volv
Global healthcare solutions

Finding GEP-NET Patients Earlier

AI-Driven Patient Identification in UK Primary Care

VOLV GLOBAL CASE STUDY | JUNE 2026

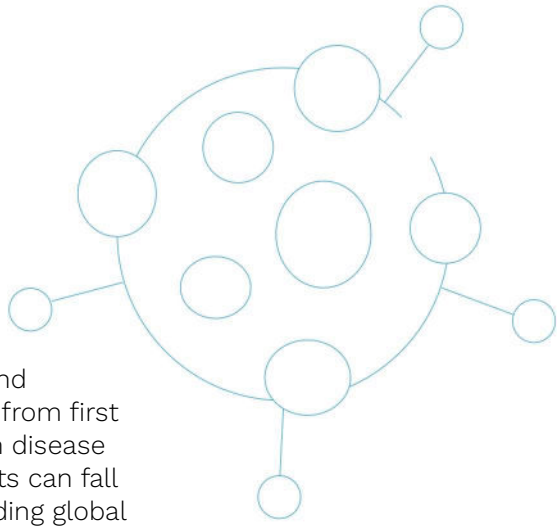

Executive Summary

Gastroenteropancreatic neuroendocrine tumours (GEP-NETs) are rare and notoriously difficult to diagnose. Patients typically wait nearly five years from first symptoms before receiving a confirmed diagnosis[1] – years during which disease can progress from lower to higher grade, and five-year survival prospects can fall from approximately 86% to as low as 25%. In a collaboration with a leading global pharmaceutical company, Volv Global applied its proprietary machine learning methodology to a large-scale UK primary care electronic health record (EHR) dataset to identify likely-undiagnosed GEP-NET patients. The work established that meaningful numbers of such patients could be identified in the UK primary care population, and that these patients are younger than diagnosed patients at the point of diagnosis – suggesting that earlier detection and intervention may be possible.

Key Findings at a Glance

~1,200

The estimated number of undiagnosed GEP-NET patients that could potentially be identified in the UK population using Volv Global's proprietary ML methodology. By comparison. there were approximately 3,813 patients with GEP-NET diagnosis codes in the dataset at that time.

5–7 yrs

Patients whose records were flagged by the model were, on average, 5–7 years younger than those currently diagnosed – suggesting the possibility of earlier-stage identification.

0.756

ROC-AUC achieved against clinically similar mimic conditions rather than the general population – good predictive performance given the substantially harder discrimination task

~24M

De-identified UK primary care patient records analysed, drawn from approximately 1,100 GP practices via Optimum Patient Care Research Database

Value Created

For patients

Earlier identification of those whose records are consistent with GEP-NETs, supporting the case for earlier access to appropriate care

For clinicians

Structured real-world evidence on GEP-NET symptom burden, phenotypic profile and treatment patterns – with transparent model outputs designed to support clinician review

For the pharmaceutical partner

A scalable methodology and visualisation platform (inVolv) transferable across geographies and data environments, providing the foundation for prospective validation and clinical deployment

Background

The Disease: Gastroenteropancreatic Neuroendocrine Tumours

Gastroenteropancreatic neuroendocrine tumours (GEP-NETs) are rare malignancies arising in the neuroendocrine cells of the gastrointestinal tract and pancreas. Prognosis is strongly associated with tumour grade and disease stage at the time of diagnosis. Published data indicate that five-year survival rates differ substantially across grades: patients with well-differentiated, low-grade disease (G1) can achieve survival rates of approximately 86%, while five-year survival for high-grade disease (G3) may be as low as 25%.[1] Because disease can progress from lower to higher grade during the period of diagnostic delay, the nearly five years patients typically wait before receiving a confirmed diagnosis carries direct prognostic consequences – patients are not simply waiting; their disease trajectory is worsening.

GEP-NETs are characterised by non-specific, often overlapping symptoms. These non-specific presentations are easily attributed to more prevalent conditions such as irritable bowel syndrome, inflammatory bowel disease, or diabetes. This diagnostic ambiguity – compounded by low disease awareness among non-specialist clinicians – contributes directly to the prolonged diagnostic delay that characterises GEP-NETs.

