

JuMP-Dev 2024

Time series modeling via JuMP

Professor: Davi Valladão

Students:

PhD: André Ramos, Marina Dietze, Tomás Gutierrez, Matheus Alves

MSc: Luiz Fernando, Matheus Nogueira

UG: Luisa Compasso, João Couto

- The main objective of this project is to implement traditional time series models in Julia, harnessing the power of JuMP to provide access to state-of-the-art solvers, explores different modeling choices and extend existing methodologies via novel optimization techniques.
- Contributions:
 - **SARIMAX.jl**
 - **UnobservedComponetsGAS.jl**
 - **StateSpaceLearning.jl**

Our packages




- **SARIMAX.jl**
 - Implements the auto SARIMA model via modern optimization techniques.
- **UnobservedComponentsGAS.jl**
 - Estimate a conditional time-varying dynamic for a given probabilistic distribution parameters.
- **StateSpaceLearning.jl**
 - Estimate a structural time series model via a high-dimensional regression.

- Given ARMA(p,q)

$$y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{j=1}^q \theta_j \varepsilon_{t-j} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma^2)$$

- Parameter estimation: existing approaches
- Maximum likelihood via Kalman filtering



- Under the optimization lens:

$$\min_{c, \Phi, \Theta, \varepsilon} \sum_{t=1}^T \varepsilon_t^2$$
$$s.t: \quad y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{j=1}^q \theta_j \varepsilon_{t-j} + \varepsilon_t$$

- Under the optimization lens:

$$\min_{c, \Phi, \Theta, \varepsilon} \sum_{t=1}^T \varepsilon_t^2$$
$$s.t: \quad y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{j=1}^q \theta_j \varepsilon_{t-j} + \varepsilon_t$$

- JuMP.jl formulation

```
m = Model(Ipopt.Optimizer)
@variable(m, c)
@variable(m, θ[1:q])
@variable(m, φ[1:p])
@variable(m, ε[1:T])
@constraint(m, [t = (max(p, q) + 1):T], ε[t] == y[t] - (c + sum(φ[j]*y[t-j]
    for j in 1:p) + sum(θ[j]*ε[t-j] for j in 1:q)))
@objective(m, Min, sum((ε[t])^2 for t in max(p, q) + 1:T))
optimize!(m)
```

- Package usage

```
model = SARIMA(trainingSet, 1, 0, 1;  
                seasonality=12, P=0, D=1, Q=2, allowMean=false)  
fit!(model; objectiveFunction="mse")  
forecast = predict!(model, stepsAhead)  
scenarios = simulate(model, stepsAhead, 200)
```

- Controlled experiment: 1000 random ARMA stationary processes

Model	Percentage of best estimation metrics					
	AIC	BIC	AICc	MSE	Log Likel.	Time
Statsmodels, pmdarima	0.9	0.9	0.9	0.8	0.7	0.0
StateSpaceModels.jl	1.0	1.0	1.0	0.8	0.9	20.7
Rforecast	0.9	0.9	0.9	1.2	1.2	64.1
SARIMAX.jl (MSE)	89.1	89.1	89.1	89.1	89.1	15.1
Tied models	8.0	8.0	8.0	8.0	8.0	0.0

Model	Percentage of erros during estimation (%)
	(%)
Statsmodels	0.0
Pmdarima	0.0
StateSpaceModels.jl	9.8
Rforecast	0.6
SARIMAX.jl (MSE)	0.0

Possible extensions

$$\begin{aligned} & \underset{c, \phi_t, \theta_t \epsilon_t}{\text{minimize}} && \sum_{t=1}^T \epsilon_t^2 \\ & \text{subject to} && y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \epsilon_t, \quad \forall t \in 1, \dots, T \end{aligned}$$

Add
explanatory
variables



$$\begin{aligned} & \underset{c, \Phi, \Theta, \beta, \epsilon_t}{\text{minimize}} && \sum_{t=1}^T \epsilon_t^2 \\ & \text{subject to} && y_t = c + \sum_{i=1}^m \beta_i X'_{ti} + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \epsilon_t, \quad \forall t \in \{1, \dots, T\} \end{aligned}$$

Change
objective
function



$$\begin{aligned} & \underset{c, \phi_t, \theta_t \epsilon_t}{\text{minimize}} && \sum_{t=1}^T |\epsilon_t| \\ & \text{subject to} && y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \epsilon_t, \quad \forall t \in 1, \dots, T \end{aligned}$$


