

EXTENDING DATA-LIMITED ASSESSMENT METHODS THROUGH AI-ASSISTED INTERPRETATION: THE FISHERIES INTELLIGENCE AND ADVISORY SYSTEM (FIAS)

ALISTAIR DUNN¹, KERRIE ROBERTSON²¹

¹ Ocean Environmental Ltd, Wellington, New Zealand. Corresponding author:
alistair.dunn@oceanenvironmental.co.nz

² Adira Consulting

Purpose

To inform the Working Party on Methods (WPM) of the Fisheries Intelligence and Advisory System (FIAS), a proposed AI-assisted architecture for applying accepted data-limited stock assessment methods in fisheries administrations with limited analytical capacity. The paper sets out the rationale for the system, its three-component architecture, and its current development status (Sections 1–7).

Discussion

Many coastal administrations, including a number of IOTC Contracting Parties, collect catch and sampling data that meet reporting requirements but have too few analysts to turn those data into assessments. Where assessments do occur, they are usually one-off, donor-funded projects, so they are rarely repeated at the interval needed to track change in a stock.

Many IOTC stocks, including neritic tunas, billfish, and elasmobranchs, are assessed only with data-limited methods, if they are assessed at all; the main constraint is capacity — too few specialist analysts, and too little funding and time — rather than a lack of suitable methods. FIAS does not propose a new estimator; it separates a deterministic, reproducible statistical backend from a large language model (LLM) layer that drafts interpretive text, subject to a numeric-fidelity check, a mandatory AI-assisted disclosure, and expert sign-off before any submission-grade use. The initial scope is limited to single-species, single-stock fisheries. The system is specified but has not yet been built or tested.

Feedback

The WPM is requested to NOTE the proposed FIAS architecture and to PROVIDE comment on: (i) the data-adequacy criteria proposed for Component I (Section 3.1); (ii) the pathway-selection criteria proposed for Component II, and how outputs should be weighted where more than one pathway is supported (Section 3.2); and (iii) the harvest-control-rule families proposed for Component III (Section 3.3). The WPM is also invited to advise whether the concept, or the results of a future pilot,

¹ The authors submit this paper in an independent professional capacity, at the invitation of IOTC Secretariat. They do not represent a CPC, the Secretariat or any funder, and accordingly do not seek a decision or endorsement from the Working Party. Their aim is technical comment on the design. The authors are presently exploring funding opportunities for this project.

should be presented to the Working Parties on Neritic Tunas, Billfish, or Ecosystems and Bycatch, or to the Scientific Committee.

Abstract

Many stocks in the IOTC area, including neritic tunas, billfish and elasmobranchs, are assessed only with data-limited methods, if they are assessed at all. These methods are applied intermittently by a small pool of specialist analysts. The main constraint is capacity rather than the statistical methods themselves: there are too few analysts, and too little funding and time, to select, apply and defensibly interpret them.

This paper introduces the Fisheries Intelligence and Advisory System (FIAS), a proposed AI-assisted architecture. Its aim is to make existing, accepted data-limited methods accessible to capacity-limited fisheries administrations, and to ensure that those methods can be applied rapidly and efficiently in house.

The methods include CPUE standardisation, LBSPR, and catch-only estimators such as CMSY++ and DCAC. FIAS's contribution lies in applying AI to provide a curated interpretation of the outputs from these methods. It strictly separates the deterministic, auditable statistical backend and a large language model (LLM) interpretation layer. The interpretation layer drafts the scientific narrative, diagnostic commentary, and management advice from the structured statistical output. The drafted text is subject to a numeric-fidelity check, carries a mandatory statement that it is AI-assisted, and must be reviewed and signed off by an expert before any submission-grade use.

The initial scope is limited to single-species, single-stock fisheries. Later iterations may extend the architecture once this first implementation has been built, piloted, and reviewed. This paper invites the Working Party's comment on the proposed data-adequacy criteria, pathway-selection criteria, and harvest-control-rule families, and notes the project's interest in working with IOTC CPCs and the Secretariat as future review or pilot partners.

