Message passing-based inference in an autoregressive active inference agent

Wouter Kouw, Tim Nisslbeck, Wouter Nuijten
International Workshop on Active Inference 2025

Problem statement

Navigation under *known* dynamics is easy:

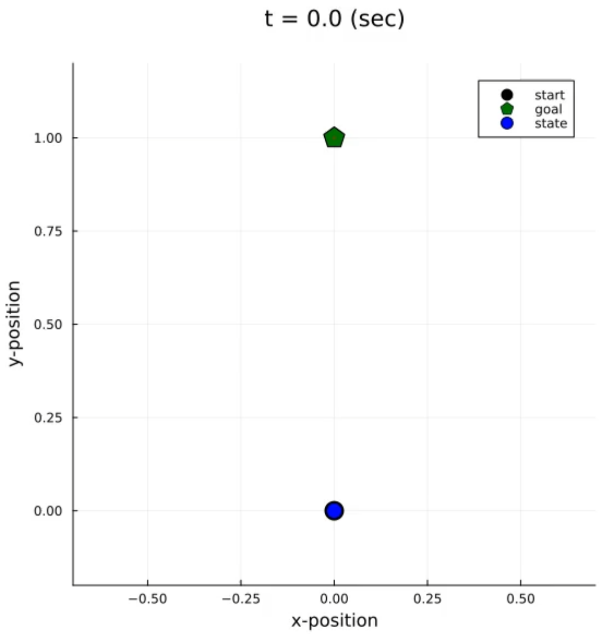

Navigation under *unknown* dynamics is hard:

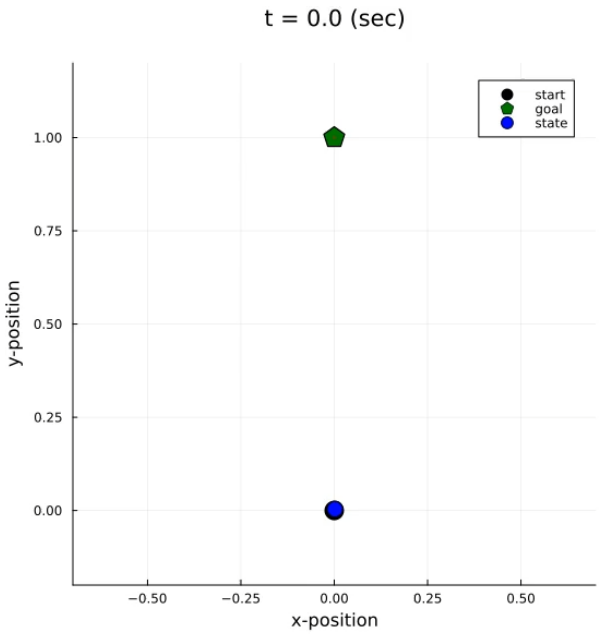

Active inference agent

We specify an agent that will cautiously learn dynamics before heading towards the goal.

- It will infer parameters by exact Bayesian filtering in an autoregressive model.
- It plans actions by minimizing expected free energy in continuous states and actions.

The agent infers by message passing in factor graphs:

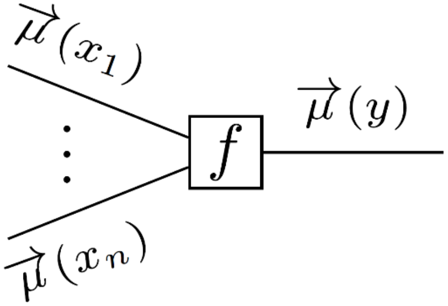

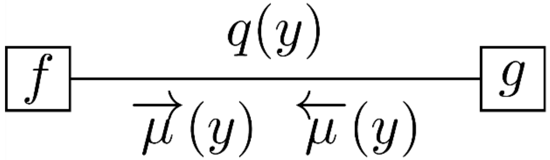

Generative model specification

Multivariate autoregressive likelihood function:

$$p(y_k \,|A, W, u_k, \bar{u}_k, \bar{y}_k) = \mathcal{N}\left(y_k \,\middle|\, A^\top \begin{bmatrix} \bar{y}_{k-1} \\ u_k \\ \bar{u}_{k-1} \end{bmatrix}, W^{-1} \right)$$

Conjugate prior over parameters:

$$p(A, W) = \mathcal{MNW}(A, W \mid M, \Lambda^{-1}, \Omega^{-1}, \nu)$$

Zero-mean Gaussian prior over controls:

$$p(u_k) = \mathcal{N}(u_k | \, 0, \Upsilon^{-1})$$

Gaussian distributed goals:

$$p(y_k | y_*) = \mathcal{N}(y_k \mid m_*, S_*)$$

TU/e

Inference

Parameter estimation by exact Bayesian filtering,

$$p(A, W|\mathcal{D}_k) = \frac{p(y_k \,|\mathrm{A}, \mathrm{W}, u_k, \bar{u}_k, \bar{y}_k)}{p(y_k \,|u_k, \mathcal{D}_{k-1})} p(A, W \,|\mathcal{D}_{k-1}) \,,$$

leading to closed-form parameter update rules.