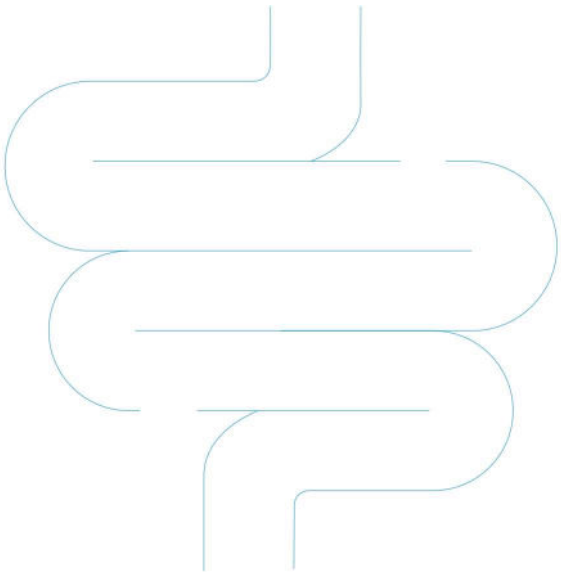

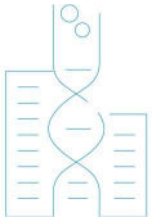

The Challenge for the Pharmaceutical Partner

The client, a leading global pharmaceutical company, approached Volv Global with a clearly defined need: to understand whether, and to what extent, there existed an identifiable population of GEP-NET patients who were currently undiagnosed or misattributed within routine records.

The challenge was multi-dimensional. Primary care coding for neuroendocrine disease is inherently imprecise: GEP-NET patients may carry a range of non-specific diagnostic codes, with organ-specific NET codes representing only a proportion of the true population. Any analytical approach would need to function reliably in the face of this real-world data fragmentation – and do so in a dataset spanning approximately 24 million de-identified patient records across approximately 1,100 UK GP practices.

The client also needed assurance that any identified population estimate would be scientifically robust, reproducible, and suitable for informing downstream strategic and medical affairs decisions.

Problem Statement

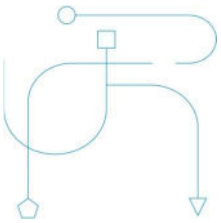

The Medical Challenge

The absence of a simple, definitive diagnostic test for GEP-NETs means that primary care clinicians must rely on a constellation of symptoms, investigations, and referral patterns to raise suspicion of the disease. In practice, this rarely happens: patients present with symptoms that are common across a wide range of conditions, and GEP-NETs are seldom at the top of the differential diagnosis. The result is that many patients remain unrecognised, sometimes for years.

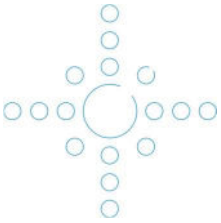

The Data Challenge

The clinical ambiguity of GEP-NETs is mirrored in how they appear – or fail to appear – in the data. A substantial proportion of patients with NET-consistent records carry unspecified-organ diagnostic codes rather than organ-specific GEP-NET codes, meaning a straightforward code-based query would systematically under-count the true population. Defining GEP-NETs in the record is therefore a methodological challenge in its own right.

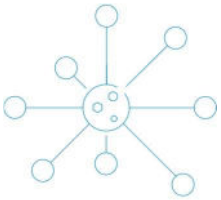

The Pharma Challenge

For a pharmaceutical company with treatments indicated in GEP-NETs, a credible estimate of the undiagnosed population size – and its demographic and clinical profile – carries direct implications for medical affairs, commercial, and clinical development functions alike. The challenge was to produce insights that were both analytically rigorous and operationally useful.

Key Questions

- Can patients with undiagnosed GEP-NETs be identified in UK primary care EHR data using a machine learning approach?
- What is the estimated size of the potentially undiagnosed GEP-NET population in the UK, and what are their demographic and clinical characteristics?
- How do the profiles of flagged undiagnosed patients compare with those of currently diagnosed patients?