1. Introduction

IOTC's mandate covers a taxonomically and operationally diverse set of stocks. A persistent constraint is that quantitative assessment coverage lags well behind the number of stocks nominally under management. The gap is widest for the six neritic tuna species considered by WPNT, for billfish, and for many elasmobranch species. It is also wide for the coastal and artisanal fisheries of a number of Contracting and Cooperating Non-Contracting Parties that have limited specialist staff, funding, or time.

This is not new, and the Working Party has already responded to it. WPM and WPNT have turned to data-limited estimators where a full integrated assessment is not feasible. For example, Fu (2024) applied CMSY and CMSY++ to Indo-Pacific king mackerel (*Scomberomorus guttatus*) using catch data alone, because the effort and biological sampling data needed for a conventional assessment were not available. That example illustrates both parts of the underlying problem. The first is a data constraint, which is well recognised. The second is a capacity constraint, which is discussed less often. The people, funding, and time needed to select, configure, and defensibly interpret these methods are limited relative to the number of stocks and administrations that could benefit from

them. The shortage is one of numbers, funding, and time, not of skill. The specialists who are available are capable, but there are too few of them, with too little time, to meet the need.

FIAS is proposed as a response to this second constraint. It is not a new scientific assessment method. Instead, it aims to make already-accepted methods usable by administrations with limited capacity to allow these administrations the opportunity to manage their stocks on their own terms.

This paper introduces the FIAS concept to the Working Party as an initial, single-species implementation. It sets out the proposed architecture and data-adequacy logic and invites methodological comment. It also describes the project's current development status, the iterative pathway by which later stages could extend beyond this initial scope, and the project's interest in working with IOTC should development proceed.

2. The FIAS Concept: Addressing the Capacity Bottleneck through AI-Assisted Interpretation

Applying a data-limited method correctly is only part of the analytical task. A specialist typically spends a comparable amount of time choosing the method that suits the available data, interpreting the diagnostic output, and turning a numerical result into text of the standard expected in a national report or an RFMO working paper. This interpretive work, rather than the statistics, is where the bottleneck lies, and it often falls on the small specialist pool of scientists. The central design proposition of FIAS is that this work is the more tractable target for AI assistance, and one that has so far received less attention. It can be automated safely if it is kept structurally separate from the numbers it describes.

The FIAS design therefore treats the following as mandatory safeguards:

- the statistical backend is deterministic and fully reproducible from its input data alone, independent of the AI model or prompt in use;
- a scripted pipeline runs the analysis, using only the pre-defined tools and methods set out in Table 1, following the fixed data-adequacy and pathway-selection logic (Sections 3.1–3.2); no AI model selects, modifies, or substitutes methods, and no AI model computes any statistic;
- a separate interpretation layer, which has no access to the data or the statistical tools, receives only the resulting structured summary of results (index values, confidence intervals, covariate summaries, and diagnostic outcomes) and generates interpretive text from that summary; it does not perform, and cannot alter, any statistical computation;
- a numeric-fidelity check traces every number appearing in AI-generated text back to a value present in the structured results object; text introducing an unsupported number is rejected and regenerated;
- standard method descriptions and caveats are drawn from a controlled template library rather than generated by the model;
- diagnostic failures and data-quality flags are passed to the AI layer as explicit structured signals that require correspondingly cautious language, rather than being treated as optional context;



- every output carries a mandatory, explicit statement that it is AI-assisted, together with the data, method, and model version used, and every output is subject to expert review before any submission-grade use;
- system performance is formally re-validated whenever the provider updates the underlying AI model, so that a change outside the project's control cannot silently degrade output quality.

Because of this separation, the paper presents FIAS as an interpretation and reporting layer applied to accepted statistical methods. It is not a new statistical assessment method, but does require careful checking and validation of the system outputs.

