



# Hyperion – A fast, versatile symbolic Gaussian Belief Propagation Framework for Continuous-Time SLAM

David Hug<sup>1</sup>, Ignacio Alzugaray<sup>2</sup>, Margarita Chli<sup>1</sup>

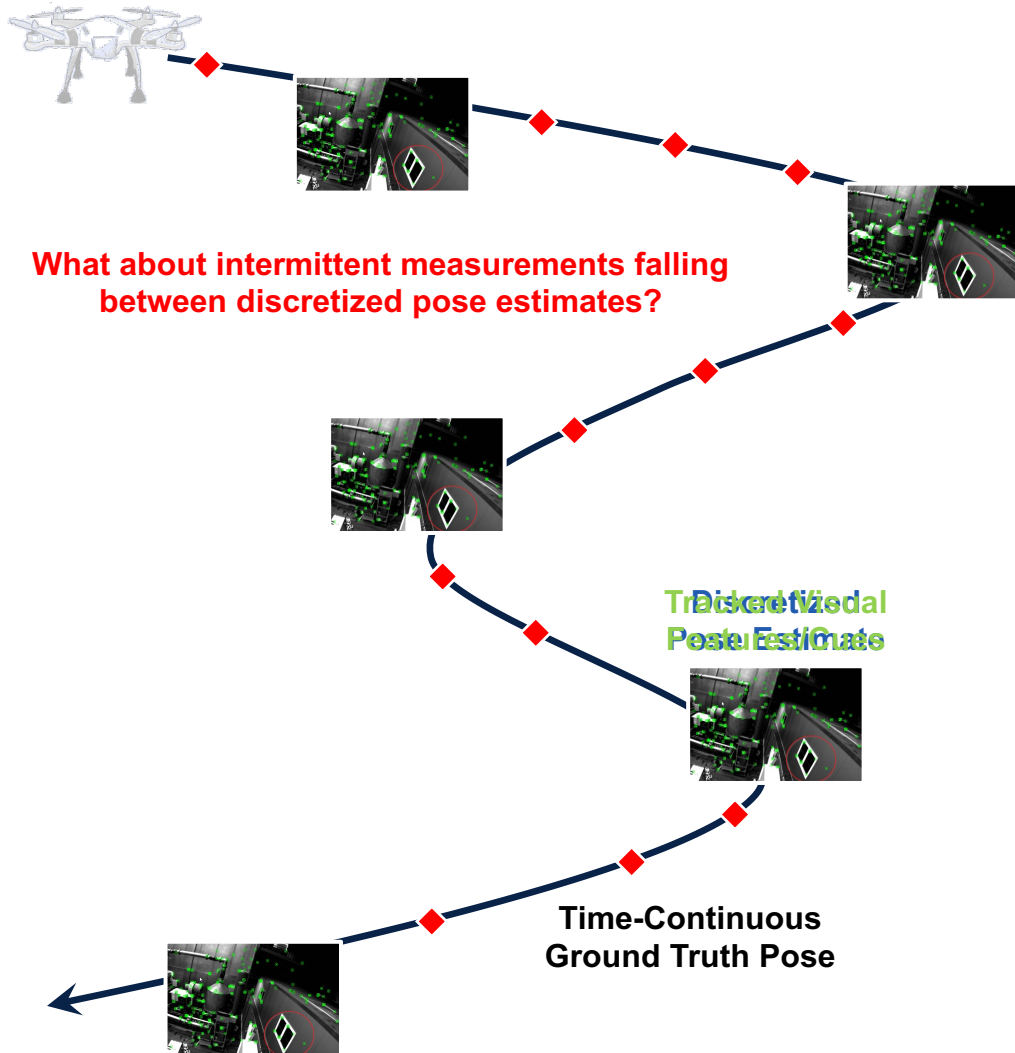
<sup>1</sup>Vision for Robotics Lab (V4RL), ETH Zurich and University of Cyprus

<sup>2</sup>Dyson Robotics Lab, Imperial College London

Presented at the European Conference on Computer Vision 2024

# Related Work

## Discrete-Time SLAM



Speed x7

Follow Camera  
Show Points  
Show Keyframes  
Show Graph  
Localization Mode  
Reset