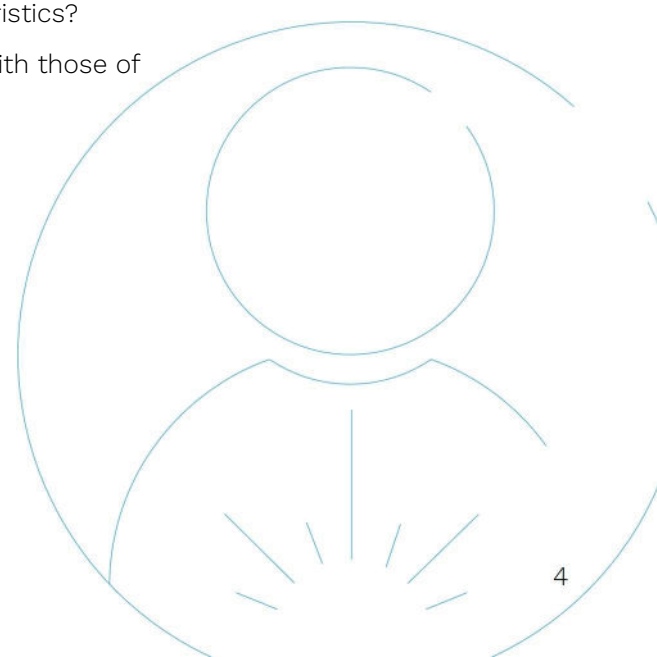

The Solution

Volv Global applied its proprietary machine learning methodology – operating through the inTrigue framework – to the Optimum Patient Care Research Database (OPCRD), a de-identified UK primary care EMR dataset.

The project was structured as a phased analytical programme, proceeding from dataset construction through feasibility assessment, initial model development, and preliminary population estimation. The inVolv platform was developed in parallel to enable geographic heatmapping of identified patient populations, providing a visual and interactive layer on top of the model and analytical outputs. All work was conducted under an ADEPT (Anonymised Data Ethics and Protocol Transparency) governance framework, administered by Optimum Patient Care (OPC), with formal protocol approval obtained prior to analysis.

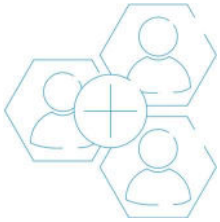

Data and Cohort Construction

The OPCRD offers an unusually large-scale view of UK primary care – approximately 24 million de-identified records from roughly 1,100 GP practices – making population-level patient identification tractable in a disease as rare as GEP-NETs. However, it carries the familiar challenges of real-world EHR data: sparsity, feature and label noise, and structural artefacts such as patient record duplication, all of which must be addressed directly before any modelling begins.

Constructing the positive cohort presented a specific challenge given the coding ambiguity described earlier: a substantial proportion of NET patients are not captured by a direct organ-specific code query. Volv Global's cohort construction methodology recovers these patients through a proprietary procedure combining coded and contextual signals from the record, producing a quality-controlled positive cohort of 1,857 GEP-NET patients suitable for supervised learning. The negative cohort was constructed from clinically relevant comparator conditions rather than the general population – a deliberate design choice that frames model performance against the discrimination task with real-world clinical value, rather than a trivially easier baseline.

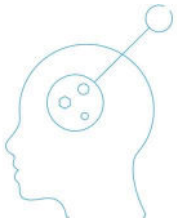

Disease Characterisation and Patient Insights

Prior to model development, Volv Global conducted comprehensive phenotypic characterisation of the diagnosed GEP-NET cohort, generating insights that carry value independently of the patient-finding model. Analysis of symptom prevalence confirmed the multi-system burden of GEP-NETs: gastrointestinal symptoms (abdominal pain, diarrhoea, gastritis) and respiratory symptoms (cough, chest pain, dyspnoea) were highly prevalent across the cohort, as were neurological symptoms including anxiety and depression. All symptom categories were more prevalent in the GEP-NET cohort compared to a matched random population from the same database.

Treatment patterns revealed the underrepresentation of specialist therapies in primary care records. Somatostatin analogues (octreotide and lanreotide) were among the most frequently recorded disease-specific treatments, consistent with established GEP-NET management.

Model Development and Performance

Volv Global trained and validated patient-finding models using 80/20 training and test data splits with three-fold cross-validation. The initial full model achieved a test set ROC-AUC of 0.756 and a PR-AUC of 0.427 – strong predictive performance given that discrimination was evaluated against clinically similar mimic conditions rather than the general population, a substantially harder task than the standard population-level benchmark.

The most predictive features were clinically coherent and consistent with the known pre-diagnostic presentation of GEP-NETs: serum biochemistry markers, patient age, referrals to colorectal and gastroenterology clinics, relevant imaging and procedural findings, and a range of gastrointestinal presentations. The interpretability of these features supports clinician review and provides a transparent basis for downstream validation.

