained for safety analysis [1]. Generic medicines may also be marketed under different trade names, making it important to distinguish between the reported product and its active ingredient when sufficient information is available.

Herbal and traditional medicines may also be reported as suspect or concomitant products. Coding can be difficult when these products are reported using local names or when their composition is not clearly described [6]. When an exact standardized match cannot be established, the available product information should be documented according to the applicable procedures while preserving the original information reported in the case.

WHO-DD is updated periodically, and organizations need to manage dictionary versions when coding cases and reviewing safety data [6]. Version changes may affect the representation and retrieval of medicinal product information and should therefore be considered during database maintenance and safety analysis. Quality control involves checking whether the selected coded product matches the information provided in the case, including the product name, active ingredient, formulation, strength, and country of

report.

Accurate drug coding is important because coded medicinal product information is used in case retrieval, signal detection, and aggregate safety analysis. An incorrect product match can affect the association between a medicine and an adverse event and may influence the interpretation of safety data.

**Causality Assessment Tools**

Causality assessment examines the extent to which the reported medicinal product may have contributed to the adverse event. It considers factors such as the timing of drug exposure and event onset, dechallenge and rechallenge, alternative causes, and other relevant clinical information [7]. The assessment helps reviewers interpret individual cases and contributes to signal detection, benefit–risk evaluation, and regulatory review [26].

The WHO-Uppsala Monitoring Centre (WHO-UMC) system is widely used for assessing suspected adverse drug reactions. It classifies drug-event relationships as certain, probable/likely, possible, unlikely, conditional/unclassified, or unassessable [7]. The assessment is based on clinical judgment rather than a numerical score. This allows the assessor to consider the overall case context, but different assessors may reach different conclusions when the available information is limited or open to interpretation.

The Naranjo algorithm uses a questionnaire and scoring system to assess causality [8]. It considers factors such as the timing of the event, previous reports of the reaction, dechallenge, rechallenge, alternative causes, and other relevant observations. The total score categorizes the relationship as doubtful, possible, probable, or definite. Its structured approach can improve consistency between assessors, but some questions may be difficult to answer when information such as rechallenge or alternative causes is unavailable.

The reliability of either approach depends on the information available in the case. Missing details about drug exposure, event timing, dechallenge, rechallenge, or concomitant medicines can limit the assessment [9, 30]. Polypharmacy can add another challenge because several medicines may have been taken during the same period, making it difficult to determine which product contributed to the event.

Organizations may also use internal causality procedures based on their products, therapeutic areas, or regulatory requirements. Regardless of the method used, the assessment should be supported by the available clinical evidence and documented clearly so that the reasoning can be understood during case review and regulatory evaluation [9].

**Medical Narrative Writing**

A medical narrative provides a chronological prose account of the case that brings together the patient's relevant medical history, the suspect and concomitant medications, the course of the adverse event, and its outcome into a single coherent account [1]. Where coded terms and structured data fields present case information as discrete elements, the narrative is what allows a medical reviewer, regulator, or inspector to understand the case as a clinical story rather than a set of unconnected data points.

A well-constructed narrative generally follows a consistent structure: it opens with brief patient demographics and relevant medical history, describes the suspect drug's indication, dose, and start date, presents the adverse event in chronological order including any diagnostic findings, and closes with the treatment given, the outcome, and the reporter's or company's causality assessment [3]. This structure mirrors the order in which a reviewer needs the information to interpret the case and evaluate causality without having to cross-reference other fields.

Balancing completeness against length is a recurring difficulty in narrative writing. Omitting relevant clinical detail, such as a relevant lab value or a missed dechallenge outcome, can weaken the case's interpretability, while including excessive irrelevant history can obscure the clinically significant information a reviewer actually needs. Narratives are generally expected to include all information necessary to support the causality assessment and coding decisions already made for the case, without introducing new interpretation not supported by the source documentation [1].

Common deficiencies identified during audits and inspections include narratives that omit the temporal relationship between drug administration and event onset, fail to mention concomitant medications with their own potential to explain the event, or state a causality conclusion without describing the clinical reasoning behind it. Inconsistencies between the narrative and the coded terms or structured fields in the same case are also frequently flagged, since a narrative describing one event while the case is coded to a different term undermines confidence in the case as a whole [1].

Because the narrative is often the first element a reviewer reads, its clarity and internal consistency with the rest of the case shape how the case is understood well before any structured data is examined in detail.