3. System Architecture

FIAS is structured around three components that follow the sequence of a fisheries assessment: a data-adequacy assessment, a scientific analysis stage, and a management advice stage. Each component produces a usable output on its own, and together they form an end-to-end advisory workflow (Figure 1). The architecture described here is for the initial, single-species implementation discussed in Section 7. Extending it to multispecies or spatially structured fisheries is left to later development stages. This section summarises the elements most relevant to the Working Party's methodological remit.




Figure 1. FIAS information flow. Fishery data pass through Component I (data-adequacy assessment and pathway determination), Component II (one or more analytical pathways, applied according to data adequacy), and Component III (integration, stock-status classification, and harvest-control-rule selection), producing a structured advice object. The AI-assisted interpretation layer drafts narrative text from that structured object only, subject to a numeric-fidelity check, before expert review and sign-off. Successive assessments feed back into Component I as data accumulate.

3.1 Component I: Monitoring Agent

Component I assesses the adequacy of a fishery's data collection programme. This covers observer coverage against relevant RFMO or user-specified thresholds, data quality screening, and the adequacy of biological sampling. Component I then determines which of the Component II analytical pathways the available data can support. It must be completed, and any shortfall documented, before Component II proceeds. The assessment can be re-run as data accumulate, so that successive assessments form a documented time series rather than a set of unconnected outputs.

3.2 Component II: Scientific Analysis Engine

This initial version of FIAS defines four analytical pathways, selected according to data adequacy. More than one pathway may apply to a single fishery, in which case their outputs are integrated in Component III. The set is a starting point and additional pathways, covering other established data-

limited methods or classes of data not addressed in Table 1, are expected to be added as the concept is developed. Any new pathway would be subject to the same data-adequacy gating and interpretation-layer safeguards described in Section 2.

Table 1. Component II analytical pathways defined in this initial version of FIAS, their data-adequacy application criteria, and representative methods.

Pathway	Applied when	Representative methods
A. CPUE standardisation	Catch–effort records carry adequate stratification and temporal coverage to estimate year effects	Regression-based standardisation (response distribution, covariate structure, and treatment of zeros determined by the data) with a full diagnostic suite (residual structure, temporal autocorrelation, index stability, nominal-vs-standardised comparison)
B. Age/length-based methods	Biological length or age composition data of adequate sample size are available	LBSPR (Hordyk et al. 2015, 2016); standard per-recruit SPR calculation; Chapman–Robson total mortality estimate
C. Catch-based methods	A usable effort series is absent or cannot be standardised, but a sufficient catch time series exists	Catch-MSY / CMSY++ and the index-informed Bayesian Schaefer model (Froese et al. 2017); DB-SRA (Dick & MacCall 2011); OCOM (Zhou et al. 2018); DCAC (MacCall 2009); Martell & Froese (2013); index-informed SPiCT ((Pedersen & Berg 2017) or JABBA (Winker et al. 2018) where a partial index can be recovered
D. Qualitative evidence synthesis	Quantitative data are insufficient for the other pathways, or as a supplementary evidence stream	Structured, source-weighted synthesis of interview evidence from fishers, vessel operators, and fisheries officers, consistent with ICES Category 5/6 practice (ICES 2012)

Every method listed in Table 1 is an existing, published, and peer-reviewed approach that is already in use or under active discussion within IOTC working parties. Fu (2024) is a direct precedent for Pathway C applied to an IOTC neritic species. Here, FIAS’s contribution lies in the data-adequacy gating that selects among pathways, and in the interpretation layer described in Section 2. The Working Party’s views on which other established methods would be useful additions to this initial set would be welcome.

