



Fisheries and Oceans  
Canada

Pêches et Océans  
Canada

Science

Sciences

## **CSAS**

**Canadian Science Advisory Secretariat**

**Research Document 2020/XXX**

Not to be cited without  
permission of the author\*

A framework for the assessment of snow crab (*Chionoecetes opilio*) in Maritimes Region (NAFO Div 4VWX)

## **SCCS**

**Secrétariat canadien de consultation scientifique**

**Document de recherche 2020/XXX**

Ne pas citer sans  
l'autorisation des auteurs\*

Un cadre évaluation de stock du crabe des neiges (*Chionoecetes opilio*) du plateau continental de la Nouvelle-Écosse (NAFO Div 4VWX)

Choi, Jae S.

Population Ecology Division  
Department of Fisheries and Oceans  
Bedford Institute of Oceanography  
P.O. Box 1006, Dartmouth, Nova Scotia B2Y 4A2

\*This series documents the scientific basis for the evaluation of fisheries resources in Canada. As such, it addresses the issues of the day in the time frames required and the documents it contains are not intended as definitive statements on the subjects addressed but rather as progress reports on ongoing investigations.

Research documents are produced in the official language in which they are provided to the Secretariat.

This document is available on the Internet at:

<http://www.dfo-mpo.gc.ca/csas/>

\*La présente série documente les bases scientifiques des évaluations des ressources halieutiques du Canada. Elle traite des problèmes courants selon les échéanciers dictés. Les documents qu'elle contient ne doivent pas être considérés comme des énoncés définitifs sur les sujets traités, mais plutôt comme des rapports d'étape sur les études en cours.

Les documents de recherche sont publiés dans la langue officielle utilisée dans le manuscrit envoyé au secrétariat.

Ce document est disponible sur l'Internet à:

ISSN 1499-3848 (Printed / Imprimé)  
©Her Majesty the Queen in Right of Canada, 2020  
©Sa Majesté la Reine, Chef du Canada, 2020

**Canada**



## **Abstract**

This is a review of spatiotemporal modeling in the context of stock assessment of snow crab in Maritimes Region. We discuss the rationale for using autocorrelation function and conditional autoregressive areal unit models. We demonstrate the utility of the latter to estimate abundance in sampling programs prone to incomplete sampling or logistical failure. These considerations suggest significant utility in basing future assessments on areal unit models. They also suggest that a systematic application of such methods across all species would provide meaningful comparisons of interspecific use of space and ecosystem requirements/sensitivities. They are also easily adapted to prognostic climate change scenarios. These are simple but necessary steps towards a truly ecosystem-based management of fisheries.

## Résumé

To be completed

---

## Contents

|       |  |    |
|-------|--|----|
| 1     | Introduction . . . . .                                     | 3  |
| 1.1   | Snow crab biology . . . . .                                | 4  |
| 1.2   | Data collection . . . . .                                  | 6  |
| 1.3   | Data analysis . . . . .                                    | 7  |
| 1.3.1 | Spatiotemporal aggregation . . . . .                       | 7  |
| 1.3.2 | Synthesis . . . . .  | 11 |
| 2     | Abundance index estimation (data aggregation) . . . . .    | 13 |
| 2.1   | Ad hoc smoothing models . . . . .                          | 14 |
| 2.2   | Static models . . . . .                                    | 14 |
| 2.3   | Continuous autocorrelation models . . . . .                | 15 |
| 2.3.1 | Spatial autocorrelation . . . . .                          | 15 |
| 2.3.2 | Temporal autocorrelation . . . . .                         | 17 |
| 2.3.3 | Spectral estimation of autocorrelation . . . . .           | 17 |
| 2.3.4 | Spatial model . . . . .                                    | 19 |
| 2.3.5 | Hierarchical spatial models . . . . .                      | 20 |
| 2.3.6 | Stationary spatiotemporal models . . . . .                 | 21 |
| 2.3.7 | Complex spatiotemporal models . . . . .                    | 23 |
| 2.4   | Hybrid Cartesian-Lagrangian models . . . . .               | 25 |
| 2.4.1 | Spatial Conditional autoregressive models . . . . .        | 26 |
| 2.4.2 | Spatiotemporal Conditional autoregressive models . . . . . | 28 |
| 3     | Comparison of results . . . . .                            | 29 |
| 3.1   | Non-spatial, non-temporal models . . . . .                 | 29 |
| 3.2   | Continuous autocorrelation models . . . . .                | 31 |
| 3.3   | Conditional autoregressive models . . . . .                | 38 |
| 3.3.1 | Poisson-CAR model . . . . .                                | 39 |

---

|       |   |    |
|-------|---|----|
| 3.3.2 | Ecosystem variability of static variables . . . . .     | 39 |
| 3.3.3 | Ecosystem variability of dynamic variables . . . . .    | 43 |
| 4     | Synthesis . . . . .                                     | 58 |
| 5     | Concluding thoughts . . . . .                           | 58 |
| 5.1   | Coda: Habitat (“species distribution”) models . . . . . | 58 |
| 6     | References . . . . .                                    | 66 |
| 7     | Appendix . . . . .                                      | 71 |
| 7.1   | Sampling bias . . . . .                                 | 71 |
| 7.2   | Residual plots . . . . .                                | 71 |

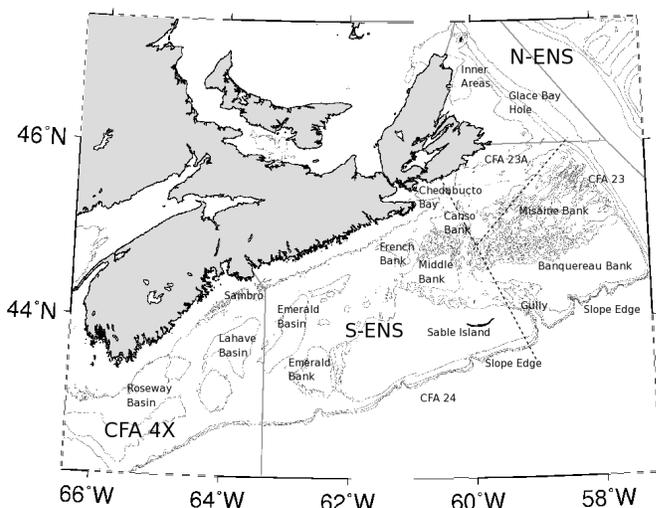


Fig. 1. Snow crab fishing areas in Maritimes region.

Version: 19 March 2020

## 1 Introduction

The goal of this document is to provide a framework for snow crab stock assessment resident in the Maritimes Region of Atlantic Canada (NAFO Statistical Division 4VWX; Figure 1; herein, *snow crab assessment*). Most formal stock assessments share the following key elements:

- data collection
  - surveys for numbers, weight, kind, environmental variables
  - quality assurance and quality control (QA/QC)
  - assimilation into long term storage
- data analysis
  - *spatiotemporal aggregation* to some fixed area to estimate temporal trends using experimental (survey) design
  - *synthesis* of fishery performance, survey indices, ecosystem variability and biological constraints to provide projections and risk analysis

More specifically, the *data analysis (spatiotemporal aggregation)* step, above, has presented a challenge in previous snow crab assessments. We will, in the following, focus upon this element by describing some of these difficulties and propose a solution going forward that is both robust and operational.

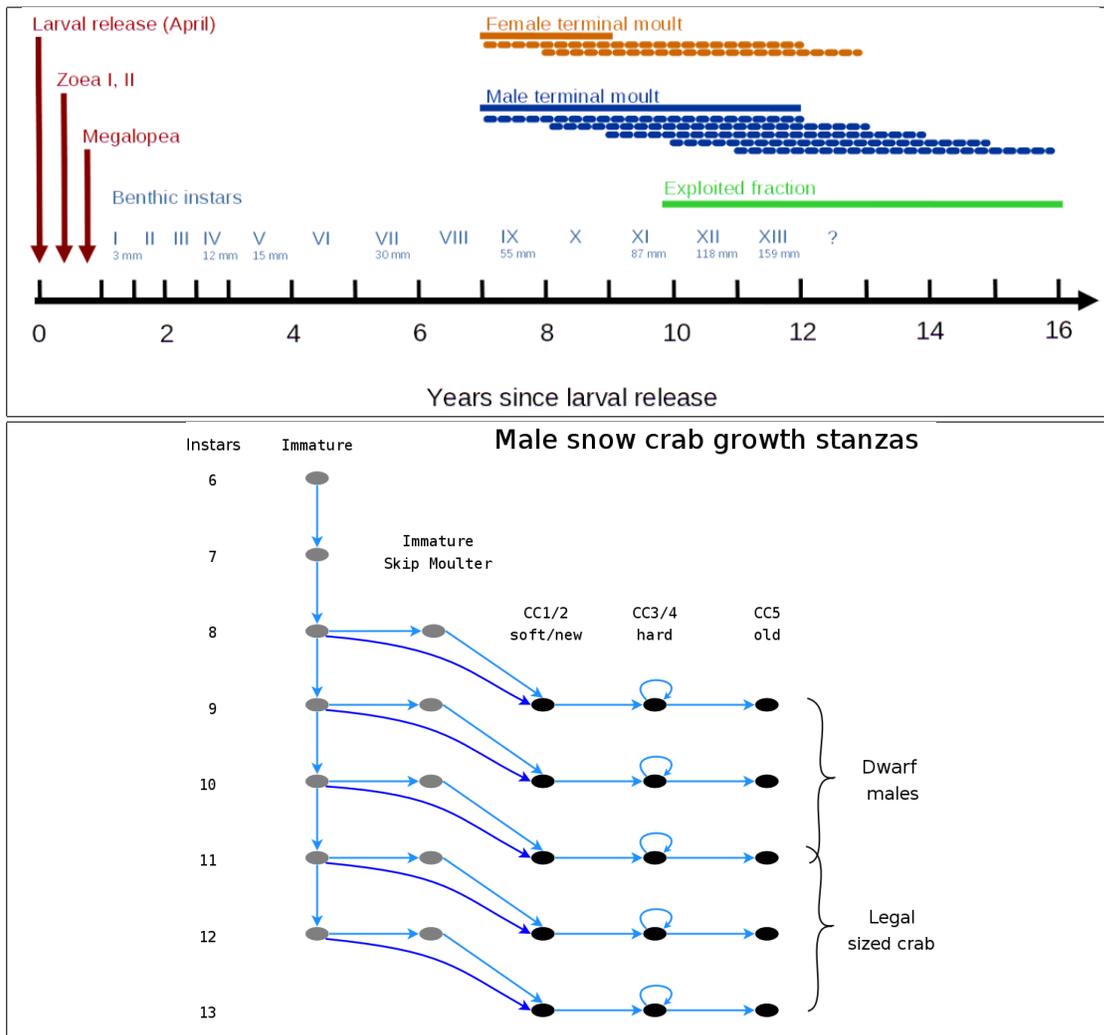
## 1.1 Snow crab biology

Snow crab have a complex life history. This complexity causes spatiotemporal complexity in their abundance distributions and so is the primary cause of the analytical difficulties of *spatiotemporal aggregation*. They are, first and foremost, cold-water *stenotherms*. In the Maritimes region, commercially fished snow crab are generally observed between depths of 50 to 300 m and between temperatures of -1 to 7°C. Snow crab are thought to avoid temperatures above 7°C, as metabolic costs have been shown to match metabolic gains near this temperature (Foyle et al. 1989). This means that their spatial distributions can fluctuate seasonally and annually due to the thermally complex nature of the Maritimes Region. Smaller crab and females also have differences in thermal *preferenda*. Further, snow crab are generally observed on soft mud bottoms. Small-bodied and moulting crabs are also found on more complex (boulder, cobble) substrates presumably as they afford more shelter (Sainte-Marie and Hazel 1992; Comeau et al. 1998). There are of course many other factors that regulate their distributions and local densities (discussed below). The end result of these interactions is that their environmental niche will vary depending upon the time of year and location.

Snow crab eggs are brooded by their mothers for up to two years. This also depends directly upon ambient temperatures as well as food availability (which also depends upon temperature via primary and secondary production) and the health condition and maturity status of the mother (up to 27 months in primiparous females – first breeding event; and up to 24 months in multiparous females – second or possibly third breeding events; Sainte-Marie 1993). More rapid development of eggs (from 12 to 18 months) has been observed in other systems (Elner and Beninger 1995; Webb et al. 2007). Over 80% of the female snow crab on the Scotian Shelf are estimated to follow an annual, rather than the bi-annual cycle possibly due to exposure to warmer temperatures (Kuhn and Choi 2011). A primiparous female of approximately 57.4 mm carapace width produces between 35,000 to 46,000 eggs, which are extruded between February and April (Sainte-Marie 1993). Multiparous females are thought to be more fecund, with more than 100,000 eggs being produced by each female. Eggs are hatched from April to June when larvae are released. This pelagic larval stage lasts for three to five months (zoea stages 1 and 2) during which snow crab feed upon zooplankton. Pelagic stages seem to have highest abundance in October and so may begin settling in January. Thereafter, they settle closer to the bottom in their megalopea stage. Very little is known of survival rates at these early life stages, but these also are thought to depend highly upon temperature.

Once settled to the bottom (benthic phase), snow crab moult approximately twice a year (Sainte-Marie et al. 1995; Comeau et al. 1998; Figure 2), the frequency and size increments of the moult also depending upon temperature and food availability. The first inter-moult stage (instar 1) is approximately 3 mm carapace width. After the 5th instar (15 mm carapace width), the moulting frequency declines to annual spring moults until they reach a terminal maturity moult. Growth is allometric with weight increasing approximately 250% with each moult. Terminal moult has been observed to occur between the 9th to the 13th instar in males and the 9th to 10th instar in females. Just prior to the terminal moult, male crab may skip a moult in one year to moult in the next (Conan et al. 1992). Male snow crab generally reach legal size ( $\geq 95$  mm carapace width) by the 12th instar; however, a variable fraction of instar 11 snow crab are also within legal size. Male instar 12 snow crab represent an age of approximately nine years since settlement to the bottom and 11 years since egg extrusion. Thereafter, the life expectancy of a male is approximately five to six years. Up to ten months are required for the shell to harden, and up to one year for meat yields to be commercially viable. After hardening of the carapace, the male is able to mate. Near the end of the lifespan of a snow crab, the shell decalcifies and softens, and may have heavy epibiont growth.

Females reproducing for the first time (primiparous females) generally begin their moult to maturity at an average size of 60 mm carapace width and mate while their carapace is still soft (early spring: prior to the fishing season in ENS, and during the fishing season in 4X). A second mating period later in the year (May to June) has also been observed for multiparous females (Hooper 1986). Pair formation, a mating embrace where the male holds the female, may occur up to three weeks prior to the mating event. During mating, complex behavioral patterns have also been observed: the male helps the primiparous female moult, protects her from



**Fig. 2.** Life history patterns of snow crab. Top: the timing of the main life history stages of snow crab. Bottom, the growth stanzas of the male component and "decision paths" to maturity and terminal moult.

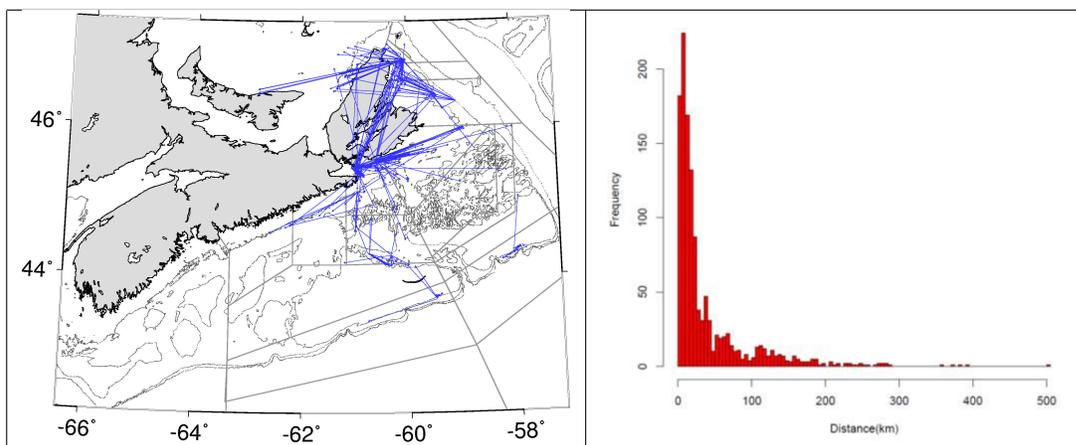


Fig. 3. Snow crab movement and distances travelled.

other males and predators and even feeds her indirectly. Upon larval release, males have been seen to wave the females about to help disperse the larvae (i.e., prior to a multiparous mating). Females are selective in their mate choice, as is often the case in sexually dimorphic species, and they have been seen to die in the process of resisting mating attempts from unsolicited males (Watson 1972; Hooper 1986). Males compete heavily for females and often injure themselves (losing appendages) while contesting over a female. Larger males with larger chela are generally more successful in mating and protecting females from harm.

Snow crab have a complex life history; are environmentally sensitive with ontogenetic variations in this sensitivity; long-lived and exposed to ecosystem variability (biotic and abiotic) that operate at a range of spatiotemporal scales; and an ability to move large distances (on average < 10's of km but up to 100's of km a year) for foraging and reproduction and complex and wide ranging interspecific interactions (Figure 3). This make static spatial aggregation techniques nonviable (see below, Section 1.3.1).

## 1.2 Data collection

The first records of snow crab activity in eastern Canada began in 1960 with incidental by-catches from ground-fish draggers near Gaspé, Quebec. In Maritimes Region, the snow crab fishery has been active since the mid-1970s. A scientific assessment of the status of snow crab in the region began in the mid-1990's on the fishing vessel *Marco Michel* (Biron et al. 1997). This was funded by industry and developed by the Gulf Fisheries Centre (GFC; Moncton, New Brunswick). The assessment mandate was moved to the Bedford Institute of Oceanography (BIO; Dartmouth, Nova Scotia) in 2004 (Choi et al. 2005b) and continues to be industry funded to the present day. The industry see this as an investment towards the long term sustainability of the resource. Prior to 2004, surveys were undertaken from April to July, however, this transitioned to the September to December period, primarily to reduce overlap with the fishing season (April-September; with the exception of area CFA4X in October to March). From 2004 to 2013, surveys were conducted on the fishing vessel, *The Gentle Lady* until its sinking in December 2013. The surveys from 2014 to the present were conducted on the fishing vessel *Ms. Jessie*. Throughout these periods of transition, every effort was made to ensure vessel characteristics, crew and sampling protocol were as consistent as possible. But of course, vessel-dependent bias is possible and likely.

A *Bigouden Nephrops* trawl, a net originally designed to dig into soft sediments for the capture of European lobsters, was used to sample the snow crab and other benthic/demersal fauna (Biron et al. 1997). It has a

headline of 20 m, with a mesh size of 80 mm in the wings and 60 mm in the belly and 40 mm in the cod-end. Net configuration varies with local conditions and so it was recorded with wireless trawl monitoring sensors; depth and temperature were recorded with wireless sensors; and positional information was recorded with global positioning systems. Actual duration of bottom contact was assessed from the trawl monitoring and temperature/depth profiles, with a nominal target of a 5 minute tow duration. Swept area of the net was computed from swept distance and monitored net width. The ship speed was maintained at approximately two knots. The warp length was approximately three times the depth. Every effort has been made to make this protocol consistent, however, there has been an evolution of the net structure to make performance more precise and accurate (chafing gear, usage of floats and weights to keep the net structure more consistent, angling of the doors, etc.) and improved electronics precision and accuracy. Some bias can be introduced here, however, the improvement in accuracy and precision was considered worthwhile.

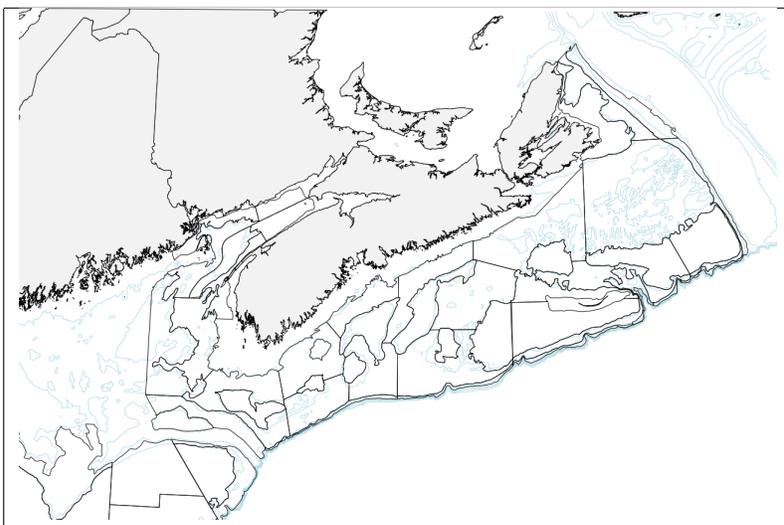
All crab were enumerated; measured with calipers; shell condition determined; claw hardness measured with a durometer; and weighed with motion-compensated scales. The latter permitted direct biomass measurement rather than relying upon allometric relationships between body parts to approximate biomass (the latter was the approach prior to 2004). The maturity status of males was determined from a combination of biological staging and morphometric analysis. While physiological maturity is not directly co-incident with the onset of morphometric maturity, the latter is more readily determined. This is determined by a disproportionate increase of chela height relative to carapace width in males. The maturity status of females is assessed from direct visual inspection of eggs or gonad development. Where maturity status was ambiguous, it was determined morphometrically, as the width of abdomen (measured by the width of the fifth abdominal segment) increases rapidly relative to carapace width at the onset of morphometric maturity, facilitating the brooding of eggs. Other bycatch were identified and numerical abundance and biomass assessed, with subsampling where required. Morphometrics were not included for non-targetted species due to time/personnel constraints at sea.

In the snow crab assessment, an open-ended, open sourced and seamless solution has evolved that assimilates the information from electronic and paper records into modern data systems. Survey data are entered directly into the At-sea-observer relational database system, held at DFO, BIO (Bedford Institute of Oceanography, Dartmouth, Nova Scotia) with double key punching of data and database range limitations to reduce input errors. Further post-processing through more finely tuned QA/QC processes are encoded in an extensible and open-sourced R-package (<https://github.com/jae0/bio.snowcrab>) such as the identification of outliers and extreme data for manual follow-up and correction in the storage data base. Data derived from electronic sensors are stored locally and post-processed for QA/QC in R and prepared for assimilation into the assessment process. Some of this data gets shared with other assessments (e.g., temperatures). All data processing, imputation and further analyses are conducted at local workstations using the above R-library. Data collection methods are re-evaluated continuously and open to evolve; however, changes in survey protocol over time have now normalized.

## 1.3 Data analysis

### 1.3.1 Spatiotemporal aggregation

The core challenge in stock assessments is tuning the limited resources and time required to sample a very large spatiotemporal domain, often with poorly defined spatial bounds and imperfect tools prone to wear and tear, in order to obtain a representation of an unobservable but real (*latent*) process(es) which are able to inform upon the status of a species! In the snow crab survey, there is a high station density relative most other directed surveys, with about 400 stations being visited annually, and taking up to four months to complete. The spatial domain of the snow crab assessment is approximately  $109,120 \text{ km}^2$  which means that each station, on average, is representing  $273 \text{ km}^2$ ! The problem is that each station samples only  $0.0039 \text{ km}^2$ ; a ratio of 1:70,000. Similarly, in terms of time, a station is sampled for on average about 5 minutes but it is supposed to represent a whole year (525,600 minutes); a ratio of 1:105,120. In terms of space-time units, this is a ratio of  $1 : 7.36 \times 10^9$ , or 1



**Fig. 4.** An example of the areal units used in the Atlantic Maritimes Region of eastern Canada (black lines) for the "Ecosystem survey" (formerly known as Ground fish survey) sampling for the purposes of stock abundance assessment of many species. Overlaid are isobaths (light blue) that show that most of these areal units actually follow simple bathymetric contours.

part in 10 billion! If the ocean bottom environment were "well mixed" (homogenous chemostat) then just a few samples would suffice. Unfortunately it is not, and so this rather sparse sampling of the ocean forces us to make assumptions, and the estimation of any feature of interest becomes a balance between the information gained by sampling intensity, that is, precise (low variability) and accurate (unbiased) estimates *vs.* the costs in resources and time associated with obtaining the information.

In fisheries assessments, a survey's *experimental design* encodes these assumptions and as such are used *religiously* as a basis for aggregation of samples across space. Advice on stock status is usually demanded on an annual basis for some arbitrarily determined administrative spatial areas (herein, "*areal units*" or *AUs*). These AUs are usually associated with historical precedents relating to access and allocation and/or attempts to rationalize sampling based upon some additional information. Sampling strategies of such AUs can range from completely random sampling in the absence of additional information, to some form of stratified random design (see also, Appendix 1). In the latter, samples are chosen randomly from AUs that are characterized, *a priori*, by *informative factors*. An example for groundfish is shown in Figure 4 where it is clear that depths were the main consideration. Analysis of (co)variance (AN(C)OVA) is a common application of such stratification, numerically *blocking* variability of such *informative factors* that are not the focus of the study (in an attempt to be more cost-effective). The lower the variability within AUs (relative to between-AU variability) of the focal variable, the more "successfully" stratification has controlled for these "nuisance" effects.

The hope is that a sample within an AU is representative of the AU (*i.e.*, "well mixed" inside AUs). If such "nuisance" factors were slow changing (e.g., depth changes due to tectonic processes) or very fast changing (e.g., biochemical nutrient levels due to bacterial re-mineralization) relative to the time scale of the biological feature of interest (abundance change due to population dynamics), then this may be a useful assumption and approximation. The problem of course is that these "nuisance" factors change at space-time scales that are similar to the feature of interest which makes samples not always representative of the area that they are supposed to represent. This can cause poor precision (elevated variability) and poor accuracy (bias). As the number of "nuisance" factors increase, the number of AUs required to adequately "block out" such factors, statistically in the sense of classical ANOVA, increases to non-viable levels (in terms of required time and resources). This increase is exponential in terms of covariates, as more than one sample is required in each covariate block and in the case of two spatial dimensions, the number of units increases by a factor of two relative to any linear

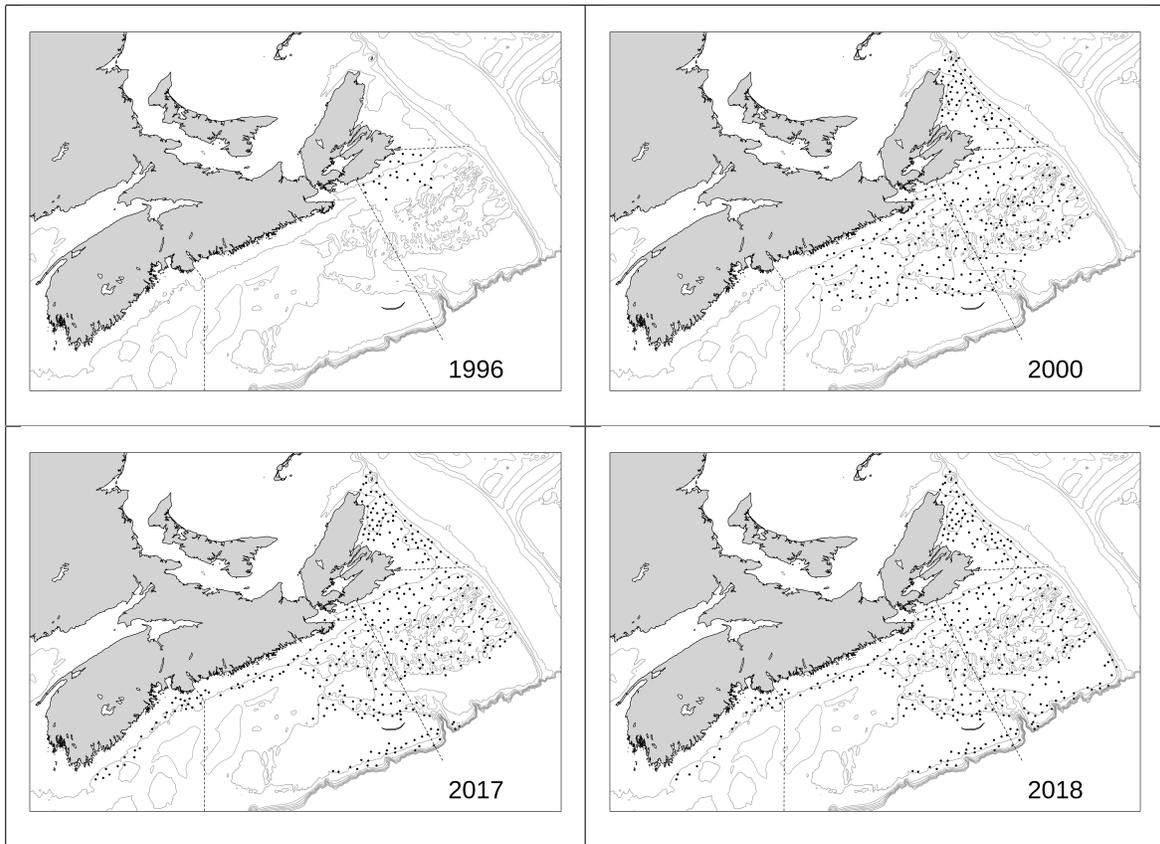
reduction of scale and with the addition of time, by an additional factor or more depending upon if seasonal and or diurnal discrimination is required.

For example, when the informative features are dynamic, their relevance in static AUs can cause a mismatch with the presumed factors of importance. That is, there can be spatial and/or temporal *aliasing* (sometimes heuristically referred to as "upscaling" and "downscaling" issues) in that a sample is taken at a very different temperature (for example) or time of year than the overall average temperature of the areal unit that it is supposed to represent. In the snow crab survey, the sampling period spans over four months during which environmental conditions on the surface range from summer conditions to winter conditions. This blocking or *factorial* approach, only crudely blocks out the influence of a handful of these extraneous factors (in the example, only depth) and otherwise ignores them as *nuisance* factors, as they are usually not even measured. Herein, we will refer to this approach of applying a *static* experimental design as a "*Cartesian*" perspective. This *Cartesian perspective* dominates in stock assessments and fisheries-related literature in general, and troubling for reasons we will see below.

In reality, these "nuisance" factors are, actually highly informative and facilitate understanding the focal process(es) of interest. If carefully measured and treated, they can support more precise and accurate predictions. Recall, that in the case of snow crab, sampling is 1 part in 10 billion. Additional information helps to improve the scale of this sampling. Further, when they operate on spatiotemporal scales similar to those of the focal process(es) they can become highly influential such that ignoring them can cause poor precision and accuracy. As such, in the snow crab assessment, we embrace these factors rather than try to ignore all the variability in the processes that influence biota; herein we will refer to this as *ecosystem variability*. Examples of *ecosystem variability* in the marine context include the interactions of organisms with variations in ambient temperatures, light levels, pH, redox potential, salinity, food/nutrient availability, shelter space, predator abundance, disease prevalence, etc. Characterizing *ecosystem variability* is indeed difficult and time-consuming due to their numerous and interacting nature; however, as they, by definition, directly and indirectly influence a population's status, they cannot and should not be ignored. This is especially the case if this information pre-exists or is cheaper to sample extensively than the population density.

Indeed, this *ecosystem variability* induces something very important: complex spatiotemporal autocorrelations in the abundance of the organism of interest. As stated by *Tobler's First law of geography*: "*everything is related to everything else, but near things are more related than distant things*" (Tobler 1970). If it were not the case, then everywhere we look would be completely chaotic and discontinuous, without localized areas of homogeneity. Similar arguments can be made for time. As such, experimental design must pay careful attention to *ecosystem variability* and the spatiotemporal autocorrelation that they induce. If the strategy is to "block out" nuisance factors, then each such AU must also be independent of each other, that is, without spatial (and temporal) autocorrelation. If a survey's experimental design is inadequate to guarantee such independence between AUs, then an appropriate spatiotemporal model can be used to attempt to rectify these biases and begin the process of (1) iteratively improving the experimental design and/or (2) improving the data collection required to analytically correct any biases induced by *ecosystem variability*. This point of view, we will herein refer to as a *Lagrangian* perspective, one that perceives AUs as being more fluid in definition than the static *Cartesian* view identified previously.

In hydrodynamically complex areas such as NAFO Statistical Division 4VWX (Maritimes Region of Atlantic Canada) where a number of oceanic currents converge and there has also been ongoing rapid climate change and historically significant and rapid ecological change (Choi et al. 2005a), the influence of *ecosystem variability* can not be safely ignored. Snow crab surveys, initially in the mid-1990s were exploratory and occurred in near-shore areas of known fishing activity (Figure 5). Surveys rapidly expanded to shelf-wide scales when it became evident that their spatial distribution was much larger and more heterogeneous than previously understood. This expansion of the survey only stabilized in the early 2000s. As such, spatial bias exists, precluding any simple analysis or naive aggregation scheme. Additional bias is also possible with the use of a new fishing vessel in 2004 and again in 2014. Further, in the southern-most area of snow crab distribution (Crab fishing area 4X), trawl survey coverage has been historically sporadic but have stabilized since 2004 (with the exception of 2014).