**Duplicate Detection**

A duplicate occurs when reports concerning the same patient and adverse event enter the pharmacovigilance system more than once—whether through different reporters, different channels, or as a follow-up to an earlier report that was not correctly linked to the original case [10]. Left unidentified, duplicates inflate the apparent frequency of an adverse event within a safety database, which can distort disproportionality analyses and other statistical methods used for signal detection [18].

Manual duplicate search relies on case processors comparing key identifying details—patient initials, date of birth, gender, country, suspect drug, and event terms—against existing cases in the database to judge whether two reports likely describe the same event [10]. Such an approach allows for clinical judgment in cases where identifying details are only partially reported, but it is time-consuming and depends on the reviewer noticing a potential match, which becomes harder as database size grows.

Algorithmic duplicate detection uses defined matching rules or statistical similarity scoring across case fields to flag potential duplicates for review, reducing reliance on a processor happening to recall or search for a similar case [10, 19]. These systems typically generate a ranked list of possible matches rather than an automatic determination, since even a strong field-level match does not confirm that two reports describe the same patient and event; a human reviewer generally makes the final decision.

Cross-database and cross-organization duplicate detection presents additional difficulty, since the same event may be reported to a regulatory authority, a manufacturer, and a clinical trial sponsor independently, with each party holding only partial identifying information about the patient [10, 20]. Variations in how patient identifiers are recorded—different date formats, transliterated names, or incomplete initials—can prevent even algorithmic matching from identifying cases that describe the same underlying event.

When a duplicate is confirmed, one case is typically designated as the primary record and the other is linked to it or closed, with any additional information from the duplicate incorporated into the primary case rather than discarded [10]. Maintaining accurate duplicate linkage is treated as an ongoing task rather than a one-time check, since new reports of an already-known event can continue to arrive after the original case has been processed.

**Quality Control & Quality Assurance**

Quality control and quality assurance have complementary but different functions within ICSR processing. Quality control refers to case-level checks performed on individual reports—verifying that data entry, coding, causality assessment, and the narrative are accurate and consistent with the source documentation—while quality assurance refers to the broader systems, procedures, and periodic audits that confirm the case-processing function as a whole is operating as intended [11].

Case-level quality control typically involves reviewing a defined sample of processed cases against their source documents, checking that structured fields were entered correctly, that MedDRA and WHO-DD terms were appropriately selected, that the causality assessment is supported by the case information, and that the narrative is consistent with the coded data [11, 24]. Discrepancies identified during this review are corrected in the case and, where a pattern emerges, fed back to the individual processor or coding team as part of ongoing training.

Organizations track quality metrics to monitor case-processing performance over time, including error rates by processing stage, timeliness against regulatory reporting deadlines, and the rate at which cases require follow-up queries to reporters. These metrics allow a pharmacovigilance department to identify whether errors are concentrated in a particular stage, such as coding or narrative writing, or associated with a particular case source or product.

When a quality issue is identified—whether through routine sampling, an internal audit, or a regulatory inspection finding—organizations generally conduct a root-cause analysis to determine why the error occurred, followed by a corrective and preventive action (CAPA) plan addressing both the specific case and the underlying process gap that allowed it [12]. A CAPA might involve additional processor training, a revision to internal coding conventions, or a change to how a particular data field is validated at entry.

Quality assurance activities extend beyond individual case review to periodic audits of the overall pharmacovigilance system, assessing whether procedures are being followed consistently across processing teams and whether the quality control process itself remains effective as case volume, product portfolio, or regulatory requirements change [12].

**Electronic Submission (ICH E2B(R3))**

After coding, causality assessment, narrative preparation, and quality review have been completed, it is transmitted to regulatory authorities in a standardized electronic format defined by the ICH E2B(R3) implementation guide [3]. This standard specifies how case information—patient details, drug information, event terms, and the narrative—should be structured into defined data elements so that safety databases across different organizations and regulators can exchange case information in a consistent, machine-readable format.

E2B(R3) organizes case data into sections covering the case identification, patient characteristics, drug information for each suspect and concomitant product, reaction or event information coded to MedDRA, and the case narrative, among other elements [3]. Each data element has defined formatting and, in many cases, defined value sets, meaning a field such as country of the reaction must draw from a standardized list rather than free text. Cases are validated against these formatting rules before submission, and a case failing validation is returned for correction rather than accepted into the receiving system.

Submission requirements vary somewhat by region despite the shared E2B(R3) foundation. In the European Union, cases are submitted to EudraVigilance under GVP Module VI, while other regions maintain their own receiving systems and regional business rules [1]. Each system may apply its own

additional validation rules on top of the core E2B(R3) structure, meaning a case accepted by one regulator's system is not guaranteed to pass validation in another without adjustment.