The posterior predictive distribution is also exact:

$$p(y_{k+1}|u_{k+1}, \mathcal{D}_k) = \int p(y_{k+1} \,|\mathrm{A}, \mathrm{W}, u_{k+1}, \bar{u}_{k+1}, \bar{y}_{k+1}) p(A, W|\mathcal{D}_k) d(A, W)$$

$$= \mathcal{T}_{\eta_{k+1}}(y_{k+1} \mid \mu_{k+1}(u_{k+1}), \Sigma_k(u_{k+1}))$$

Nisslbeck & Kouw (2025). Online Bayesian system identification in multivariate autoregressive models via message passing. *ECC.*

Inference

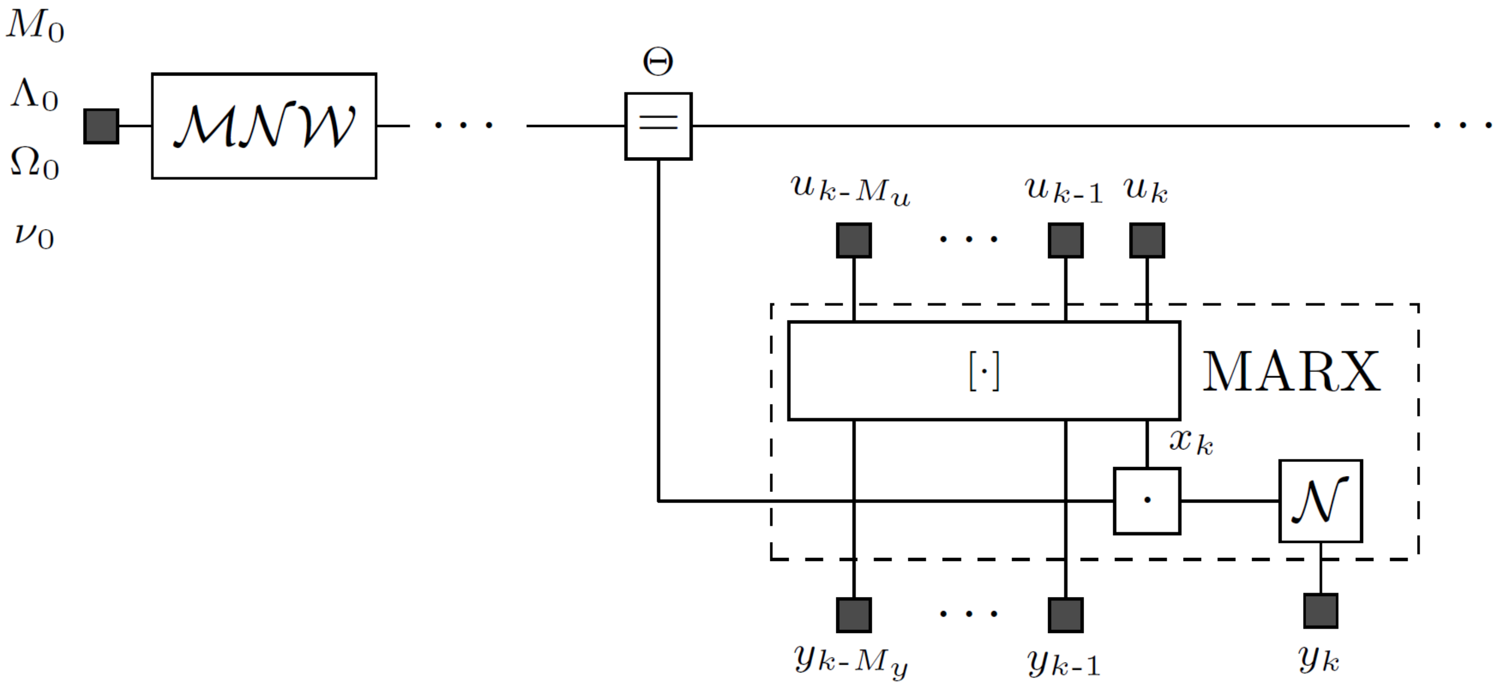

TU/e

Inference

Planning by expected free energy minimization, for $t = k + 1$:

$$F[q] = \int q(y_t, u_t, \Theta) \ln \frac{q(u_t, \Theta)}{p(y_t, u_t, \Theta | y_*, \mathcal{D}_k)} d(y_t, u_t, \Theta)$$

where $q(y_t, u_t, \Theta) = p(y_t \mid \Theta, u_t, \bar{u}_t, \bar{y}_t)\, p(\Theta \mid \mathcal{D}_k)\, q(u_t)$.