**Fig. 5.** Variability of sampling locations in the snow crab survey.

Historically, the statistical approach used to aggregate/integrate sampled observations was "kriging" (Biron et al. 1997), a method of data interpolation that borrows from the information found in the manner in which variability changes with distance from a given location (spatial covariance/autocorrelation; Cressie 1993). The optimal survey design for this approach is generally considered to be a uniform grid. To this end, the snow crab survey locations were randomly chosen locations in a  $10 \times 10$  minute areal grid (Biron et al. 1997). The rationale for the choice grid size is not clear as 10 minutes in longitude is not the same distance as 10 minutes in latitude, and also the distance of 10 minutes in longitude varies with latitude. This renders each sampling grid unequal in surface area and therefore an unequal influence upon the integration in space. Sampling within each such "grid" is/was also pseudo-random at best. This is because many locations are not directly accessible to trawls. The substrate in Maritimes Region is notoriously rugged, ragged, rocky with deep trenches in many areas; features that do not permit fishing nets to pass unharmed or force the nets to become part of the substrate. There was, therefore, bias introduced in bottom type selection, one that preferred "trawlable" locations, these locations also tend to be uniform depths, softer, gravel or mud substrates that coincide with preferred snow habitats. This resulted in preferential sampling of deeper locations, colder temperatures, and very different species associations (Figure 6).

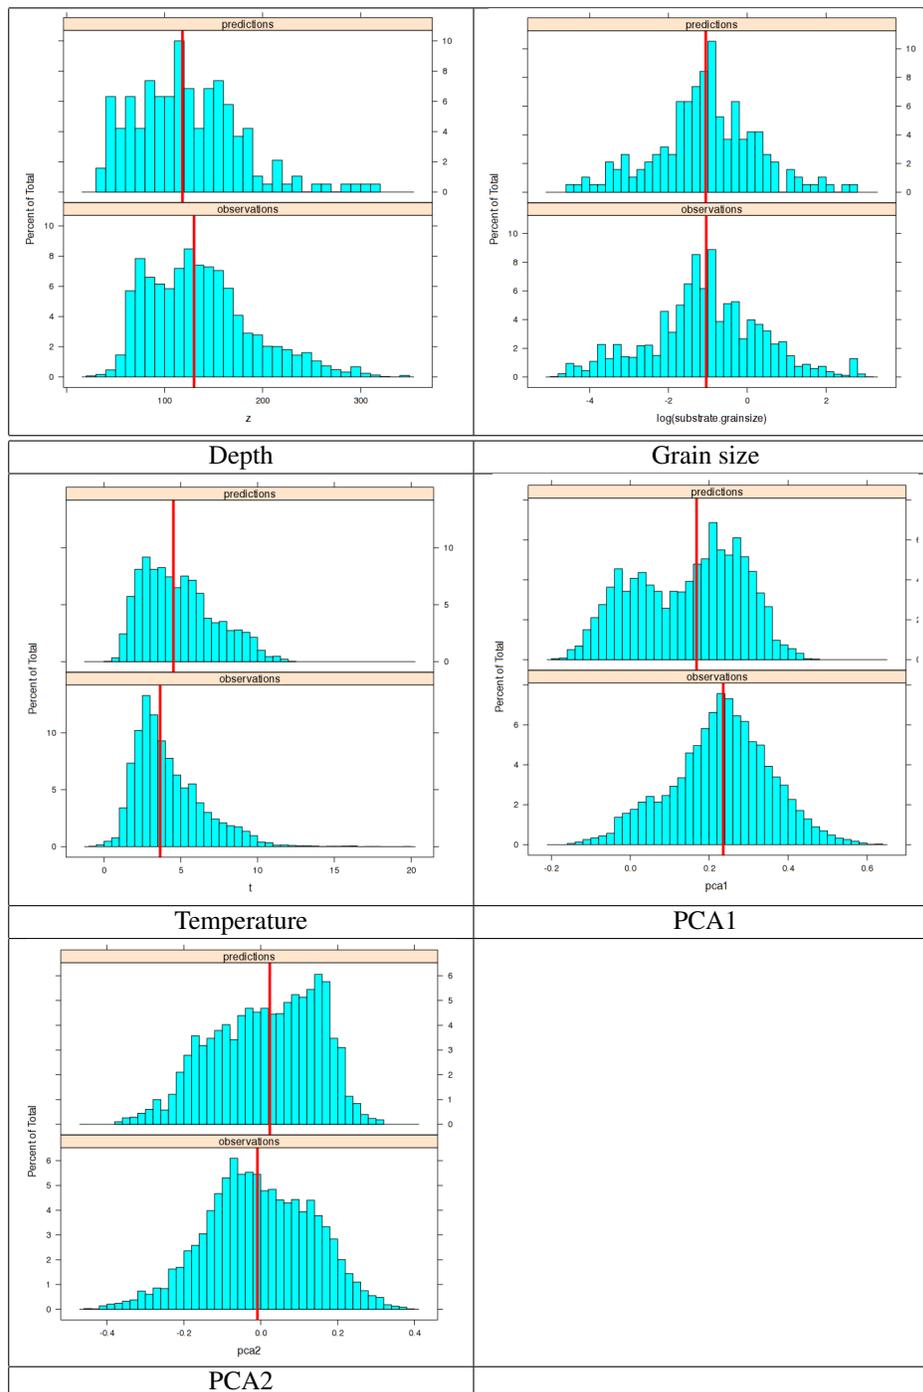
Survey locations are clustered in complex ways due to the historicity of fisheries management designations of fishing areas, number of licenses and quotas, and the industry funded nature of the survey with varying levels of contributions. Ultimately, this resulted in aggregations of survey stations in core fishing areas, no sampling in non-core areas, again introducing bias towards areas of snow crab habitat (Figure). Indeed, sampling of locations known to be ephemeral or unlikely to carry snow crab were considered by many in industry to be a waste of their funding. Even to this day, we still do not know the full limits of their spatial distribution and movement due to operational and logistical constraints (survey trawls cannot reasonably go beyond 350 m depth), and instead focus upon the locations that are accessible to the fishery that have their own depth, distance and operational constraints.

Currently, sampling is as extensive and intensive as possible, limited by trawl-accessible depths (350 m), to permit an objective determination of the spatial bounds of the snow crab population; information that must be known if reliable estimates of biomass and population structure (e.g., size, sex, maturity) are to be made. The spatial distribution of snow crab is quite dynamic and so can rapidly shift to areas where they are not "traditionally" found. In addition, the distributional patterns of immature, soft-shelled, very old and female crabs do not correspond completely to those of legal size males. The former are considered to be less competitive and more susceptible to predation (Hooper 1986) and usually observed in environments or substrates with greater cover (gravel, rocks; Comeau et al. 1998). Sampling that focused upon only those areas where large hard-shelled males occur in high frequency would preclude the reliable estimation of the relative abundance of these other important segments of the crab population.

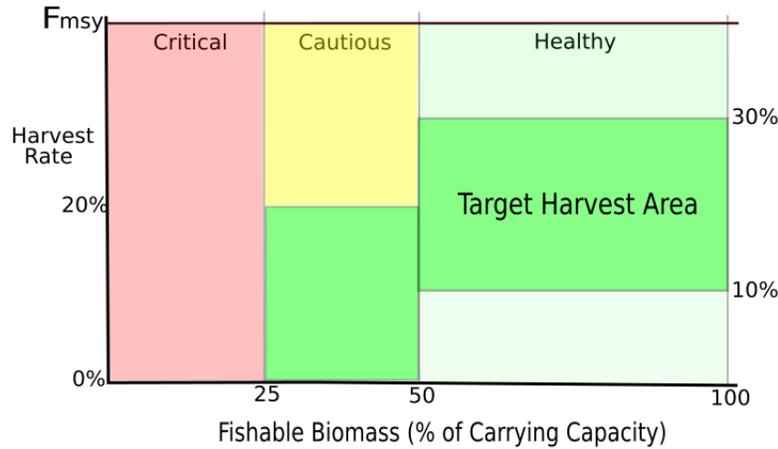
### 1.3.2 Synthesis

In most stock assessments, sampled observations are aggregated across space. The resulting aggregate temporal indices of abundance are used as is; and if there is sufficient information, they are subjected to time-series (state-space) modelling that *synthesizes* information on fisheries performance and biological/structural constraints in the form of natural mortality, growth, stock-recruitment patterns. The end result is an assessment of status (usually focusing only upon "spawning stock biomass", *i.e.*, mature reproductive females) relative to some biological reference points estimated from the stock assessment model.

Following a number of international agreements, a departmental position was formulated (DFO 2006) to move towards a "precautionary approach" to fisheries assessment and management. Here "precautionary" was used in a very particular sense, to use such *synthetic fisheries models* to define biological reference points and harvest control rules that used such reference points. In the snow crab assessment, a simple biomass dynamics model



**Fig. 6.** Sampling bias in ecosystem variability between surveyed locations (“observations”) and the overall domain (“predictions”). See Section 3.2 for explanations of the variables.



**Fig. 7.** Harvest control strategy framed in the context of the "precautionary approach" for snow crab (Choi and Zisserson 2011).

is used address this requirement and estimate biological reference points such as carrying capacity and intrinsic rates of increase. Harvest control rules (Figure 7) were framed within this context in 2010 (Choi and Zisserson 2011) and remains the current management approach.

## 2 Abundance index estimation (data aggregation)

The estimation of the state of a system  $y_{st}$  is an intermediate goal in a stock assessment. When the domain size in time and space of  $y_{st}$  is large, it cannot be *directly* observed. *Indirect* observations at some small subset of locations and times  $Y_{st}$  are measured in a coordinate space  $\{(s,t) \in D \in \mathfrak{R}^d \times \mathfrak{R}^1\}$  in the domain  $D$  of dimensionality  $d+1$  with  $\{s = 1, \dots, S\}$  spatial and  $\{t = 1, \dots, T\}$  temporal locations, are used to infer  $y_{st}$ . Here, we focus upon the simple case of  $d = 2$  spatial dimensions. In the case of aggregated areal and temporal units, these observations exist in a set of  $u = 1, \dots, U$  non-overlapping areal units and  $v = 1, \dots, V$  temporal units. To connect  $Y_{st}$  to  $y_{st}$ , we assume that the observations  $Y_{st}$  are derived from some spatiotemporal stochastic distributional *function* that (statistically) "generates" the observations (i.e., a spatiotemporal random field):

$$y_{st} \xrightarrow{\text{spatiotemporal process}} Y_{st} \quad (1)$$

and so the observations  $Y_{st}$  are used to infer the real state  $y_{st}$ .

Stock assessments, almost always, focus only upon a purely *temporal process*  $\Upsilon_t$  that "generates" a spatially integrated spatiotemporal process  $\int y_{st} ds$ :

$$\Upsilon_t \xrightarrow{\text{temporal process}} \int y_{st} ds \quad (2)$$

This integration is encoded in the *experimental design* and is almost always a purely *spatial process*. What this means is that spatial processes are being treated as being nested within temporal processes. The problem

is that a *spatial experimental design* can represent a spatiotemporal process only if the temporal component is stationary (at least, first and second order). This is generally not true. (Indeed, even the assumption of a single *spatial process* being valid throughout the spatial domain of interest at some time is itself problematic, as we will see below.) A naive application of a spatial design to spatiotemporal problem can, therefore, be problematic; this assumption needs to be examined.

To facilitate examination of this assumption, we begin with the formalism of a Generalized linear model (Banerjee et al. 2004):

$$\begin{aligned} Y &\sim f(\mu, {}^\varepsilon\sigma^2), \\ g(\mu) &= x^T \beta + O + \varepsilon, \end{aligned} \tag{3}$$

with constant offsets  $O = (o_1, \dots, o_S)$ , if any; a  $S \times V$  matrix of covariates  $x = (x_{sv}) \in \mathfrak{R}^{S \times V}$ ; the  $V$  covariate parameters  $\beta$  with a multivariate normal prior with mean  $\mu$  and diagonal variance matrix  $\Sigma$ ; and  $\varepsilon = (\varepsilon_1, \dots, \varepsilon_S)$  residual error. (Notationally, we are using left superscripts to denote variable types and reserve right subscripts for indexing). The  $f(\cdot)$  indicates an exponential family (Binomial, Normal, Poisson) and  $E(Y) = \mu$  with an invertible link function  $g(\cdot) = f^{-1}(\cdot)$ . In the Normal case,  $Y \sim \text{Normal}(\mu, {}^\varepsilon\sigma^2)$  and  $\mu = x^T \beta + O + \psi$  with standard deviation  ${}^\varepsilon\sigma$  associated with the residual error  $\varepsilon$ . In the Binomial case,  $Y \sim \text{Binomial}(\eta, \theta)$  and  $\ln(\theta/(1-\theta)) = x^T \beta + O + \psi$ , where  $\eta$  is the vector of number of trials and  $\theta$  the vector of probabilities of success in each trial. In the Poisson case,  $Y \sim \text{Poisson}(\mu)$  and  $\ln(\mu) = x^T \beta + O + \varepsilon$ . It is worthwhile emphasizing that there is a strong assumption that *residuals are independent and identically distributed*, that is:  $\varepsilon \stackrel{iid}{\sim} f(\cdot)$ . Therefore, if any autocorrelation exists, such as when there is spatial or temporal, or spatiotemporal autocorrelation, then model parameter estimation can become biased, including estimates of abundance.

## 2.1 Ad hoc smoothing models

Rapid visualization methods such as linear, or inverse-square distance weighted interpolation or simple spline-based smoothing attempt to address *Tobler's First law* in an *ad hoc* manner. However, these methods run a high risk of producing biased spatial patterns as they focus solely upon magnitude effects (first order) and/or make strong *ad hoc* assumptions such as the patterns of residual error or the location and order of inflection points. As producing biased estimates is anathema in stock assessments, naive application of these methods are not recommended as a tool for assessment and should only be used as quick visualization tools. Usually, as a result of the simplicity of the methods, time variation is also treated independently (i.e., totally ignored or at least the temporal autocorrelation is ignored).

## 2.2 Static models

Ignoring spatiotemporal autocorrelation altogether is perhaps the most frequently encountered approach. This might seem to be the "simplest" and assumption free. But it is actually the approach with the strongest assumptions on data; it assumes *Tobler's First law* is not applicable such that survey experimental design is static, always correct and unbiased and most precise. What this means is that snow crab are expected to be distributed inside of areal (i.e.,  $10 \times 10$  minute areal grid) and temporal units (5 minutes) in a regular manner while distributions between adjacent areal and temporal units are completely random and chaotic. These are, however, untenable assumptions as they are not supported by observations. Adjacent areal units and temporal units tend to have some form of autocorrelation. Of course, one can still invoke the Central limit theorem, and some

measure of central tendency of a random sample would provide an estimate of density and therefore overall abundance even in this situation. But the implicit assumptions are actually not "simple", though the explicit part of the model may be.

In a regression model context, this fundamentally *Cartesian perspective* is equivalent to an interaction only model between two "fixed" factors (areal unit and temporal unit, encoded in the model matrix  $x^T$ ) with some variability around each mean:

$$Y_{uv} \sim N(\mu_{uv}, \varepsilon \sigma_{uv}^2)$$

$$g(\mu_{uv}) = x^T \beta + \varepsilon_{uv} \tag{4}$$

$$\varepsilon_{uv} \sim N(0, \varepsilon \sigma_{uv}^2) \tag{5}$$

That is, all samples inside each areal ( $u$ ) and temporal ( $v$ ) unit is assumed to be equally representative (carries the same amount of information, "well mixed"). The residual error inside each space-time unit are independent and identically distributed (*iid, i.e. exchangeable*). The usual assumption is that they follow a Normal distribution:  $\varepsilon_{uv} \stackrel{iid}{\sim} N(0, \varepsilon \sigma_{uv}^2)$  with an invertible link function  $g(\cdot)$ , in this case, identity, and variance  $\varepsilon \sigma_{uv}^2$ . It must also be positive valued as we are dealing with organisms. The relationships between spatial and temporal units are fully ignored as each space-time unit is assumed to operate *completely independently* of the other. One relies strictly upon the survey's experimental design being able to provide a *precise* and *accurate* representation of reality. The random stratified model is the default, "standard" approach in most stock assessments. A significant weakness of this approach is that if a survey is incomplete for any reason, then there is no information for that location at that time. No inference can be made as there is no sharing of information between areal units nor temporal units.

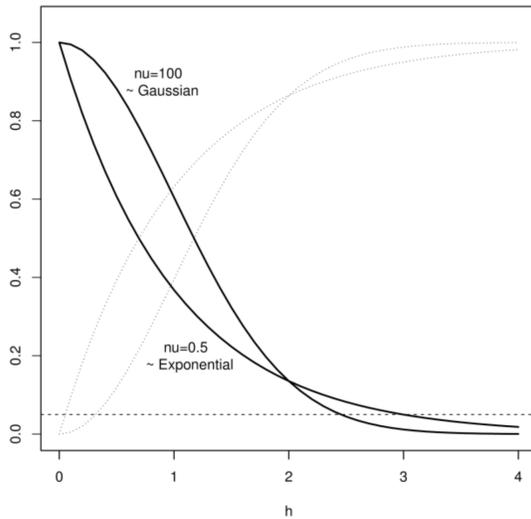
## 2.3 Continuous autocorrelation models

A *Lagrangian* perspective ignores the boundaries between areal ( $u$ ) and temporal ( $v$ ) units and instead focuses upon the spatiotemporal relationships and autocorrelations between all data in space  $s$  and time  $t$ .

### 2.3.1 Spatial autocorrelation

A spatial autocorrelation function  $\rho$  indicates how the proportion of the spatial variance  $C(h=0) = \omega \sigma^2$  drops as distance increases  $h$ . It is the covariance function  $C(h)$  scaled by the total variance  $C(0)$ , that is,  $\rho(h) = C(h)/C(0)$ . The spatial covariance function  $C(h) = C(s, s'; \theta)$  expresses the tendency of observations closer together to be more similar to each other than those further away (Tobler's First law). Here, the distance,  $h = \|s - s'\|$ , where  $\|\cdot\|$  indicates a spatial norm which in  $d = 2$  spatial dimensions is simply the Euclidean distance,  $h = (\Delta \text{northing}^2 + \Delta \text{eastng}^2)^{1/2}$ . Commonly used forms include:

$$\begin{aligned}
 C(h)_{\text{Spherical}} &= \begin{cases} \omega \sigma^2 (1 - \frac{3}{2}h/\phi + \frac{1}{2}(h/\phi)^3); & 0 < h \leq \phi \\ 0; & h > \phi, \end{cases} \\
 C(h)_{\text{Exponential}} &= \omega \sigma^2 e^{-h/\phi}, \\
 C(h)_{\text{Gaussian}} &= \omega \sigma^2 e^{-(h/\phi)^2}, \\
 C(h)_{\text{Powered exponential}} &= \omega \sigma^2 e^{-|h/\phi|^p}, \\
 C(h)_{\text{Matérn}} &= \omega \sigma^2 \frac{1}{2^{v-1}\Gamma(v)} (\sqrt{2vh}/\phi)^v K_v(\sqrt{2vh}/\phi).
 \end{aligned} \tag{6}$$



**Fig. 8.** Matérn autocorrelation function,  $\rho(h) = C(h)/C(0)$ , the covariance function  $C(h)$  scaled by the total variance  $C(0)$ , for two values of  $\nu$  (dark lines). As  $\nu$  increases ( $\nu = 100$ ), it approaches the Gaussian curve (upper dark curve on the left side) while at smaller values ( $\nu = 0.5$ ) the curve is exponential (lower dark curve on the left side). This flexibility has made it a popular choice in geostatistics. The associated semivariograms (scaled to unit variance)  $\gamma(h)$  are shown in light stippled lines. Spatial scale is defined heuristically as the distance  $h$  at which the autocorrelation falls to a low value (dashed horizontal line). The semivariance (also called "semivariogram"),  $\gamma(h)$ , is more commonly used in the geostatistics literature, and is simply the covariance function  $C(h)$  reflected on the horizontal axis of the global variance  $C(0)$  such that  $\gamma(h) = C(0) - C(h) = \frac{1}{2} \text{Var}[Y_s - Y_s'] = \omega \sigma^2 [1 - \rho(h)]$ .

At zero distance,  $C(0) = \text{Cov}(Y_s, Y_s) = \text{Var}(Y_s) = \varepsilon \sigma^2 + \omega \sigma^2$  (i.e., global spatial variance, also called the "sill"), where  $\varepsilon \sigma$  is the non-spatial, unstructured error (also called the "nugget"),  $\omega \sigma$  is the spatially structured error (also called the "partial sill"), and  $\omega \theta = \{\phi, \nu, p, \dots\}$  are function-specific parameters including  $\phi$  the range parameter.  $\Gamma(\cdot)$  is the Gamma function and  $K_\nu(\cdot)$  is the Bessel function of the second kind with smoothness  $\nu$ . The Matérn covariance function is frequently used in the more recent literature as the shape of this function is more flexible and known to be connected to a Gaussian spatial random process (Figure 8). The *stmv* approach (see 2.3.7) permits a local estimation of these autocorrelation parameters and the spatial scale.

Use of this approach permits a consistent and unambiguous definition of the spatial scale of observations (Choi et al. 2018). A sense of the spatial scale of some process is imperative for the development of any empirical assessment. This spatial scale can be defined as the distance  $h$  at which the autocorrelation falls to some level beyond which there is minimal relationship. Here, we use  $\rho(x) = 0.1$  as this threshold; that is, where spatial variance drops to 10% of the total spatial variance. This spatial scale is informative. When this *spatial scale* is small, this means short-range processes dominate (relative to the scale of the whole domain) and vice versa. For example, when a spatial feature (e.g., abundance distribution in space) demonstrates short characteristic *spatial scales* (i.e., a lot of spatial variability at short distances), sampling approaches must respect this and similarly operate at such shorter scales or even smaller ones if we are to be able to resolve the patterns and describe properly the subject of interest. Similarly, if a spatial feature is long-ranged and one wishes to resolve such patterns properly, then a sampling protocol must be similarly long-ranged to resolve the larger pattern. A sampling program much smaller than the characteristic spatial scale would be beneficial, but the accrued benefits relative to cost of sampling would diminish rapidly, in that time, effort and resources requirements generally increase more rapidly than any benefit (e.g., in the simplest case, if one is looking only naively at standard error as a measure of benefit, then it would increase asymptotically with increased effort with a power of  $-1/2$ ).

Autocorrelation function-based methods assume first and second order stationarity, that is, the mean and variance is stable and constant across the whole spatial domain. They are also quite slow to compute.

### 2.3.2 Temporal autocorrelation

Ecological systems also exist in a temporal frame. As such, similar to the above spatial considerations, there also exists some characteristic temporal scale upon which the processes internal to an area of interest and time period of interest, operate. The canonical example is how some quantity changes from one discrete-time period to another. This discrete-time notion of temporal autocorrelation is the slope parameter  $\rho$  from a plot of a variable as a function of itself with an offset of one time unit:

$$v_{t+1} = \rho v_t + \eta_t,$$

with  $\eta_t \sim N(0, \tau \sigma^2)$  and a temporal (linear) autocorrelation parameter  $\rho$ . This is known as an AR(1) process, where the 1 indicates 1 unit time lag. Note that the expectation is 0 for  $\eta_t$ , that is, *first order stationary*. More complex models with moving averages and additional time-lags can also be specified. Collectively these are known as AR, ARMA, and ARIMA models. The difficulty with these autocorrelation time-series formulations is the requirement of a complete data series without missing data, in the standard form.

In a completely identical approach to the spatial case, a temporal autocorrelation function can be used to describe the form of the temporal autocorrelation pattern (Choi et al. 2018). More specifically, we define a *temporal stochastic process*,  $y_t$ , that is, some latent, unobservable but real, *stochastic function* that generates observations  $y_t \rightarrow Y_t$ , where  $Y_t$  are any temporally referenced observation at some time  $t$ , measured in a coordinate space whose domain  $D$  has dimensionality 1 such that  $\{t \in D \in \mathfrak{R}\}$  with  $\{l = 1, \dots, L\}$  temporal "locations". The manner in which the variability of  $Y_t$  changes as a function of the norm (distance),  $h = \|t - t'\|$ , is the temporal autocorrelation function  $\rho(h)$ . The latter can take any form including the same as the spatial autocorrelation functions.

The covariance function, for example, when expressed as an exponential decay model controlled by time range parameter  $\phi_t$  is:

$$C(t, t'; \tau \theta) = \tau \sigma^2 e^{-|h|/\phi_t}.$$

At zero time difference,  $C(0) = \text{Cov}(Y_t, Y_t) = \text{Var}(Y_t) = \varepsilon \sigma^2 + \tau \sigma^2$  (*i.e.*, global temporal variance), where  $\varepsilon \sigma^2$  is the nontemporal, unstructured error,  $\tau \sigma^2$  is the temporally structured error. The *temporal autocorrelation function* is defined as the covariance function scaled by the global variance:  $\rho_t(h) = C(h)/C(0)$ . Heuristically, the *temporal autocorrelation scale* is defined as the time difference at which the temporal autocorrelation decreases asymptotically to some low level. In this document, we will use the same threshold as the "practical spatial range",  $\rho_t(h) = 0.1$ , and refer to it as the "temporal scale" at which temporal autocorrelation approaches 0.1 (that is 10% of the total temporal variance). The *stmv* approach (see 2.3.7) permits estimation of these local autocorrelation parameters and the temporal scale.

### 2.3.3 Spectral estimation of autocorrelation

Fourier transforms decompose any function in a "continuous" domain (e.g., time, space) as an infinite sum of sine and cosine functions (Fourier 1822). The sine and cosine functions are interpreted as amplitudes and

phase shifts associated with an infinite range of frequencies in the "spectral" domain. Computationally efficient algorithms for Fast Fourier Transforms (FFT) were initially developed by Gauss in 1805 (Heideman et al. 1985), and elaborated upon by Yates (1937), Danielson and Lanczos (1942), and fully generalized by Cooley and Tukey (1965). This enabled operations in the spectral domain that are orders of magnitude faster than their equivalent operations in the continuous time and/or space domains. For example, for a problem with  $n$  data, the FFT has a computational complexity of order  $\mathcal{O}(n \cdot \log_2(n))$ . In contrast, an operation of importance in the context of spatiotemporal modelling is inversion of a covariance matrix  $\Sigma_{n \times n}$  that has a computational complexity of order  $\mathcal{O}(n^3)$ . This represents an improvement by a factor of  $n^2 / \log_2(n)$ , which even for a simple problem with  $n = 10^3$  data locations, can represent up to a  $10^2$  fold improvement in computational speed. Parenthetically, the Discrete Fourier Transform (DFT) has the intermediate computation complexity of order  $\mathcal{O}(n^2)$ .

Beyond computational complexity, there exist two additional features of the Fourier Transform that are especially significant in the context of spatiotemporal modelling. The first is known as the Wiener-Khinchin (-Einstein, -Kolmogorov) theorem (Wiener 1930; Khintchine 1934; Robertson and George 2012), which connects the autocorrelation function of a stationary random process with the power spectral density (also known as a power spectrum) of the process. That is, a rapid estimation of the autocorrelation (and cross-correlation) of a process can be obtained from the power spectrum. The second aspect of significance is the **convolution** theorem: the combination of two functions in the continuous domain becomes simple multiplication in the spectral domain. The convolution of an autocorrelation function with a spectral representation of the spatiotemporal process of interest amounts to a kernel-based smoothing interpolator respecting the temporal/spatial/spatiotemporal autocorrelation. These two aspects, respectively, permit fast and unbiased representations of the autocorrelation function and rapid estimation/prediction without having to invert and solve the covariance matrix  $\Sigma_{n \times n}$ .

In what follows, we will focus upon a one-dimensional problem for simplicity, with the awareness that this can be simply extended to higher dimensions, including space and space-time dimensions. Specifically, any measurements along the one dimensional  $d = 1$  coordinate space  $\{(x) \in D \in \mathfrak{R}^d\}$ , of domain  $D$ , generated from the process of interest  $g(x)$  can be represented in the frequency domain as a series of complex trigonometric coefficients  $G(k)$  and vice versa. The forward and backward Fourier transforms are respectively:

$$G(k) = \mathcal{F}_x[g(x)](k) = \int_{-\infty}^{\infty} g(x) \cdot e^{-(2\pi ki)x} dx,$$

$$g(x) = \mathcal{F}_k^{-1}[G(k)](x) = \int_{-\infty}^{\infty} G(k) \cdot e^{(2\pi xi)k} dk.$$

The above make use of Euler's formula,  $e^{2\pi\theta i} = \cos(2\pi\theta) + i \cdot \sin(2\pi\theta)$ , to compactly represent the amplitudes and phase shifts of the sine and cosine functions with amplitudes  $g(x)$  and  $G(k)$ , also called Fourier pairs. The  $i$  represent by convention, imaginary numbers.

In an applied setting, the discrete form of the transform is particularly useful as measurements are usually discrete at some fundamental level (such as a sampling event). The discrete Fourier transform and it's inverse are as follows:

$$G_k = \mathcal{F}_x[g_x](k) = \sum_{n=0}^{N-1} g_x \cdot e^{-(2\pi ki)(x/N)},$$

$$g_x = \mathcal{F}_k^{-1}[G_k](x) = \frac{1}{N} \sum_{k=0}^{N-1} G_k \cdot e^{(2\pi xi)(k/N)}.$$

The  $g_x = \{g_0, g_1, \dots, g_{N-1}\}$  are vector of values in time or space of the data and the  $G_k = \{G_0, G_1, \dots, G_{N-1}\}$  are discrete frequency bands. For each frequency band  $k$ , the associated amplitude is:

$$|G_k|/N = \sqrt{\text{Re}(G_k)^2 + \text{Im}(G_k)^2}/N$$

and the phase is the angle  $\arg()$  between the real and imaginary line:

$$\arg(G_k) = \text{atan2}(\text{Im}(G_k), \text{Re}(G_k)) = -i \cdot \ln\left(\frac{G_k}{|G_k|}\right).$$

The utility of the spectral representation is that the autocorrelation function  $\rho(x)$  of some **stationary** function  $g(x)$  is equivalent to the inverse Fourier transform of the power spectral distribution.

$$\rho(x) = \mathcal{F}_k^{-1}[|G_k|^2](x)$$

where,  $|G_k|^2$  is the modulus squared power spectral density derived from:

$$\mathcal{F}[g(x) * g(-x)](k) = G_k \cdot G_k^* = |G_k|^2$$

and the asterisk denotes a complex conjugate.

This relation is the result of the Wiener-Khinchin theorem (Robertson and George 2012). The autocorrelation function is, of course, directly related to the covariance function used in temporal, spatial and spatiotemporal interpolation methods. The utility is, however, not simply a matter of the significantly reduced computational complexity from  $\mathcal{O}(n^3)$  to  $\mathcal{O}(n \cdot \log_2(n))$ . Semivariogram/autocorrelation estimation methods generally make repeated use of data for variance estimation in various distance groups, causing spurious autocorrelation and, therefore, biased parameter estimates. The autocorrelation function, being derived from a spectral decomposition are, by definition, independent and therefore parameter estimates are unbiased!

As an aside, these FFT-based methods require additional handling due to "missing" data being common in most empirical applications. This can be handled by a locally weighted average scheme (kernel weights) which in the spectral domain is known as a Nadaraya-Watson kernel convolution smoother (see e.g., the documentation in the R-library, "fields", Nychka et al. 2017). And it is worth re-iterating that such FFT-based approaches face the same challenges as the covariance-based approaches in that stationarity (first and second order) is assumed/required (Wiener-Khinchin theorem). The *stmv* approach (see 2.3.7) uses these FFT-based approaches to facilitate estimation of these autocorrelation parameters and estimate the spatial and temporal scales.