Reporting timelines are tied to the seriousness and expectedness determination made earlier in the case lifecycle: serious, unexpected adverse events generally require submission within a fixed number of calendar days from when the organization first became aware of the case, while other cases are submitted through periodic reports covering a defined reporting interval [2]. Missing a submission deadline is treated as a compliance finding independent of the quality of the case itself, which places practical pressure on every earlier processing stage to be completed within the time available.

Once submitted, an acknowledgment message confirms whether the receiving system accepted or rejected the case, and rejected cases must be corrected and resubmitted, making the submission stage a feedback point that can surface earlier data entry or coding errors not caught during quality control [3].

**AI & NLP in ICSR Processing**

Artificial intelligence (AI) and natural language processing (NLP) are being increasingly incorporated across the early stages of ICSR management, particularly tasks involving large volumes of unstructured text such as intake screening, literature monitoring, and medical coding [13]. NLP models can screen incoming reports, literature articles, or social media content to identify text that meets the minimum criteria for a valid ICSR, reducing the manual effort needed to review sources that ultimately contain no reportable case [16].

In medical coding, NLP-based tools can suggest MedDRA or WHO-DD terms by matching free-text descriptions in a case against the standardized dictionaries, presenting a coder with candidate terms rather than requiring the coder to search the dictionary manually for each concept [16,27]. Similar approaches are applied to narrative writing, where models can generate a draft narrative from structured case data as a starting point for a medical reviewer to revise, rather than requiring the narrative to be composed entirely from scratch.

Duplicate detection has also benefited from AI-based approaches, with machine learning models trained to identify likely duplicate cases based on patterns across multiple fields simultaneously, rather than relying on fixed matching rules alone [14,29]. These models can surface potential matches that differ in wording or formatting but describe the same underlying event, which fixed-rule systems may miss.

Despite these applications, human review remains essential rather than merely an optional safeguard. Regulatory guidance generally expects that AI-suggested coding, narrative drafts, or duplicate matches are reviewed and confirmed by a qualified individual before being finalized in a case, since current models can misinterpret clinical context, generate plausible but inaccurate text, or fail on cases involving rare terminology or unusual presentations not well represented in their training data [17]. Validating model performance, documenting how a tool reached a given output, and maintaining an audit trail of human review are treated as necessary components of using these tools in a regulated environment, rather than as separate from the core processing workflow [13,28].

Across these applications, the prevailing model intake, coding, narrative writing, and duplicate detection is one of augmentation rather than replacement: AI and NLP tools reduce the manual burden of processing large case volumes, while causality judgment, final coding decisions, and case approval remain the responsibility of trained personnel [15].

## Challenges & Future Directions

A number of difficulties extend across the entire ICSR lifecycle rather than being confined to a single stage. Causality assessment methods such as WHO-UMC and Naranjo remain dependent on the completeness of case information, and neither fully resolves the inter-assessor variability that arises when cases involve polypharmacy or limited rechallenge data [9]. Medical and drug coding face a related difficulty: MedDRA and WHO-DD are updated on fixed schedules, and organizations must continually manage version changes to avoid inconsistency in how the same clinical concept or product is represented over time [5,6]. Duplicate detection, too, remains incomplete even with algorithmic support, since variation in how patient identifiers are recorded across sources can prevent matching systems from linking reports that describe the same event [10].

A common thread across these challenges is that later-stage processes cannot fully compensate for gaps introduced earlier in the pipeline. An incomplete case at intake limits the causality assessment that follows; inconsistent coding fragments signals that quality control may not catch until aggregate analysis; and a narrative inconsistent with its coded terms can undermine reviewer confidence regardless of how thoroughly the case was otherwise processed. Addressing these gaps calls for closer integration between stages, rather than treating each as an independent checkpoint.

The growing use of AI and NLP tools offers a partial response to some of these challenges, particularly for coding suggestion, narrative drafting, and duplicate matching at scale, but introduces its own open questions [13,17]. Regulatory expectations for validating and documenting these tools are still developing, and consistent standards for how much human review is sufficient for a given task have not been fully established across regions. Future work will likely need to address how AI-assisted processing is validated for regulatory acceptance, how performance is monitored as case volume and reporting sources continue to diversify, and how smaller organizations without extensive in-house data science resources can adopt these tools without compromising case quality.

Continued refinement of causality assessment tools to reduce inter-assessor variability, closer harmonization of regional electronic submission requirements under E2B(R3), and clearer regulatory guidance on AI validation in pharmacovigilance represent the areas most likely to shape how ICSR processing evolves in the coming years.