Add
regularization




$$\begin{aligned} & \underset{c, \phi_t, \theta_t, \epsilon_t}{\text{minimize}} && \sum_{t=1}^T \epsilon_t^2 + \lambda \left(\sum_{i=1}^N |\Psi_i| \right) \\ & \text{subject to} && y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \epsilon_t, \quad \forall t \in 1, \dots, T \\ & && \Psi = \{\phi, \theta\} \end{aligned}$$

Possible extensions

$$\begin{aligned} & \underset{c, \phi_t, \theta_t \epsilon_t}{\text{minimize}} && \sum_{t=1}^T \epsilon_t^2 \\ & \text{subject to} && y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \epsilon_t, \quad \forall t \in 1, \dots, T \end{aligned}$$




Specification: Auto-ARIMA (Hyndman-Khandakar)



Optimal Specification/Estimation: MINLP

Possible extensions

$$\begin{aligned} & \underset{c, \phi_t, \theta_t \epsilon_t}{\text{minimize}} && \sum_{t=1}^T \epsilon_t^2 \\ & \text{subject to} && y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \epsilon_t, \quad \forall t \in 1, \dots, T \end{aligned}$$



Optimal Specification/Estimation: MINLP

$$\begin{aligned} & \underset{c, \phi_t, \theta_t \epsilon_t}{\text{minimize}} && \sum_{t=1}^T \epsilon_t^2 \\ & \text{subject to} && y_t = c + \sum_{i=1}^p \phi_i y_{t-i} + \sum_{i=1}^q \theta_i \epsilon_{t-i} + \epsilon_t, \quad \forall t \in 1, \dots, T \\ & && \|\Psi\|_0 \leq K \end{aligned}$$

Advantages

- State-of-the-art optimization algorithms through modern solvers
 - Gurobi;
 - Ipopt;
 - SCIP.
- Possibility to use global solvers, like
 - Alpine;
 - Eago;
- Access to a wide range of optimization resources, including:
 - Flexibility to add necessary restrictions;
 - Regularization terms in the objective function;
 - Initialization techniques, e.g. warm-start;
 - Mixed Integer Programming (MIP);
 - Robust optimization techniques.

- Estimates a conditional time-varying dynamic for a given probabilistic distribution parameters

$$y_t \sim p(y_t | \mu_{t|t-1}; \mathbf{Y}_{t-1}, \sigma, \theta) = N(\mu_{t|t-1}, \sigma)$$

$$\mu_{t+1|t} = m_{t+1|t}$$

$$m_{t+1|t} = m_{t|t-1} + b_{t|t-1} + \kappa_m S_t$$

$$b_{t+1|t} = b_{t|t-1} + \kappa_b S_t$$

- Optimization problem

$$\max_{\substack{\mu, m, b, \kappa_m, \kappa_b, \sigma, \\ s(\mu, \sigma, y)}} \sum_{t=1}^T \left(-\frac{1}{2} \ln(2\pi\sigma^2) - \frac{(y_t - \mu_{t|t-1})^2}{2\sigma^2} \right)$$

$$s. t. : \mu_{t+1|t} = m_{t+1|t} \quad \forall t = 1, \dots, T$$

$$m_{t+1|t} = m_{t|t-1} + b_{t|t-1} + \kappa_m S_t \quad \forall t = 1, \dots, T$$

$$b_{t+1|t} = b_{t|t-1} + \kappa_b S_t \quad \forall t = 1, \dots, T$$

$$\sigma > 0$$

$$LB \leq \kappa_m \leq UB$$

$$LB \leq \kappa_b \leq UB$$

- Currently available distributions:
 - Gaussian
 - LogNormal
 - t-Location-Scale
- More continuous and discrete distributions will be implemented.
- Currently available components structures:
 - Level: random walk, random walk with slope, $AR(1)$
 - Seasonality with trigonometric functions: deterministic and stochastic
 - Autoregressive of order p
 - Explanatory variables

- Package usage

```
dist = UnobservedComponentsGAS.NormalDistribution()
d     = 0.0

params      = [true, false]
level      = ["random walk slope", ""]
seasonality = ["deterministic 12", ""]
ar         = [missing, missing]
gas_model  = UnobservedComponentsGAS.GASModel(dist, params, d, level,
                                              seasonality, ar)

fitted_model = UnobservedComponentsGAS.fit(gas_model, y)
forec        = UnobservedComponentsGAS.predict(gas_model, fitted_model, y,
                                              steps_ahead, num_scenarios)
```

- The Basic Structural Model

$$y_t = \mu_t + \gamma_t + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2)$$

$$\mu_t = \mu_{t-1} + \nu_t + \xi_t, \quad \xi_t \sim N(0, \sigma_\xi^2)$$

$$\nu_t = \nu_{t-1} + \zeta_t, \quad \zeta_t \sim N(0, \sigma_\zeta^2)$$

$$\gamma_t = - \sum_{j=1}^{s-1} \gamma_{t-j} + \omega_t, \quad \omega_t \sim N(0, \sigma_\omega^2)$$