The optimal variational posterior is $q(u_t) \propto p(u_t) \exp\big(-G(u_t)\big)$ where

$$\begin{aligned} G(u_t) &= -\mathbb{E}_{p(y_t, \Theta \mid u_t, \mathcal{D}_k)} \left[\ln \frac{p(y_t, \Theta \mid u_t, \mathcal{D}_k)}{p(y_t \mid u_t, \mathcal{D}_k)\, p(\Theta | \mathcal{D}_k)} \right] - \mathbb{E}_{p(y_t \mid u_t, \mathcal{D}_k)} [\ln p(y_t | y_*)] \\ &= \text{constants} - \frac{1}{2} \ln |\Sigma_t(u_t)| + \frac{1}{2} \mathrm{Tr} \left[S_*^{-1} (\Sigma_t(u_t) \frac{\eta_t}{\eta_t - 2} + (\mu_t(u_t) - m_*)(\mu_t(u_t) - m_*)^\top \right] \end{aligned}$$

Inference

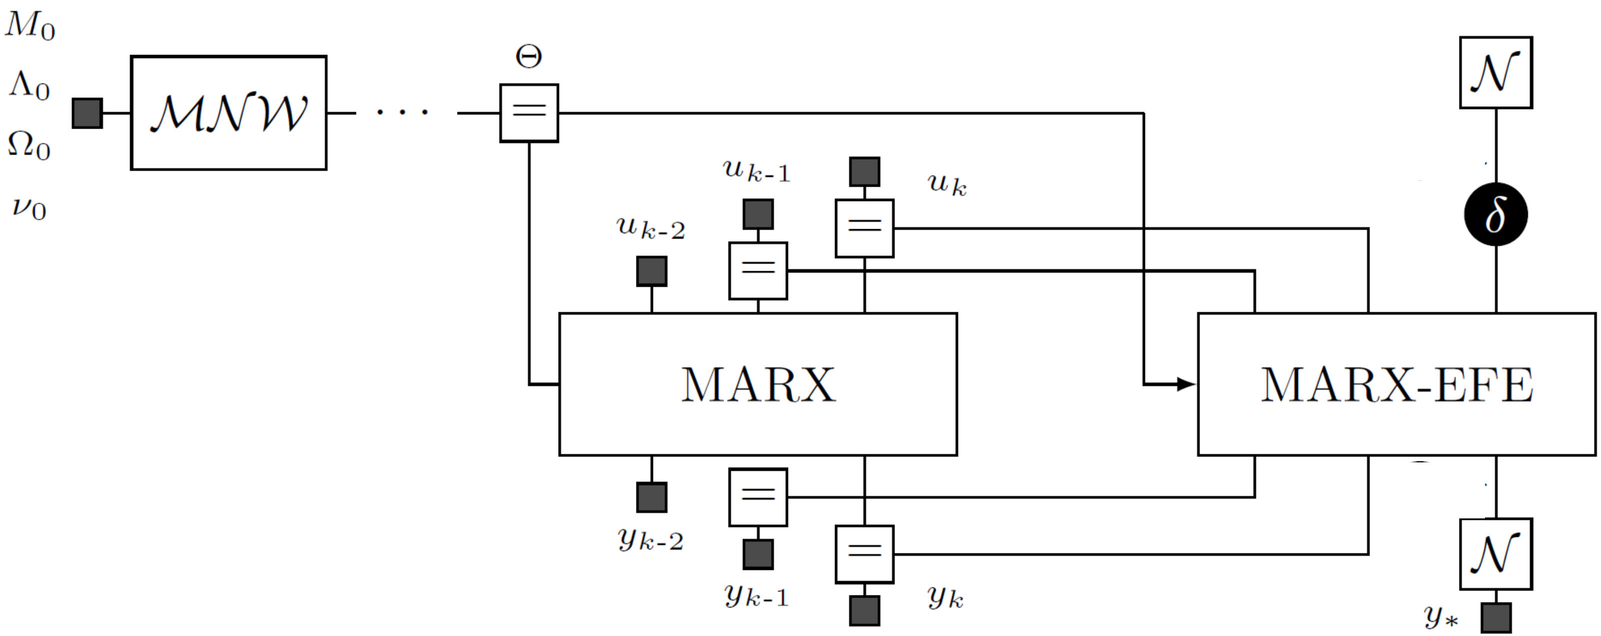

M_0
Λ_0
Ω_0
ν_0
$\mathcal{MNW}$
Θ
$u_{k\text{-}2}$
$u_{k\text{-}1}$
u_k
MARX
MARX-EFE
$\mathcal{N}$
δ
$\mathcal{N}$
$y_{k\text{-}2}$
$y_{k\text{-}1}$
y_k
y_*

Inference

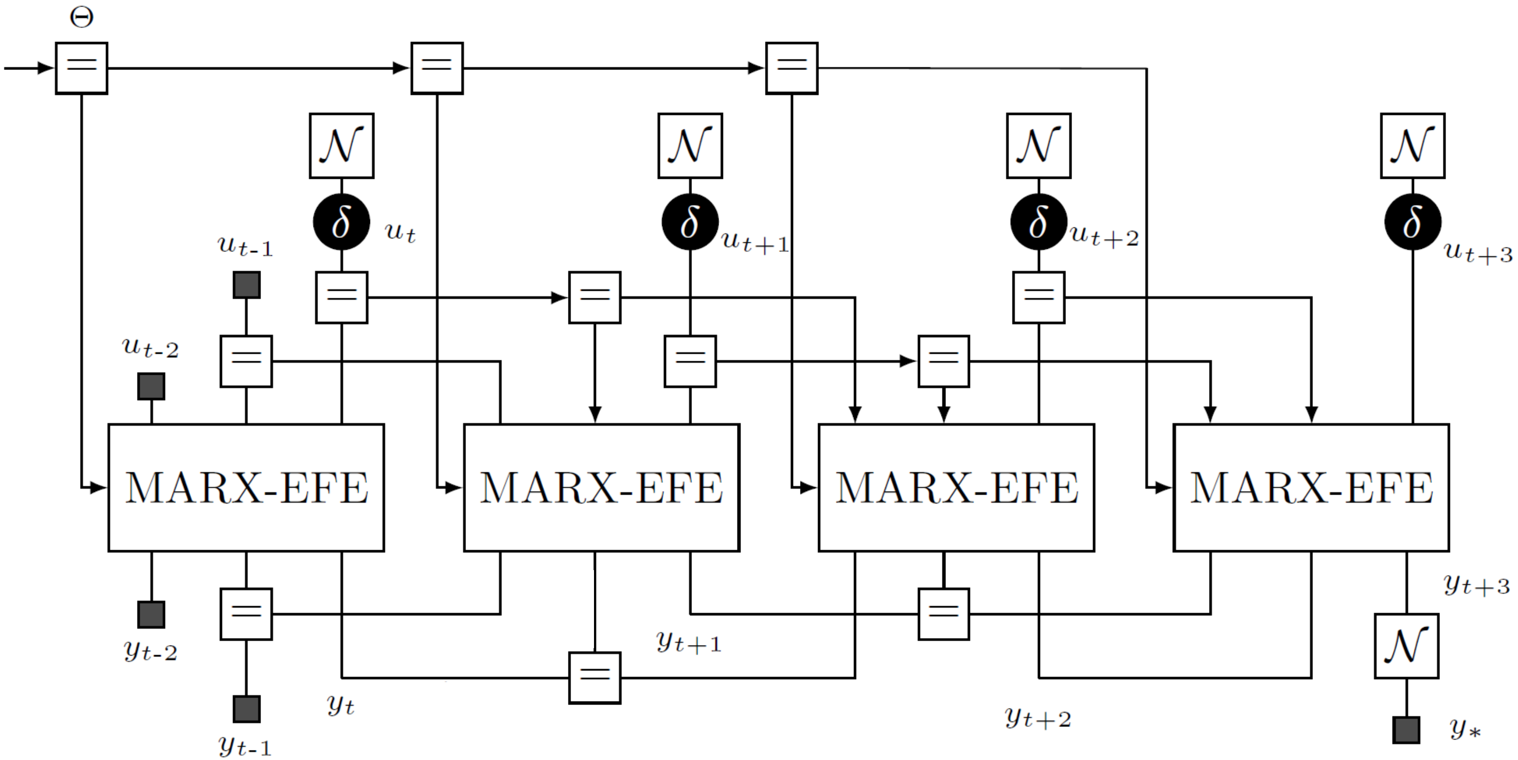

Θ
𝒩
δ
u_t
u_{t+1}
u_{t+2}
u_{t+3}
u_{t-1}
u_{t-2}
MARX-EFE
MARX-EFE
MARX-EFE
MARX-EFE
y_{t-2}
y_{t-1}
y_t
y_{t+1}
y_{t+2}
y_{t+3}
y_*