## Conclusion

ICSR processing is best understood as an interconnected sequence rather than as one isolated task, each of which depends on the accuracy of the stage before it. A gap at intake carries through to triage; an inconsistency in coding can fragment a signal that should have been detected as one; an incomplete causality assessment limits how confidently a case can be interpreted; and a narrative that does not match its coded terms undermines the case as a whole, regardless of how carefully each individual field was completed. Viewing these stages together, rather than as separate technical topics, makes this interdependence clear in a way that studying any single stage in isolation does not.

Standardized tools such as MedDRA, the WHO Drug Dictionary, the WHO-UMC causality system, and the Naranjo algorithm provide the shared vocabulary and structured judgment that make case comparison possible across organizations and regulators. Their value depends on consistent application, which in turn depends on the quality control, duplicate detection, and audit processes that check whether that consistency is actually being achieved in practice. Electronic submission under the ICH E2B(R3) framework then carries this structured, quality-checked information to regulators in a common format,

closing the loop between individual case processing and the aggregate safety data that regulatory decisions rely on.

Artificial intelligence and natural language processing are beginning to reduce the manual burden of processing large case volumes, particularly in intake screening, coding suggestion, and duplicate detection. These tools do not remove the need for human judgment in causality assessment, final coding decisions, or case approval, and their reliable use in a regulated setting depends on validation, documentation, and consistent human review rather than treating automation as a replacement for trained personnel.

As case volumes continue to grow and reporting sources diversify, the reliability of pharmacovigilance data will continue to depend on how well each stage of ICSR processing is executed and how well these stages are understood as parts of a single connected system, rather than as isolated technical tasks.

**References**

1. European Medicines Agency, Guideline on good pharmacovigilance practices (GVP) – Module VI (Rev 2), EMA/873138/2011 Rev 2, 2017. https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-good-pharmacovigilance-practices-gvp-module-vi-collection-management-submission-reports-suspected-adverse-reactions-medicinal-products-rev-2_en.pdf
2. International Council for Harmonisation, ICH E2A: Clinical Safety Data Management: Definitions and Standards for Expedited Reporting, 1994. https://database.ich.org/sites/default/files/E2A_Guideline.pdf
3. International Council for Harmonisation, ICH E2B(R3): Guideline on Electronic Transmission of Individual Case Safety Reports (ICSRs). https://www.ich.org/page/e2br3-individual-case-safety-report-icsr-specification-and-related-files
4. E.G. Brown, L. Wood, S. Wood, The medical dictionary for regulatory activities (MedDRA), Drug Safety, 1999, 20 (2), 109–117. https://doi.org/10.2165/00002018-199920020-00002
5. MedDRA MSSO, Introductory Guide MedDRA Version 27.0, 2024. https://admin.meddra.org/sites/default/files/guidance/file/intguide_27_0.pdf
6. Uppsala Monitoring Centre, WHO Drug Dictionary Enhanced (WHO-DDE) – User Guide. https://who-umc.org/whodrug/whodrug-global/
7. The Uppsala Monitoring Centre, The use of the WHO-UMC system for standardised case causality assessment. https://www.who.int/docs/default-source/medicines/pharmacovigilance/whocausality-assessment.pdf
8. C.A. Naranjo, U. Busto, E.M. Sellers, et al., A method for estimating the probability of adverse drug reactions, Clinical Pharmacology and Therapeutics, 1981, 30 (2), 239–245. https://doi.org/10.1038/clpt.1981.154
9. S. More, S. Atal, P. Solanki, Inter-rater agreement between WHO-Uppsala Monitoring Centre system and Naranjo algorithm for causality assessment of adverse drug reactions, ScienceDirect, 2024. https://www.sciencedirect.com/science/article/abs/pii/S1056871924000248
10. European Medicines Agency, GVP Module VI Addendum I – Duplicate management of suspected adverse reaction reports, EMA/405655/2016, 2017.

https://www.ema.europa.eu/en/documents/regulatory-procedural-guideline/guideline-good-pharmacovigilance-practices-gvp-module-vi-addendum-i-duplicate-management-suspected-adverse-reaction-reports_en.pdf