### 2.3.4 Spatial model

Snow crab demonstrate clustered aggregations in space and time. This means that the residual error inside each area-time unit will not be *iid*, *a priori*. Furthermore, areal units closer together tend to be more similar than those more distant; and samples in a given time unit will also show correlation. The first attempts at snow crab abundance estimation were developed with a focus upon the spatial component of the residual error. In terms of the model formulation, it is an incremental extension to the *simple model* (Eq. 4) for  $\mathbf{Y}_t = (Y_{1t}, \dots, Y_{St})^T$  that adds a spatial error ( $\omega_t$ ) to the basic model, thus treating each time slice  $t$ , *independently*. The index  $t$  could be

dropped for brevity, however, we leave it explicitly to remind ourselves that it is a spatiotemporal process where time is assumed to be independent:

$$\begin{aligned} Y_t &\sim \text{MVN}(\mu_t, \Sigma_t) \\ g(\mu_t) &= x_t^T \beta_t + \omega_t + \varepsilon_t \end{aligned} \quad (7)$$

Here, the non-spatial component is still the independent and identically distributed component, but now applicable to the full domain and not a single area unit, while the spatial component has an error that changes as a function of distance (the autocorrelation function):

$$\begin{aligned} \varepsilon_t &\sim \text{N}(0, \varepsilon \sigma_t^2) \\ \omega_t &\sim \text{GP}(0, C(s_t, s'_t; {}^\omega \theta_t)) \end{aligned}$$

The non-spatial component for each year  $\varepsilon_t$  represents measurement and/or micro-scale process variability. In its simplest (historical) form, it is usually assumed to have a Normal distribution with a mean of zero and standard deviation  $\varepsilon \sigma$  (also called a *nugget* error). The spatial error (also called a "partial sill") is assumed to follow a **Gaussian process** with mean 0 and a spatial covariance function  $C(s_t, s'_t; {}^\omega \theta_t)$  that describes the form of the variance as a function of distance between data, controlled by the spatial parameters  ${}^\omega \theta_t$  and spatially structured standard deviation  ${}^\omega \sigma$  (see below). The covariance matrix is the sum of these two variances:  $\Sigma_t = [C(s_{it}, s_{jt}; {}^\omega \theta_t)]_{i,j=1}^K + \varepsilon \sigma_t^2 I_S$ , with  $I_S$  an identity matrix of size  $S$ .

This *spatial model* (eq. 7, also known as *kriging*) was used in the early snow crab assessments. However, it necessitated the repeated calculation of the autocorrelation function for each year (actually, the covariance matrix; see below). What was notable was that the form of the autocorrelation function was not constant. It varied due to variations in spatiotemporal abundance and distributions. Especially *alarming* was that it was not possible to obtain a reliable estimate of an autocorrelation function when densities and spatial distributions contracted (e.g. in 2005), a condition where in fact one needs more precise and accurate estimate, rather than less. Indeed, it was not clear if one should model the spatial autocorrelation function itself as a time-autocorrelated feature or to embed time autocorrelation to the model or simply to default to a time-averaged autocorrelation function (all were attempted).

### 2.3.5 Hierarchical spatial models

The addition of spatial autocorrelation improves our understanding of the influence a station has over its immediate environment and how it decays with distance based on variance characteristics. In this way, it is an incremental improvement over a non-spatial model. However, the assumption of the former model is that the unstructured error is  $\varepsilon \stackrel{iid}{\sim} \text{N}(0, \sigma_\varepsilon^2)$ . Unfortunately, snow crab data distribution suggest that the error is non-Gaussian; it is perhaps closer to a logarithmic distribution. This is due to in a large part, to spatial aggregations arising from behavioral factors such as mating clusters, predator avoidance, foraging upon patchy food distributions and habitat constraints such as temperature limitations and substrate availability, what we naively call *ecosystem variability*. The outcome of such spatial aggregations is that the *numerical* distribution of abundance sampled in a grid-like or any naive experimental design will also be *non-Gaussian*. Most samples will miss the aggregations leading to many samples being low density locations and a rarer and rarer number of stations with high density aggregations.

The data distributional model is important as computing the arithmetic average of a distribution that is logarithmic is not the best estimate of central tendency; due to the shape of the distribution it would overestimate it. However, most historical forms of kriging generally operates upon the assumption that the error distribution of data is Normal (Gaussian). In other words, the clustered nature of the abundance of snow crab will result in spatial estimates of abundance that is generally greater than likely exists, simply due to the data and spatial distribution. This would not be a precautionary approach.

So it is critical to alter our data distributional assumption. This is readily implemented by altering the distributional assumption of the *spatial model* (eq. 7):

$$\begin{aligned}
 Y_t &\sim \text{MVN}(\mu_t, \Sigma_t) \\
 \log(\mu_t) &= x_t^T \beta_t + \omega_t + \varepsilon_t \\
 \varepsilon_t &\sim \text{N}(0, \varepsilon \sigma_t^2) \\
 \omega_t &\sim \text{GP}(0, C(s_t, s'_t; \omega \theta_t))
 \end{aligned} \tag{8}$$

for the positive valued biomass. It is attractive as the results are less prone to distributionally induced bias. However, as mentioned above, this adds another form of bias to the model in that all zero values are ignored (dropped) as  $\log(0)$  is not defined. One way around this situation is to have a separate binomial model to describe the zero and non-zero values.

$$\begin{aligned}
 H_t &\sim \text{Binomial}(\eta_t, \Theta_t) \\
 \ln(\Theta_t / (1 - \Theta_t)) &= x_t^T H \beta_t + H \omega_t + H \varepsilon_t \\
 H \varepsilon_t &\sim \text{N}(0, H \varepsilon \sigma_t^2) \\
 H \omega_t &\sim \text{GP}(0, C(s_t, s'_t; \omega \theta_t))
 \end{aligned} \tag{9}$$

where the left superscript  $H$  indicates the parameter is obtained from the binomial (habitat) model,  $\eta_t$  is the number of trials and  $\Theta_t$  the probabilities of success in each trial in a given year  $t$  and  $H_t$  is the proportion of samples with snow crab present in a given location (i.e. habitat probability) and year. This latter model becomes an excellent habitat probability model (also called species distribution models in the current literature) that can provide many insights into the habitat space of snow crab on its own. When combined with the the abundance model becomes an hierarchical zero-inflated model. However, it comes at the cost of additional computations and a doubling of the number of parameters as well as a requirement to conduct posterior simulations to properly propagate errors and no simple way to assess overall model performance. [ASIDE: this hierarchical approach is also known as “Delta models” or “Hurdle models”].

### 2.3.6 Stationary spatiotemporal models

So far, the models have used only a representation of the spatial autocorrelation. Most biological features such as abundance are, however, variable not only in space but also in time. Just as space can be modelled as an autocorrelated process, so too can time. However, due to the brevity of most biological time-series, a full and reliable temporal autocorrelation function cannot normally be defined/estimated. Instead, an operational extraction of a few components, namely, an annual and possibly a seasonal component can provide some information and constraints upon the residual temporal errors. The simplest such model utilizes a single spatial autocorrelation

function of the full domain and a single (separable, additive) temporal autocorrelation parameter or function (if there is sufficient time-series data).

Spatiotemporal models can be seen as a simple extension of the spatial regression model. The space-time regression model can then be specified as:

$$\begin{aligned}
 Y &\sim \text{MVN}(\mu, \Sigma) \\
 \log(\mu) &= x^T \beta + \omega + \varepsilon \\
 \varepsilon &\sim \text{N}(0, \varepsilon \sigma^2) \\
 \omega &\sim \text{GP}(0, C(s, s'; \omega \theta))
 \end{aligned} \tag{10}$$

where the error process is decomposed into a spatiotemporal structured component  $\omega$  and an unstructured component  $\varepsilon$ . The complexity arises in how these *spatiotemporal* parameters are parameterized ( $\omega, \varepsilon, \beta, \theta$ , etc.). For example, the *spatially structured error* can be parameterized as:

- $\omega$  – completely fixed
- $\omega_{-t}$  – spatial effects nested in time (spatial autocorrelation at each time slice, no temporal autocorrelation);
- $\omega_{s-}$  – temporal effects nested in sites (temporal autocorrelation at each site, no spatial autocorrelation);
- $\omega_{s-} + \omega_{-t}$  – *separable* (spatial and temporal autocorrelations are independent and additive (or multiplicative if on log scale) with  $\omega_{-t} \sim \text{GP}(0, C(s_t, s'_t; \omega \theta_t))$  and  $\omega_{s-} \sim \text{GP}(0, C(t_s, t'_s; \tau \theta_s))$ );
- $\omega_{st}$  – non-separable (both time and space structure evolve in a non-simple manner).

Each of the various errors and parameters can have similar complexities which leads to rather extraordinary expressiveness and equally extraordinary computational requirements. While conceptually coherent and elegant, the evaluation of the likelihoods in these models requires the repeated computation of the inverse of the covariance matrix  $\Sigma_{n \times n}$ , an operation that scales with  $\mathcal{O}(n^3)$  operations. This has been a bottleneck to further development and use of these covariance-based methods in large scaled problems of space and space-time. Approximations have been suggested to overcome this computational limit: modelling the spatial process  $\omega$  with a lower dimensional process via kernel convolutions, moving averages, low rank splines/basis functions and predictive processes (projection of spatial process onto a smaller subset; Sølna and Switzer 1996, Wikle and Cressie 1999, Hung et al. 2004, Xu et al. 2005, Banerjee et al. 2004); approximating the spatial process as a Markov random field with Laplace and SPDE Approximations (Lindgren and Rue 2015); and approximating the likelihood of the spatiotemporal SPDE process with a spectral domain process (Sigrist et al. 2012).

Separable models are almost always used for the sake of computational speed as this treats space and time independently, reducing the problems crudely from  $\mathcal{O}((ST)^3)$  to  $\mathcal{O}(S^3) + \mathcal{O}(T^3)$  operations; where  $S$  is the number of spatial locations and  $T$  the number of time slices. In reality, however, such separable models are usually inappropriate unless the study area is homogeneous and truly first and second order constant (i.e., constant mean, variance) across both time and space, a fact that is seldom true in most ecological systems. This is notoriously the case with biology where aggregation and behavior is highly clustered and context (location and time) dependent (non-linear/non-additive). There are discontinuities in space and time due to landscape, species collapse, invasions, etc.

### 2.3.7 Complex spatiotemporal models

All these approaches, including most operational spatiotemporal models, share a core assumption, that the mean and variance is stable throughout the domain (first and second order stationarity) and that a single autocorrelation function in space and a single autocorrelation function in time is sufficient to describe the full spatiotemporal pattern. The correctness of these rather important assumptions are seldom verified in the literature (I have never encountered them). When they are examined, there is little evidence to suggest this to be an appropriate assumption (see Section 3.2).

One way to deal with such non-stationarity is to induce autocorrelation through covariates. By expanding the model to include time-varying and static covariates that can account for some of these discontinuities, the error distributions can become more stationary. That is, expanding the contents of  $x^T \beta$  to include processes that represent *ecosystem variability*. The range of possibilities are of course quite large and as each component of *ecosystem variability* usually requires their own modelling and assessment, it can become very quickly an onerous exercise.

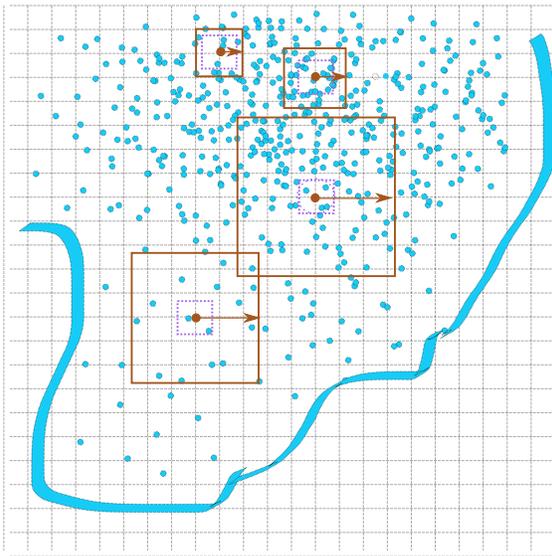
Another solution is to identify a local sub-domain where they actually are stationary in space and potentially stationary in time as well. This reduces the problem into small manageable subdomains where assumptions of local stationarity are valid and modelling of spatiotemporal processes and associated parameters become computationally feasible. This approach is implemented in the *stmv* library (Choi et al. 2018; <https://github.com/jae0/stmv>). There is some conceptual similarity of this approach to "geographically weighted regression" (e.g., Fotheringham et al. 2002) in that each subdomain can have their own model parameters,  $\beta_{st}$ . However, geographically weighted regression permit only the model parameters  $\beta_{st}$  to vary; however, here both the model parameters  $\beta_{st}$  and the spatiotemporal errors  $\omega_{st}$  and  $\varepsilon_{st}$  are also varying in space.

To be more precise, we define statistical nodes  $\{N_{m=(1,\dots,M)} | m \in \mathfrak{R}^d\}$  in a spatial lattice that are the centres of areal units (Figure 9). The norm (distance) of data from each node is  $h_m = \|s_m, s_Y\|$ . A local subdomain of a given node  $m$  is  $\{S_{m=(1,\dots,M)} \in D | h_m < h_u\}$  or more briefly as  $\{S_m\}$  which represents all locations within some distance to the statistical node  $\{h_u | \rho(h_u)_{\text{Matérn}} > 0.1\}$ ; that is, the distance at which the local spatial autocorrelation drops to a negligible value (arbitrarily taken as  $\rho > 0.1$ ) and associated parameter values are "supported". The spatiotemporal data found within this subdomain  $m$  is  $\{Y_{st} | (s, t) \in D | h_m < h_u\}$  and is herein, notationally abbreviated as  ${}^m Y_{st}$ .

Operating upon all components of the regression model simultaneously is computationally prohibitive. Even with very simplistic Generalized Additive Model (GAM) or Generalized Additive Mixed effects Model (GAMM) parametrizations of spatial and temporal structure, the solutions require weeks on fast machines (5 GHz CPU, 64GB RAM in 2016), depending of course also upon the amount of data and resolution and model complexity. As a compromise between model complexity and computational speed, a global covariate model  $x^T \beta_{st} + \varphi_{st}$  can be used to parameterize the covariate effects using a linear, generalized linear or generalized additive model. A low order penalized basis splines of the covariate predictors of 3 knots or less are biologically plausible as modality can often be expected:

$$\begin{aligned} Y &\sim \text{N}(\mu_{st}, \varphi \sigma^2) \\ g(\mu_{st}) &= x^T \beta_{st} + \varphi_{st}, \\ \varphi_{st} &\sim \text{N}(0, \varphi \sigma^2). \end{aligned}$$

The resultant residual spatiotemporal error  $\varphi_{st}$  can be decomposed into other components in an approach similar to that of "regression kriging" and universal "kriging with external drift" (Hengl et al. 2004). First a local spatial autocorrelation scale is derived from a rapid (coarse grained) fit of the local residuals based upon an iteratively increasing guess of  ${}^m \varphi_{st}$  to a Matérn autocorrelation function.



**Fig. 9.** Spatial distribution of data (blue dots) overlaid by a statistical grid. The  $m$  nodes represent the centers of each local subdomain  $S_m$  which extends to a distance (right-facing arrows; solid squares) that varies depending upon the underlying spatial variability of the data, defined as the distance at which the spatial autocorrelation drops to some small value (e.g.,  $\rho_s = 0.1$ ). Data within this distance and parameters obtained from the local analysis are, under the assumption of second order stationarity, used to complete the local model of the spatial or spatiotemporal processes and then predict/interpolate to some fraction of the distance between statistical grid nodes (stippled square). Every statistical node is visited for spatiotemporal modelling. Any overlapping predictions are locally averaged (weighted by number of predictions and prediction variance). As grid size decreases the number of computations increase while reducing computational load and RAM requirements; however, the utility of the model also declines due to small sample sizes entering analyses. Judicious choice of statistical grid density as well as maximum and minimum number of data points and upper and lower bounds of spatial bounds must be balanced.

To treat time in a similar manner, one would also need to determine temporal nodes and define appropriate temporal autocorrelation scales. In practice, temporal data are often sparse and limiting in survey data and so data from all time periods are used to determine a crude scaling, essentially amounting to a temporally averaged spatial autocorrelation. Once the approximate bounds of the subdomain (support) are estimated, the  ${}^m\varphi_{st}$  are modelled as a Fourier series with two harmonics, one inter-annual and one sub-annual (seasonal):  $f_m(\cdot)$  (interannual, seasonal). In other words, a full temporal autocorrelation (covariance) model is not used but rather one that uses only a subset of the components at fixed wavelengths:

$$\begin{aligned} {}^m\varphi_{st} &= f_m(\cdot) + {}^m\zeta_{st}, \\ {}^m\zeta_{st} &\sim N(0, \zeta\sigma_m^2). \end{aligned}$$

The temporal autocorrelation is, therefore, carried by the individual temporal processes at each spatial datum and the temporally structured error  ${}^\tau\sigma_m^2$  is the variance component of the model  $f_m(\cdot)$ , that is,  ${}^\tau\sigma_m^2 = \varphi\sigma_m^2 - \zeta\sigma_m^2$ .

The spatial autocorrelation function is parameterized as being derived from the subdomain mean Gaussian process with a Matérn covariance function with parameters  ${}^\omega\theta_m = \{\phi_m, \nu_m\}$  and a time-varying spatially structured standard error  ${}^\omega\sigma_m$ . As the data used to estimate the spatial autocorrelation structure are often sparse, the data are augmented by temporal predictions of the residual error process at each spatial datum (and notationally designated by an asterisk). These temporally "augmented" residual processes are modelled spatially at each time slice  ${}^m\varphi_{st}^*$  as the sum of a time-varying spatial **Gaussian process**  ${}^m\omega_{st}$  parameterized as a Matérn spatial covariance function:

$${}^\omega\sigma_{mt}^2 \frac{1}{2^{\nu_{mt}-1}\Gamma(\nu_{mt})} (\sqrt{2\nu_{mt}h}/\phi_{mt})^{\nu_{mt}} K_{\nu_{mt}}(\sqrt{2\nu_{mt}h}/\phi_{mt})$$

with a local spatial error  ${}^\omega\sigma_{mt}$ ; and a spatially and temporally unstructured error process assumed to be derived from a Normal error process with mean zero and error  $\sigma_{\varepsilon|m}$ :

$$\begin{aligned} \varphi_{mt}^* &= \omega_{mt} + \varepsilon_{mt}, \\ \omega_{mt} &\sim \text{GP}(0, C(\mathbf{s}, \mathbf{s}'; {}^\omega\theta_{mt} = \{\nu_{mt}, \phi_{mt}, {}^\omega\sigma_{mt}\})), \\ \varepsilon_{mt} &\sim \text{Normal}(0, {}^\varepsilon\sigma_m^2). \end{aligned}$$

This approach represents a practical balance between computational time and model complexity/realism. For additional speed, an FFT-based Matérn convolution implementation is used. This solution focuses upon prediction using a mosaic of solutions in space, overall likelihood or AIC evaluation is a challenge. At present, predictive success is the only means to evaluate utility and eventually some form of Expectation-Maximization approaches through iteration once computational speeds improve would provide less biased parameter estimates.

## 2.4 Hybrid Cartesian-Lagrangian models

Spatial models without fixed spatial boundaries (i.e., the *Lagrangian* models) are intuitive and elegant. However, practical application in data poor situations invites difficulties and challenges requiring model complexities that can escalate exponentially when attempting to do justice to the natural non-stationary nature of these spatiotemporal processes. An alternative approach is to treat spatiotemporal autocorrelation in a discrete but local manner. This approach begins with the AUs of the Cartesian approach, however, extends them by permitting local interactions. The nature of these interactions are simply determined by either proximity (distance) or real connectivity (roads, pathways, ocean currents), which must therefore be determined carefully. This permits a

combination of both the *Cartesian* model's focus upon AUs and the changing underlying fields of *ecosystem variability* of the *Lagrangian* models. As the size of the AUs decrease, the solutions of this hybrid approach will in theory converge with those of the *Lagrangian* models, due to the Hammersly-Clifford Theorem (see below). Such conditional autoregressive models can be implemented in Bayesian form using INLA (<http://www.r-inla.org/>). We use the R-library "*carstm*" (<https://github.com/jae0/carstm>) to facilitate modeling by acting as a simple front end to INLA that manages the data required for these models and the resultant predictions.

### 2.4.1 Spatial Conditional autoregressive models

The CAR model is an extension of the generalized regression model and so initially very similar to the setup of the *Lagrangian* models. The study region  $D = D_{11}, \dots, D_{U1}, \dots, D_{UV}$  is a set of  $u = 1, \dots, U$  areal units and  $v = 1, \dots, V$  temporal units. Observations  $Y$  on such a set  $D$ , generally demonstrate spatial and temporal autocorrelation and can be modelled as the same extension of a Generalized linear model (Eq: 3):

$$Y \sim f(\mu, \varepsilon\sigma^2)$$

$$g(\mu) = x^T \beta + O + \varepsilon$$

As with the autocorrelation function approach, the CAR-structured approach also decomposes the residual error structure into two additive random components. These components are a spatial "CAR-structured" error  $\phi$  and a non-spatial (unstructured) *iid* error  $\varepsilon$ :

$$g(\mu) = x^T \beta + O + \phi + \varepsilon$$

Where it differs is that there is attention paid to the connectivity between AUs. This connectivity enters into the estimation of  $\phi$  and is by convention, represented by the elements  $w_{ij}$  of an adjacency matrix  $W_{U \times U}$  such that, connectivity is designated by  $w_{ij} = 1$ , non-connectivity is designated by  $w_{ij} = 0$ , and self-connectivity, i.e.,  $\text{diag}(W)$ , is designated as  $w_{ii} = 0$ . The residual spatial autocorrelation error  $\phi$ , is treated as a random variable that is assumed to have a *local conditional distribution* between neighbouring lattices  $\{p(\phi_i | \phi_{i \sim j})\}$ , excluding the self ( $i \sim j$ ) which approximates the *global joint distribution*  $p(\phi_1, \phi_2, \dots, \phi_U)$ , on the strength of the Hammersly-Clifford Theorem (Besag 1974, Besag et al. 1991, Leroux et al. 2000). The *local conditional distribution* (Banerjee et al. 2004) can be expressed as a weighted sum of the random quantity, where the weights are derived from the *local neighbourhood connectivity* in the adjacency matrix:

$$p(\phi_i | \phi_{i \sim j}, \tau) = N(\alpha \sum_{i \sim j}^K w_{ij} \phi_j, \tau^{-1}).$$

Here,  $\tau$  is a precision parameter; and  $\alpha \in [0, 1]$  controls the strength of spatial dependence/association such that when  $\alpha = 0$  there is no spatial structure and  $\alpha = 1$  indicates a fully intrinsic (conditional) autoregressive model, I(C)AR. The above essentially amounts to assuming a single global spatial autocorrelation  $\alpha$ , tempered by the local average of spatial errors to approximate the spatial structure of the area of interest. Any difference in autocorrelation behaviour in different areas must therefore be associated with covariate variations. This assumption is unrealistic in many systems as covariates may not properly exist, are unknown, or be able to express these differences between subareas, especially when real discontinuities exist (but see, *stmv*).

Under the assumption that  $\phi$  is a Markov Random Field, the *joint distribution* of the random variable is, via Brook's Lemma (Besag 1974):

$$\phi \sim N(0, Q^{-1}),$$

where,

$$Q = \tau D(I - \alpha B)$$

is the precision matrix, and  $D$  is a diagonal matrix of the same shape as  $W$ , and  $\text{diag}(D) = d_{ii} = \text{no. of neighbours for each AU}_{i_i}$ . Further,  $B = D^{-1}W$  is the scaled adjacency matrix that quantifies the relative strength/influence of the connectivity for each neighbour;  $\text{diag}(B) = b_{ii} = 0$ . This simplifies to:

$$Q = [\tau(D - \alpha W)]^{-1}.$$

The spatial dependence parameter is usually difficult to estimate and so is often assumed to be  $\alpha = 1$ . This simplifies the precision matrix to:

$$Q = [\tau(D - W)]^{-1}.$$

The  $\phi$  are non-identifiable (adding a constant to  $\phi$  leaves the joint distribution unchanged) rendering it an improper prior. In other words, the covariance matrix is singular and additional constraints are required to make it computationally tractable. By convention, this constraint is that they sum to zero ( $\sum_u \phi_u = 0$ ). This model is a well known and frequently utilized simplification of the CAR model, and is known as an I(C)AR (Intrinsic conditional autoregressive) model. When applied to a Poisson process, it is also known in the epidemiological literature as a convolution model or a Besag-York-Mollie (BYM) model (Besag et al. 1991).

Setting sensible priors on the precision of the structured and unstructured random effects has been found to be challenging as they are dependent upon the connectivity of the area of interest (Bernardinelli 1995). Simpson et al. (2017) and Riebler et al. (2016), following Leroux (1999), approach this issue by a clever parametrization the problem to what they call the "bym2" model in INLA:

$$\psi = \psi(\sqrt{1 - \rho}\varepsilon^* + \sqrt{\rho}\phi^*)$$

where  $\psi$  is the overall standard deviation;  $\rho \in [0, 1]$  models the relative balance of spatial and non-spatial heterogeneity;  $\varepsilon^* \sim N(0, I)$  is the unstructured random error with fixed standard deviation of 1; and the ICAR random variable scaled so that  $\text{Var}(\phi^* = \phi/\nu) \approx 1$ . These assumptions ensure that  $\text{Var}(\varepsilon) \approx \text{Var}(\phi) \approx 1$  which in turn ensures that  $\psi$  can be interpreted as an overall standard deviation. The priors can then be more conventionally understood and assigned as some prior of the standard deviation such as a half-normal or a half-t or exponential distribution and a Beta(0.5, 0.5) for  $\rho$ . INLA's bym2 model by default uses the type 2 Gumbel distribution for  $\rho$ , that shrinks the distribution towards zero (Simpson et al. 2017). For the spatial process ("bym2") we use the recommended "pc" prior (rho=0.5, alpha=0.5). For the AR1 process, the "cor0" PC prior is used which has a "base" model of 0 correlation, that is, no correlation.

As abundance is generally modal along environmental gradients, the covariates are here modelled as a Random walk process of second order (i.e., 2nd derivative; called "rw2" in INLA); it is discretized along its quantiles (9

discretization points). This permits a smooth functional representation of the local relationship of the covariates which in the case of a uniformly spaced index  $i$  is:

$$\Delta^2 x_i = (x_i - x_{i+1}) - (x_{i+1} - x_{i+2}) \sim N(0, \tau^{-1})$$

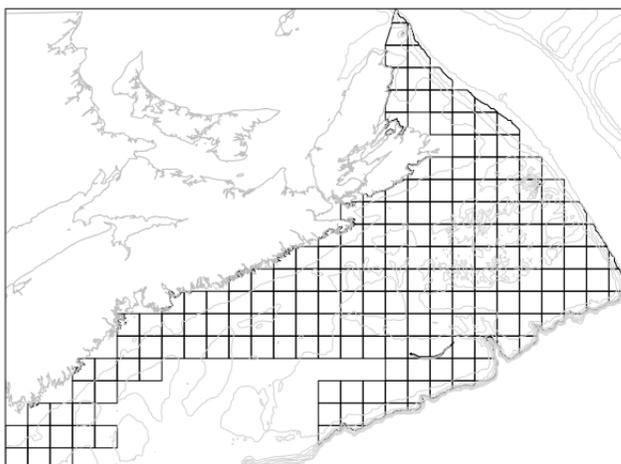
This is a local perspective and so philosophically akin to the CAR approach. The recommended “pc” prior for precision is used, an exponential function that decays to a “base” model of 0 precision. Though additive in the logarithmic link space, these are multiplicative for the Poisson models that follow.

## 2.4.2 Spatiotemporal Conditional autoregressive models

Time variation of the focal process of interest can also be treated as an autoregressive process. The argument, similar to the spatial case, is that there is a continuity in time for many processes and so an inherent temporal autocorrelation. If covariates exist and exhibit this temporal autocorrelation, then temporal autocorrelation is induced and enters mechanistically into the model. Any unaccounted or induced autocorrelation must, however, be explicitly accounted to ensure that  $\varepsilon \stackrel{iid}{\sim} f(\cdot)$ . In the time-series literature, these are known as the AR(I)MA type models (AutoRegressive Integrated Moving Average), often used when deterministic models do not exist or difficult to parameterize.

In a fisheries stock-assessment context, the stock assessment model is usually a constrained time-series model, constrained by assumptions of biological rate processes such as mortality and growth. Current practice is to aggregate of all AUs and treat it as a single integral time-series process. This is, of course, a stop-gap solution used to avoid the difficulties of model complexity and historical limitations of computational capacity. When all/most AUs act in an homogeneous manner, this approximation may be sufficient to express the focal process of interest. However, this is not always the case, especially in the oceanographic context as temporal and spatial variability is occurring at all scales.

At present, in keeping with the focus upon simple and operational autoregressive processes in CAR models and obtaining computationally viable solutions, we extend the CAR model by the simple addition of a temporal AR(1) process.



**Fig. 10.** Areal units in lattice structure (25 km x 25 km) used for snow crab assessment.

### 3 Comparison of results

There is flexibility in the choice of the areal units used in areal-unit-based models. The management regions could have been used (Figure 1); however, these would have been too coarse to permit identification of local aggregations and gaps in distributions, information that is useful in the assessment of stock status. An alternative is to use the groundfish areal units (Figure 4); however, they are really depth-based areal units, that do not account for other environmental covariates. Many of these latter units are also quite large and so not completely appropriate to the spatial scale of snow abundance patterns. Other alternatives are to alter any of the above to reduce the size of the larger polygons where required or even simply to use Voronoi-based polygons or meshing algorithms (e.g., see INLA). In the snow crab survey, the original design was lattice based (10 min X 10 min) and so, here, we have adopted a simple lattice structure, where areal units are defined as 25 km X 25 km units, slightly larger than the 10 min X 10 min original grid. This was found to be a reasonable balance between computational time and RAM required to model (important in more complex spatiotemporal CAR models) vs reasonably representative solutions sufficient to describe areal patterns. Of course, to expect an area of 625 km<sup>2</sup> to be homogenous in terms of ecosystem variability and snow crab densities is still a strong assumption (in the simple no-covariate case). AU's were constrained to those where there have been at least 3 surveys conducted since 1996. Polygon boundaries were further constrained to follow the 0 and 350 m depth contours, the operational limits of survey trawls and commercial fishing.

Aside:

All code for running the areal unit models presented in this section are found in:

<https://github.com/jae0/bio.snowcrab/blob/master/inst/scripts/>

03.abundance\_estimation\_carstm.r

99.snowcrab.framework.2020.R.

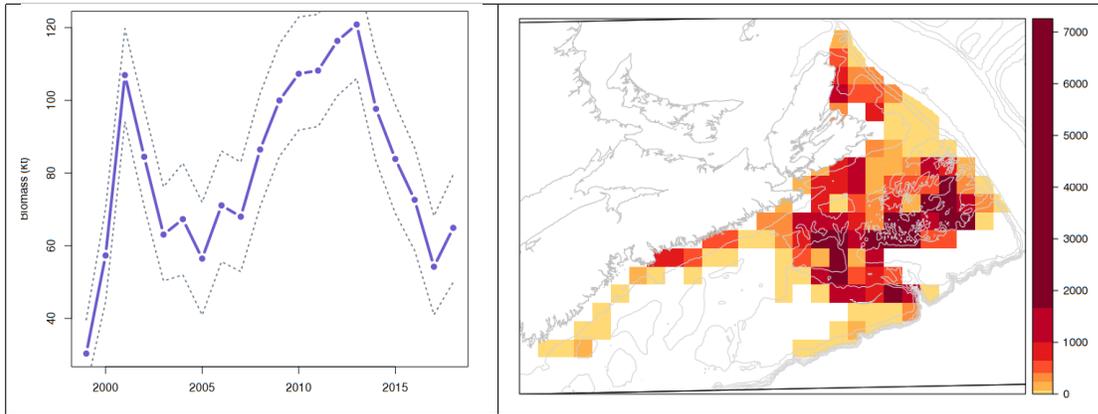
All code for running the stmv-based models are found in:

<https://github.com/jae0/stmv/blob/master/inst/scripts/>

[https://github.com/jae0/aegis.\\*/blob/master/inst/scripts/](https://github.com/jae0/aegis.*/blob/master/inst/scripts/)

#### 3.1 Non-spatial, non-temporal models

For the sake of simplicity of comparison with CAR-based models, in this section, we use the same AUs that were developed for CAR model-based computations (Figure 10). As these non-spatial and non-temporal representations are minimally parametrized; this means, essentially that they make strong assumptions upon the



**Fig. 11.** “Factorial crossed Gaussian (biomass)” model biomass (kt) variations over years assuming a non-spatial, non-temporal model (left) with 95% CI from marginal simulations. Note the abundance estimates for 1999, 2000, 2014 and 2017 are likely underestimated due to incomplete surveys not being accounted in the model. The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in  $\text{kg}/\text{km}^2$ ).

data: that autocorrelations and environmental variability induced spatiotemporal bias do not exist. Instead, each sample is assumed to be *perfectly representative* of the AU from which it is derived and importantly, that there are no relationships between areal units. Though unrealistic, this *random-stratified* design represents a simple basis from which to compare the appropriateness of an *experimental design* to represent the spatial process, especially as this is also a very frequently encountered assumption.