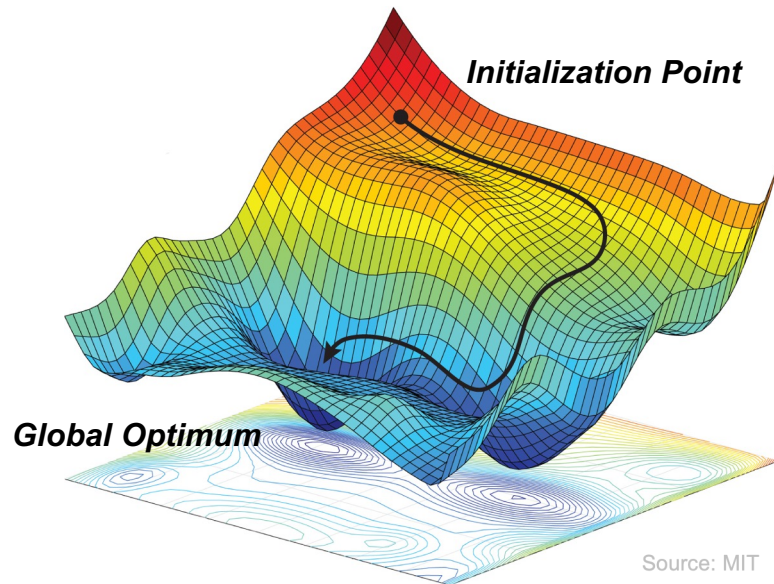
Conventional pipelines assume synchronized inputs and steady sensor acquisition rates

Source: Mur-Artal et al.

RGB-D

# Related Work

## Non-Linear Least Squares Optimization (Single Agent)



Surface Plot of  
Cumulative Residuals

**Minimization Problem<sup>[1]</sup>**

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \left[ \frac{1}{2} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \|\bar{\mathbf{r}}(t, \theta)\|^2 \right] \text{ with } \theta \subseteq \Theta$$

*Optimal Parameters*    *Sum over square of weighted residuals stemming from all sensors and measurement times*

**Weighted Residual**

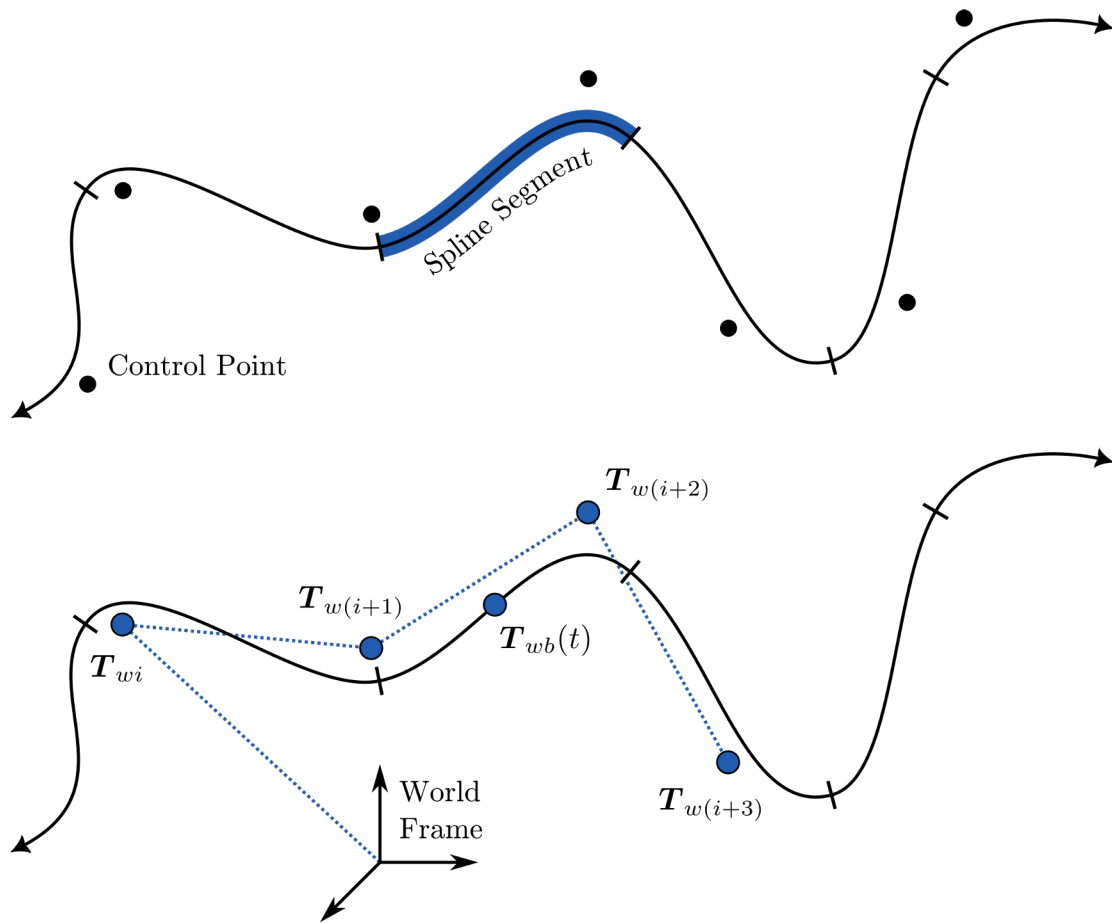
$$\|\bar{\mathbf{r}}(t, \theta)\|^2 = \bar{\mathbf{r}}^\top \bar{\mathbf{r}} = \mathbf{r}^\top \underset{\text{Precision}}{\Sigma_m^{-1}} \mathbf{r}$$