- Can be written as a high-dimensional regression:

$$y_t = \mu_0 + t\nu_0 + \sum_{\tau=1}^{t-1} (t - \tau + 1)\zeta_\tau + \sum_{\tau=1}^{t-1} \xi_\tau + \gamma_{\tau-s}M_t + \sum_{j=1}^{M_t} (\omega_{t-js} - \omega_{t-js-1}) + \varepsilon_t^*$$

- For an illustrative example, consider the local level model:

$$y_t = \mu_t + \epsilon_t$$
$$\mu_t = \mu_{t-1} + \eta_t$$

- By iterating the level component we obtain:

$$y_t = \mu_1 + \sum_{\tau=2}^t \eta_\tau + \epsilon_t$$

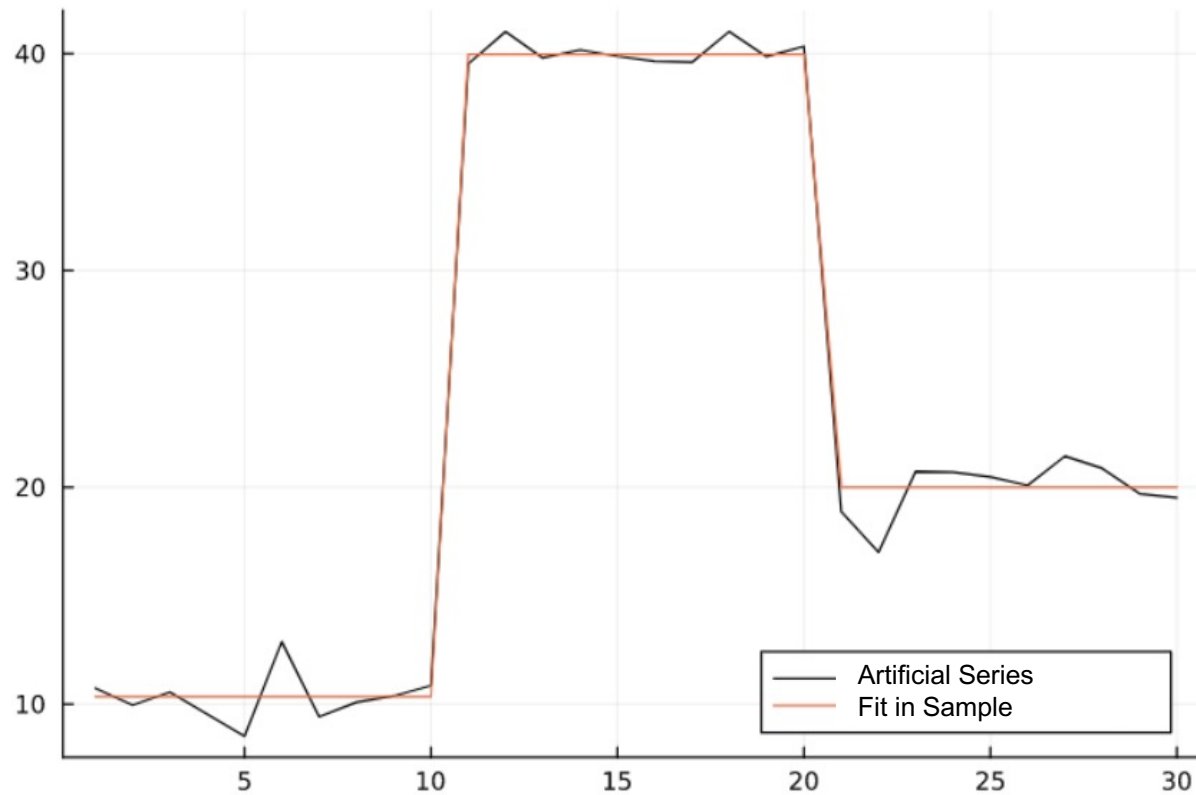
- This equation can be written in a vectorized form:

$$\mathbf{y} = \mu_1 + \mathbf{X}\boldsymbol{\eta} + \boldsymbol{\epsilon},$$

- Where $\mathbf{y} = (y_1, \dots, y_T)$, $\boldsymbol{\eta} = (\eta_2, \dots, \eta_T)$ and:

$$X = \begin{bmatrix} 0 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ 1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & 1 & \dots & 0 \\ 1 & 1 & \dots & 1 \end{bmatrix}$$

- Result of a Lasso estimation for the local level model in regression form :



- In principle we fit a regularized regression (“AdaLasso”) using “coordinate descent” algorithm.
- To explore other regularization techniques such as the norm L0, we will rely on JuMP to solve the optimization model.