The most common distributional assumption is a Gaussian/Normal assumption:

$$Y_{st} \sim \text{Normal}(\mu_{st}, \psi \sigma_{st}^2)$$

$$\mu_{st} = \beta_{st} + \psi_{st}$$

The biomass variations in this model, which we will refer to as, “**Factorial crossed Gaussian (biomass)**” (Figure 11; Table 12) only represent aggregate estimates *where there are samples*. When AUs are not sampled due to logistical limitations (2017) or boat loss (2014), spatial bias will be induced as they must be *ad hoc* assumed to have some value (usually zero). Thus the low values in the years 1999 and 2000 are due to the small extent of the survey. Similarly, when visited at differing times of the year (pre-2004 vs post-2004), temporal (seasonal) bias can be induced. All these factors make using a factorial crossed approach nonsensical for interannual comparisons and are only shown here for heuristic value. See also Appendix 1. In the following we will focus especially upon the year 2017, as due to logistical issues the survey was not completed, resulting in some areas not being sampled (offshore; Figure 11).

As animal abundance distributions are generally logarithmically scaled, the following model form is potentially useful:

$$Y_{st} \sim \text{Normal}(\mu_{s,t}, \psi \sigma_{st}^2)$$

$$\ln(\mu_{s,t}) = \beta_{st} + \psi_{st}$$

Unfortunately, while fitting the positive-valued information reasonably well, there exist numerous zero-valued biomass density locations that must necessarily be ignored. This is due to  $\log(\text{zero})$  being undefined. It was (and still is) common practice in numerical ecology to add a small constant to these zeros, however, this is ill-advised as the magnitude of this constant can still have large effects and bias estimates tremendously. These zero-valued locations are, however, informative: some are due to real absence and others due to observation errors such as detection limits of sampling gear in low number situations or incomplete surveys. Zero-value locations would need to be treated as part of a second stage model that defines locations of zero-values. So again, a seemingly simple model but with a doubling of the number of parameters would be required. This model is not examined numerically in this report, though a similar approach is used in the *stmv*-based analysis. The additional computational complexity is counter to our desire to obtain a simple operational solution to data integration.

An alternative that is more functional is a Poisson assumption that operates upon numerical densities rather than biomass densities. It is also the simplest of these models as there is only an intensity parameter (vs a mean and variance in the Gaussian and Lognormal):

$$Y_{s,t} \sim \text{Poisson}(\mu_{s,t})$$

$$\ln(\mu_{s,t}) = \beta_{0|s,t} + O_{s,t} + \psi_{s,t}$$

where  $O$  are the offsets (swept area). Further, zero-values are permissible. If the fishery *synthesis* operates upon numerical abundance, this is an advantage. If biomass is required, the numbers need to be converted to biomass via some assumptions related to mean weight of individuals in a sample; this is straightforwardly accomplished as both weight and numbers are recorded. If absent it can also be imputed. As such, we favor the Poisson distributional assumption.

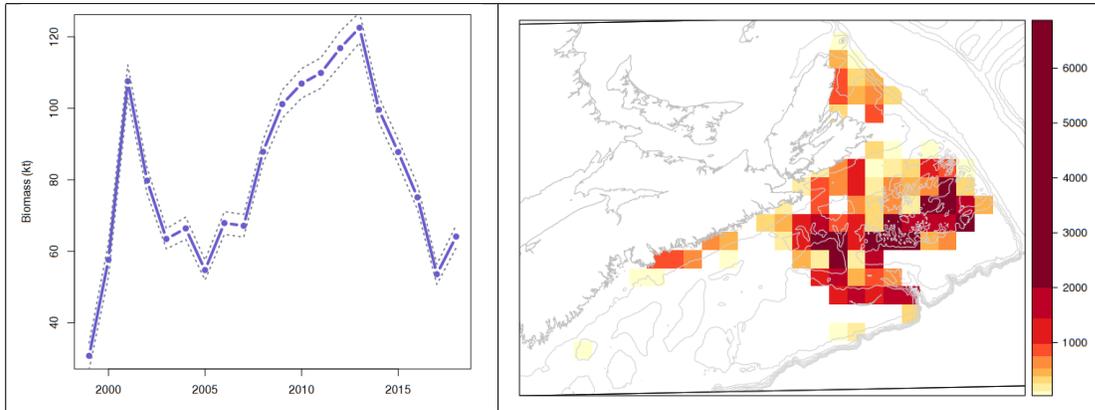
The aggregate temporal solutions derived from this Poisson distributional assumption which we will refer to as a “**Factorial crossed**” model (Figure 12; Table 1) are indeed very similar to those from a direct operation upon biomass and a Gaussian distributional assumption (“**Factorial crossed Gaussian (biomass)**”; Figure 11), though the variability in estimates are reduced. The differences are clearer in the spatial representation where we see the spatial aggregation is more clustered (less diffuse) in the Poisson model vs the Gaussian model. Again, this suffers from lack of assumptions on how to treat unsampled locations, rendering interannual comparisons meaningless.

It is worth emphasis again that in factorial models, when there are AUs that are not sampled, there is no way to estimate abundance in such areas as there are no assumptions to permit this. This in contrast to the CAR-based and the autocorrelation methods that do permit such imputation/interpolation, to which we now turn.

## 3.2 Continuous autocorrelation models

*Before continuing in this section, the reader is warned that the results of this approach are currently not available. The computational time required to complete this approach has now grown to 2-3 months of computing time. This is not operational and so is actually the reason for this new framework. What is presented in this section for the spatial-only portions is up to date, but the spatiotemporal results are from prior analyses and only here to provide context.*

For pragmatic reasons, namely computational time and computational resources, predictions are discretized to a 1 km X 1 km resolution and a temporal resolution to 10 sub-annual units. As the approach is continuous, these



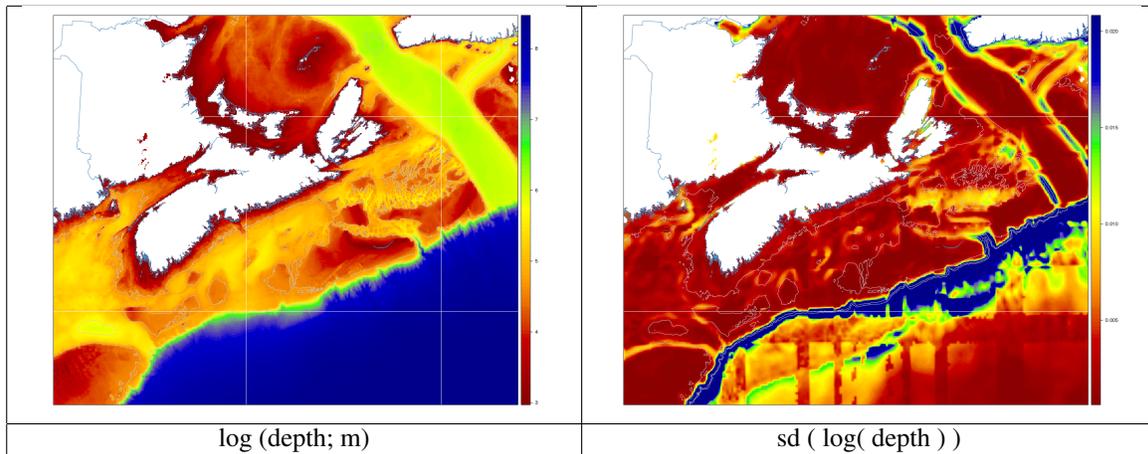
**Fig. 12.** “Factorial crossed” model biomass (kt) variations over years assuming a non-spatial, non-temporal model with a Poisson distributional assumption (left) with 95% CI from marginal simulations. Note the abundance estimates for 1999, 2000, 2014 and 2017 are likely underestimated due to incomplete surveys not being accounted in the model. The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in  $\text{kg}/\text{km}^2$ ).

**Tab. 1.** “Factorial crossed” model results. This is GLM using INLA for abundance estimation with a Poisson distributional assumption, with space and time being treated as factors. There is no autocorrelation.

```

Formula:
  inla(
formula = totno ~ -1 + year_factor:aid + offset(log(data_offset)),
family = "poisson",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 0.01, tolerance = 1e-04),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor = TRUE),
control.fixed = list(prec.intercept = 1))
Time used:
  Pre = 32.6, Running = 112, Post = 2.39, Total = 147
Fixed effects:
              mean      sd 0.025quant 0.5quant 0.975quant  mode kld
year_factor1999:aid1 -0.986 22.483   -54.061    3.115    31.004 15.538  0
...
year_factor2018:aid190 -1.094 21.889   -52.781    2.906    30.010 15.062  0
Expected number of effective parameters(stddev): 2590.09(0.00)
Number of equivalent replicates : 2.94
Deviance Information Criterion (DIC) .....: -
Deviance Information Criterion (DIC, saturated) ....: -
Effective number of parameters .....: -
Marginal log-Likelihood: -34659

```



**Fig. 13.** Variation of mean and total standard deviation of depth from a local (sub-domain, i.e., *stmv*) analysis. Note that mean and variance is not uniform across space.

discretizations are arbitrary and readily altered depending upon available computational resources and problem requirements. For snow crab assessments, this resolution was deemed a good balance between information gain, resolving *ecosystem variability* and data (RAM)/speed (CPU) requirements, costs and viability within the time frame of assessments (~1 month from data availability to assessment).

Each of the covariates used in this approach also requires modelling and prediction to obtain a consistent space-time basis. For snow crab, these covariates include: depth, bottom slope, bottom curvature, substrate grain size, bottom temperatures, and bottom species composition. The first two have temporal scales that are long relative to the time scale of snow crab and so can be treated as a pure spatial problem.

**Bathymetry** (depth; m; =-elevation) is one such feature that is highly informative in that it determines ambient light levels, surface complexity/rugosity, hydrodynamic stability and overall environmental stability. There are no covariates and so it is an extremely simple model. However, data density is extremely high and also variable. Sources are varied including, snow crab surveys, groundfish surveys, CHS, etc. Here, we model it as a localized Lognormal process:

$$\begin{aligned}
 \log(Y_s) &= F(\text{constant offset}) + \varphi_s, \\
 \varphi_s &\sim \text{Normal}(0, \sigma_\varphi^2), \\
 \varphi_{s|m} &= \omega_{s|m} + \varepsilon_{s|m}, \\
 \omega_{s|m} &\sim \text{GP}(0, C(\mathbf{s}, \mathbf{s}'; \boldsymbol{\theta}_m = \{\mathbf{v}_m, \phi_m, \sigma_m\})), \\
 \varepsilon_{s|m} &\sim \text{Normal}(0, \sigma_{\varepsilon|m}^2).
 \end{aligned}$$

The residual error process  $\varphi_{s|m}$  is readily decomposed into a local spatial process  $\omega_{s|m}$  and a local unstructured error  $\varepsilon_{s|m}$ . An FFT-based Matérn convolution implementation is used to express the spatial process for computational speed improvements. Data tapering is used to resolve short-scale variations and granular details; this tapering depends upon data density (90% to 50% similarity in spatial autocorrelation level tend to work well). Importantly, one observes that even in a "simple" variable such as depth, the mean and standard deviation vary in complex ways. That is, bathymetry is not first and second order stationary and further, that in more complex spatiotemporal models, simple separable models are likely to be inappropriate (Figure 13).

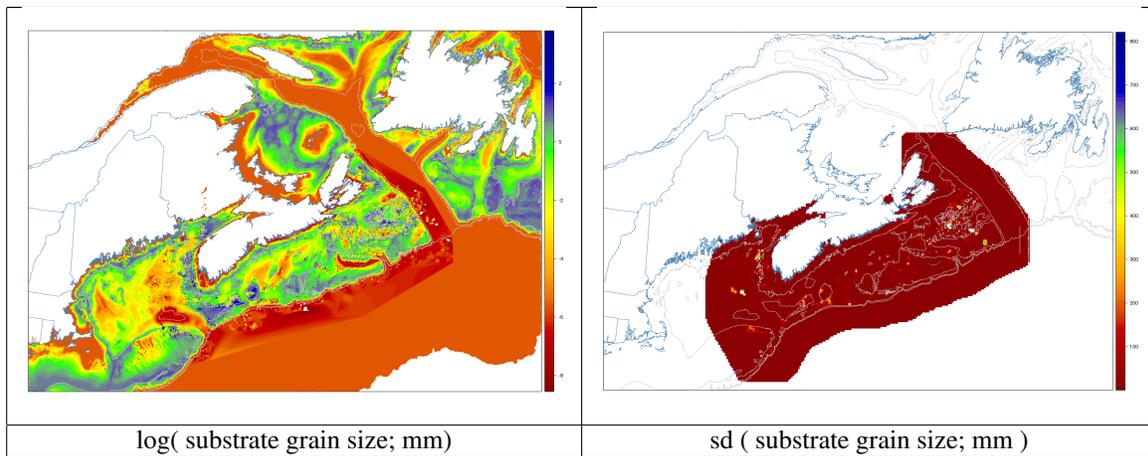
Similar to the case of bathymetry, **substrate grain size** (mm) can be considered a purely spatial model as it also varies at geological time scales. Of course, catastrophically rapid changes can and do occur but these can be still

**Tab. 2.** *Stmv* global model  $F()$  results for substrate grain size derived from Generalized additive models.

```

Family: gaussian
Link function: log
Formula:
substrate.grainsize ~ s(b.sdSpatial, k = 3, bs = "ts") + s(b.localrange,
  k = 3, bs = "ts") + s(log(z), k = 3, bs = "ts") + s(log(dZ),
  k = 3, bs = "ts") + s(log(ddZ), k = 3, bs = "ts")
Parametric coefficients:
      Estimate Std. Error t value Pr(>|t|)
(Intercept) -0.9053906  0.0111006 -81.5626 < 2.22e-16
Approximate significance of smooth terms:
      edf Ref.df      F    p-value
s(b.sdSpatial) 1.99628      2 304.92667 < 2.22e-16
s(b.localrange) 1.99992      2 4461.65359 < 2.22e-16
s(log(z))      1.99666      2 18144.34273 < 2.22e-16
s(log(dZ))     1.98167      2  26.51703 2.2327e-12
s(log(ddZ))    1.97150      2   7.92948 0.00031715
R-sq.(adj) =  0.14  Deviance explained = 13.8%
GCV = 5.6951  Scale est. = 5.695    n = 713021
AIC = 3263839.33

```

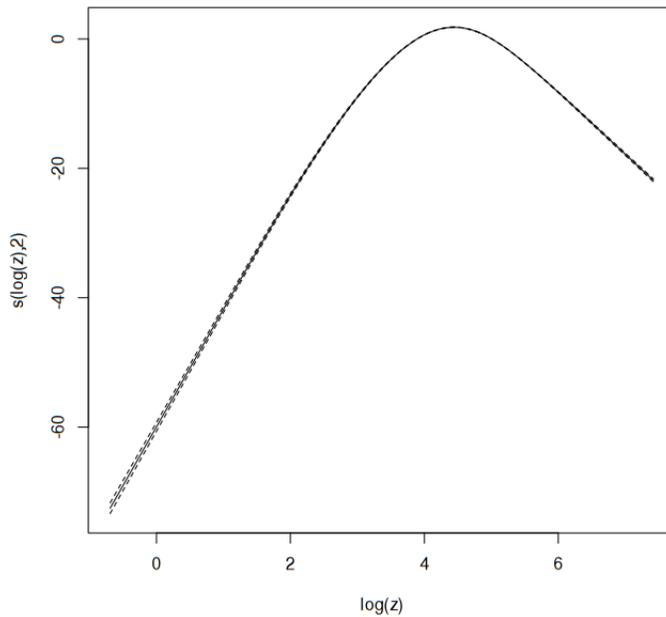


**Fig. 14.** Variation of mean and total standard deviation of substrate grain size from a local (sub-domain, i.e., *stmv*) analysis. Note that mean and variance is not uniform across space. The data are constrained to a smaller area as seen in  $sd(\text{substrate grain size})$  and that beyond this area it is extrapolated from the mean model which is driven by depth.

be considered statically as we are not focused upon the geological processes. Grain size is ultimately a proxy measure of the type of substrate (mud, sand, gravel, rock, etc.) and so can be informative for benthic, demersal and infaunal organisms. Unfortunately, the only available data currently available is an (over-)smoothed surface (Kostylev and Hannah 2007). Some data have been added from the snow crab surveys. It is also modelled as a Lognormal process, FFT based Matérn covariance convolution problem:

$$\begin{aligned}
 \log(Y_s) &= F(\text{depth, slope, curvature}) + \varphi_s, \\
 \varphi_s &\sim \text{Normal}(0, \sigma_\varphi^2), \\
 \varphi_{s|m} &= \omega_{s|m} + \varepsilon_{s|m}, \\
 \omega_{s|m} &\sim \text{GP}(0, C(\mathbf{s}, \mathbf{s}'; \theta_m = \{v_m, \phi_m, \sigma_m\})), \\
 \varepsilon_{s|m} &\sim \text{Normal}(0, \sigma_{\varepsilon|m}^2).
 \end{aligned}$$

Substrate grain size is clearly associated with depth. However, surprisingly, it is a modal relationship. From observational studies, one expected a decline in abundance with grain size. Instead, grain size in the 50 (=exp(4)) to 250 (=exp(5.5)) mm range seem to be optimal. Irrespective of the cause, this relationship therefore permits

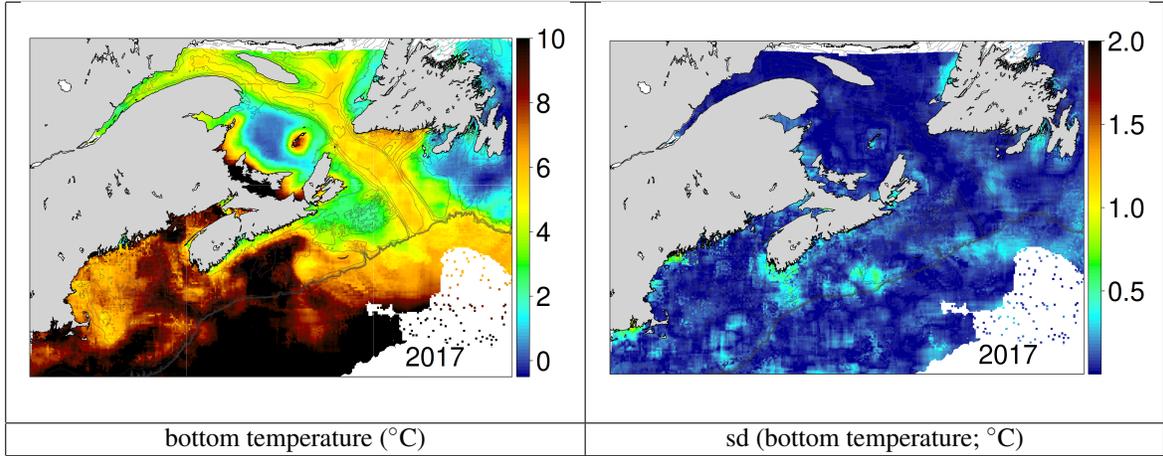


**Fig. 15.** Substrate grain size is highly associated with depth (Table 2). The form of the relationship is modal with a maximum grain size found near  $\exp(4.5) = 90$  m.

us to extrapolate well past the spatial domain of the data for substrate, by "borrowing" from the information content of *ecosystem variability* (in this case, depth-related variables).

**Temperature** is a fundamental modulator of metabolism, growth, reproduction, predator and prey distribution and abundance, disease incidence, species composition, *etc.* Bottom temperatures, in particular, are the focus due to their relevance to snow crab and modelled as an hierarchical, spatiotemporal, "nonseparable" spatiotemporal process. As they have high frequency variations, some additional complexity is required in modelling their spatiotemporal patterns. Here, the temporal effects are nested in spatial subdomains  $S_m$ . This is done using a two-harmonic approach where sinusoids are used to represent an inter-annual and a sub-annual variations in time-series at each statistical node using data assessed to be part of the local sub-domain. Predictions in space are then estimated as a pure spatial problem under the now correct assumption of local stationarity, for each time slice. In other words, it is a non-separable (hierarchical) model.

The global covariate model is simply an intercept model with an identity link such that  $\varphi_{s,t}$  are centered upon zero. Salinity or water density data can conceivably enter to delineate water masses and origins, however, this data does not exist at sufficient quality, density and coverage to be informative enough to merit the additional computational load (at present). Instead, the residuals errors are modelled locally in each subdomain as a time-series with two Fourier harmonics in time (an inter-annual and a sub-annual/seasonal component). Additional penalized thin-plate spline smooth terms for local depth and location are used to resolve local spatial trends and aliasing to third order or less (via shrinkage). Temporal predictions at each spatial datum are then used to "augment" the modelling of the spatial processes  $\varphi_{s,t|m}^*$  which are treated independently for each time slice as a **Gaussian process**. The temporal autocorrelation is, therefore, carried only indirectly by the individual temporal processes centered at each spatial datum. For faster computations, a Fast Fourier Transform (FFT) based convolution method is used to approximate the spatial Gaussian process. The model specification is, therefore:



**Fig. 16.** Variation of mean and total standard deviation of bottom temperature for 2017, from a local (sub-domain, i.e., *stmv*) analysis. Note that mean and variance is not uniform across space.

$$\begin{aligned}
 Y_{s,t} &= F(\text{identity}) + \varphi_{s,t}, \\
 \varphi_{s,t} &\sim \text{Normal}(0, \sigma_{\varphi}^2), \\
 \varphi_{s,t|m} &= f_m(\text{interannual,seasonal,northing,easting,depth}) + \zeta_{s,t|m} \\
 \zeta_{s,t|m} &\sim \text{Normal}(0, \sigma_{\zeta|m}^2), \\
 \varphi_{s,t|m}^* &= \omega_{s,t|m} + \varepsilon_{s,t|m}, \\
 \omega_{s,t|m} &\sim \text{GP}(0, C(\mathbf{s}, \mathbf{s}'; \boldsymbol{\theta}_{t|m} = \{v_{t|m}, \phi_{t|m}, \sigma_{t|m}\})), \\
 \varepsilon_{s,t|m} &\sim \text{Normal}(0, \sigma_{\varepsilon|m}^2).
 \end{aligned}$$

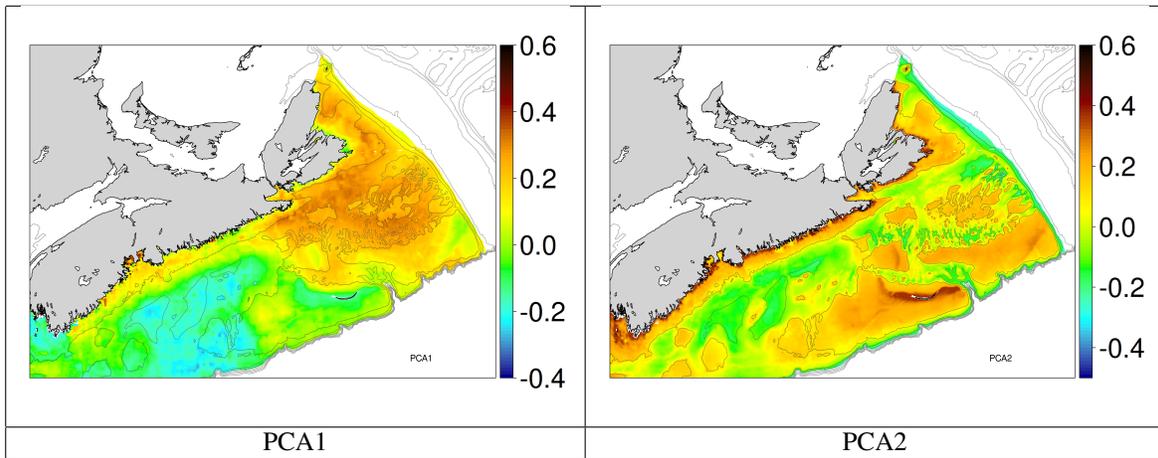
The temperature data originate from a number of different sources, including the snow crab survey, groundfish survey, AZMP survey, FSRS survey, scallop survey and many other opportunistic samples maintained and kindly provided by Roger Petitpas (OSD, DFO).

The results are not presented in this report as computational time is significant (> 2 months) given the computing resources available. The depictions are from earlier versions of the model (Figure 16).

Bottom species composition is modelled in a similar manner to bottom temperature. "Species composition" is simply the first two eigen-decompositions (Principal components analysis eigenvectors) of the correlation matrix of species abundance.

$$\begin{aligned}
 Y_{s,t} &= F(\text{identity}) + \varphi_{s,t}, \\
 \varphi_{s,t} &\sim \text{Normal}(0, \sigma_{\varphi}^2), \\
 \varphi_{s,t|m} &= f_m(\text{interannual,seasonal,northing,easting,depth, temperature}) + \zeta_{s,t|m} \\
 \zeta_{s,t|m} &\sim \text{Normal}(0, \sigma_{\zeta|m}^2), \\
 \varphi_{s,t|m}^* &= \omega_{s,t|m} + \varepsilon_{s,t|m}, \\
 \omega_{s,t|m} &\sim \text{GP}(0, C(\mathbf{s}, \mathbf{s}'; \boldsymbol{\theta}_{t|m} = \{v_{t|m}, \phi_{t|m}, \sigma_{t|m}\})), \\
 \varepsilon_{s,t|m} &\sim \text{Normal}(0, \sigma_{\varepsilon|m}^2).
 \end{aligned}$$

The data for this analysis is derived from the snow crab surveys as well as the groundfish surveys. Time period is from 1999 to present as prior to this period in the groundfish surveys, species identification was inconsistent for



**Fig. 17.** Variation of mean and total standard deviation of species composition for 2017, from a local (sub-domain, i.e., *stmv*) analysis. Note that mean is not uniform across space.

most invertebrates. The first axis of representation is primarily associated with temperatures while the second is primarily associated with depths.

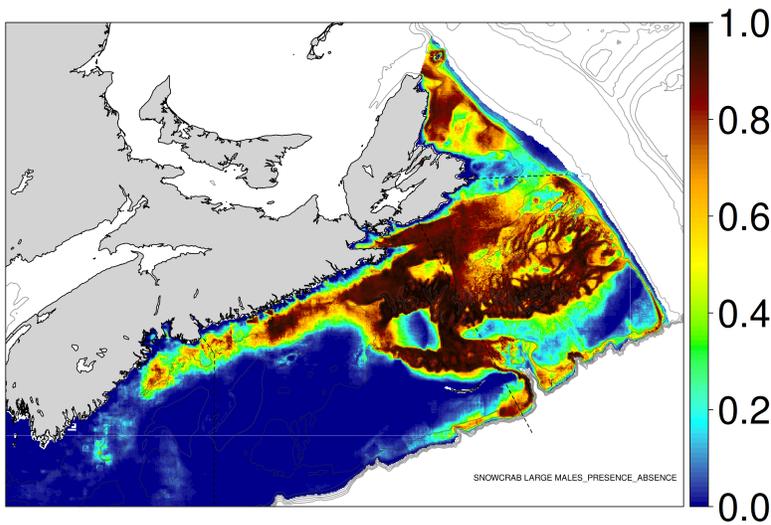
The snow crab estimation process uses the above and other covariates interpolated to the same spatial support of the snow crab assessment for prediction. This is required to refine predictions of snow crab abundance and habitat while avoiding issues of bias due to aliasing.

Additional covariates that express the ecosystem state at a given time and location ("indicators") are informative in delineating spatiotemporal processes that are structured from those that are random. Their model formulation is similar in that they follow a similar model structure with temporal effects nested in spatial subdomains and the use of link functions in a Generalized Linear Model/Generalized Additive Model setting where the covariates used to model these indicators rely upon spatial predictions of depth and substrate grain size and the spatial derivatives of the former (slope and curvature). The spatiotemporal error process is modelled locally in each subdomain as a space-time "nonseparable" model, using time-varying covariates related to bottom temperature variations and associated statistics.

For the estimation of habitat preferences and the creation of species distribution maps that rely upon presence-absence data. The data  $Y$  are assumed to come from a Bernoulli binomial process with a logit link function  $g(\cdot)$ :

$$\begin{aligned}
 \text{logit}(Y_{s,t}) &= F(\text{depth, slope, curvature, substrate grainsize}) + \varphi_{s,t} \\
 \varphi_{s,t} &\sim \text{Normal}(0, \sigma_{\varphi}^2), \\
 \varphi_{s,t|m} &= f_m(\text{ecosystem indicators}) + \zeta_{s,t|m} \\
 \zeta_{s,t|m} &\sim \text{Normal}(0, \sigma_{\zeta|m}^2), \\
 \varphi_{s,t|m}^* &= \omega_{s,t|m} + \varepsilon_{s,t|m}, \\
 \omega_{s,t|m} &\sim \text{GP}(0, C(\mathbf{s}, \mathbf{s}'; \theta_{r|m} = \{v_{t|m}, \phi_{t|m}, \sigma_{r|m}\})), \\
 \varepsilon_{s,t|m} &\sim \text{Normal}(0, \sigma_{\varepsilon|m}^2).
 \end{aligned}$$

For the estimation of abundance, the positive valued data  $Y$  are assumed to come from a lognormal process:



**Fig. 18.** Probability of observing snow crab for 2017, from a local analysis (sub-domain, i.e., *stmv*).

$$\begin{aligned}
 \log(Y_{s,t}) &= F(\text{depth, slope, curvature, substrate grainsize}) + \varphi_{s,t} \\
 \varphi_{s,t} &\sim \text{Normal}(0, \sigma_{\varphi}^2), \\
 \varphi_{s,t|m} &= f_m(\text{ecosystem indicators}) + \zeta_{s,t|m} \\
 \zeta_{s,t|m} &\sim \text{Normal}(0, \sigma_{\zeta|m}^2), \\
 \varphi_{s,t|m}^* &= \omega_{s,t|m} + \varepsilon_{s,t|m}, \\
 \omega_{s,t|m} &\sim \text{GP}(0, C(\mathbf{s}, \mathbf{s}'; \boldsymbol{\theta}_{t|m} = \{\nu_{t|m}, \phi_{t|m}, \sigma_{t|m}\})), \\
 \varepsilon_{s,t|m} &\sim \text{Normal}(0, \sigma_{\varepsilon|m}^2).
 \end{aligned}$$

The results of this analytical approach estimated  $y_{st}$  as shown in Figure (19) as well as a direct measure of the habitat in terms of probability of observing a snow crab (Figure 18). The results using updated methodology are unfortunately not available and currently not operational due to the computational power/time demands of the approach.

### 3.3 Conditional autoregressive models

The CAR-based results are based upon a discretization of space to a 25 km X 25 km scale and 10 sub-annual units in time (Figure 10). Neighbourhood structure was determined based upon a Queen-adjacency rule (any contact along all borders, including diagonals). Smaller units could have been chosen to help identify local processes, however, this would make computations slower and also run the risk of imputing data for a very large proportion of areal units where there would be insufficient sampling and neighbours with no sampling. Conversely, larger units would result in fewer imputations but areal units would not be able to resolve important spatial features. Irrespective, different polygon formulations can be explored in future to see if any additional benefit in precision and accuracy could be accrued relative to the costs of computational time, solution instability.

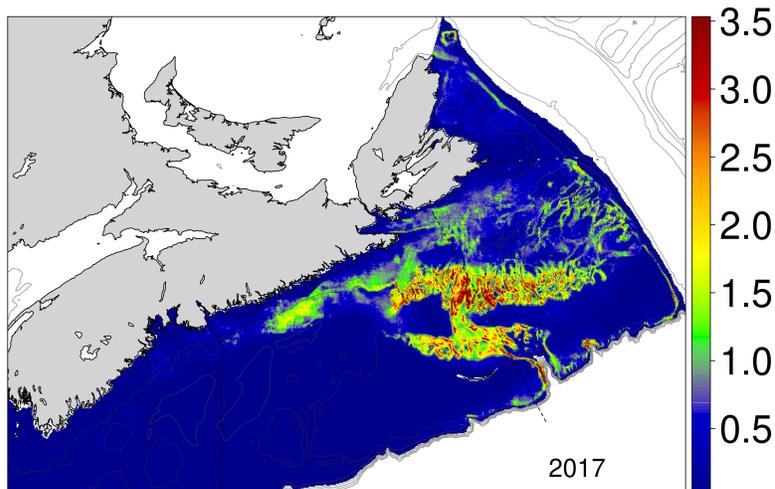


Fig. 19. Snow crab biomass density for 2017, from a local analysis (sub-domain, i.e., *stmv*) analysis.