**Residual**                      **Distance Metric**

$$\mathbf{r}(t, \theta) = \underset{\text{Predicted Measurement}}{\hat{\mathbf{m}}(t, \theta)} \ominus_{\mu} \underset{\text{Actual Measurement}}{\mathbf{m}(t, \theta)}$$

# Related Work

## Continuous-Time SLAM<sup>2</sup>



$$T_{wb}(t) = T_{wi} \prod_{j=1}^k \left( T_{w(i+j-1)}^{-1} T_{w(i+j)} \right)^{\lambda_{i,j}(t)} \in \text{SE}(3)$$

*Spline Degree*
*Interpolation Coefficient*

*Base Control Point*
*Incremental Elements in the Lie Group*

$$= \exp \left( \lambda_j(t) \log \left( T_{w(i+j-1)}^{-1} T_{w(i+j)} \right) \right)$$

*Geodesic Curves*

$$\Lambda_i(t) = \begin{bmatrix} \lambda_{i,0}(t) \\ \lambda_{i,1}(t) \\ \lambda_{i,2}(t) \\ \lambda_{i,3}(t) \end{bmatrix} = \frac{1}{3!} \begin{bmatrix} 6 & 0 & 0 & 0 \\ 5 & 3 & -3 & 1 \\ 1 & 3 & 3 & -2 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u^0(t) \\ u^1(t) \\ u^2(t) \\ u^3(t) \end{bmatrix}$$

*B-Spline Interpolation Coefficients<sup>[1]</sup>*

*Vector of Normalized Times*

[1] Qin, PCCGA 1998  
 [2] Hug et al., RA-L 2022

# Methodology

## Approach



Mathematical expressions for spline-related *residuals remain tedious and error-prone*, leading to *suboptimal performance*



Standard NLLS optimizers *neither model uncertainties* nor *do they (trivially) extend to distributed computations* across multiple agents

### Approach

Leverage Gaussian Belief Propagation (GBP) for *distributed, stochastic inference* along with *automating code generation*


# Methodology

## Code Generation

**Approach**  
Extend SymForce<sup>[1]</sup> to *delegate the generation of performance-critical code* within the framework