Experiment

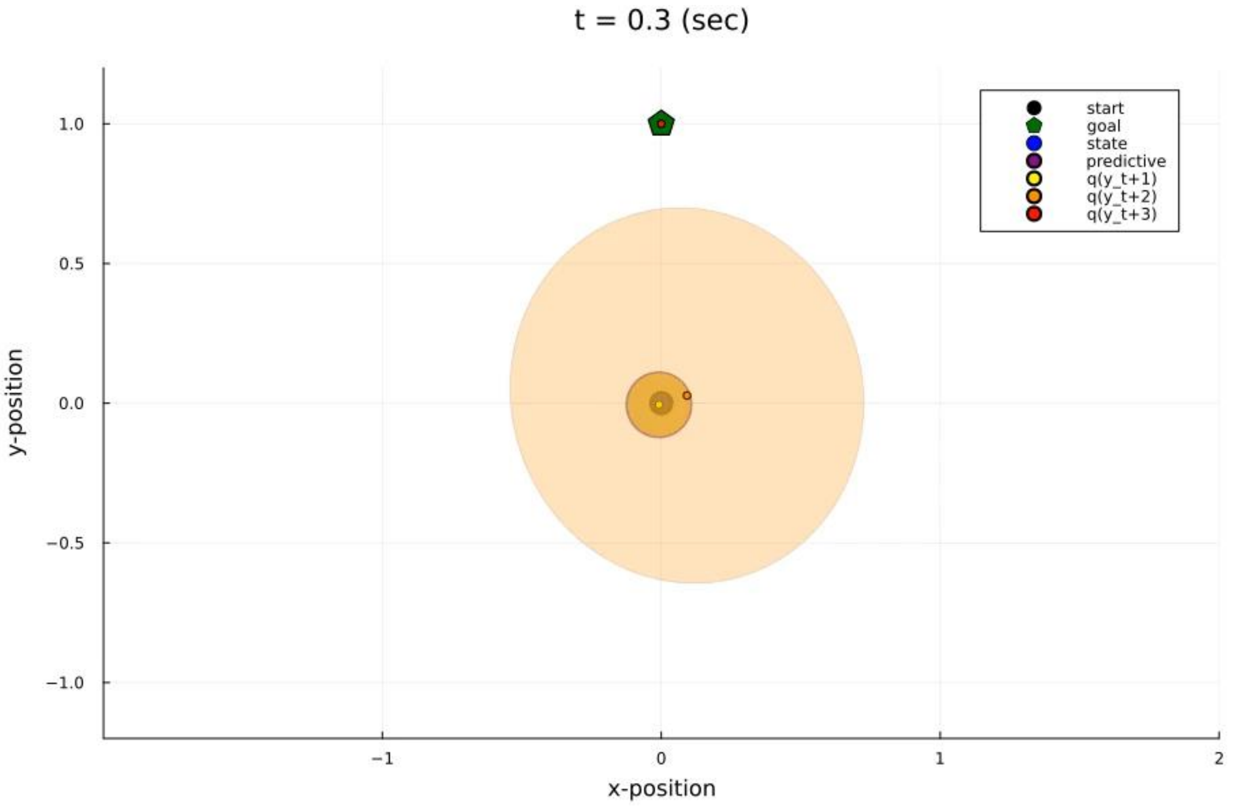
t = 0.3 (sec)
1.0
0.5
0.0
−0.5
−1.0
y-position
−1
0
1
2
x-position
start
goal
state
predictive
q(y_t+1)
q(y_t+2)
q(y_t+3)

Discussion

Take-aways

- Forward + backward pass generates sequence of intermediate goal priors.
- Each time step in planning horizon solves a small EFE minimization problem.
- Local computation and modular design to ensure scalability.

Future work

- Hierarchies over timescales.
- A higher-order discrete agent that switches between lower-order agents.
- Adding a memory that captures obstacles in space.

Thank you

Wouter Nuijten

Tim Nisslbeck

BIASlab rxinfer

TU/e