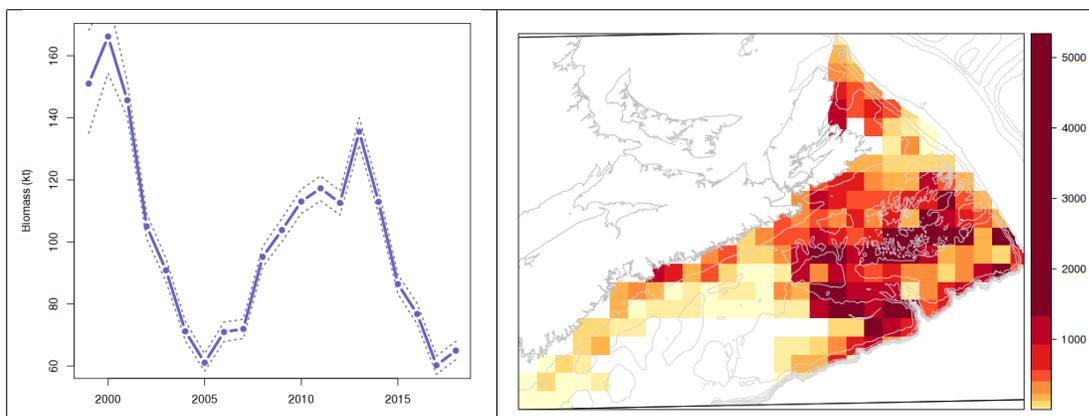
### 3.3.1 Poisson-CAR model

As we favor the Poisson distributional model of numerical abundance (Figure 12), we add to it a single CAR spatial autocorrelation as a random effect to see the incremental effect (Table 3; Figure 20). We will refer to this as the “**Mixed effects simple**” model, as space is treated as a random effect and year as a fixed effect. This assumption permits estimation in unsampled areal units. The temporal trend is similar to the non-CAR model, however, the aggregate magnitudes are higher as can be expected relative to the fixed effects models. In the spatial representation, we can see that this is due to a rather optimistic expectation (extrapolation/imputation) of unsampled neighbouring AUs. Further, as sampling is not completely representative of the domain and in fact preferential (Figure 6), these estimates are expected to be positively biased.

### 3.3.2 Ecosystem variability of static variables

Ecosystem variability needs to be accounted to try to adjust for this preferential sampling. This is readily done in a CAR-based inference and prediction schema as covariates (and year is still treated as a factor, so this would be a “mixed effects” model). These factors, however, need to be expressed upon the same space and time scales as the focal process of interest. This can be accomplished in an *ad hoc* manner from available data. However, such attempts would have the same basic problems associated with reduced accuracy and precision as with any *Cartesian* approach. It would be prudent to model them. Though it is possible to use *stmv*-modelled fields, it is assumption-wise, and methodologically more self-consistent to use CAR-based methods for these covariate fields. Of note, however, is that due to the very large size of these AUs relative to the discretization in *stmv*, some variables such as bottom slope and curvature are not as informative relative and so not computed. At present only a limited number of ecosystem variables are available for systemic analysis. Here, we present a summary of some of these results.

First and foremost is bottom depth due to its correlation with a large number of biological, physical and chemical processes with depth. The CAR-based solution summarizes depths upscaled at each AU, appropriate to the autocorrelation and unstructured error assumptions (Figure 21 and 22). The spatial variations are consistent



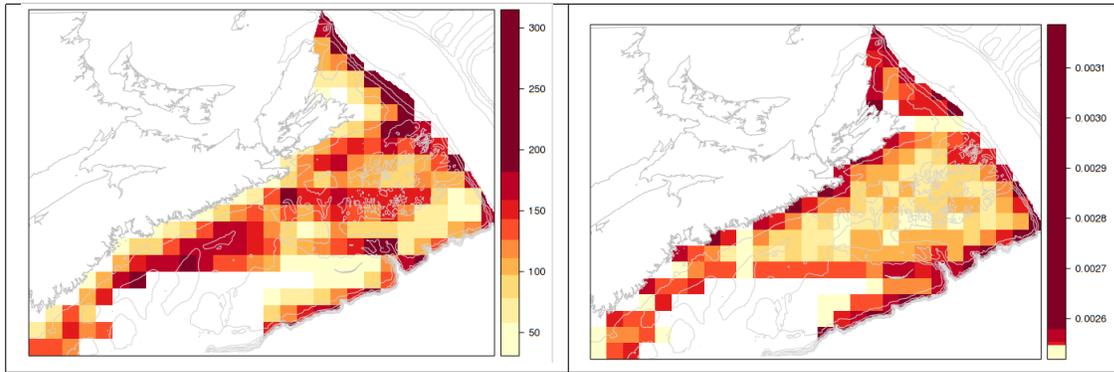
**Fig. 20.** “Mixed effects simple” model. Biomass (kt) variations over years assuming a single CAR random effect, and years as a fixed effect model with a Poisson distributional assumption (left). The 95% CI are based upon posterior simulations from the joint distributions. The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in  $\text{kg}/\text{km}^2$ ).

**Tab. 3.** Results for the “Mixed effects simple” model.

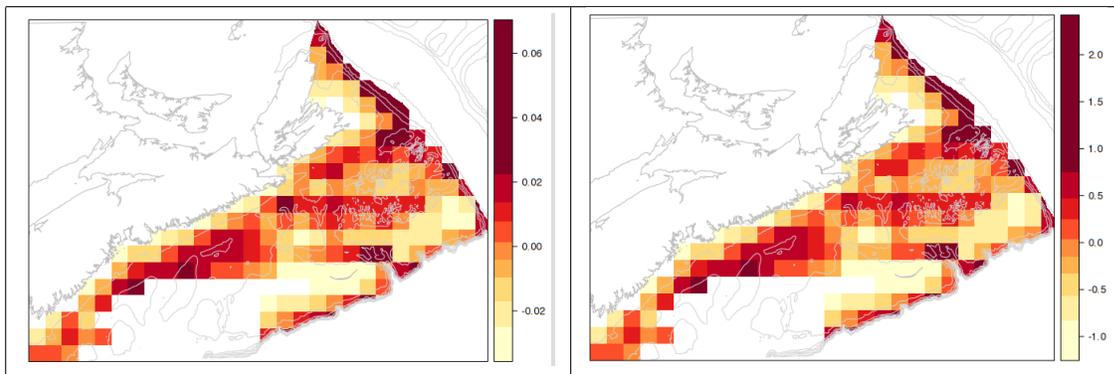
```

Formula:
inla(formula = totno ~ 1
+ year_factor
+ f(auid, model="bym2", graph=sppoly@nb, scale.model=TRUE, constr=TRUE, hyper= H$bym2)
+ offset(log(data_offset)),
family="poisson",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 1e-06, tolerance = 1e-12),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor = TRUE),
control.fixed = H$fixed)
Fixed effects:
      mean      sd 0.025quant 0.5quant 0.975quant   mode kld
(Intercept)  6.721 0.088      6.546   6.722      6.894  6.723  0
year_factor2000 0.101 0.068      -0.031   0.101      0.235  0.100  0
...
year_factor2018 -0.808 0.061      -0.926  -0.808      -0.687 -0.809  0
Model hyperparameters:
      mean      sd 0.025quant 0.5quant 0.975quant   mode
Precision for auid 0.313 0.059      0.201   0.313      0.428 0.316
Phi for auid      0.836 0.095      0.621   0.849      0.972 0.906

```



**Fig. 21.** Bathymetry (m) CAR-analysis: spatial predictions (left), and standard errors (right).



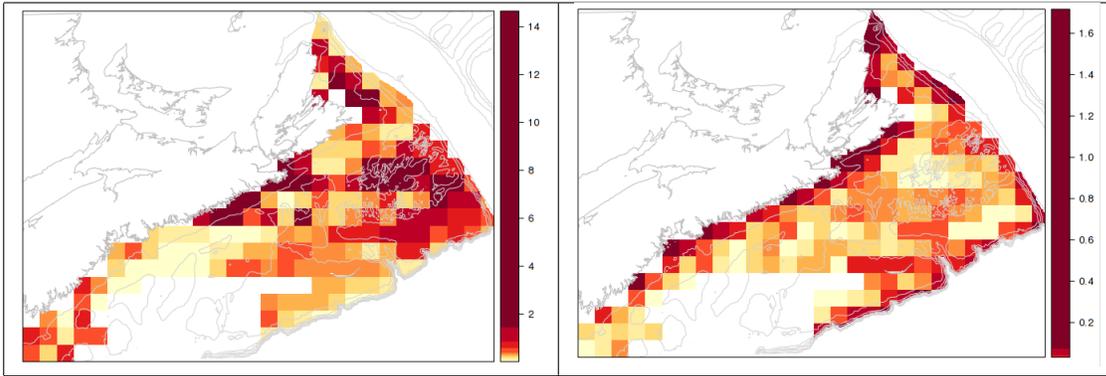
**Fig. 22.** Bathymetry (m) CAR-analysis: unstructured errors  $\varepsilon$  (left), and spatial error components  $\phi$  (right).

with known bathymetry in the region (Figure 13). The representation is least satisfactory in areas of rapidly changing depths in short distances (e.g., Misaine bank, near shore areas) where the AUs due to their large size were not able to resolve the variability resulting in a reduction of the mean to an intermediate magnitude (Figure 21). Of note, the spatial correlation components dominate over the unstructured error, indicating that between AU correlations are high (Figure 22). Indeed, the posterior distribution of the proportionality factor  $\rho$  that indicates the relative importance of spatial error is much higher than unstructured errors: mode of 0.996 and a 95% Bayesian Credible Interval (BCI): 0.922, 0.999).

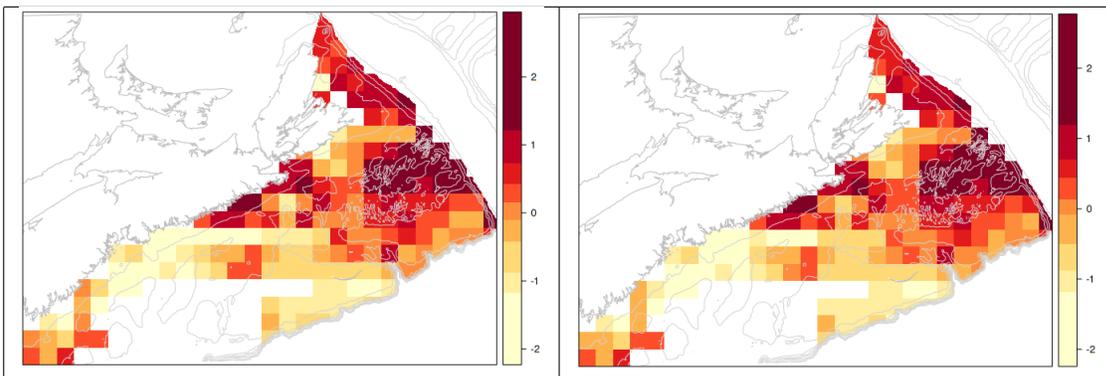
Substrate grain size is also an important factor for benthic and demersal organisms as they indicate the predominant geological and hydro-dynamic processes in a region. For the area of interest, grain size is associated with depth. Spatial autocorrelated error processes dominate over unstructured errors, with a posterior  $\rho$  mode of 0.959 and a 95% BCI of (0.867, 0.995).

Note that there are some differences between the relationship of grain size and depth in the stmv (Figure 15) vs CAR (Figure 25) approaches. Though maximum grain size is attained near the same depths (50-100 m), the modality is less pronounced. This demonstrates the effect of spatial resolution and aggregation of data (often referred to as the “Map area unit problem”) that reduces the ability of the CAR approach to resolve small-scaled variations, especially in topographically highly variable areas, nearshore, from 10 to 30 m (Figure 15).

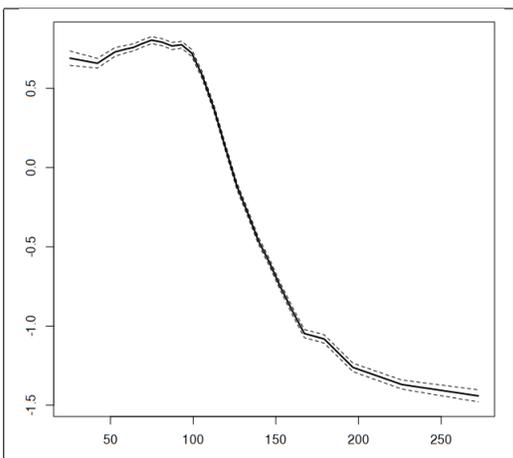
We will refer to the addition of these static covariates to the “**Mixed effects simple**” model as a “**Mixed effects static**” model. We observe that spatial random effects still dominated over non-spatial errors in the “**Mixed effects static**” model, at  $\rho = 0.89$  (Table 4), a marginal increase from  $\rho = 0.84$  in the “**Mixed effects simple**”



**Fig. 23.** Substrate grain size (mm) CAR-analysis: spatial predictions (left), and standard errors (right).



**Fig. 24.** Substrate grain size (mm) CAR-analysis: unstructured errors  $\epsilon$  (left), and spatial error components  $\phi$  (right).



**Fig. 25.** Substrate grain size's (log: mm) relationship with depth (m). Most variation in substrate grain size is seen from 100 m to 200 m, beyond which it stabilizes.

**Tab. 4.** Results for the “Mixed effects static” model.

```

Formula:
inla(formula = totno ~ 1
+ year_factor
      + f(inla.group(z, method="quantile", n=9), model="rw2", scale.model=TRUE, hyper=H$rw2)
      + f(inla.group(substrate.grainsize, method="quantile", n=9 ), model="rw2",
          scale.model=TRUE, hyper=H$rw2)
      + f(auid, model="bym2", graph=sppoly@nb, scale.model=TRUE, constr=TRUE, hyper= H$bym2)
+ offset(log(data_offset)),
family="poisson",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 1e-06, tolerance = 1e-12),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor = TRUE),
control.fixed = H$fixed)
Fixed effects:

```

|                 | mean   | sd    | 0.025quant | 0.5quant | 0.975quant | mode   | kld |
|-----------------|--------|-------|------------|----------|------------|--------|-----|
| (Intercept)     | 6.519  | 0.086 | 6.349      | 6.520    | 6.686      | 6.521  | 0   |
| year_factor2000 | 0.160  | 0.068 | 0.027      | 0.160    | 0.295      | 0.159  | 0   |
| ...             |        |       |            |          |            |        |     |
| year_factor2018 | -0.706 | 0.062 | -0.826     | -0.707   | -0.584     | -0.707 | 0   |

```

Model hyperparameters:

```

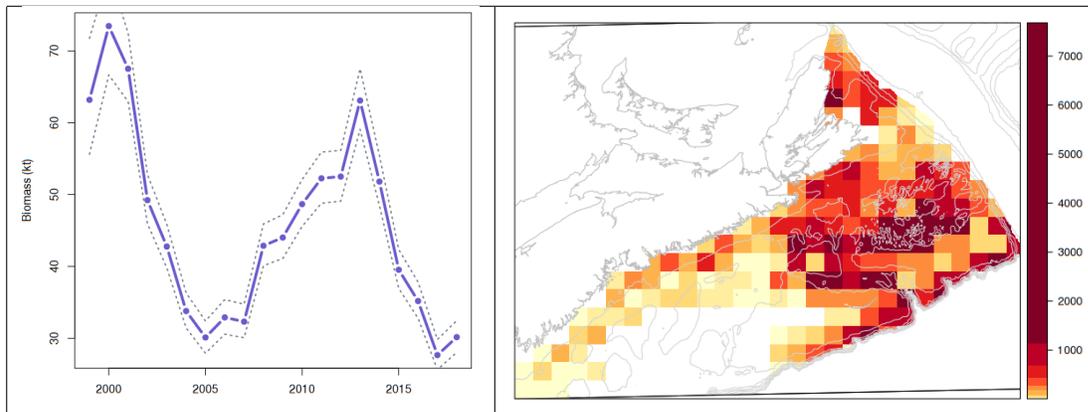
|  | mean  | sd    |
|--|-------|-------|
| Prec inla.group(z, method = "quantile", n = 9)                   | 4.449 | 1.779 |
| Prec inla.group(substrate.grainsize, method = "quantile", n = 9) | 0.211 | 0.056 |
| Prec auid  | 0.298 | 0.047 |
| Phi for auid   | 0.889 | 0.070 |

model (Table3). Based upon various information theoretic criteria (DIC, WAIC, CPO; Table 12), it is as improvement relative to the “Mixed effects simple” model. But this came at a cost of an increase in the number of effective parameters, associated with the covariate parametrizations. Overall, the magnitudes are reduced with a shifting of densities to areas that are considered to be positively associated with numerical density, based upon covariate information, mostly depth (Figure 27). Substrate grain size has a weaker effect, more variable effect/relationship in part as it is itself strongly associated with depth; however, a marginal increase in density is suggested with increase in grain size which is counter to empirical observations. This is likely due to empirical observation of substrate grain size being confounded with depth (Hooper 1986), and potentially not having adjusted for this effect.

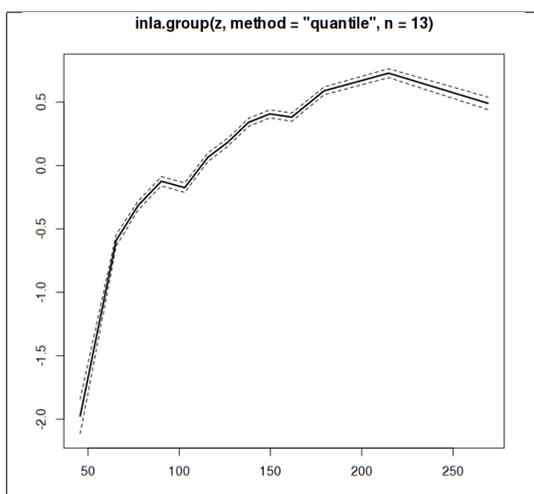
### 3.3.3 Ecosystem variability of dynamic variables

**Temperature** is a pervasive biological factor that influences snow crab at many levels from indirectly via food availability, predator prevalence, and disease prevalence to directly via growth, reproduction and basic metabolic processes. To add this factor requires that it too be modelled to a coherent state. Here, we focus upon bottom temperatures as it dominates the life time of snow crab. The highly spatiotemporal variability of bottom temperature can be simplistically modelled as a non-separable CAR-AR1 (Table 5; Figure 29). For the sake of computational tractability, we nest the CAR effects within each year and have an AR1 process to model yearly variations. Additionally, we add an additive (separable) AR1 process for seasonality and a depth-related RW2 process (second order random walk, exponential) to account for depth-related temperature variability. Spatial autocorrelations accounted for 42.7% of the variability (Table 5; Figure 30). Temporal autocorrelations at a yearly scale was 0.841 while on an sub-annual scale, the autocorrelation was 0.725 (Table 5).

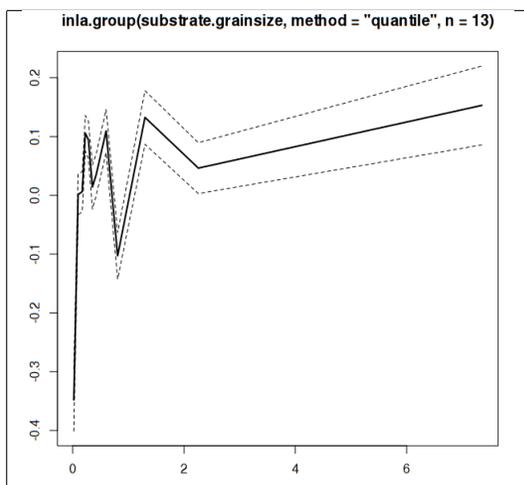
The relationship of bottom temperatures with other environmental covariates are complex. Relative to bottom depth, a pot-and-handle type relationship is seen with minimal temperatures being regularly observed at intermediate depths in the 70 to 90m range (Figure 31). This is mostly associated with the Nova Scotia current, a cold water mass flowing along those depth, close to the coastline.



**Fig. 26.** “Mixed effects static” model results. Poisson model-based biomass (kt) variations over years assuming a single CAR random effect and substrate grain size and depth as covariates with a Poisson distributional assumption (left). The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in  $\text{kg}/\text{km}^2$ ).



**Fig. 27.** Numerical density vs depth (m) after adjustment for CAR-based errors (centered on the global mean:  $6.723/\text{km}^2$ ).



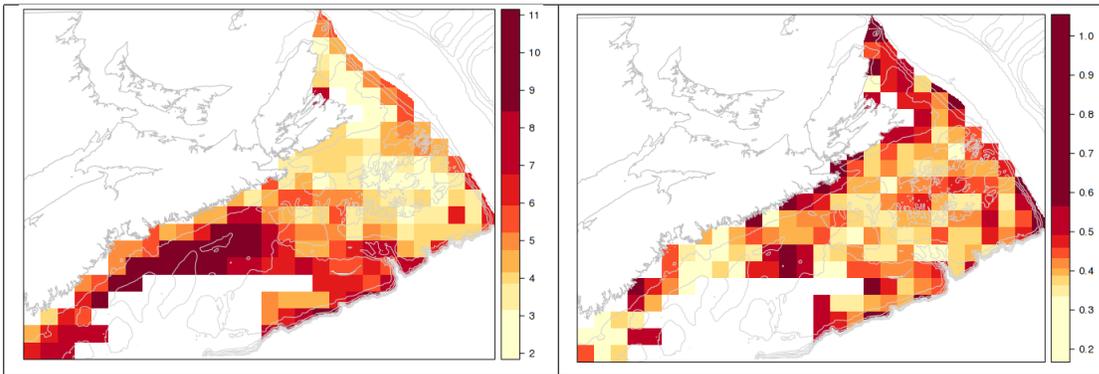
**Fig. 28.** Numerical density vs substrate grain size (mm) after adjustment for CAR-based errors (centered on the global mean:  $6.723/\text{km}^2$ ).

**Tab. 5.** Results for bottom temperature estimation CAR at each year.

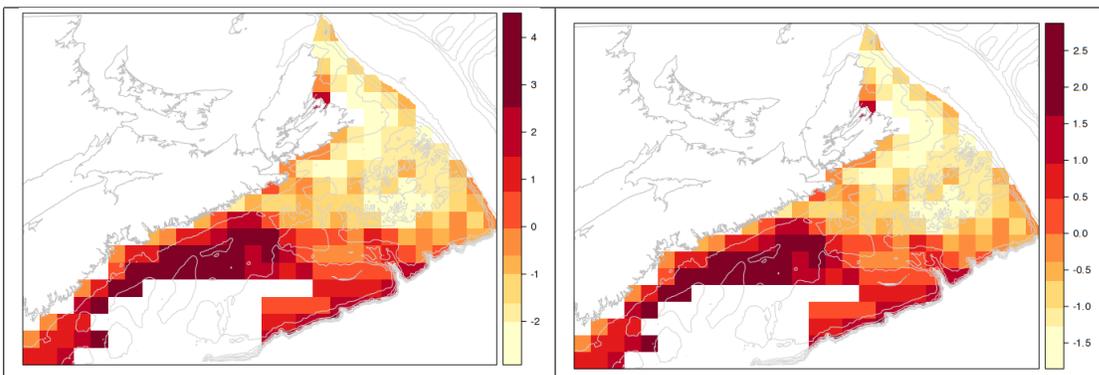
```

Formula:
inla(formula = t ~ 1
+ f( dyri, model="ar1", hyper=H$ar1 )
+ f( inla.group( z, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( auid, model="bym2", graph=spoly@nb, group=year_factor, scale.model=TRUE, constr=TRUE,
hyper=H$bym2, control.group=list(model="ar1", hyper=H$ar1_group)),
family = "normal",
data= M,
control.compute=list(dic=TRUE, config=TRUE),
control.results=list(return.marginals.random=TRUE, return.marginals.predictor=TRUE ),
control.predictor=list(compute=FALSE, link=1 ),
control.fixed=H$fixed, # priors for fixed effects, generic is ok
)
Fixed effects:
      mean   sd 0.025quant 0.5quant 0.975quant  mode kld
(Intercept) 4.52 0.285      3.937   4.522   5.089 4.524  0
Model hyperparameters:
                mean   sd
Precision for the Gaussian observations      0.504 0.005
Precision for dyri                          6.302 2.551
Rho for dyri                                0.725 0.121
Precision for inla.group(z, method = "quantile", n=9) 0.413 0.111
Precision for auid                          0.427 0.024
Phi for auid                                0.427 0.034
GroupRho for auid                           0.841 0.011
Deviance Information Criterion (DIC) .....: 96494.96
Deviance Information Criterion (DIC, saturated) ....: 28729.33
Effective number of parameters .....: 1862.94
Marginal log-Likelihood: -47560.07

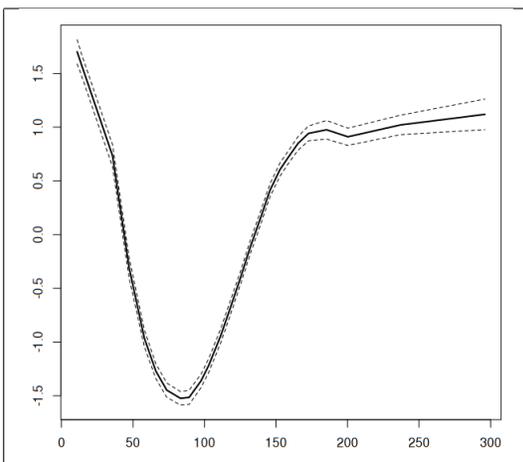
```



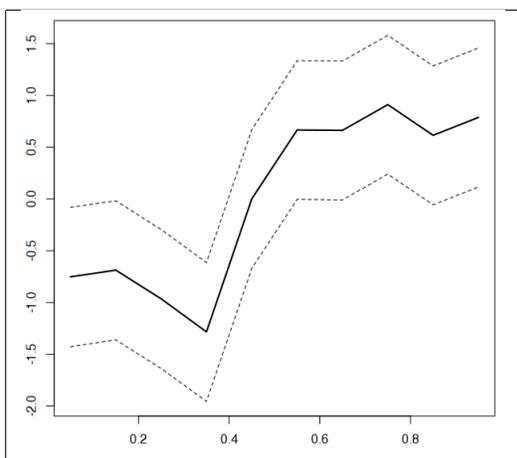
**Fig. 29.** Bottom temperature ( $^{\circ}\text{C}$ ) CAR-analysis: for early October 2000, spatial predictions (left), and standard errors (right).



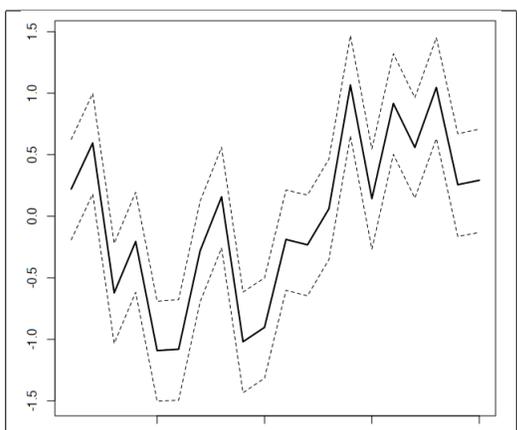
**Fig. 30.** Bottom temperature ( $^{\circ}\text{C}$ ) CAR-analysis: for the year 2000, unstructured errors  $\varepsilon$  (left), and spatial error components  $\phi$  (right). The posterior mode for  $\rho$ , the proportion of error attributable to spatial effects is 0.983 (95% BCI: 0.983, 1.00).



**Fig. 31.** Bottom temperature ( $^{\circ}\text{C}$ ) vs depth (m) after adjustment for CAR-based errors (centered on the global mean:  $4.52^{\circ}\text{C}$ ).



**Fig. 32.** Bottom temperature variations – seasonal trends (centered on the global mean: 4.52°C).



**Fig. 33.** Bottom temperature variations – annual trends from 1999 to 2018 (centered on the global mean: 4.52°C). This indicates that temperatures have been elevated for the last eight years..

In terms of seasonality in the domain of the snow crab assessment, average bottom temperatures are lowest from Jan to May and then rise significant for the remainder of the year, on average, about 2°C (Figure 32).

Inter-annual variability is constrained to the time span of the snow crab series, although it does not have to be. In that time frame, the average changes have been of the same magnitudes as that of the seasonal variations (about 2°C; Figure 33). From the year 2010 to 2018, average bottom temperatures have been higher, relative to the colder period from 2001 to 2009.

**Species composition/assemblage** at a given location and time is a highly informative marker for ambient environmental conditions. The first two (dominant) eigenvectors summaries information related to the temperature and depth of a given location (Figures 17) and are each subjected to the same model are temperature, with the addition of temperature as an additional covariate. For the first axis (PCA1), approximately 51.9% of the error can be attributed to spatial structure (Table 6) while for the second, this was 53.8% (Table 7).

The overall predictions were similar to the *stmv*-based results; however, the effects of aggregation is to spatially smear the patterns (Figures 34, 36). The interpretations of the axes are similar, with PCA1 being mostly a temperature-driven pattern of species association (Table 6; Figure 7) and PCA2 being one associated with

**Tab. 6.** Results for species composition axis 1 – CAR at each year.

```

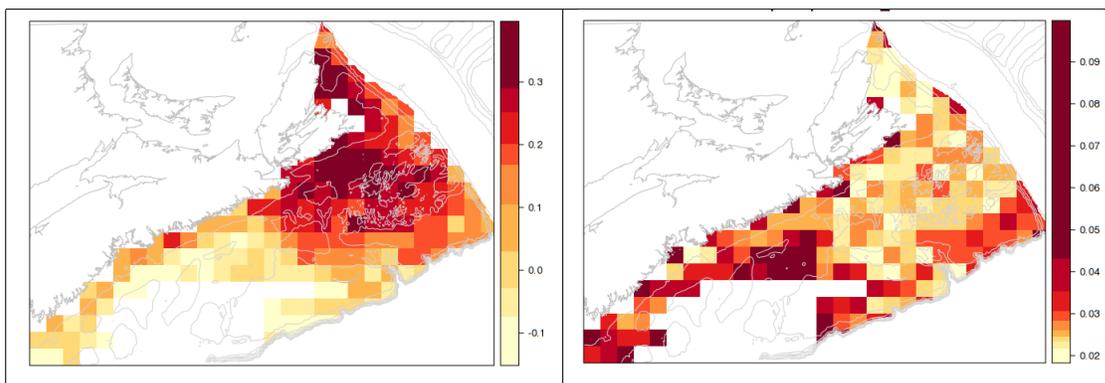
Formula:
inla(formula = pca1 ~ 1
+ f( dyri, model="ar1", hyper=H$ar1 )
+ f( inla.group( t, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( z, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( substrate.grainsize, method="quantile", n=9 ), model="rw2", scale.model=TRUE,
hyper=H$rw2)
+ f( auid, model="bym2", graph=sppoly@nb, group=year_factor, scale.model=TRUE, constr=TRUE,
hyper=H$bym2, control.group=list(model="ar1", hyper=H$ar1_group)),
family = "normal",
data= M,
control.compute=list(dic=TRUE, config=TRUE),
control.results=list(return.marginals.random=TRUE, return.marginals.predictor=TRUE ),
control.predictor=list(compute=FALSE, link=1 ),
control.fixed=H$fixed, # priors for fixed effects, generic is ok
)
Fixed effects:
              mean      sd 0.025quant 0.5quant 0.975quant  mode  kld
(Intercept) 0.143 0.015          0.11   0.143          0.173 0.144  0
Model hyperparameters:
              mean      sd
Prec the Gaussian observations      1.91e+02 2.88e+00
Prec dyri                          1.03e+03 6.60e+02
Rho for dyri                        5.50e-02 2.47e-01
Prec inla.group(t, method = "quantile", n=9) 1.70e+04 1.45e+04
Prec inla.group(z, method = "quantile", n=9) 1.61e+02 7.50e+01
Prec inla.group(substrate.grainsize, method = "quantile", n=9) 7.29e+02 2.72e+03
Prec auid                            1.71e+02 1.21e+01
Phi for auid                          5.19e-01 3.80e-02
GroupRho for auid                     9.19e-01 7.00e-03
Expected number of effective parameters(stdev): 1129.49(30.74)
Number of equivalent replicates : 9.58
Deviance Information Criterion (DIC) ..... -24998.14
Deviance Information Criterion (DIC, saturated) .... 11953.32
Effective number of parameters ..... 1132.66
Marginal log-Likelihood: 13670.07

```

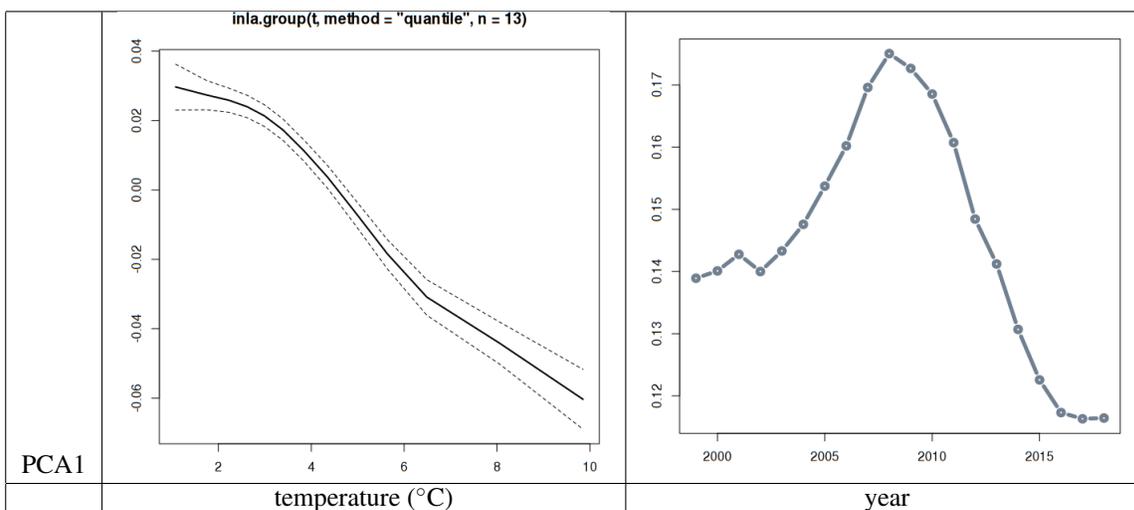
depth. Temporal autocorrelation at annual scales is high at 0.919 for PCA1 and 0.914 for PCA2. Sub-annual autocorrelations were lower at 0.055 and 0.601, respectively. The latter was stronger as “depth”-related associations changed less than the “temperature”-related associations that fluctuated more in seasonally. Note that the timeseries of PCA1 and PCA2 indicate that species assemblages have been changing strongly from a cold-water dominated one in the late 2000’s to a warmer water dominated assemblages since 2015. The changes in PCA2 are also similar with deeper water assemblages becoming more prominent early in 2000 and changing into a dominance of shallower water assemblages since 2010. The latter may be associated with pycnocline changes due to changes in temperature, salinity and the hydrodynamics of the mixing of cold low salinity water masses with the warm salty Gulf stream water.