3.3 Component III: Management Advice Module

Component III integrates the outputs of Component II across pathways. Where more than one pathway applies, it describes how consistent their signals are rather than averaging conflicting results. Stock status is classified against documented proxy reference points, with uncertainty conditioned explicitly on the data-quality findings from Component I. For example, low or non-representative observer coverage or a short series widens the range placed on an index or SPR estimate, and that wider uncertainty is carried through into the harvest recommendation. The size of any recommended change is derived from an explicit, auditable harvest control rule matched to

the available evidence. Index-based rules of the ICES rfb or 2-over-3 type are used where a standardised or nominal index exists. Where only catch-based estimates are available, a rule conditioned on status relative to an MSY proxy is used instead. The advice is accompanied by a risk assessment covering overfishing risk, monitoring adequacy, and compliance-related risk. Component III produces a first-pass analytical product for national use. Its outputs are intended for expert review and, where appropriate, for submission as a working paper to the relevant IOTC Working Party. They are not IOTC advice. They do not replace Working Party review, the Scientific Committee's advice to the Commission, or the Commission's decisions on management measures. The harvest control rule families described here are proposed for data-limited stocks for which the Commission has not adopted a rule.

4. Relevance to IOTC-Area Fisheries

The four initial pathways map onto recognisable gaps in IOTC-area assessment coverage. Neritic tuna and mackerel fisheries that have a length-frequency sampling programme, but for which full CPUE standardisation is not feasible, are natural candidates for Pathway B. Where appropriate, these results could be supported by Pathway C, following the approach used by Fu (2024). Billfish and elasmobranch bycatch in artisanal fleets, landed by coastal states with limited observer coverage, are candidates for Pathway D. This pathway provides a documented and auditable way of using fisher and observer knowledge that would otherwise go unused. More generally, administrations that can currently only submit data to the Secretariat, and cannot analyse the data themselves, would be able to produce a first-pass scientific narrative. Nominated experts would review this narrative before it was submitted to WPNT, the Working Party on Billfish, the Working Party on Ecosystems and Bycatch, or WPM.

FIAS is intended to complement IOTC's existing capacity-building mechanisms, not to duplicate them. The IOTC–OFCF Project, run by the IOTC Secretariat and Japan's Overseas Fisheries Cooperation Foundation since 2002, strengthens statistical data collection, processing, and reporting in developing coastal states. FIAS addresses the next step. Its purpose is to turn data that are already collected, and that already meet IOTC reporting requirements, into a first-pass scientific interpretation for expert review.

5. Illustrative Output (Hypothetical)

To illustrate the interpretation-layer concept described in Section 2, this section presents a hypothetical example of FIAS output. The scenario, input data (Table 2), and values are entirely hypothetical. They are not derived from IOTC data holdings, do not represent an assessment of any real stock, and should not be cited as such.

Table 2. Simulated input data used in the illustrative example presented in this section.

Simulated input	Value
Catch time series	22 years, species-specific, no directed effort series
Length-frequency sampling	~300 fish/year from port sampling, available in 8 of 22 years
Biological parameters	Regional L_{∞} , K, M, and L50 estimates from published sources (not stock-specific)

In operation, each component, and each method or pathway applied within Component II, would produce its own detailed report. Each report would include full diagnostics, a description of uncertainty, and the structured results object described in Section 2. For brevity, only the main conclusion from each component is shown here. These are followed by the combined narrative that the Component III interpretation layer would generate from them.

Illustrative AI-generated interpretation excerpt (simulated, not a real assessment):

The excerpt below shows how the interpretation layer would combine the three structured outputs into a single narrative for a scientific report:

- **Component I, Monitoring Agent (data-adequacy output):** Length-frequency sampling was adequate in five of the eight sampled years. It fell below the sample-size threshold in the remaining three years, and this is flagged as a data-quality caveat. There is no usable effort series, so Pathway A (CPUE standardisation) is not supported. Pathway B (LBSPR) is supported as the primary pathway, and Pathway C (CMSY++) is supported as a supplementary, catch-only cross-check.
- **Component II, Scientific Analysis Engine (statistical output):** Pathway B (LBSPR) returns a selectivity-corrected SPR of 0.34 (bootstrap 95% CI: 0.21–0.48). This uses life-history ratios from a regional M/K proxy rather than stock-specific estimates. Pathway C (CMSY++, supplementary) estimates current biomass at 41–68% of BMSY, with much wider uncertainty than the SPR estimate.
- **Component III, Management Advice Module (integrated advice output):** The SPR estimate places the stock in the uncertain-status band (proxy reference range 0.20–0.40). The CMSY++ result is broadly consistent, but it is given less weight because of its wider uncertainty. The Component I data-quality flag widens the confidence bounds carried into this classification. The recommended action is a precautionary management response, together with continued or expanded length-frequency sampling to reduce uncertainty in future assessments.
- **Interpretation layer**