*High-level, symbolic mathematical expressions and residuals*



*Translation of complex expressions and symbolic optimization*




*High-performance machine code lacking interpretability*

# Experiments

## Code Generation



*High-level, symbolic mathematical expressions and residuals*



*Translation of complex expressions and symbolic optimization*



*High-performance machine code lacking interpretability*

$\mathcal{L}$	Setup		Pose [s]		Velocity [s]		Acceleration [s]		Avg. Speedup
	$k$	$\partial/\partial\mathcal{B}$	Ours	Basalt	Ours	Basalt	Ours	Basalt	
SO(3)	4	✗	<b>1.64e-7</b>	3.16e-7	<b>9.49e-8</b>	2.90e-7	<b>1.12e-7</b>	3.28e-7	2.64x
SO(3)	4	✓	<b>5.03e-7</b>	6.70e-6	<b>4.05e-7</b>	7.82e-6	<b>5.11e-7</b>	9.46e-6	17.05x
SO(3)	5	✗	<b>1.96e-7</b>	4.14e-7	<b>1.32e-7</b>	3.67e-7	<b>1.39e-7</b>	3.96e-7	2.58x
SO(3)	5	✓	<b>6.89e-7</b>	1.08e-5	<b>5.82e-7</b>	1.27e-5	<b>7.87e-7</b>	1.58e-5	19.19x
SO(3)	6	✗	<b>2.17e-7</b>	4.78e-7	<b>1.82e-7</b>	4.42e-7	<b>1.81e-7</b>	4.83e-7	2.43x
SO(3)	6	✓	<b>8.19e-7</b>	1.57e-5	<b>7.51e-7</b>	1.87e-5	<b>1.01e-6</b>	2.43e-5	22.71x
SE(3)	4	✗	<b>1.62e-7</b>	7.03e-7	<b>1.38e-7</b>	7.46e-7	<b>1.34e-7</b>	6.88e-7	4.96x
SE(3)	4	✓	<b>5.71e-7</b>	4.69e-5	<b>9.11e-7</b>	5.35e-5	<b>1.12e-6</b>	6.25e-5	65.56x
SE(3)	5	✗	<b>1.96e-7</b>	7.38e-7	<b>1.91e-7</b>	8.73e-7	<b>1.70e-7</b>	9.15e-7	4.57x
SE(3)	5	✓	<b>7.32e-7</b>	9.64e-5	<b>1.27e-6</b>	9.44e-5	<b>1.47e-6</b>	1.14e-4	94.53x
SE(3)	6	✗	<b>2.53e-7</b>	9.11e-7	<b>2.34e-7</b>	1.12e-6	<b>2.23e-7</b>	1.12e-6	4.47x
SE(3)	6	✓	<b>9.29e-7</b>	1.25e-4	<b>1.54e-6</b>	1.67e-4	<b>1.99e-6</b>	1.75e-4	110.31x

**Performance comparison between our auto-generated and optimized B-Spline implementation and the hand-crafted implementation from Sommer et al. [1]**

# Experiments

## Code Generation



High-level, symbolic mathematical expressions and residuals



Translation of complex expressions and symbolic optimization



High-performance machine code lacking interpretability

$\mathcal{L}$	Setup		Pose [s]		Velocity [s]		Acceleration [s]		Avg. Speedup
	$k$								
SO(3)	4								2.64x
SO(3)	4								17.05x
SO(3)	5								2.58x
SO(3)	5								19.19x
SO(3)	6	✗	2.17e-7	4.78e-7	1.82e-7	4.42e-7	1.81e-7	4.85e-7	2.43x
SO(3)	6	✓	8.19e-7	1.57e-5	7.51e-7	1.87e-5	1.01e-6	2.43e-5	22.71x
SE(3)	4	✗	1.62e-7	7.03e-7	1.38e-7	7.46e-7	1.34e-7	6.88e-7	4.96x
SE(3)	4	✓	5.71e-7	4.69e-5	9.11e-7	5.35e-5	1.12e-6	6.25e-5	65.56x
SE(3)	5	✗	1.96e-7	7.38e-7	1.91e-7	8.73e-7	1.70e-7	9.15e-7	4.57x
SE(3)	5	✓	7.32e-7	9.64e-5	1.27e-6	9.44e-5	1.47e-6	1.14e-4	94.53x
SE(3)	6	✗	2.53e-7	9.11e-7	2.34e-7	1.12e-6	2.23e-7	1.12e-6	4.47x
SE(3)	6	✓	9.29e-7	1.25e-4	1.54e-6	1.67e-4	1.99e-6	1.75e-4	110.31x

**Insight**  
Residual and Jacobian evaluations represent *a major performance-limitation*, rendering 110x speedups momentous

Performance comparison between our auto-generated and optimized B-Spline implementation and the hand-crafted implementation from Sommer et al. [1]



# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)

*Minimization Problem*<sup>[1]</sup>

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \left[ \frac{1}{2} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \|\bar{\mathbf{r}}(t, \theta)\|^2 \right] \text{ with } \theta \subseteq \Theta$$

# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)

**Minimization Problem<sup>[1]</sup>**

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \left[ \frac{1}{2} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \|\bar{\mathbf{r}}(t, \theta)\|^2 \right] \text{ with } \theta \subseteq \Theta$$

**Simplification**

**Stochastic Optimization Problem<sup>[2,3]</sup>**


$$\Theta^* = \underset{\Theta}{\operatorname{argmax}} \log(p(\Theta)) = \underset{\Theta}{\operatorname{argmin}} \sum_i \|\bar{\mathbf{r}}_i(t_i, \theta_i)\|^2$$

*Probability Distribution*


*Residuals*

$$p(\Theta) = \prod_i f_i(t_i, \theta_i) \propto \prod_i \exp(-E_i(t_i, \theta_i))$$

**Assume Gaussian Distributions**



**Continuous-Time Factor Graph (Visual)**



**Stochastic Continuous-Time Factor Graph (Visual)**

[1] Agarwal et al., Ceres Solver

[2] Ortiz et al., CVPR 2020

[3] Ortiz et al., ICRA 2022

# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)

Minimization Problem<sup>[1]</sup>

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \left[ \frac{1}{2} \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \|\bar{\mathbf{r}}(t, \theta)\|^2 \right] \text{ with } \theta \subseteq \Theta \leftarrow$$

*Simplification*

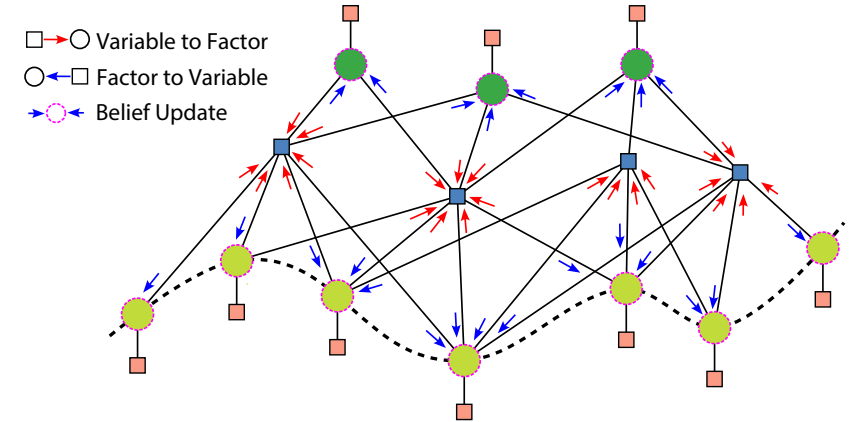
Stochastic Optimization Problem<sup>[2,3]</sup>

$$\Theta^* = \underset{\Theta}{\operatorname{argmax}} \log(p(\Theta)) = \underset{\Theta}{\operatorname{argmin}} \sum_i \|\bar{\mathbf{r}}_i(t_i, \theta_i)\|^2$$

*Probability Distribution*      *Residuals*

$$p(\Theta) = \prod_i f_i(t_i, \theta_i) \propto \prod_i \exp(-E_i(t_i, \theta_i))$$

*Assume Gaussian Distributions*




**Local Message Passing in the Factor Graph**

[1] Agarwal et al., Ceres Solver  
 [2] Ortiz et al., CVPR 2020  
 [3] Ortiz et al., ICRA 2022

# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)

*Node Update*




$$\eta_{n_j} = \sum_{f_i \in N(n_j)} \eta_{f_i \rightarrow n_j}$$

*“What do my neighbors believe about me?”*

# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)


*Node Update*



$$\eta_{n_j} = \sum_{f_i \in N(n_j)} \eta_{f_i \rightarrow n_j}$$

*“What do my neighbors believe about me?”*

**Node-to-Factor Messages**




***“Pass the updated belief to neighboring factors”***

# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)

*Node Update*



$$\eta_{n_j} = \sum_{f_i \in N(n_j)} \eta_{f_i \rightarrow n_j}$$


*“What do my neighbors believe about me?”*

*Node-to-Factor Messages*



*“Pass the updated belief to neighboring factors”*

*Factor Update*




$$\eta_i^0 = -\bar{\mathbf{J}}_i^{0,\top} \bar{\mathbf{r}}_i^0$$

*“Reevaluate the Residual and Jacobian”*

# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)

Node Update



$$\eta_{n_j} = \sum_{f_i \in N(n_j)} \eta_{f_i \rightarrow n_j}$$


“What do my neighbors believe about me?”