The addition of these spatiotemporal covariates to the model of snow crab density is simple, once the data are spatially and temporally aligned. We will refer to this as the “**Mixed effects dynamic**” model (Table 8). The spatial effect represents 76% of the extra Poisson variability. The aggregate timeseries (Figure 38) indicates a reduction a the overall magnitudes as the covariates help to refine overly optimistic imputation in unsampled locations.

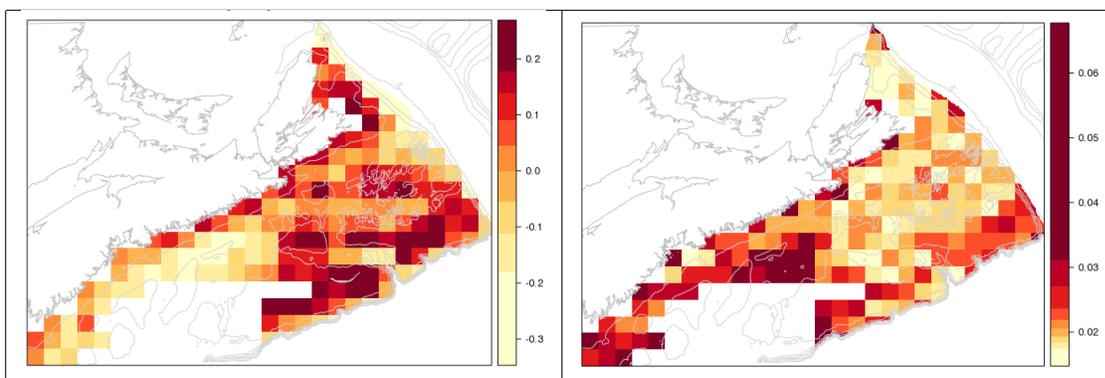
The manner in which year is treated is important. In the mixed effects models, it is treated as a factor. In reality there exists temporal autocorrelation. This can be represented through an AR1 process that is additive which means a “**Separable**” spatiotemporal model (Table 9). The model suggests that there is a significant 0.86 autocorrelation between years. The spatial effect represents 85% of the extra-Poisson variation. The overall fit is again improved, based upon DIC. However, the effect upon the aggregate timeseries of biomass and spatial predictions are minimal (Figure 39). The limits of using a single CAR effect across all years seem to have been reached.



**Fig. 34.** Species composition axis 1 for Sept 2017; mean (left) and sd (right).



**Fig. 35.** Species composition axis 1 vs temperature and time.



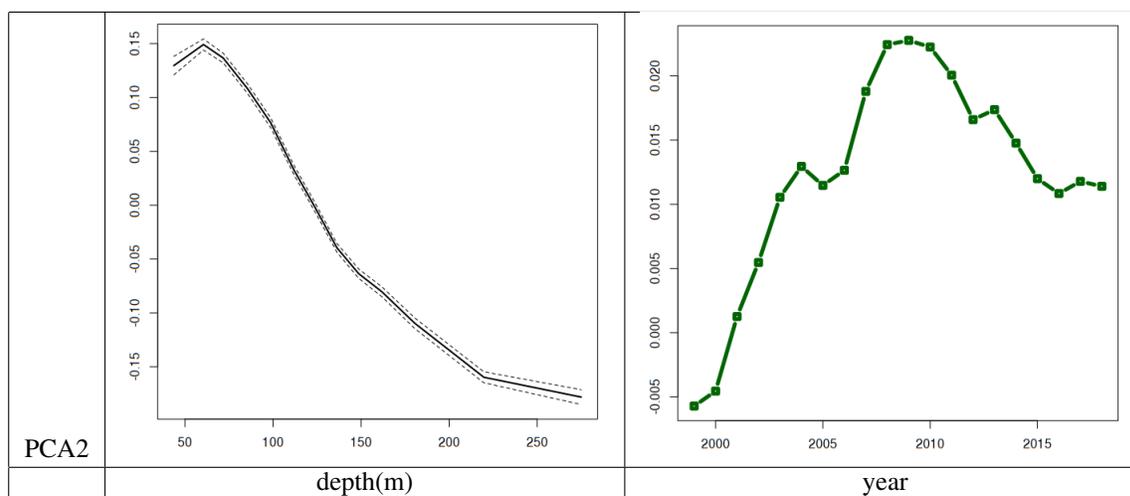
**Fig. 36.** Species composition axis 2 for Sept 2017; mean (left) and sd (right).

**Tab. 7.** Results for species composition axis 2 – CAR at each year.

```

Formula:
inla(formula = pca2 ~ 1
+ f( dyri, model="ar1", hyper=H$ar1 )
+ f( inla.group( t, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( z, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( substrate.grainsize, method="quantile", n=9 ), model="rw2",
scale.model=TRUE, hyper=H$rw2)
+ f( auid, model="bym2", graph=sppoly@nb, group=year_factor, scale.model=TRUE, constr=TRUE,
hyper=H$bym2, control.group=list(model="ar1", hyper=H$ar1_group)),
family = "normal",
data= M,
control.compute=list(dic=TRUE, config=TRUE),
control.results=list(return.marginals.random=TRUE, return.marginals.predictor=TRUE ),
control.predictor=list(compute=FALSE, link=1 ),
control.fixed=H$fixed, # priors for fixed effects, generic is ok
)
Time used:
Pre = 2.82, Running = 512, Post = 5.39, Total = 520
Fixed effects:
      mean      sd 0.025quant 0.5quant 0.975quant mode kld
(Intercept) 0.02 0.021    -0.024     0.02     0.065 0.02  0
Model hyperparameters:
                                mean      sd
Prec the Gaussian observations    255.139 3.86e+00
Prec dyri                          972.573 5.21e+02
Rho for dyri                       0.601 2.19e-01
Prec inla.group(t, method = "quantile", n=9) 7071.547 1.04e+04
Prec inla.group(z, method = "quantile", n=9) 340.448 2.10e+02
Prec inla.group(substrate.grainsize, method = "quantile", n=9) 1267.298 1.65e+03
Prec auid                           398.917 3.31e+01
Phi for auid                         0.538 7.00e-02
GroupRho for auid                    0.914 8.00e-03
Expected number of effective parameters(stdev): 921.94(40.80)
Number of equivalent replicates : 11.74
Deviance Information Criterion (DIC) .....: -28335.04
Deviance Information Criterion (DIC, saturated) ....: 11769.79
Effective number of parameters .....: 928.42
Marginal log-Likelihood: 15491.21

```

**Fig. 37.** Species composition axis 2 vs depth and time.

Tab. 8. Results for the “Mixed effects dynamic” model.

```

Formula:
inla(formula = totno ~ 1
+ year_factor
  + f( dyri, model="ar1", hyper=H$ar1 )
  + f(inla.group( z, method="quantile", n=9), model="rw2", scale.model=TRUE, hyper=H$rw2)
  + f(inla.group( substrate.grainsize, method="quantile", n=9 ), model="rw2",
    scale.model=TRUE, hyper=H$rw2)
  + f( inla.group( t, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
  + f( inla.group( pca1, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
  + f( inla.group( pca2, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
  + f(auid, model="bym2", graph=sppoly@nb, scale.model=TRUE, constr=TRUE, hyper= H$bym2)
+ offset(log(data_offset)),
family="poisson",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 1e-06, tolerance = 1e-12),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor = TRUE),
control.fixed = H$fixed
)
Fixed effects:
              mean      sd 0.025quant 0.5quant 0.975quant   mode kld
(Intercept)  6.162 0.084    5.996    6.162    6.325  6.163    0
year_factor2000 0.245 0.069     0.111    0.245    0.382  0.244    0
Model hyperparameters:
              mean      sd
Prec inla.group(z, method = "quantile", n = 9)  15.418  7.577
Prec inla.group(substrate.grainsize, method = "quantile", n = 9)  0.532  0.200
Prec inla.group(t, method = "quantile", n = 9)  169.717 111.319
Prec inla.group(pca1, method = "quantile", n = 9)  8.211  3.751
Prec inla.group(pca2, method = "quantile", n = 9) 20.566 10.102
Prec auid  0.551  0.083
Phi for auid 0.757  0.082

```

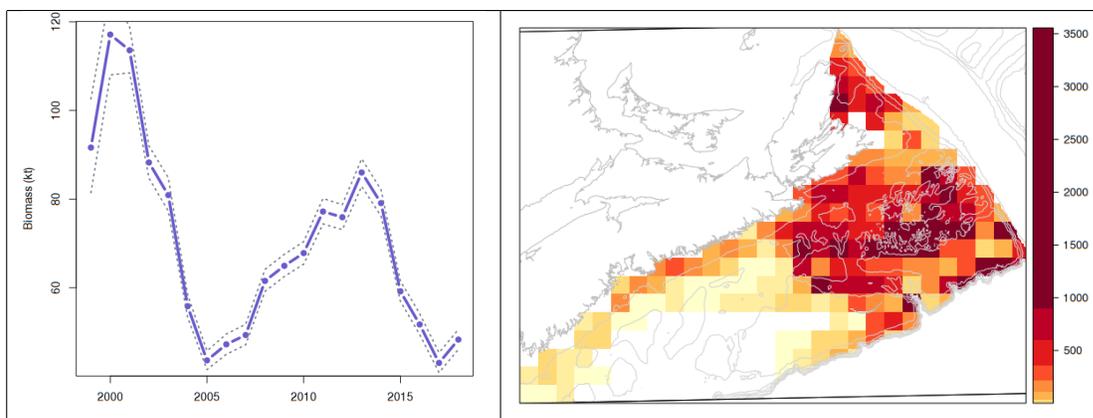


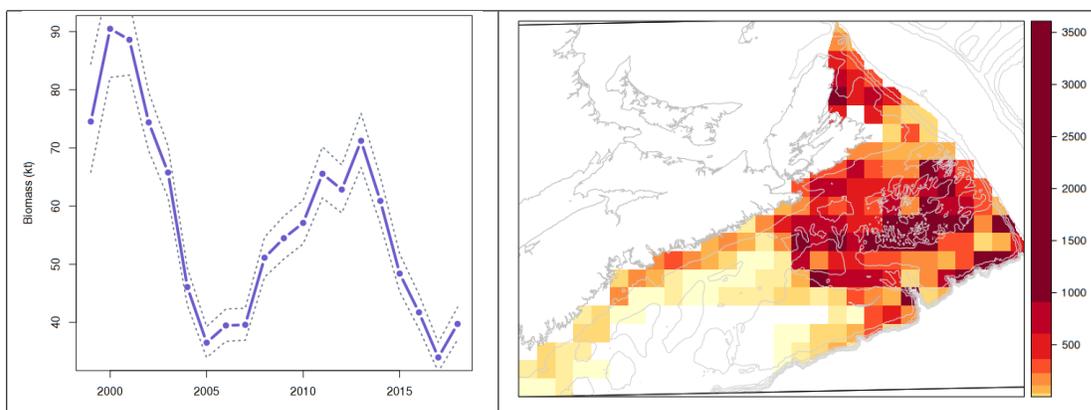
Fig. 38. “Mixed effects dynamic” model: biomass (kt) variations over years treated as a factorial variable and assuming a single additive CAR random effect and substrate grain size, depth, bottom temperature and species composition as covariates with a Poisson distributional assumption (left). The 95% CI is based upon posterior simulations from the joint distributions of model estimates. The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in  $\text{kg}/\text{km}^2$ ).

**Tab. 9.** Results for the “Separable” model.

```

Formula:
inla(formula = totno ~ 1
+ f( year_factor, model="ar1", hyper=H$ar1 )
+ f( dyri, model="ar1", hyper=H$ar1 )
+ f(inla.group( z, method="quantile", n=9), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f(inla.group( substrate.grainsize, method="quantile", n=9 ), model="rw2",
scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( t, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( pca1, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( pca2, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f(auid, model="bym2", graph=sppoly@nb, scale.model=TRUE, constr=TRUE, hyper= H$bym2)
+ offset(log(data_offset)),
family="poisson",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 1e-06, tolerance = 1e-12),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor = TRUE),
control.fixed = H$fixed
)
Fixed effects:
              mean      sd 0.025quant 0.5quant 0.975quant mode  kld
(Intercept) 5.955 0.149      5.663      5.955      6.248 5.955 0.001
Model hyperparameters:
Prec year_factor              mean      sd
Rho for year_factor           0.862    0.051
Prec inla.group(z, method = "quantile", n = 9) 16.377 12.415
Prec inla.group(substrate.grainsize, method = "quantile", n = 9) 0.522 0.194
Prec inla.group(t, method = "quantile", n = 9) 173.327 113.614
Prec inla.group(pca1, method = "quantile", n = 9) 8.143 3.786
Prec inla.group(pca2, method = "quantile", n = 9) 19.974 12.119
Prec auid                      0.491 0.098
Phi for auid                    0.849 0.092

```



**Fig. 39.** “Separable” model: biomass (kt) variations over years treated as an AR1 process and assuming a single additive CAR random effect and substrate grain size, depth, bottom temperature and species composition as covariates with a Poisson distributional assumption (left). The 95% CI is based upon posterior simulations from the joint distributions of model estimates. The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in  $\text{kg}/\text{km}^2$ ).

**Tab. 10.** Results for the “Nonseparable simple” model.

```

Formula:
inla(formula = totno ~ 1
      + f( auid, model="bym2", graph=sppoly@nb, group=year_factor, scale.model=TRUE,
          constr=TRUE, hyper=H$bym2, control.group=list(model="ar1", hyper=H$ar1_group))
      + offset(log(data_offset)),
family="poisson",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 1e-06, tolerance = 1e-12),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor
                      = TRUE),
control.fixed = H$fixed
)
Time used:
  Pre = 1.52, Running = 912, Post = 0.638, Total = 914
Fixed effects:
      mean   sd 0.025quant 0.5quant 0.975quant mode kld
(Intercept) 5.838 0.083      5.672   5.839      5.997 5.841  0
Model hyperparameters:
      mean   sd 0.025quant 0.5quant 0.975quant mode
Precision for auid 0.283 0.024      0.238   0.282      0.334 0.280
Phi for auid      0.622 0.051      0.525   0.621      0.723 0.615
GroupRho for auid 0.886 0.009      0.866   0.887      0.904 0.888

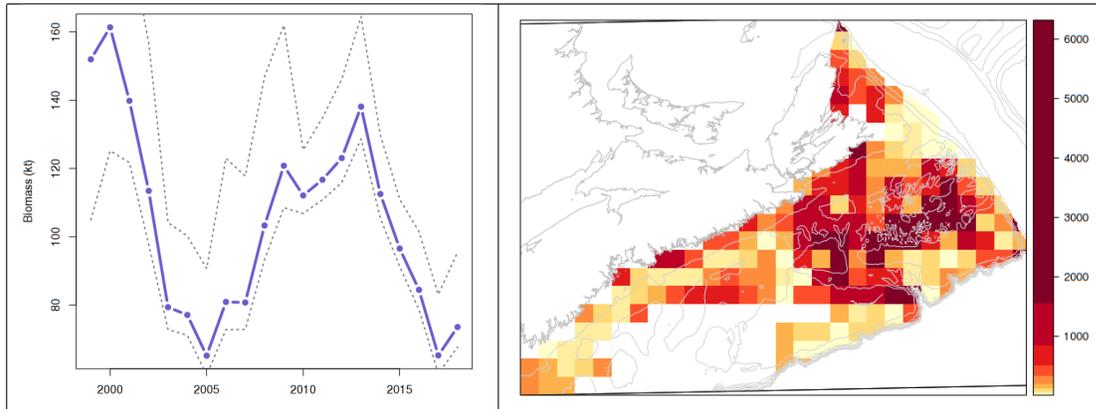
```

The next reasonable question is, therefore, what if space and time are non-separable? One way to proceed is to have a CAR effect for each year (“**Nonseparable spacetime**”). The alternative being to have a single CAR with multiple AR1 processes for each AU, “**Nonseparable timespace**”.) This is a reasonable expectation as there is no strong reason to suggest that spatial relationships in total number of snow crab persist across years. Implementation is again quite straightforward. First without any covariates, the “**Nonseparable simple**” model (Table 10) is found to be similar to the other “simple” models without ecosystem variability as covariates: the magnitudes are high, however, the timeseries is smoother. And as with the other “simple” models, the absence of ecosystem variability results in an optimistic prediction of abundance in AUs without samples (Figure 40).

The addition of covariates that incorporate ecosystem variability acts to refine these overly optimistic predictions (“**Nonseparable spacetime**” model; Table 11). The spatial effect is significant at 66% of the extra-Poisson variation and an inter-annual autocorrelation strength of 0.80. The alternate, “**Nonseparable timespace**” (not shown) had similar performance based upon various goodness of fit criteria (Table 12).

The manner in which the covariates interact with snow crab numerical densities using this “**Nonseparable spacetime**” model are as follows (Figure 42). Depth has a definitively strong relationship with maximal densities of snow crab being found at 150m and greater. Similarly, substrate grain size has a positive relationship with density. This is counter to the literature-based expectation of preference for mud and suggests that this was a confounded observation. Greater density may be related to habitat complexity/heterogeneity and is a novel finding. Temperature relationships have been known to be important and the range from 1 to 4 °C seem to be associated with highest densities. However, they are also found to persist in temperatures upto 10°C, though this is more variable. The seasonality observed likely associated with increased numbers due to moulting in spring/summer and then subsequent decline with natural and fishing mortality causing the declines in the autumn and winter. The relationship with PCA1 is very striking in that the species assemblages are associated with a temperature-gradient. The presence of cold-water assemblages is a strong predictor for higher snow crab densities. The PCA2 is associated with a depth-gradient and suggests that intermediate depths are most favourable for finding high densities of snow crab. Note that these relationships are slightly different from those derived from simpler models (e.g., the “Static” models in Figures 27, 28).

All models show a similar overall pattern: snow crab have fluctuated through at least one cycle with a periodicity of approximately 12 years, which roughly corresponds to the approximate longevity of a male snow crab. The



**Fig. 40.** “Nonseparable simple” model: biomass (kt) variations over years treated as an AR1 process and assuming a CAR random effect for each year; no covariates with a Poisson distributional assumption (left). The 95% CI is based upon posterior simulations from the joint distributions of model estimates. The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in kg/km<sup>2</sup>).

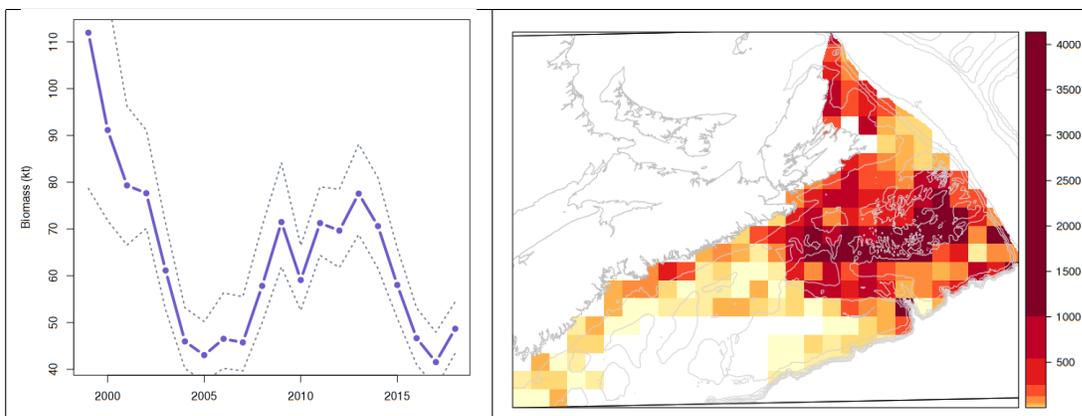
**Tab. 11.** Results for the “Nonseparable spacetime” model.

```

Formula:
inla(formula = totno ~ 1
      + f( dyri, model="ar1", hyper=H$ar1 )
      + f(inla.group( z, method="quantile", n=9), model="rw2", scale.model=TRUE, hyper=H$rw2)
      + f(inla.group( substrate.grainsize, method="quantile", n=9 ), model="rw2",
          scale.model=TRUE, hyper=H$rw2)
      + f( inla.group( t, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
      + f( inla.group( pca1, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
      + f( inla.group( pca2, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
      + f( auid, model="bym2", graph=sppoly@nb, group=year_factor, scale.model=TRUE,
          constr=TRUE, hyper=H$bym2, control.group=list(model="ar1", hyper=H$ar1_group))
      + offset(log(data_offset)),
family="poisson",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 1e-06, tolerance = 1e-12),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor = TRUE),
control.fixed = H$fixed
)
Fixed effects:Fixed effects:
              mean sd 0.025quant 0.5quant 0.975quant mode kld
(Intercept) 5.547 0.11      5.321    5.55    5.759 5.554 0
Model hyperparameters:

              mean      sd
Precision for dyri      28.692 16.806
Rho for dyri            -0.012  0.143
Precision for inla.group(t, method = "quantile", n = 9)      256.897 330.921
Precision for inla.group(z, method = "quantile", n = 9)       8.575  3.873
Precision for inla.group(substrate.grainsize, method = "quantile", n = 9) 1.460  0.859
Precision for inla.group(pca1, method = "quantile", n = 9)    8.077  3.837
Precision for inla.group(pca2, method = "quantile", n = 9)   23.812 11.599
Precision for auid      0.435  0.035
Phi for auid            0.661  0.060
GroupRho for auid      0.805  0.015

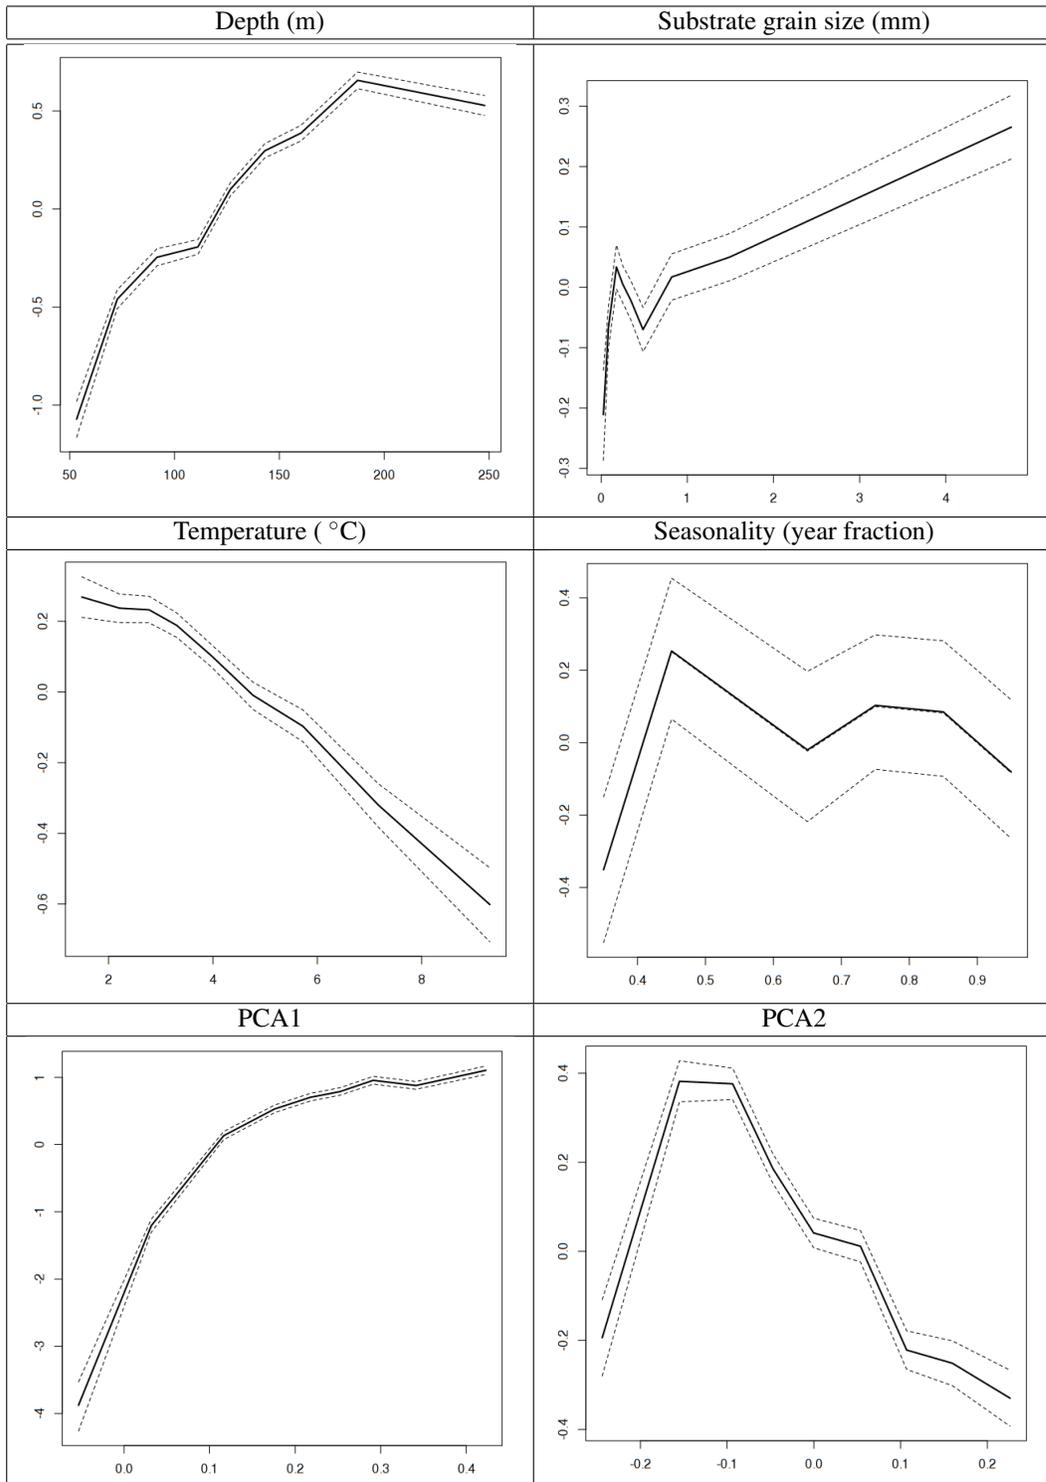
```



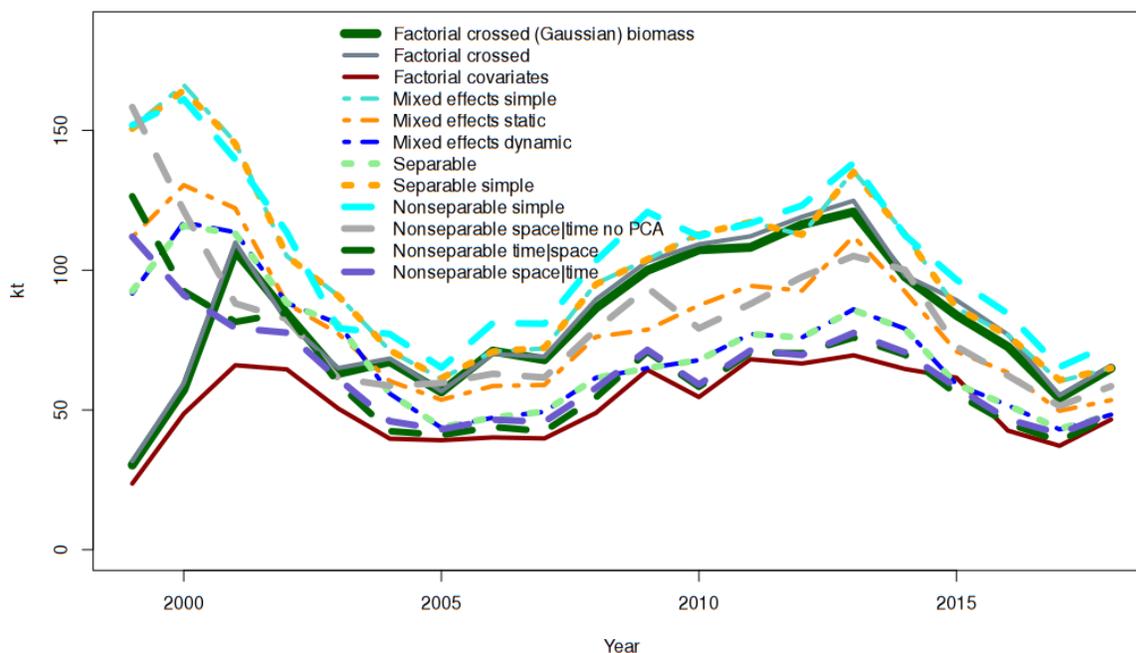
**Fig. 41.** “Nonseparable spacetime” model: biomass (kt) variations over years treated as an AR1 process and assuming a CAR random effect for each year; with all covariates and a Poisson distributional assumption (left). The 95% CI is based upon posterior simulations from the joint distributions of model estimates. The associated biomass estimates in space for the year 2017 (red is high and yellow is low on a quantile scale of spatial densities in  $\text{kg}/\text{km}^2$ ).

**Tab. 12.** Summary of various model representations of numerical abundance of snow crab in Maritimes Region in this report. All are based upon a Poisson distributional assumption and modelled with INLA, with the exception of the “Factorial crossed Gaussian” which operates upon a Gaussian assumption directly upon biomass densities using GLM. Here, “crossed” means that only the interaction terms are used, without marginal (main) effects. The “|” notation indicates “grouped by”, as commonly used for mixed effects modelling.

| Model                                 | AUs       | Year   | Season | Environmental variability (rw2) | - Log likelihood | DIC   | WAIC  | CPO    |
|---------------------------------------|-----------|--------|--------|---------------------------------|------------------|-------|-------|--------|
| Factorial crossed Gaussian (biomass ) | fixed     | fixed  | -      | -                               | -                | -     | -     | -      |
| Factorial crossed                     | fixed     | fixed  | -      | -                               | 34659            | -     | -     | 34590  |
| Factorial covariate                   | fixed     | fixed  | ar1    | depth+substrate+temp+PCA1+PCA2  | 30409            | -     | -     | 122636 |
| Mixed effects simple                  | bym2      | fixed  | -      | -                               | 36105            | 71454 | 73156 | 36458  |
| Mixed effects static                  | bym2      | fixed  | -      | depth+substrate                 | 33426            | 65926 | 67663 | 33701  |
| Mixed effects dynamic                 | bym2      | fixed  | -      | depth+substrate+temp+PCA1+PCA2  | 30115            | 59302 | 61066 | 30368  |
| Separable simple                      | bym2      | ar1    | -      | -                               | 36045            | 71449 | 73144 | 36464  |
| Separable                             | bym2      | ar1    | ar1    | depth+substrate+temp+PCA1+PCA2  | 30052            | 59301 | 61054 | 30368  |
| Nonseparable simple                   | bym2 year | ar1    | ar1    | -                               | 26684            | 53557 | 62148 | 32952  |
| Nonseparable spacetie no PCA          | bym2 year | ar1    | -      | depth+substrate+temp            | 23713            | 47310 | 55052 | 29009  |
| Nonseparable timespace                | bym2      | ar1 AU | ar1    | depth+substrate+temp+PCA1+PCA2  | 30571            | 44476 | 50380 | 26577  |
| Nonseparable spacetime                | bym2 year | ar1    | ar1    | depth+substrate+temp+PCA1+PCA2  | 22092            | 44295 | 50308 | 26940  |



**Fig. 42.** Covariate relations with numerical density for the “Nonseparable spacetime” model.



**Fig. 43.** Comparison of biomass trajectories  $\int y_{st} ds$  for all models for the full Maritimes region domain (kt).

question is how high were these amplitudes and their relative sizes. Three main clusters of solutions are observed (Figure 43):

- (1) those that ignore ecosystem-based information such as the “\* Simple” models, “Factorial crossed” models (without environmental covariates);
- (2) those that incorporate only static environmental features (depth, substrate) – intermediate magnitude estimates; and,
- (3) those that include time varying ecosystem-based information – lower magnitude estimates.