“Length-composition data from the eight sampled years indicate a selectivity-corrected SPR of 0.34 (bootstrap 95% CI: 0.21–0.48). This places the stock in the uncertain-status band (SPR 0.20–0.40) under the proxy reference points applied in this analysis. The estimate should be treated with caution. Sample sizes in three of the eight years fall below the adequacy threshold applied in Component I, and the assumed M/K ratio is taken from a regional life-history proxy rather than from species- and stock-specific estimates. The supplementary catch-based analysis (CMSY++) is broadly consistent with this result, estimating current biomass at 41–68% of BMSY. That estimate is much more uncertain, however, and is given less weight because a length-based signal is available. A precautionary management response is recommended, together with continued or expanded length-frequency sampling to narrow the SPR confidence interval in future assessments.”

The interpretive style, the structured statement of uncertainty, and the explicit link between data adequacy and precaution shown in this excerpt are the intended final product. No such analysis has

yet been run on real data and the example here is illustrative only and are not outputs of a working system.

6. Current Status and Development Pathway

FIAS has been specified in detail but has not yet been built or tested. Subject to funding, development would be a staged, iterative programme built around the initial single-species scope set out in Section 7. The programme has three stages.

- The first is a foundation stage, in which a prototype is built for the single-species, single-stock case and reviewed by independent experts before any pilot use.
- The second is a pilot stage, in which the single-species system is tested with a small number of volunteer administrations against a pre-agreed evaluation framework. This framework would include simulation-based recovery testing against known operating models and defined thresholds for agreement with expert judgement.
- The third is a review stage, which would determine whether and how the system should be extended. Depending on the findings of that review, this could include multispecies, spatially structured, or other more complex fisheries beyond the initial scope.

The authors are two individuals and have developed the concept and specification without external funding. Building and testing the system needs development funding, technical collaborators (stock assessment scientists to review the method implementations and test them against simulated and real data, plus software and data science specialists), and administrations willing to pilot the system on their own data. The pilot is what shows whether the approach works in practice. Without it the system could only be tested on simulated and historical data. The authors welcome contact from WPM participants and CPC administrations interested in any of these roles. Development funding is currently being explored.

7. Discussion and Conclusion

The initial scope of FIAS is deliberately focused and limited. Although future iterations can build on these ideas, the system is designed to be implemented first for single-species or single-stock fisheries in which catch, effort, or biological data can be attributed to a defined stock. The underlying methods carry the usual assumptions of a single stock initially at or near equilibrium, with approximately constant recruitment. These assumptions are poorly met in mixed multispecies fisheries and species where recruitment varies strongly. Spatiotemporal modelling is excluded from the initial implementation by design, not because of a technical limitation of the platform, but these analyses require considerably greater initial investment in standardised method development. Spatiotemporal modelling, together with multispecies and spatially structured extensions, is a candidate direction for later development once the initial single-species implementation has been built, piloted, and reviewed.

The specification identifies two risks explicitly. First, the interpretation layer depends on a single LLM API. The architecture abstracts this interface and documents a substitution process, but the dependency itself remains. Second, the qualitative evidence pathway (Pathway D) is, by its nature, the hardest of the initial four analysis pathways to validate against ground truth. It is proposed to test this pathway by, for example, simulation and independent expert concordance, rather than assuming it is reliable.