Node-to-Factor Messages



“Pass the updated belief to neighboring factors”


Factor Update



$$\eta_i^0 = -\bar{\mathbf{J}}_i^{0,\top} \bar{\mathbf{r}}_i^0$$

“Reevaluate the Residual and Jacobian”

Factor-to-Node Messages



$$\eta'_{f_i \rightarrow n_a} = \eta_a^0 - \mathbf{\Lambda}'_{aa}{}^\top \mathbf{\Lambda}'_{bb}{}^{-1} \eta'_b$$

“Marginalize the probability distribution for a neighboring node”

# Methodology

## Continuous-Time Gaussian Belief Propagation (GBP)

Node Update




$$\eta_{n_j} = \sum \eta_{f_i \rightarrow n_j}$$

“What do n

### Approach


Evaluate the suitability of the **combined GBP-based CT-SLAM method** in absolute and localization setups

Node-to




“Pass the updated belief to neighboring factors”

Factor Update



$$\eta^0 = -\bar{J}_i^0, \bar{r}_i^0$$

“Jacobian”



$$\eta'_{f_i \rightarrow n_a} = \eta_a^0 - \Lambda'_{aa}{}^\top \Lambda'_{bb}{}^{-1} \eta'_b$$

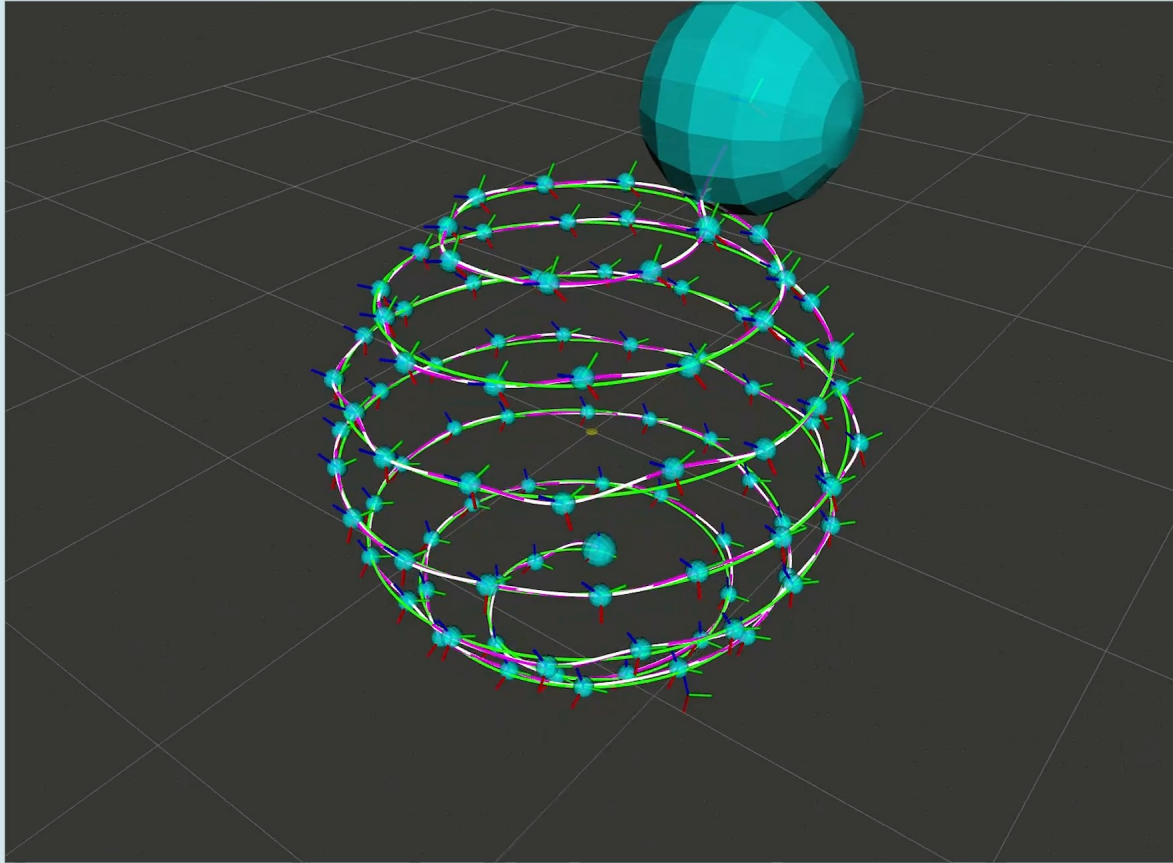
“Marginalize the probability distribution for a neighboring node”