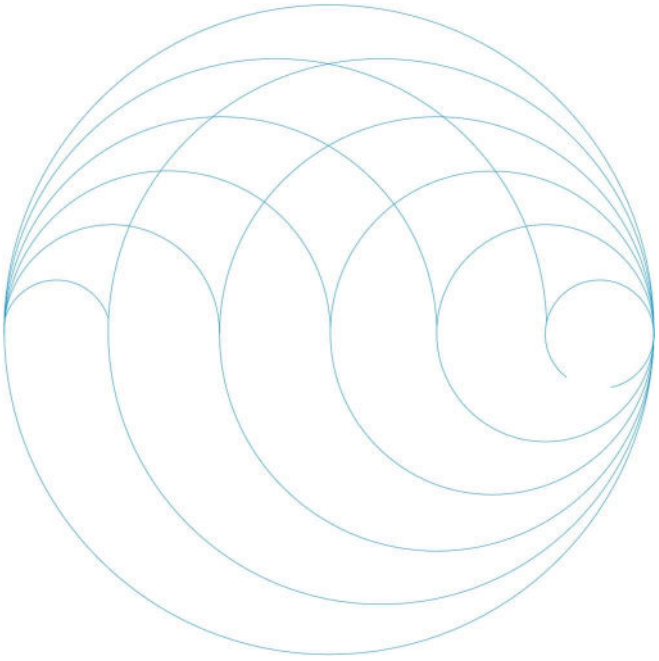

Results

Estimated Identifiable Population

At a model precision of 0.85, approximately 1,200 undiagnosed patients in the UK population have records consistent with GEP-NETs and could potentially be identified by the model.

This estimate was generated by applying the trained model to an available subset of 6.8 million patients from the OPCRD, with results then extrapolated to estimate the potentially identifiable undiagnosed population across the full database. The 0.85 precision threshold was selected to balance clinical actionability against false-positive burden; alternative thresholds would shift the estimated population size accordingly.

To contextualise this figure: there are currently approximately 3,813 patients with GEP-NET diagnosis codes in the dataset. An additional 1,200 likely-undiagnosed patients – approximately an additional one-third again of the existing diagnosed population – is a clinically and commercially meaningful estimate of unmet diagnostic need.

Earlier Identification Signal

Analysis of the demographic profile of patients flagged as likely undiagnosed by the model revealed that these patients were, on average, 5–7 years younger than diagnosed patients at the point of diagnosis – with female patients approximately 5 years younger and male patients approximately 7 years younger.

Volv Global interprets this finding carefully. The age difference is a demographic observation: the model is flagging patients who are younger in age. The hypothesis arising from this – that these patients may be at an earlier stage of their disease course at the time of flagging – is scientifically plausible and is consistent with the broader objective of earlier identification. It is also consistent with the prognostic data: given that five-year survival for GEP-NET ranges from approximately 86% for low-grade disease (G1) to as low as 25% for high-grade disease (G3), identifying patients earlier in their disease trajectory carries meaningful clinical stakes.[1] However, this hypothesis would require prospective clinical validation to confirm. No claim is made regarding clinical outcomes, disease stage, or treatment benefit on the basis of this observation.

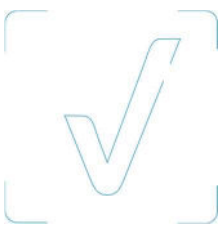

inVolv Platform: Geographic Visualisation

Alongside the analytical work, Volv Global developed and demonstrated the inVolv dashboard for GEP-NETs, providing geolocation heatmapping and analytical insights of both diagnosed and model-flagged undiagnosed patients across UK GP practices. The platform was built and validated internally during Phase 1, with GP-level deployment scoped as a Phase 2 activity.

Conclusion

This project demonstrated that Volv Global's proprietary machine learning methodology can be applied effectively to large-scale UK primary care EHR data to identify patients whose records are consistent with undiagnosed GEP-NETs. The analytical work addressed real-world data challenges – sparse coding, missing records, patient duplicates, and clinically ambiguous presentations – and produced a model with good discriminative performance against clinically similar mimic conditions.

The results deliver three outputs of value: a structured, data-grounded estimate of the potentially identifiable undiagnosed GEP-NET population in the UK; a comprehensive phenotypic characterisation of GEP-NET patients in primary care – including symptom burden, treatment patterns, and demographics – directly usable by medical affairs teams; and a geographic visualisation capability through the inVolv platform, mapping diagnosed and flagged patients to GP practice level.