The Working Party's methodological advice would be valuable on three points:

- i. whether the categories of data-adequacy check proposed for Component I — observer coverage, data-quality screening, and biological-sampling adequacy (Section 3.1) — are an appropriate basis for gating pathway selection;
- ii. whether the pathway-selection criteria set out in Table 1 (Section 3.2) are consistent with the standards WPM applies when judging whether a data-limited method is fit for use, and how outputs should be weighted where more than one pathway is supported for the same fishery; and
- iii. whether the choice of index-based (ICES rfb or 2-over-3 type) and catch-based (MSY-proxy-conditioned) harvest-control-rule families proposed for Component III (Section 3.3) are reasonable starting points for data-limited stocks of the kind assessed in the IOTC area.

FIAS is specified but not yet built or tested, and the engagement opportunities described depend on development funding being secured. The Working Party's advice on whether this concept, or the results of a future pilot, would also be of interest to the Working Party on Neritic Tunas, the Working Party on Billfish, the Working Party on Ecosystems and Bycatch, or the Scientific Committee would help direct the work to the appropriate bodies.

Declaration on the Use of AI

During the preparation of this work, the authors used Claude (Anthropic; Claude Opus 5.5) to help draft and structure the manuscript text and to format the manuscript to the IOTC working paper template. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of this publication.

References

- Dick, E.J.; MacCall, A.D. (2011). Depletion-based stock reduction analysis: A catch-based method for determining sustainable yields for data-poor fish stocks. *Fisheries Research* 110, 331–341.
<https://doi.org/10.1016/j.fishres.2011.05.007>
- Froese, R.; Demirel, N.; Coro, G.; Kleisner, K.M.; Winker, H. (2017). Estimating fisheries reference points from catch and resilience. *Fish and Fisheries* 18, 506–526. <https://doi.org/10.1111/faf.12190>
- Fu, D. (2024). Assessment of Indian Ocean Indo-Pacific king mackerel (*Scomberomorus guttatus*) using data-limited methods. IOTC Working Party on Neritic Tunas working paper IOTC-2024-WPNT14-19. Indian Ocean Tuna Commission.
- Hordyk, A.; Ono, K.; Valencia, S.; Loneragan, N.; Prince, J. (2015). A novel length-based empirical estimation method of spawning potential ratio (SPR), and tests of its performance, for small-scale, data-poor fisheries. *ICES Journal of Marine Science* 72, 217–231.
<https://academic.oup.com/icesjms/article/72/1/217/815772>.
<https://doi.org/10.1093/icesjms/fsu004>
- Hordyk, A.R.; Ono, K.; Prince, J.D.; Walters, C.J. (2016). A simple length-structured model based on life history ratios and incorporating size-dependent selectivity: Application to spawning potential ratios for data-poor stocks. *Canadian Journal of Fisheries and Aquatic Sciences* 73, 1787–1799.
<https://doi.org/10.1139/cjfas-2015-0422>

ICES (2012). ICES implementation of advice for data-limited stocks in 2012 in its 2012 advice. ICES CM 2012/ACOM 68. International Council for the Exploration of the Sea, 42 p.

MacCall, A.D. (2009). Depletion-corrected average catch: A simple formula for estimating sustainable yields in data-poor situations. ICES Journal of Marine Science 66, 2267–2271.
<https://doi.org/10.1093/icesjms/fsp209>

Martell, S.; Froese, R. (2013). A simple method for estimating MSY from catch and resilience. Fish Fish 14, 504–514. <https://doi.org/10.1111/j.1467-2979.2012.00485.x>

Pedersen, M.W.; Berg, C.W. (2017). A stochastic surplus production model in continuous time. Fish and Fisheries 18, 226–243. <https://doi.org/10.1111/faf.12174>

Winker, H.; Carvalho, F.; Kapur, M. (2018). JABBA: Just Another Bayesian Biomass Assessment. Fisheries Research 204. <https://doi.org/10.1016/j.fishres.2018.03.010>

Zhou, S.; Punt, A.E.; Smith, A.D.M.; Ye, Y.; Haddon, M.; Dichmont, C.M.; Smith, D.C. (2018). An optimized catch-only assessment method for data poor fisheries. ICES Journal of Marine Science. <https://doi.org/10.1093/icesjms/fsx226>