# Experiments


## Motion Capture Setups

### Absolute Setup (Simulation)



In slow motion: The proposed continuous-time GBP solver (in magenta) and the conventional Ceres solver (in white) converge to identical solutions close to the ground truth (in green) even under poor initialization ( $\pm 1.00$  m/rad) and substantial pose measurement noise ( $\pm 0.05$  m/rad).


### Absolute Setup (ChArUco)



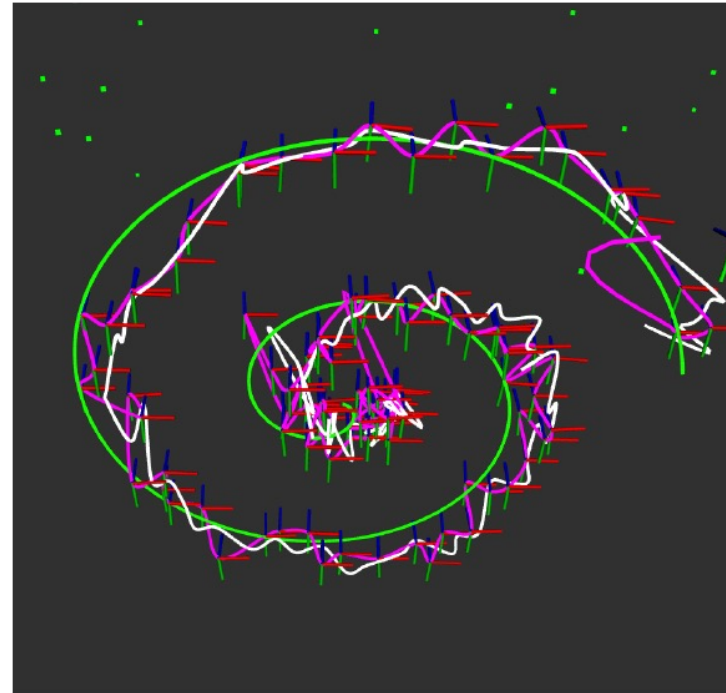
In slow motion: The proposed GBP solver (in magenta) and Ceres (in white) converge to identical solutions close in the ChArUco experiment (with initialization at identity).

# Experiments


## Localization Setup



(a) At initialization



(b) After the 1<sup>st</sup> iteration




(c) Absolute Error


**Left: Estimated motion yielded by the proposed GBP-based framework (in magenta) and Ceres (in white) across iterations and plotted against ground truth (in white). Right: Resulting errors from Hyperion and Ceres are identical.**

# Experiments

## Ablation on Message Dropouts



(a) Dropout Convergence: Absolute




(b) Dropout Convergence: Localization


**Graph energy vs. number of iterations conditioned on the message dropout percentage.**

# Experiments

## Ablation on B- and Z-Splines



**(a)** Convergence: Absolute



**(b)** Convergence: Localization

**Graph energy conditioned on the spline and solver type across iterations.**

# Conclusions

1.) Presents the *first open-source GBP-based continuous-time optimization framework* with *symbolic code generation*

# Conclusions

1.) Presents the *first open-source GBP-based continuous-time optimization framework* with *symbolic code generation*

2.) Implements the *fastest, Ceres-interoperable<sup>[1]</sup> B- and Z-Spline implementations* to date, further alleviating computational limitations

# Conclusions

1.) Presents the *first open-source GBP-based continuous-time optimization framework* with *symbolic code generation*

2.) Implements the *fastest, Ceres-interoperable<sup>[1]</sup> B- and Z-Spline implementations* to date, further alleviating computational limitations

3.) ***Demonstrates the efficacy*** of the proposed framework ***in absolute and localization setups***

# Find us on GitHub!