The *non-spatial and non-temporal* (Cartesian) methodology of the “**Factorial crossed (Gaussian) biomass**” comes closest to a naive application of a “random stratified” (spatial) design to a spatiotemporal process (Figure 43; dark green thick line). This is of course, as previously identified as being not particularly useful due to the numerous cases of incomplete sampling rendering interannual comparisons impossible. To make temporal comparisons possible, assumptions are required that can treat spatial and/or temporal processes.

The application of random CAR or AR1 effects facilitates prediction in areas without sampling (in contrast to the factorial models) and so useful in the context of an assessment program susceptible to logistical failures. However, this prediction can be overly optimistic when done simplistically; one needs to be careful. The addition of informative covariates that index relevant measures of ecosystem variability assist to tune these predictions to more reasonable bounds. They also suggest that even the predictions from our best model is likely an overestimate of the true state as the number of covariates that represent ecosystem variability are still limited and quite rudimentary.

Assumptions of separability of spatial and temporal autocorrelation are widespread in the applied literature. Here we note that their predictions are qualitatively similar to the nonseparable models, with overlapping 95%

credible intervals. However, the variations are sufficiently informative (based on DIC, WAIC and CPO metrics) to suggest operating upon the more complex model would be beneficial.

Fundamentally, these results suggest that estimating  $\int y_{st} ds$  through a nonseparable, spatiotemporal, ecosystem variability model-based aggregation/integration schema is advisable. Evidently, the rather common practice of using a purely non-spatial and non-temporal model or even a purely spatial schema to represent a spatiotemporal process can result in biased parameter estimates. Surprising? No, not really.

## 4 Synthesis

Fishery data *synthesis* operates upon the aggregate biomass ( $\int y_{st} ds$ ) through some fisheries model. In the snow crab assessment, this is a (logistic) biomass dynamics model (implemented in the STAN language; <https://mc-stan.org/>):

```
https://github.com/jae0/bio.snowcrab/blob/master/R/fishery\_model.r
```

with associated parameters as specified in:

```
https://github.com/jae0/bio.snowcrab/blob/master/inst/scripts/05.fisherymodel.R
```

Application of this, perhaps simplest possible state space model is straightforward and permits the balancing of temporal variability with fishery catches and biological constraints upon the intrinsic rate of increase  $r$  (Figure 47), the carrying capacity  $K$  (Figure 48), and catchability  $q$  (Figure 49).

*Assuming all catches have been accounted correctly*, the biomass dynamics model suggests that the biomass of the fishable component is in a low state in all management areas (Figure 44). All areas are tracking the target exploitation rates (20%) and only occasionally going beyond FMSY (Figure 45).

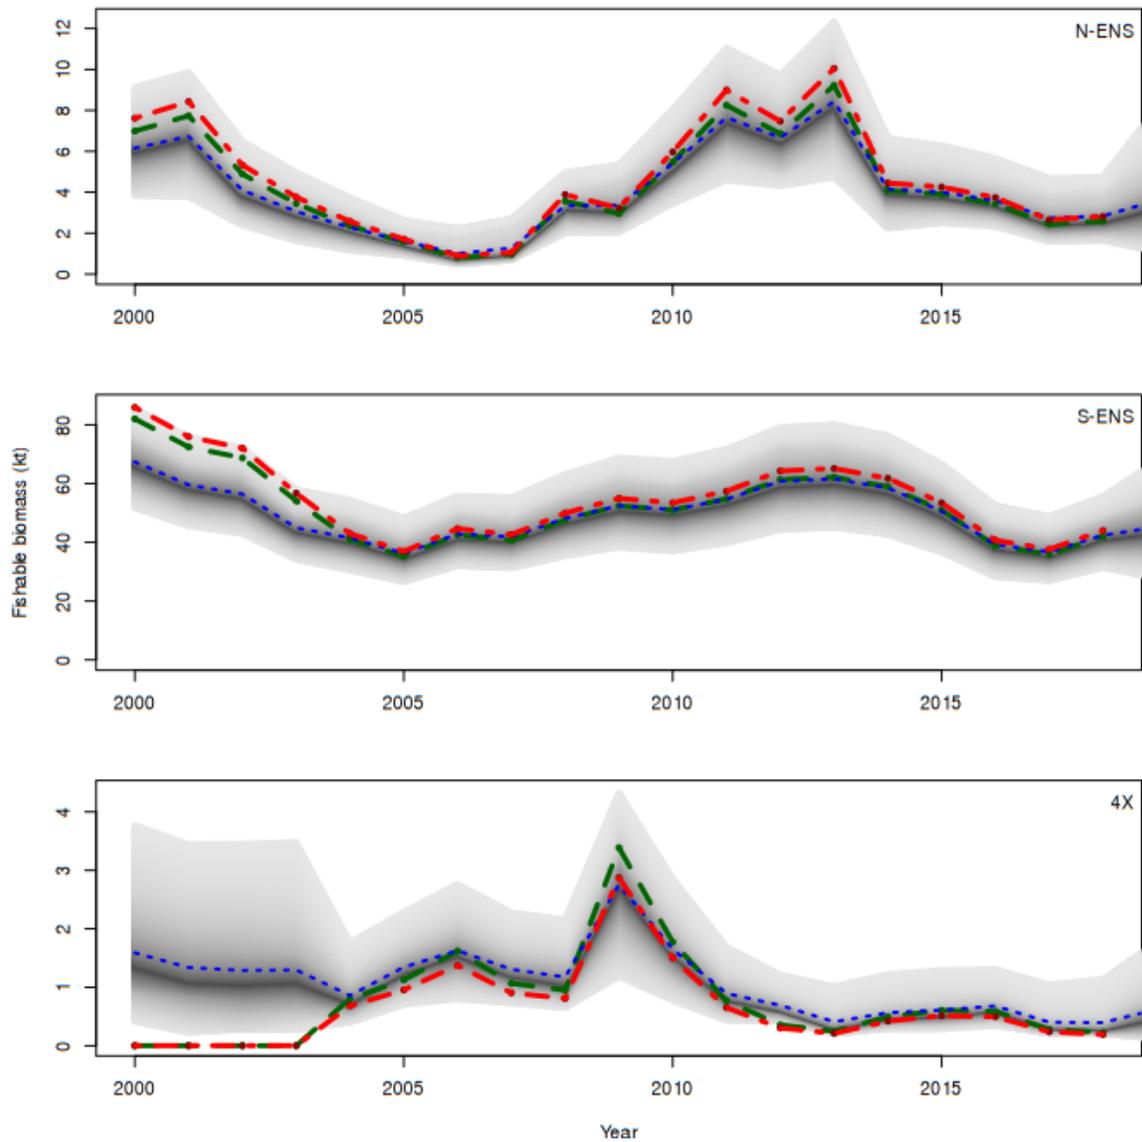
In terms of location within the harvest control rules, one area was in each zone (Figure 46).

## 5 Concluding thoughts

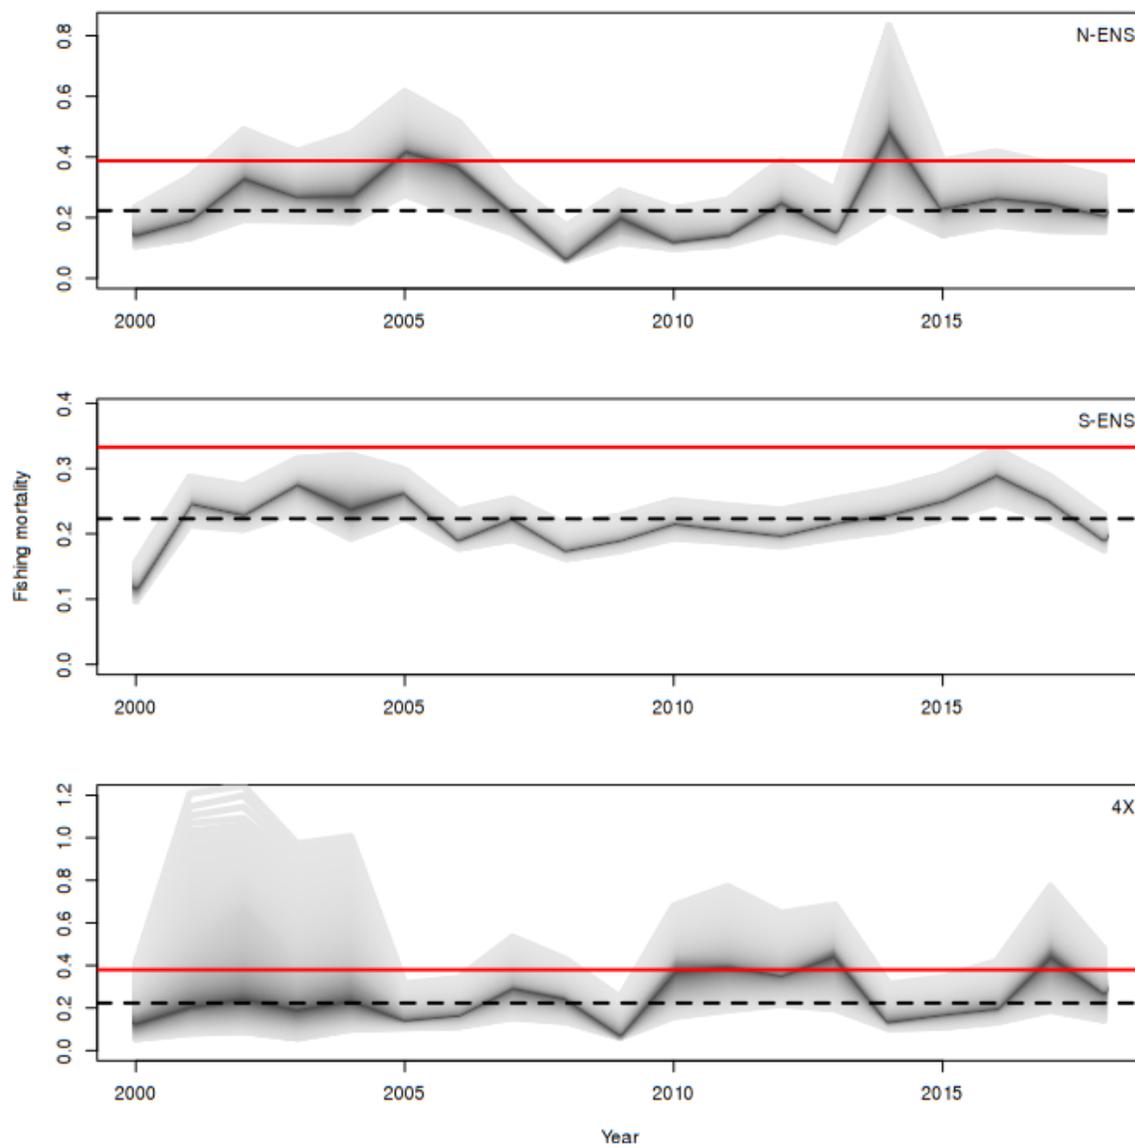
The methods presented here represent a unified, viable and coherent solution for operationally modelling the abundance of snow crab and the ecosystem variability in which they exist, and by extension, many other species and their ecological niche that urgently require assessment in the Maritimes Region; urgent due to rapid ecosystem and climate change and rapid changes in human exploitation patterns and disruption of their ecosystem.

### 5.1 Coda: Habitat (“species distribution”) models

The “**Nonseparable spacetime**” model is a viable operational model form for abundance estimation. However, utility still requires a reasonably coherent sampling of the organism of interest. For many species a coherent set of samples do not exist. For example, the Groundfish (“Ecosystem”) surveys systematically ignored noncommercial species in the mid 1980s and almost all invertebrates up until 1999). In the snow crab surveys, non-snow



**Fig. 44.** Comparison of biomass trajectories. Blue is the posterior mean of the estimate of the biomass estimate  $Y_t$ , while the dark gray is the median and light gray the 95% BCI (Bayesian Credible Interval). Red is the aggregate index of biomass  $(\int y_{st} ds)$ . Green is the catchability corrected biomass  $(q^{-1} \int y_{st} ds)$ .



**Fig. 45.** Comparison of fishing mortality estimates. Dark gray is the median and light gray the 95% BCI (Bayesian Credible Interval). Red line is the FMSY estimate. Dashed line is the target 20% exploitation rate.

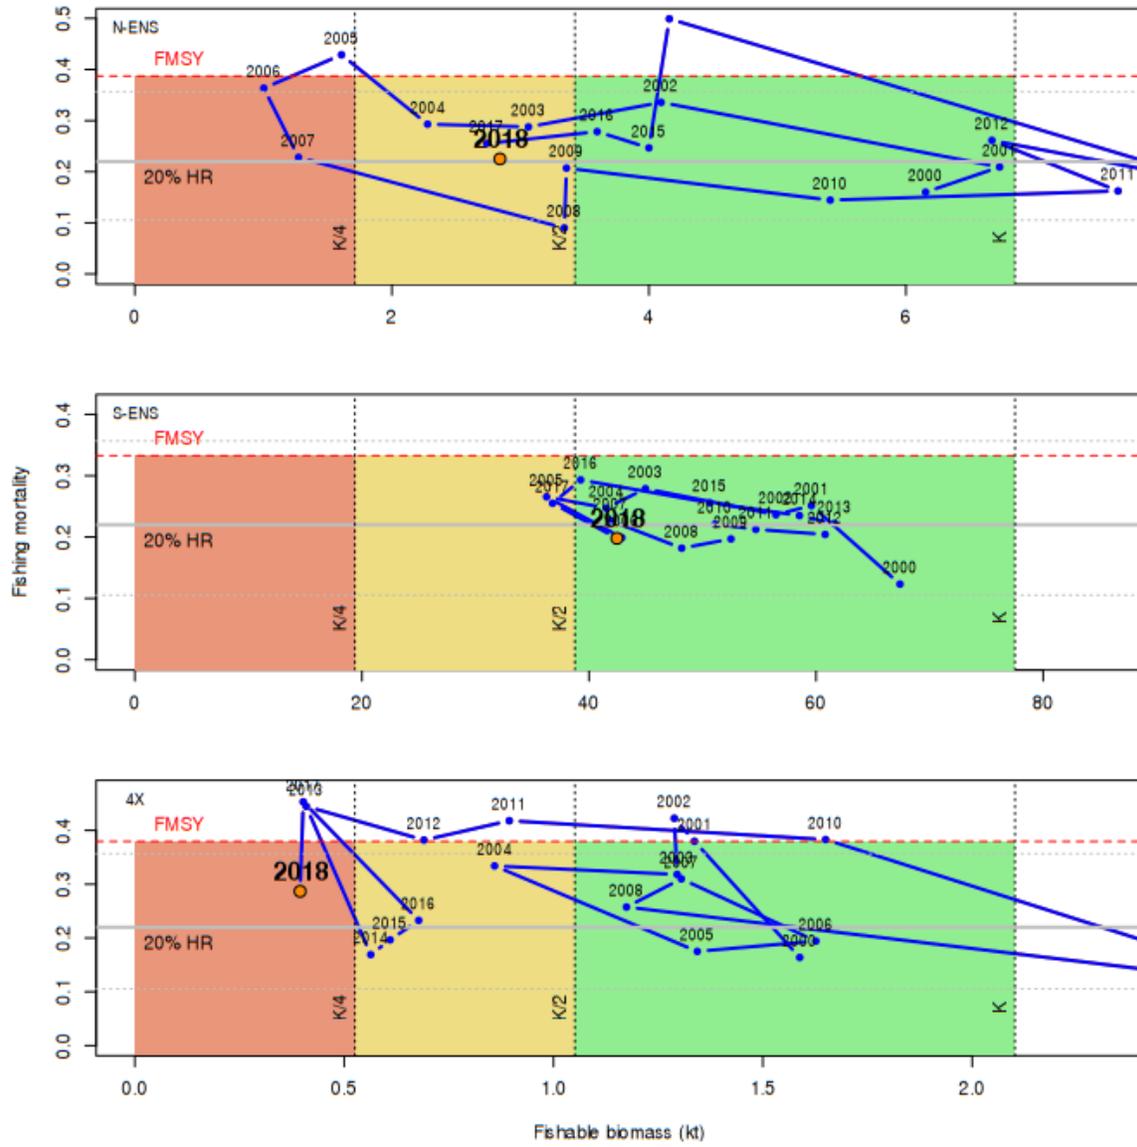
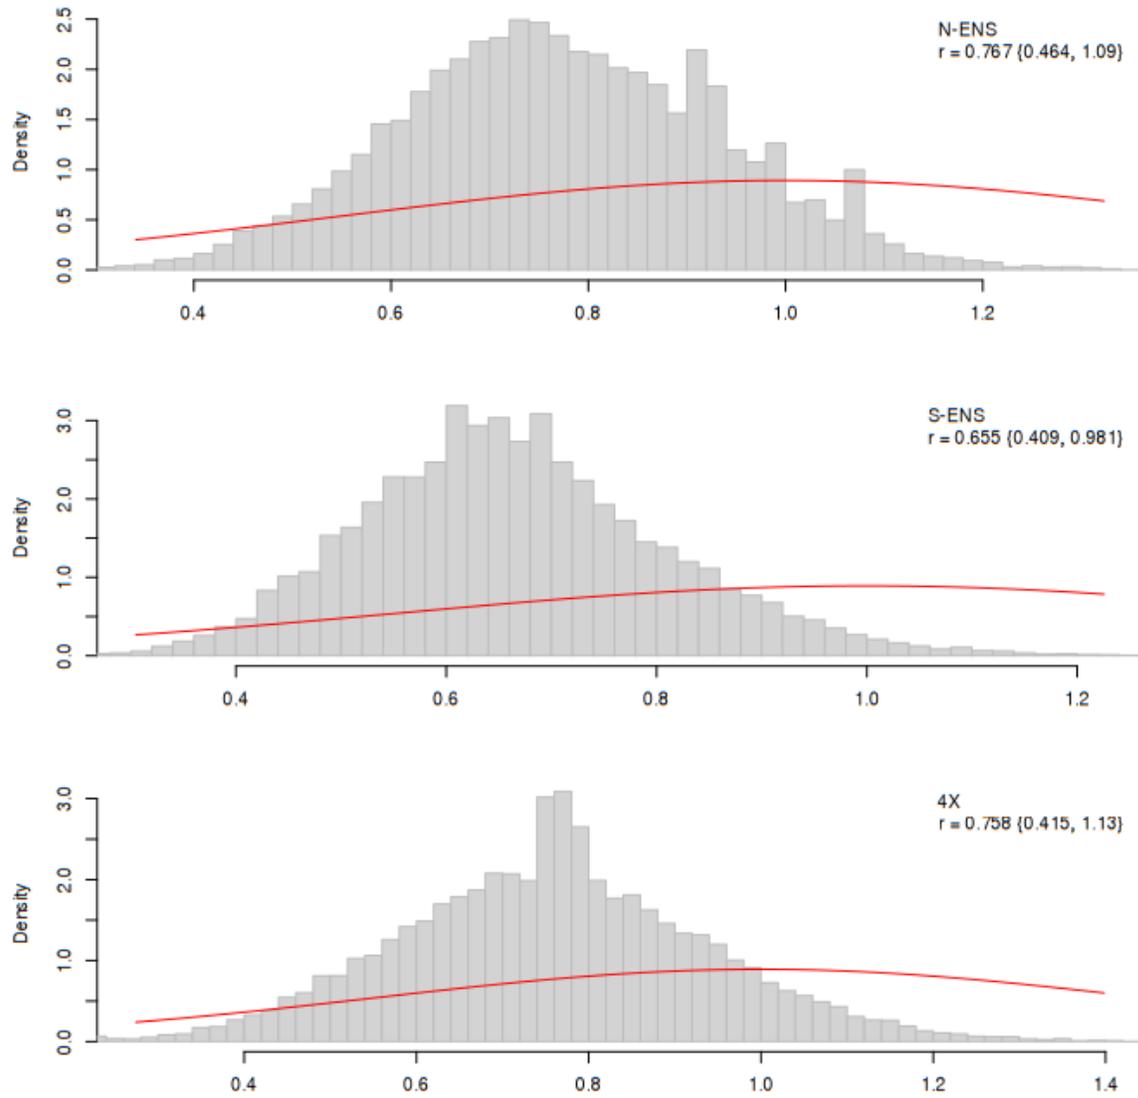
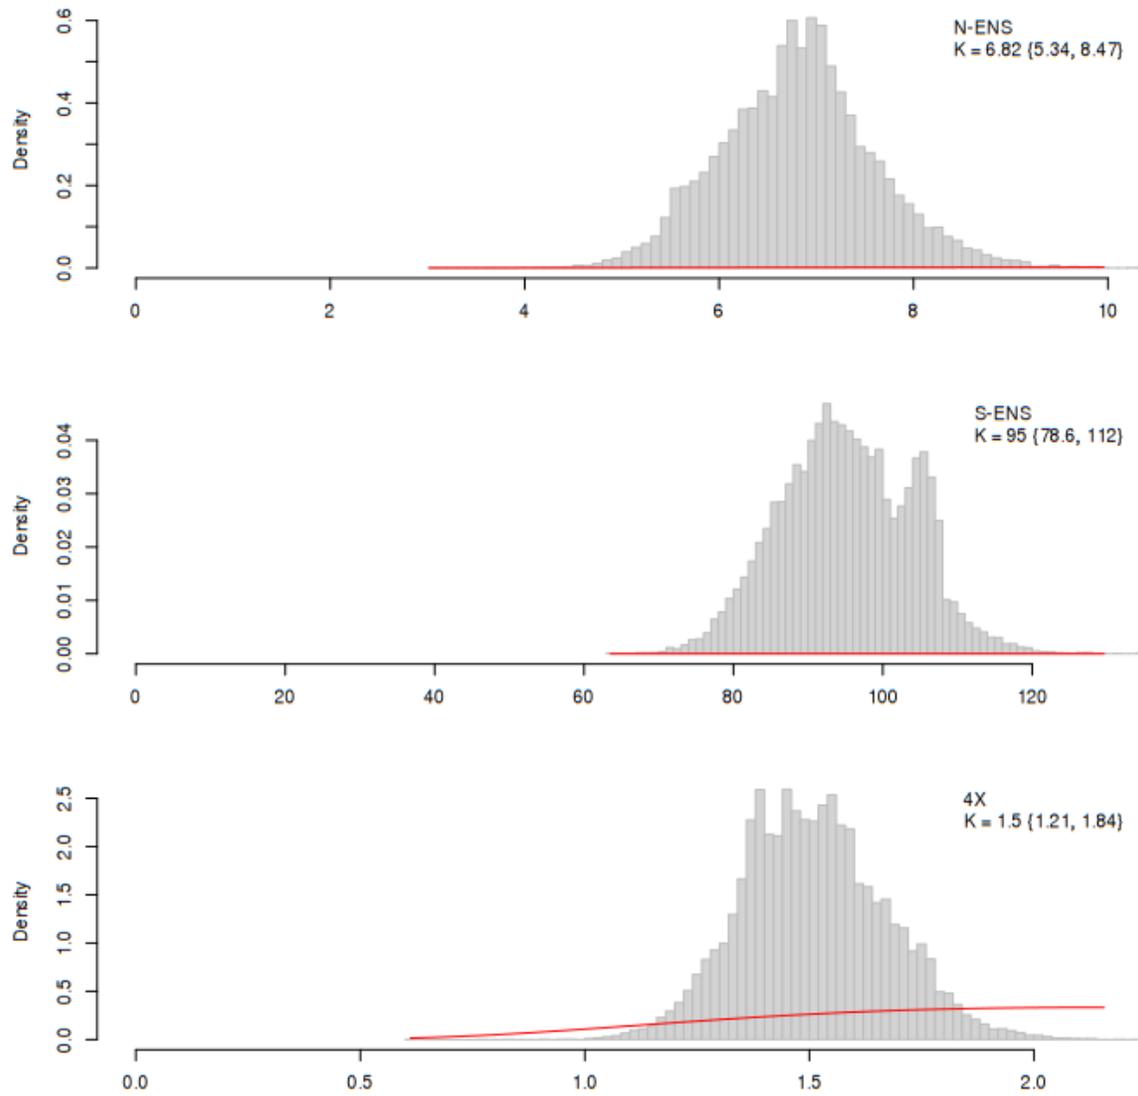


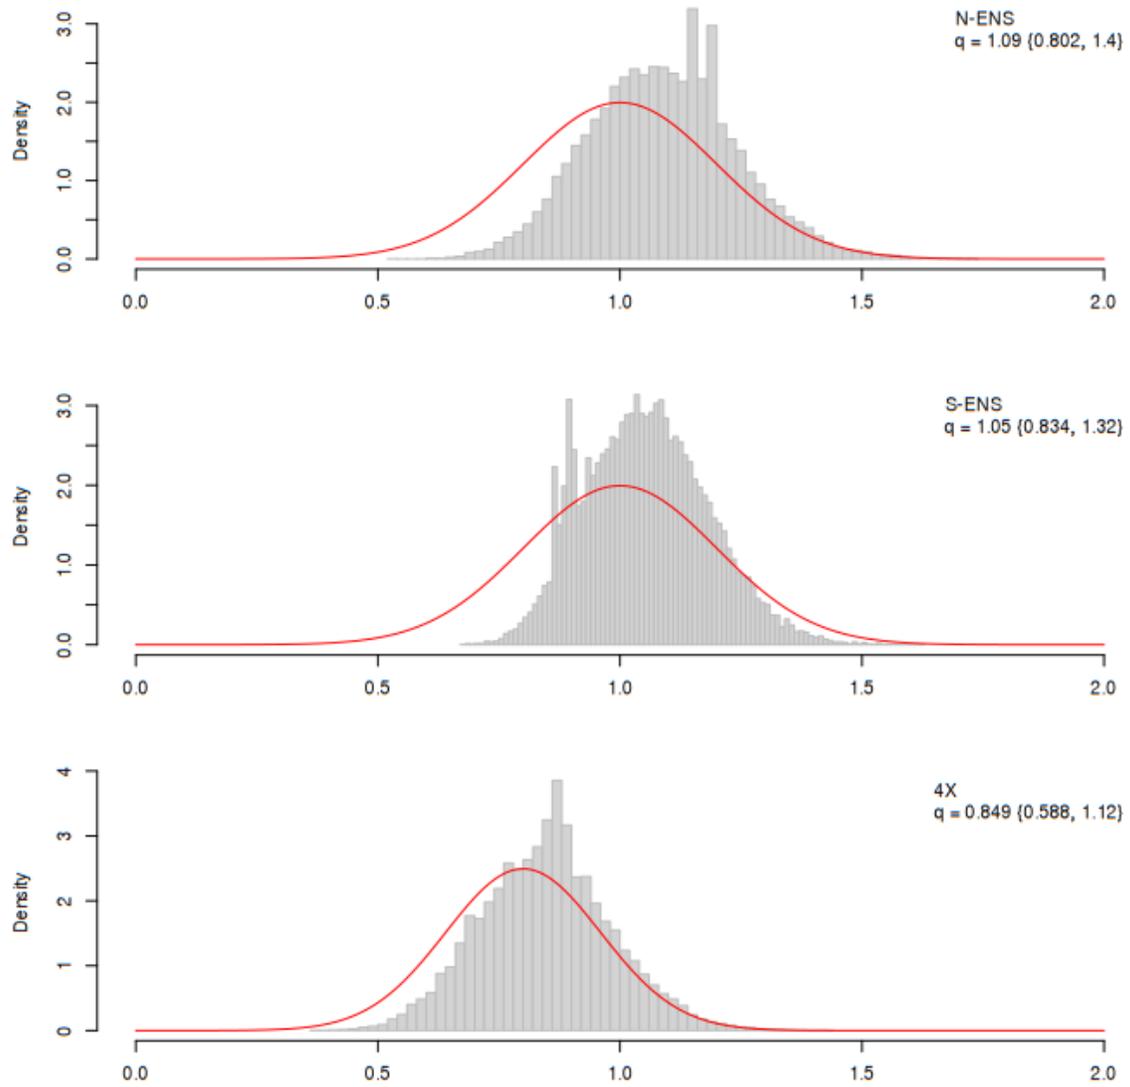
Fig. 46. Harvest control rules.



**Fig. 47.** Posterior distribution of the intrinsic rate of increase with prior in red.



**Fig. 48.** Posterior distribution of the carrying capacity with prior in red.



**Fig. 49.** Posterior distribution of the catachability coefficient with the prior in red.

**Tab. 13.** Results for the “Nonseparable spacetime habitat” model for small immature crab (<50 mm carapace width).

```

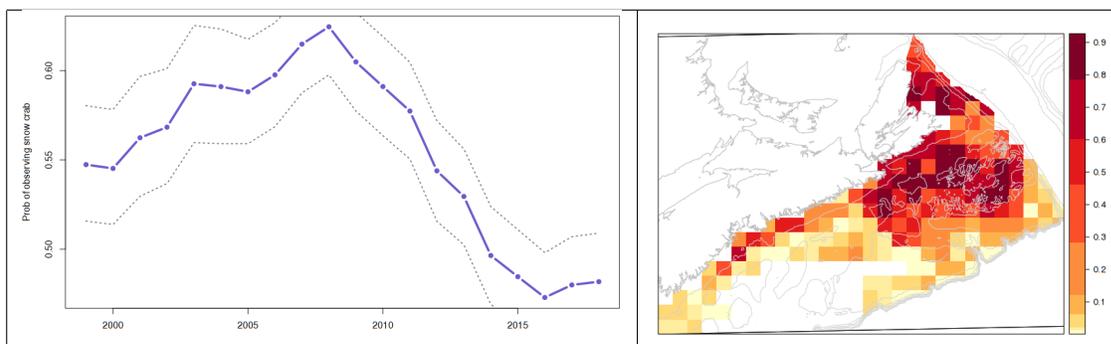
Formula:
inla(formula = pa ~ 1
+ f( dyri, model="ar1", hyper=H$ar1 )
+ f(inla.group( z, method="quantile", n=9), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f(inla.group( substrate.grainsize, method="quantile", n=9 ), model="rw2",
scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( t, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( pca1, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( inla.group( pca2, method="quantile", n=9 ), model="rw2", scale.model=TRUE, hyper=H$rw2)
+ f( auid, model="bym2", graph=sppoly@nb, group=year_factor, scale.model=TRUE,
constr=TRUE, hyper=H$bym2, control.group=list(model="ar1", hyper=H$ar1_group)),
family="binomial",
data = M,
verbose = TRUE,
control.compute = list(dic = TRUE, config = TRUE),
control.predictor = list(compute = FALSE, link = 1),
control.inla = list(cmin = 0, h = 1e-06, tolerance = 1e-12),
control.results = list(return.marginals.random = TRUE, return.marginals.predictor = TRUE),
control.fixed = H$fixed
)
Time used:
Pre = 4.47, Running = 13078, Post = 5.43, Total = 13088
Fixed effects:
      mean  sd 0.025quant 0.5quant 0.975quant  mode kld
(Intercept) -0.384 0.192      -0.792      -0.376      -0.017  -0.364  0
Model hyperparameters:
              mean      sd
Precision for dyri      13.338  13.600
Rho for dyri            0.024  0.211
Precision for inla.group(t, method = "quantile", n=9)      142.645  387.404
Precision for inla.group(z, method = "quantile", n=9)       2.163  1.742
Precision for inla.group(substrate.grainsize, method = "quantile", n=9) 16.309  53.300
Precision for inla.group(pca1, method = "quantile", n=9)    30.908  37.211
Precision for inla.group(pca2, method = "quantile", n=9)    645.305 6052.288
Precision for auid      0.426  0.057
Phi for auid            0.974  0.026
GroupRho for auid      0.921  0.014
Expected number of effective parameters(stdev): 408.60(22.44)
Number of equivalent replicates : 18.61
Deviance Information Criterion (DIC) .....: 6872.43
Deviance Information Criterion (DIC, saturated) ....: 6872.43
Effective number of parameters .....: 389.55
Marginal log-Likelihood: -1868.98

```

crab species were systematically recorded only after 2005; and due to gear mesh size, small snow crab are poorly captured. In such situations, relative abundance estimates may not be unreliable, but a delineation of habitat usage can still be estimated by modifying the distributional assumptions of the basic Generalized linear model (eq. 3) to a binomial case:  $Y \sim \text{Binomial}(\eta, \theta)$ , where  $\ln(\theta/(1-\theta)) = x^T \beta + O + \psi$ , and  $\eta$  is the vector of number of trials with  $\theta$  the vector of probabilities of success in each trial. Using such a distributional assumption and applying all available data, we can delineate the habitat preferences of small immature snow crab (< 50 mm carapace width) and so estimate the spatial and temporal changes of this important component of the population (Table 13; Figure 50).