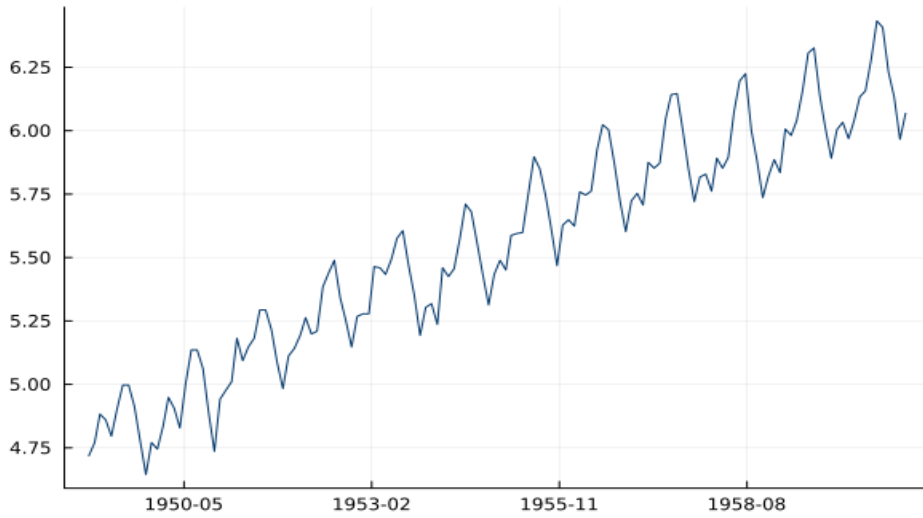
$$\min_{\beta} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|$$

Package usage:


```
output = StateSpaceLearning.fit_model(y)
prediction = StateSpaceLearning.forecast(output, steps_ahead)
```

Comparison

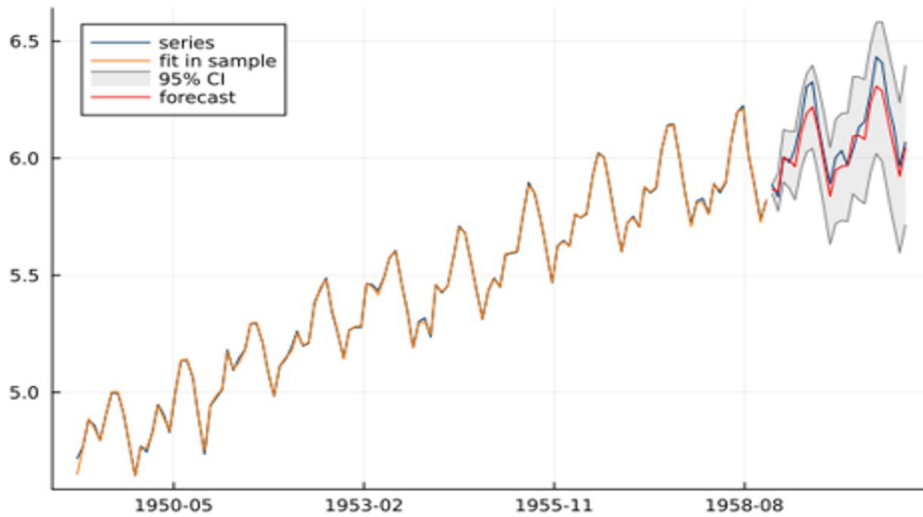
Airline Time Series




Fit and Forecast - Airline - State Space Learning



Fit and Forecast - Airline - UnobservedComponentsGAS



Fit and Forecast - Airline - Sarimax



Numerical Results

Results for out of sample forecasting accuracy for 48000 monthly M4 time series competition

MASE

Models	Mean				Median			
	Short	Medium	Long	Total	Short	Medium	Long	Total
autoARIMA Python	0.63	0.969	1.24	0.947	0.488	0.719	0.909	0.745
StateSpaceModels.jl	0.662	1.031	1.34	1.011	0.504	0.741	0.929	0.765
SARIMAX.jl	0.632	0.963	1.229	0.942	0.478	0.708	0.891	0.733
UnobservedComponentsGAS.jl	1.717	4.161	6.738	4.205	0.679	0.883	1.106	0.924
StateSpaceLearning.jl	0.654	0.96	1.194	0.936	0.501	0.695	0.845	0.72

Thank You!!