11. European Medicines Agency, GVP Module I – Pharmacovigilance systems and their quality systems. https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-good-pharmacovigilance-practices-module-i-pharmacovigilance-systems-their-quality-systems_en.pdf
12. European Medicines Agency, GVP Module IV – Pharmacovigilance audits (Rev 1). https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-good-pharmacovigilance-practices-module-iv-pharmacovigilance-audits-rev-1_en.pdf
13. Council for International Organizations of Medical Sciences (CIOMS), Artificial Intelligence in Pharmacovigilance: CIOMS Working Group report, 2025. https://doi.org/10.56759/cdob6397
14. D.J. Lewis, J.F. McCallum, Utilizing advanced technologies to augment pharmacovigilance systems: challenges and opportunities, Therapeutic Innovation & Regulatory Science, 2020, 54, 888–899. https://doi.org/10.1007/s43441-019-00023-3
15. TransCelerate BioPharma, Industry Perspective on Artificial Intelligence/Machine Learning in Pharmacovigilance, 2022. https://doi.org/10.1007/s40264-022-01164-5
16. M. Salas, J. Petracek, P. Yalamanchili, et al., Artificial intelligence in pharmacovigilance: opportunities in ICSR processing, Perspectives in Clinical Research, 2022. https://doi.org/10.1007/s40290-022-00441-z
17. R. Ball, G. Dal Pan, "Artificial Intelligence" for pharmacovigilance: ready for prime time?, Drug Safety, 2022, 45 (5), 429–438. https://doi.org/10.1007/s40264-022-01157-4
18. G.N. Norén, R. Orre, A. Bate, I.R. Edwards, Duplicate detection in adverse drug reaction surveillance, Data Mining and Knowledge Discovery, 2007, 14 (3), 305–328. https://doi.org/10.1007/s10618-006-0052-8
19. P.M. Tregunno, D.B. Fink, C. Fernandez-Fernandez, E. Lazaro-Bengoa, G.N. Noren, Performance of probabilistic method to detect duplicate individual case safety reports, Drug Safety, 2014, 37 (4), 249–258. https://doi.org/10.1007/s40264-014-0146-y
20. R. Kiguba, G. Isabirye, J. Mayengo, et al., Navigating duplication in pharmacovigilance databases: a scoping review, BMJ Open, 2024, 14 (4), e081990. https://doi.org/10.1136/bmjopen-2023-081990
21. S. Golder, K. O'Connor, Y. Wang, A. Klein, G. Gonzalez Hernandez, The value of social media analysis for adverse events detection and pharmacovigilance: scoping review, JMIR Public Health and Surveillance, 2024, 10, e59167. https://doi.org/10.2196/59167
22. J. van Stekelenborg, J. Ellenius, S. Maskell, et al., Recommendations for the use of social media in pharmacovigilance: lessons from IMI WEB-RADR, Drug Safety, 2019, 42 (12), 1393–1407. https://doi.org/10.1007/s40264-019-00858-7

23. A. Bate, S.J.W. Evans, Quantitative signal detection using spontaneous ADR reporting, Pharmacoepidemiology and Drug Safety, 2009, 18 (6), 427–436. https://doi.org/10.1002/pds.1742
24. M. Lindquist, Data quality management in pharmacovigilance, Drug Safety, 2004, 27 (12), 857–870. https://doi.org/10.2165/00002018-200427120-00003
25. T. Bergvall, G.N. Noren, M. Lindquist, vigiGrade: a tool to identify well-documented individual case reports and highlight systematic data quality issues, Drug Safety, 2014, 37 (2), 65–77. https://doi.org/10.1007/s40264-013-0131-x
26. J. Almenoff, J.M. Tonning, A.L. Gould, et al., Perspectives on the use of data mining in pharmacovigilance, Drug Safety, 2005, 28 (11), 981–1007. https://doi.org/10.2165/00002018-200528110-00002
27. C. Combi, M. Zorzi, G. Pozzani, U. Moretti, E. Arzenton, From narrative descriptions to MedDRA: automagically encoding adverse drug reactions, Journal of Biomedical Informatics, 2018, 84, 184–199. https://doi.org/10.1016/j.jbi.2018.07.001
28. B. Ohana, J. Sullivan, N. Baker, Validation and transparency in AI systems for pharmacovigilance: a case study applied to the medical literature monitoring of adverse events, arXiv:2201.00692, 2022. https://arxiv.org/abs/2201.00692
29. M.M. Zitu, S. Zhang, D.H. Owen, C. Chiang, L. Li, Generalizability of machine learning methods in detecting adverse drug events from clinical narratives in electronic medical records, Frontiers in Pharmacology, 2023, 14, 1218679. https://doi.org/10.3389/fphar.2023.1218679
30. T.B. Agbabiaka, J. Savovic, E. Ernst, Methods for causality assessment of adverse drug reactions: a systematic review, Drug Safety, 2008, 31 (1), 21–37. https://doi.org/10.2165/00002018-200831010-00003