The project was conducted as a genuine collaboration: Volv Global worked closely with the client's clinical and medical affairs teams throughout, incorporating domain expertise into cohort design, comparator selection, and feature interpretation – ensuring outputs that are clinically grounded, not merely algorithmic.

The underlying analytical framework is directly transferable to other geographies and data environments, with geographic scale-out identified as a priority next step.

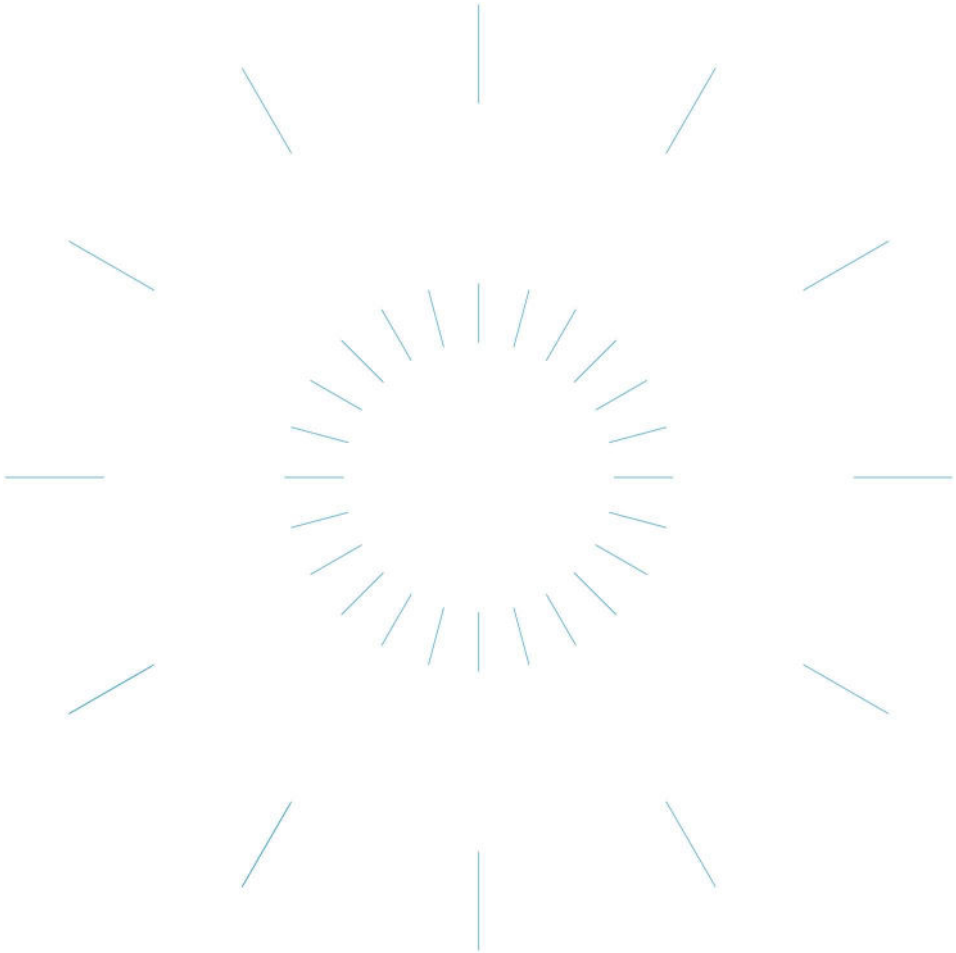

Recommended Next Steps

The Phase 1 analytical programme established a scientifically sound foundation for further development. The following steps are recommended to translate the model's outputs into clinical and commercial impact.

- Prospective clinical validation: Subjecting the model's flagged patient population to independent clinician chart review would validate precision in a real-world setting and generate the evidence needed for deployment at scale.
- Deployment at GP practice level: With a validated model and an operational inVolv platform, the architecture exists to enable GP-level alerts for patients whose records are consistent with GEP-NETs. This requires a deployment programme involving clinical governance, NHS engagement, and coordination with specialist referral centres.
- Geographic scale-out: The methodology is directly transferable to other data environments. Applying the same approach to US claims data or European EMR datasets would both extend the reach of the patient identification capability and provide cross-system validation of the model's generalisability.
- Integration with treatment access pathways: Earlier identification – once validated – creates an opportunity to examine whether earlier access to appropriate therapies improves outcomes, through a prospective study of the flagged cohort.
- Cross-functional evidence generation: The disease characterisation outputs – symptom burden, treatment patterns, referral pathways – are directly usable by medical affairs for clinical education, payer submissions, and HTA, and are suitable for peer-reviewed publication.