As can be seen, the core “nursery” areas are found in near shore, shallow and slightly warmer environments with a warmer-water species assemblage. To foster long-term sustainability of the fishery, some meaningful protection of these areas from human disturbance are worth further discussions. They also represent areas where high natural mortality due to predation is likely to occur. Ecosystem-based management would suggest that the abundance of predators (such as halibut and cod) in such areas need to be more closely monitored and perhaps targeted fishing of predators in such areas may need to be discussed with other fishery sectors.

Further, the same habitat analysis for the fishable component indicates that their habitat has declined significantly since the year 2010 (Figure 52). Mature females in contrast have slightly differing habitat preferences



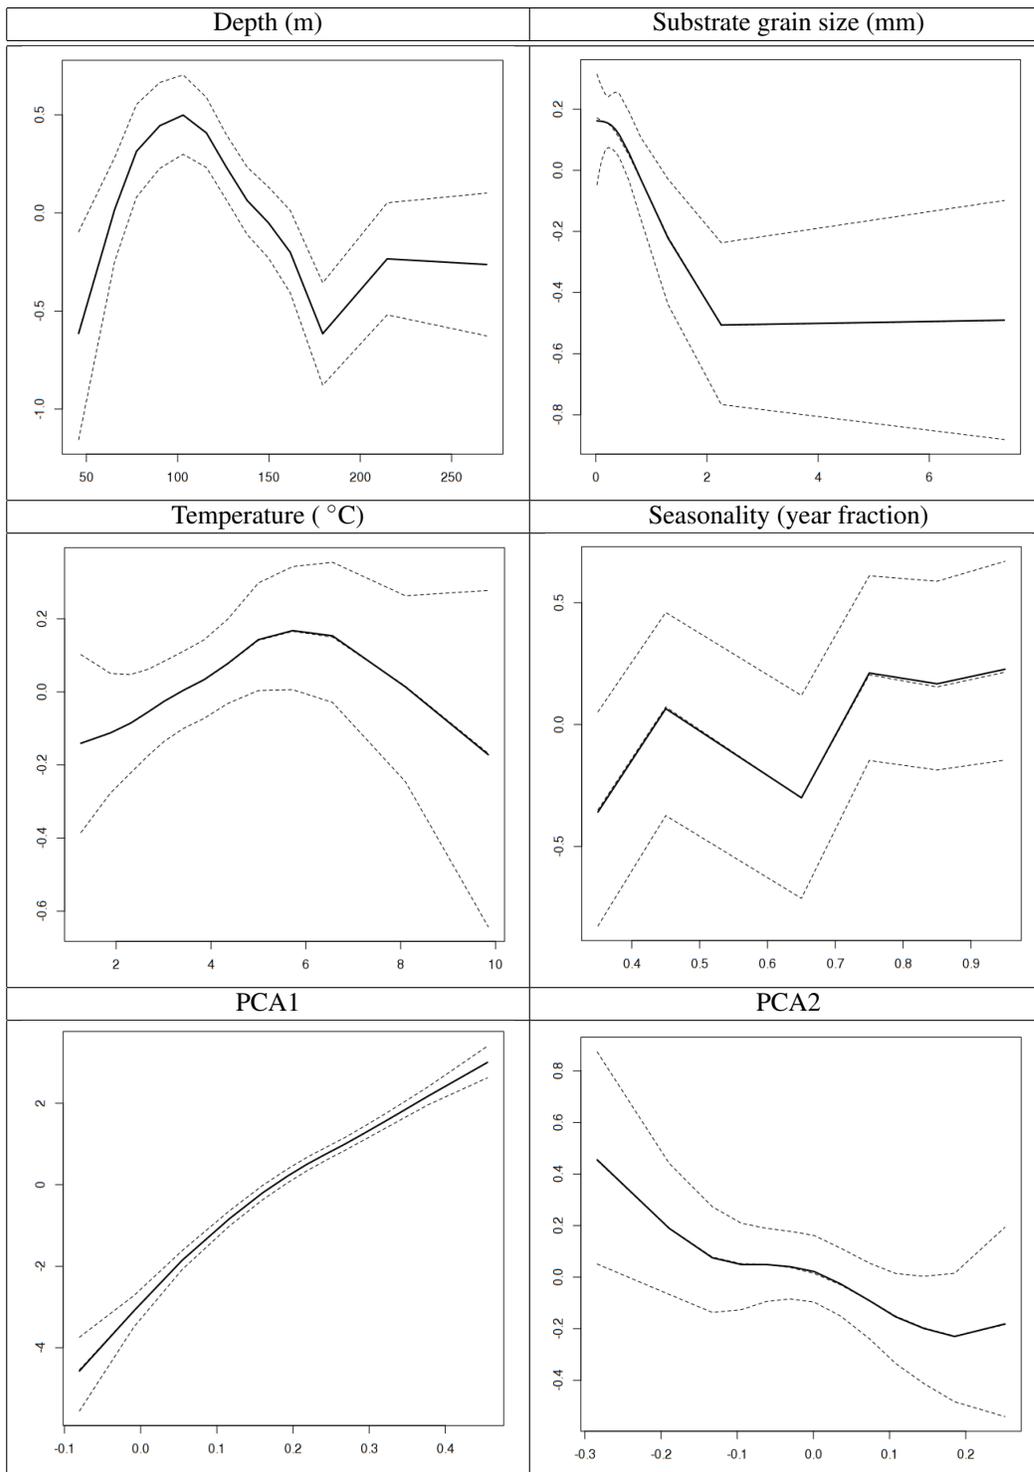
**Fig. 50.** “Nonseparable spacetime habitat” model for small immature snow crab: surface area-weighted probability of observing small snow crab as a function of time (left). The 95% CI is based upon posterior simulations from the joint distributions of model estimates. The associated probabilities of observing small snow crab in space for the year 2017 (red is high and yellow is low on a quantile scale of probability).

(warmer and shallower than large males) that result in a lower overall probability of observing them in the Maritimes Region (Figure 52). The peak in 2007 and 2008 are associated with cool bottom temperatures occurring in shallower waters (75 - 150 m) near shore.

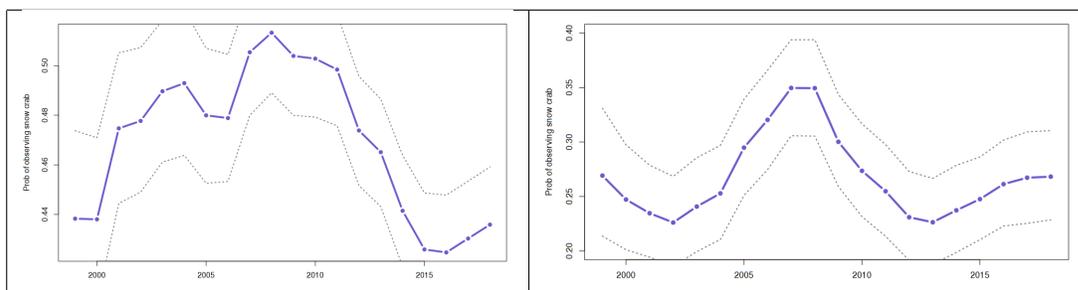
Though these assessments of habitat and abundance are extremely simplistic (e.g., relative to the *stmv*-approach), their systematic application would nonetheless provide meaningful comparisons of interspecific use of space and ecosystem requirements and/or sensitivities in general, applications that are easily adapted to prognostic climate change scenarios. These are small and simple but necessary steps towards a truly ecosystem-based management of fisheries.

## 6 References

- Banerjee, S., Carlin, B. P., and Gelfand, A. E.. 2004. Hierarchical Modeling and Analysis for Spatial Data. *Monographs on Statistics and Applied Probability*. Chapman and Hall/CRC.
- Bernardinelli, L., Clayton, D. and Montomoli, C. 1995. Bayesian estimates of disease maps: How important are priors? *Statistics in Medicine* 14: 2411–2431.
- Besag, Julian. 1974. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society Series B (Methodological)* 1974: 192-236.
- Besag, J., J. York, and A. Mollie. 1991. Bayesian image restoration with two applications in spatial statistics. *Ann Inst Stat Math* 43: 1–59.
- Biron, M., Moriyasu, M., Wade, E., DeGrace, P., Campbell, R., and Hebert, M. 1997. Assessment of the 1996 snow crab (*Chionoecetes opilio*) fishery off eastern Cape Breton, Nova Scotia (CFAs 20 to 24, and 4X). Department of Fisheries and Oceans. *Canadian Stock Assessment Secretariat. Research Document* 1997/102.
- Choi, J.S., Frank, K.T., Petrie, B., and Leggett, W.C. 2005a. Integrated assessment of a large marine ecosystem: A case study of the devolution of the eastern Scotian Shelf, Canada. *Oceanography and Marine Biology: An Annual Review* 43:47–67.



**Fig. 51.** Covariate relations with numerical density for the "Nonseparable spacetime habitat" model of immature small snow crab (<50 mm carapace width).



**Fig. 52.** “Nonseparable spacetime habitat” model for fishable snow crab (left) and mature females (right): surface area-weighted probability of observing snow crab as a function of time. The 95% CI is based upon posterior simulations from the joint distributions of model estimates.

Choi, J.S., Zisserson, B., and Reeves, A. 2005b. An assessment of the 2004 snow crab populations resident on the Scotian Shelf (CFAs 20 to 24). Department of *Fisheries and Oceans*. *Canadian Science Advisory Secretariat Research Document* 2005/028.

Choi, J.S. and Zisserson, B. 2011. Assessment of Scotian Shelf Snow Crab in 2010. *Canadian Science Advisory Secretariat Research Document* 2011/110.

Choi, J.S., Vanderlaan, A.S.M., Lazin, G., McMahon, M., Zisserson, B., Cameron, B., and Munden, J. 2018. St. Anns Bank Framework Assessment. *Canadian Science Advisory Secretariat Research Document* 2018/066.

Comeau, M., Conan, G.Y., Maynou, F., Robichaud, G., Therriault, J.C., and Starr, M. 1998. Growth, spatial distribution, and abundance of benthic stages of the snow crab, *Chionoectes opilio*, in Bonne Bay, Newfoundland, Canada. *Canadian Journal of Fisheries and Aquatic Sciences* 55:262–279.

Conan, G.Y., Comeau, M., and Robichaud, G. 1992. Life history and fishery management of Majid crabs: The case study of the Bonne bay (Newfoundland) *Chionoectes opilio* population. *International Council for the Exploration of the Seas, Council Meeting Document* 1992/K:21–21.

Cooley, J.W., and Tukey, J.W. 1965. An algorithm for the machine calculation of complex Fourier series. *Mathematics of Computation* 19 (90): 297–301.

Cressie, N. 1993. *Statistics for spatial data*. Wiley-Interscience, New York. 900p.

Danielson, G.C., and Lanczos, C. 1942. Some improvements in practical Fourier analysis and their application to x-ray scattering from liquids. *Journal of the Franklin Institute* 233: 365–380.

Dean, C. B., Ugarte, M. D. and Militino, A. F. 2001. Detecting interaction between random region and fixed age effects in disease mapping. *Biometrics* 57: 197–202.

DFO. 2006. A harvest strategy compliant with the precautionary approach. Department of Fisheries and Oceans. *Canadian Science Advisory Secretariat. Science Advisory Report* 2006/023.

Elnor, R.W., and Beninger, P. 1995. Multiple reproductive strategies in snow crab, *Chionoectes opilio*: Physiological pathways and behavioural plasticity. *Journal of Experimental Marine Biology and Ecology* 193:93–112.

Fotheringham, A.S., Brunson, C., and Charlton, M.E. 2002. Geographically Weighted Regression: The Analysis of Spatially Varying Relationships. Wiley, Chichester.

Fourier, J. 1822. *Théorie analytique de la chaleur*. Firmin Didot Père et Fils, Paris.

- Foyle, T., O'Dor, R., and Elner, R. 1989. Energetically defining the thermal limits of the snow crab. *Journal of Experimental Biology* 145:371–393.
- Fuentes, M. 2001. A high frequency kriging approach for non-stationary environmental processes. *Environmetrics*, 12: 469-483.
- Haran, M. 2011. Gaussian random field models for spatial data. *Handbook of Markov Chain Monte Carlo* 2011: 449-478.
- Heideman, M.T., Johnson, D.H., Burrus, C.S. 1985. Gauss and the history of the fast Fourier transform. *Archive for History of Exact Sciences* 34: 265–277.
- Hengl, T., Heuvelink, G.B.M., and Stein, A. 2004. A generic framework for spatial prediction of soil variables based on regression-kriging. *Geoderma* 120: 75-93.
- Hooper, R. 1986. A spring breeding migration of the snow crab, *Chionoectes opilio* (O. Fabr.), into shallow water in Newfoundland. *Crustaceana*. 50:257–264.
- Huang, H.-C., and Hsu, N.-J. 2004. Modeling transport effects on ground-level ozone using a non-stationary space-time model. *Environmetrics* 15: 251–268.
- Khintchine, Alexander. 1934. "Korrelationstheorie der stationären stochastischen Prozesse". *Mathematische Annalen* 109: 604–615.
- Kuhn, P., and Choi, J.S. 2011. Influence of temperature on embryo developmental cycles and mortality of female *Chionoectes opilio* (snow crab) on the Scotian Shelf, Canada. *Fisheries Research* 107:245-252.
- Kostylev, V. and Hannah, C. 2007. Process-driven characterization and mapping of seabed habitats. *Special Paper - Geological Association of Canada*. 47: 171-184.
- Lindgren, F., Rue, H., and Lindstrom, J. 2011. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73: 423–498.
- Nychka, D., Furrer, R., Paige, J., Sain S. 2017. Fields: Tools for spatial data. R package version 9.8-3, <URL: <https://github.com/NCAR/Fields>>.
- Riebler, A., Sørbye, S.H., Simpson D., and Rue, H. 2016. An intuitive Bayesian spatial model for disease mapping that accounts for scaling. *Statistical methods in medical research* 25: 1145-1165.
- Robertson, C., and George, S.C. 2012. Theory and practical recommendations for autocorrelation-based image correlation spectroscopy. *Journal of biomedical optics*, 17, 080801.
- Sigrist, F., Künsch, H.R., and Stahel, W.A. 2012. A dynamic nonstationary spatio-temporal model for short term prediction of precipitation. *Ann. Appl. Statist.* 6: 1452–1477.
- Simpson, D., Rue, H., Riebler, A., Martins, T.G., and Sørbye, SH. 2017. Penalising Model Component Complexity: A Principled, Practical Approach to Constructing Priors. *Statist. Sci.* 32: 1-28.
- Sølna, K., and Switzer, P. .1996. Time trend estimation for a geographic region. *Journal of the American Statistical Association* 91: 577–589.

- Sainte-Marie, B. 1993. Reproductive cycle and fecundity of primiparous and multiparous female snow crab, *Chionoecetes opilio*, in the northwest Gulf of Saint Lawrence. *Canadian Journal of Fisheries and Aquatic Sciences* 50: 2147–2156.
- Sainte-Marie, B., and Hazel, F. 1992. Molting and mating of snow crabs, *Chionoecetes opilio*, in shallow waters of the northwest Gulf of St. Lawrence. *Canadian Journal of Fisheries and Aquatic Sciences* 49: 1282–1293.
- Sainte-Marie, B., Raymond, S., and Brethes, J.-C. 1995. Growth and maturation of the benthic stages of male snow crab, *Chionoecetes opilio* (Brachyura: Majidae). *Canadian Journal of Fisheries and Aquatic Sciences* 52:903–924.
- Tobler W. 1970. A computer movie simulating urban growth in the Detroit region. *Economic Geography* 46: 234–240.
- Wakefield, J. 2007. Disease mapping and spatial regression with count data. *Biostatistics* 8: 158–183.
- Watson, J. 1972. Mating behaviour of the spider crab, *Chionoecetes opilio*. *Journal of the Fisheries Research Board of Canada* 29: 447–449.
- Webb, J.B., Eckert G.L., Shirley, T.C., and Tamone, S.L. 2007. Changes in embryonic development and hatching in *Chionoecetes opilio* (Snow Crab) with variation in incubation temperature. *Biological Bulletin*. 213:67-75.
- Wiener, N. 1930. Generalized Harmonic Analysis. *Acta Mathematica* 55: 117–258.
- Wikle, C.K., and Cressie, N. 1999. A dimension-reduced approach to space-time Kalman filtering. *Biometrika* 86: 815–829.
- Xu, K., Wikle, C. K., and N. I. Fox. 2005. A kernel-based spatio-temporal dynamical model for nowcasting weather radar reflectivities. *Journal of the American Statistical Association* 100: 1133–1144.
- Yates, F. 1937. The design and analysis of factorial experiments. *Technical Communication of the Commonwealth Bureau of Soils*, 35.

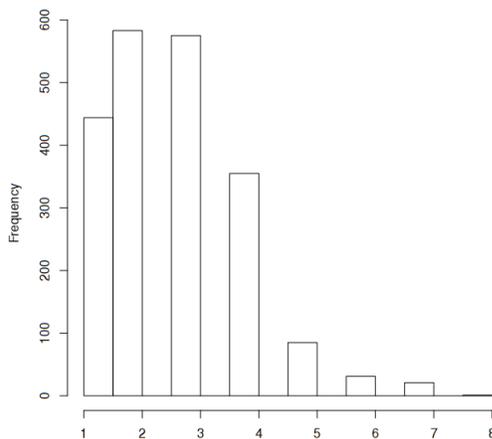


Fig. 53. Number of stations in each areal unit in each year (mean=2.6, sd=1.3).

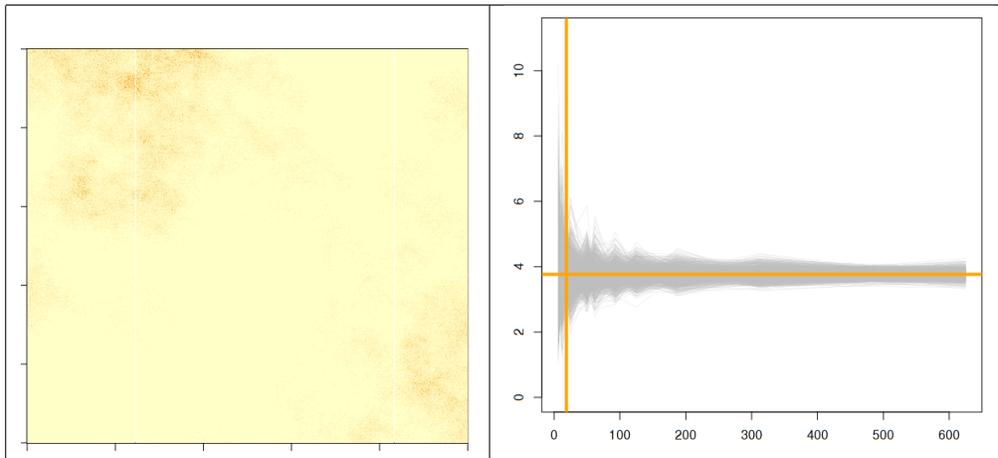
## 7 Appendix

### 7.1 Sampling bias

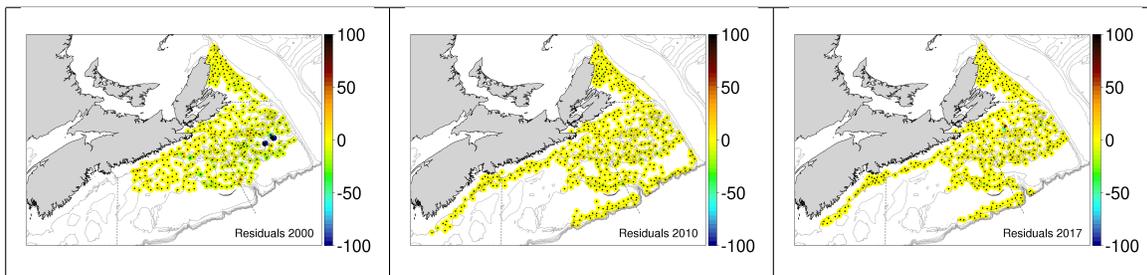
It is well understood via the Central limit theorem that multiple samples are required to estimate the overall mean and that this converges with an infinite number of samples. Not having the luxury to sample infinitely, we run into estimation issues due to small sample size. There exist heuristic rules of thumb that  $n=30$  is sufficient for beginning to estimate a stable value for central tendency. In the case of the snow crab each AU is sampled from 0 to 8 stations in each year (Figure 53). Central limit properties do not apply to such low numbers. Using a simple simulation of an AU of  $25 \times 25$  km, with cells discretized to a  $25 \times 25$  m, a single snow crab trawl is on average  $0.0039 \text{ km}^2$  which amounts to approximately 6.25 grid cells (Figure 54; left panel) out of a total of  $10^6$  cells in each AU. The simulation suggests that variability in the mean does not stabilize until  $>200$  cells are sampled which is about 32 stations (Figure 54; right panel). As a consequence, there is a high likelihood that a mean based on such limited samples will be biased relative to the (true) latent mean, especially if preferential sampling is also occurring. This reiterates the problem of using simple averages to represent a given AU; additional information is required to constrain the estimates to more realistic values which entails the use of a model-based approach to estimation.

### 7.2 Residual plots

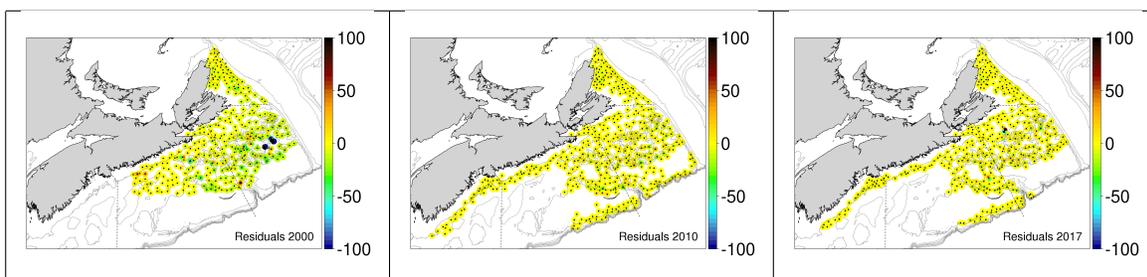
These are the model residuals (observed - expected) for a subset of the models used in this document to further evaluate model performance in a space-time context and supplement the DIC, WAIC and CPO measures of model performance, as requested by the reviewers of this document. All methods show that there are difficulties posed by hyper-aggregation of snow crab (red) and low numbers of snow crab in areas that are otherwise expected to have more (blue). The latter are likely due to localized depletion in these areas due to persistent fishing activity. The former are likely due surrounding environmental conditions being poor, forcing crab into localized areas of better conditions (“crab holes”). Both these areas tend to be in Misaine Bank, an area of high topographic complexity and nonsimple connectivity.



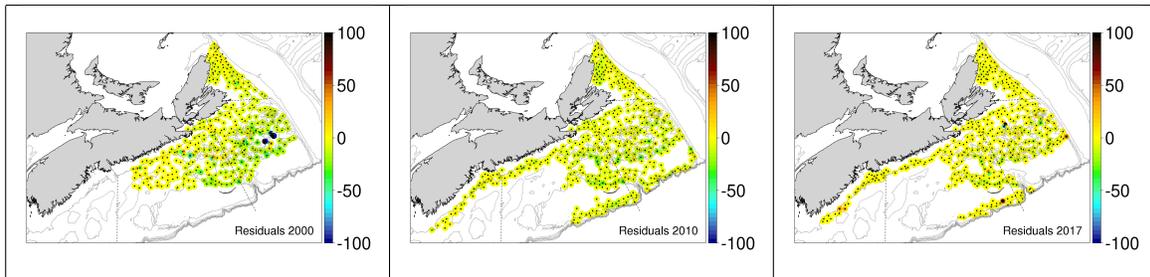
**Fig. 54.** Left: Random Gaussian field, exponentiated and truncated to mimic a Poisson process. It has an exponential variogram scale of 20 km (practical range) and a variance of 0.9 with a nugget of 0.1 and overall mean of 1. It has dimensions of 25 km x 25 km and is discretized to 1000 units in each direction. Each grid unit is therefore 25 m x 25 m. Overall mean is 5 individuals per grid cell and a standard deviation of 3.9. Right: Trajectories of mean estimates as a function of number of cells. For reference  $n=3$  stations (18.75 cells) is given as the vertical line.



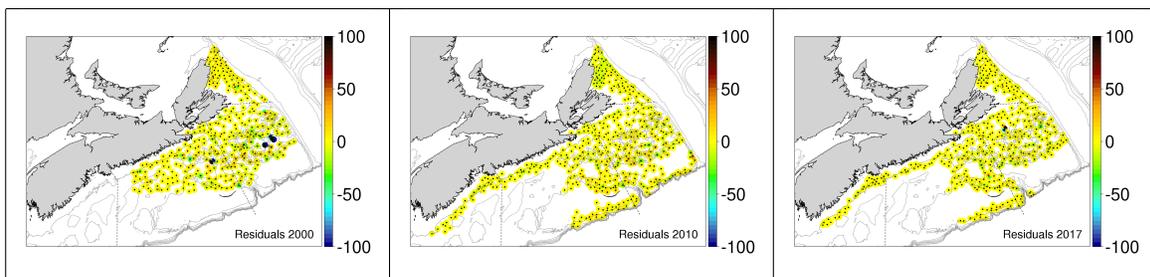
**Fig. 55.** Residuals (observed-expected) of modelled solutions in the snow crab survey, on the scale of biomass density per set (kg/rawl set) – **Factorial crossed (Gaussian) biomass.**



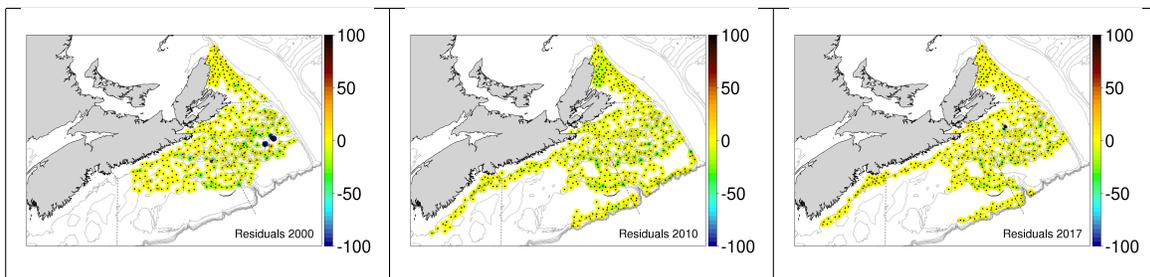
**Fig. 56.** Residuals (observed-expected) of modelled solutions in the snow crab survey (number of individuals per trawl set) – **Factorial crossed.**



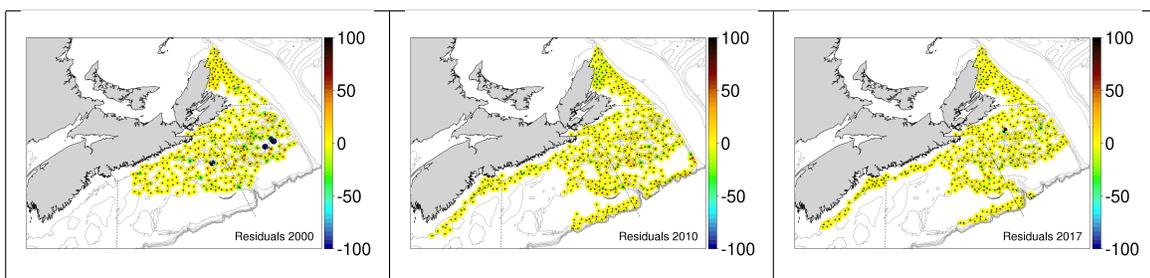
**Fig. 57.** Residuals (observed-expected) of modelled solutions in the snow crab survey (number of individuals per trawl set) – **Factorial covariates.**



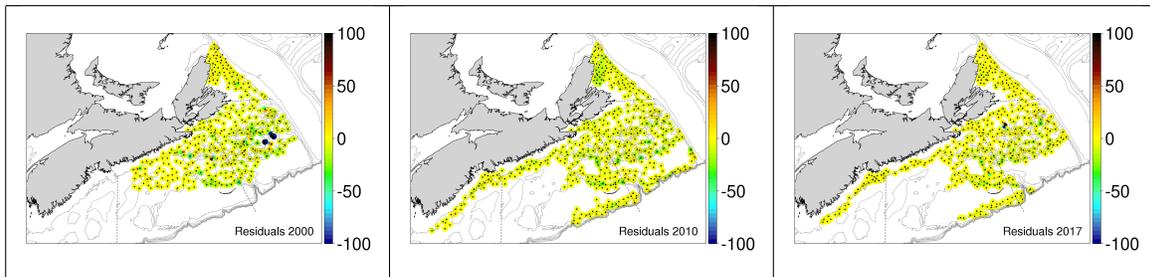
**Fig. 58.** Residuals (observed-expected) of modelled solutions in the snow crab survey (number of individuals per trawl set) – **Mixed effects simple.**



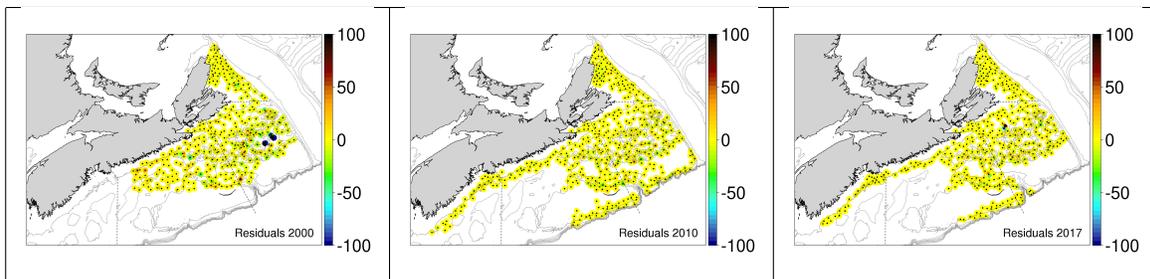
**Fig. 59.** Residuals (observed-expected) of modelled solutions in the snow crab survey (number of individuals per trawl set) – **Mixed effects dynamic.**



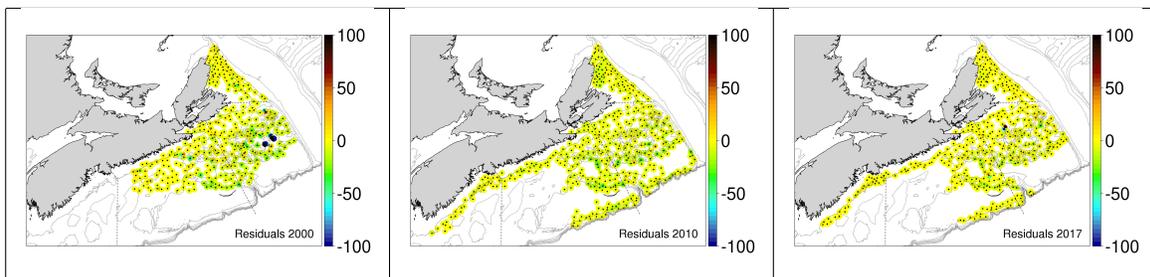
**Fig. 60.** Residuals (observed-expected) of modelled solutions in the snow crab survey (number of individuals per trawl set) – **Separable simple.**



**Fig. 61.** Residuals (observed-expected) of modelled solutions in the snow crab survey (number of individuals per trawl set) – **Separable**.



**Fig. 62.** Residuals (observed-expected) of modelled solutions in the snow crab survey, on the scale of number of individuals per trawl set – **Nonseparable simple**.



**Fig. 63.** Residuals (observed-expected) of modelled solutions in the snow crab survey, on the scale of number of individuals per trawl set – **Nonseparable spacetime**.