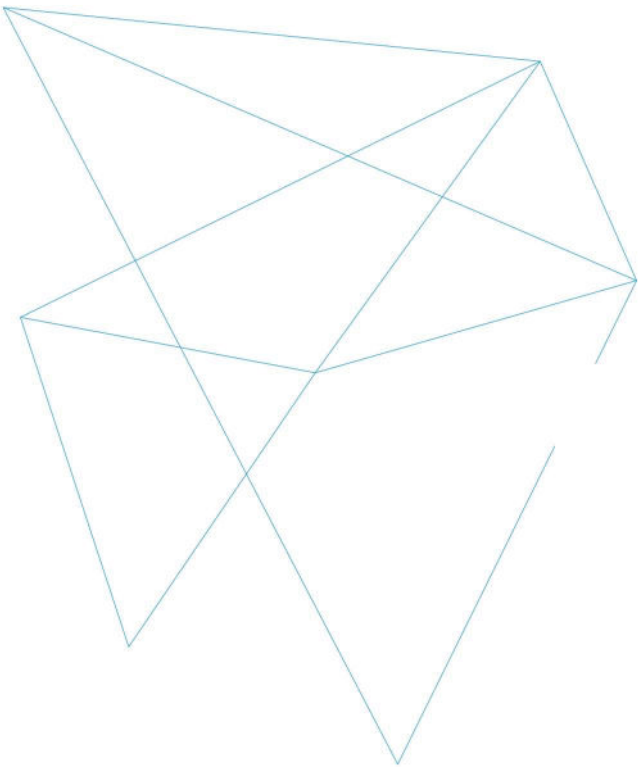

References

1. Basuroy R, Bouvier C, Ramage JK, Sissons M, Srirajaskanthan R. Delays and routes to diagnosis of neuroendocrine tumours. BMC Cancer. 2018;18:1122. doi:10.1186/s12885-018-5057-3 (PMC6240263). (Median time from first symptom to diagnosis: 53.8 months; 80% of patients visited their GP a median of 11 times prior to diagnosis.)

1. Ran S, et al. The assessment of Ki-67 for prognosis of gastroenteropancreatic neuroendocrine neoplasm patients: a systematic review and meta-analysis. Front Oncol. 2023;13:1145093. (Cumulative 5-year survival: GEP-NEN G1 ~86%, G2 ~65%, G3 ~25%.)

Disclaimer

This case study describes analytical research conducted by Volv Global in collaboration with a pharmaceutical partner. All analytical outputs are designed to support strategic and clinical decision-making and are not intended to constitute clinical diagnosis. Final clinical judgement remains with qualified healthcare professionals. This document has been prepared for informational purposes and does not represent regulatory guidance.

About Volv Global

Volv Global is a healthcare AI company founded in 2017 and headquartered in Épalinges, Switzerland. Its mission is to generate new knowledge at speed, close the diagnostic gap, as well as other gaps in the care pathway, to improve patient outcomes. Volv Global works across conditions where patients are difficult to identify, where the window for effective treatment is narrow, and where understanding which patients will progress or need therapy beyond standard care can meaningfully change outcomes. It is a trusted partner to leading pharmaceutical organisations across the USA and Europe, with solutions deployed in live clinical programmes.

Applying a proprietary machine learning methodology to population-scale real-world data – accessed through trusted data partners covering more than 400 million patients – Volv Global generates disease intelligence that enables pharmaceutical teams to de-risk clinical programmes, identify and stratify patient populations with greater precision, and build stronger real-world evidence. For clinicians, Volv Global's insights are designed to surface actionable signals within existing care pathways. For patients, they translate into earlier diagnosis, better-informed treatment decisions, and a faster path to the treatment that can help them, through a diagnostic system that too often leaves difficult-to-diagnose diseases unrecognised for years.

Volv Global's solutions each address a distinct clinical question across the patient journey, and are configured to the client's specific research question, disease area, and healthcare setting. Volv Global does not hold patient data; all work is conducted within the governed environments operating under applicable privacy and regulatory frameworks.

www.volv.global