Monte Carlo Methods for Comparing Inspection Data

Michael C Byington, Anouk van Pol, John van Pol
Ingu Solutions Inc

Pipeline Pigging and Integrity Management Conference

February 12-16, 2024

Organized by
Clarion Technical Conferences

https://doi.org/10.52202/072781-0007

https://doi.org/10.52202/072781-0007

Abstract

The development of conventional inline inspection technology has been geared towards ultra-high resolution tools requiring a completely cleaned pipeline and coming with increasing costs and data handling time. By comparing data between inspections, unconventional inline inspection tools equipped with off-the-shelf micro-electromechanical magnetometers provide actionable information on the best allocation of resources for digs and high resolutions tools as well as the location of unexpected changes such as illegal hot tap installations. To ensure an accurate comparison between inspections, the data of subsequent inspections must be perfectly aligned. This presentation will discuss a fully automated approach to align inspection data using Monte-Carlo based time warping strategies.

We will illustrate the method with a case study conducted by Pan American Energy. They installed a hot tap at an undisclosed location along a nearly 3,000 meter steel pipeline and asked INGU to locate it. The hot tap was unambiguously identified and Pan American Energy confirmed that the provided location was within the 6 meter acceptance criterion for the project.

Introduction

In 2016, a pipeline failure in North America resulted in the uncontrolled release of 2,000 metric tons of hydrocarbons (both liquid and gas) [1]. Since the 1960s, magnetic flux leakage (MFL) and ultra sound devices have been the primary tools used for pipeline integrity management since the 1980s. However, due to tight bends, non-circular valves, diameter changes or unknown geometry [1,2], approximately 70% of US gas lines built before modern inline inspection (ILI) was a viable technology are considered unpiggable [3], up to 40% of the lines in service as of 2012 [4]. In response to this issue, more recent technologies have been developed since the early 2000s to address unpiggable lines [5,6]. These devices have applications in leak detection, either by acoustic signature [7] or pressure differential [8], pipeline path reconstruction [7,9], and wall condition assessment [10,11].

Free-floating inline inspection tools have the potential to make a large impact in illegal hot tap detection. Estimates from 2012-2017 indicate Nigeria loses 30% of their total hydrocarbon output to theft, around $8 billion, largely through hot taps [12]. The EU loses 4 billion euros in revenue due to hydrocarbon theft worldwide [13] and the worldwide loss is estimated at $133 billion [14]. Any welding or cutting process involving heat or stress should change the remnant magnetic flux of the pipeline which is detectable with passive magnetometers in free-floating devices. To make such a comparison the relative error in the distance scale between the two runs must be smaller than the features in the magnetic flux data.

Free-floating tools take measurements in time and infer a distance scale later rather than directly measuring distance with odometer wheels. This time to distance mapping function is created from observed features (bends, joints, elevation changes, etc) [11]. Recent advances in automation of joint detection from remnant magnetization have lowered the costs and improved the accuracy of time to distance maps for free floating tools [15]. Errors in distance less than 10 m on a 10 km pipeline are now standard. However, this accuracy is still orders of magnitude lower than is useful for a remnant magnetization comparison. To make this comparison feasible, the signals must be aligned with respect to eachother.

https://doi.org/10.52202/072781-0007

Comparing time series data while accounting for shifts in the x axis between data sets has been addressed since the late 1950s [16] and in 1968 dynamic time warping (DTW) method was published [17]. DTW is still an active area of research today [18,19]. The main challenges in the field are physical constraints on the warp mapping and computational cost as a function of sequence length. Naive DTW is prone to chasing noise which, in extreme cases, can result in mapping the majority of one signal onto a single data point in another. Stated rigorously, the warp function $W(x) -> x'$ may have a derivative between 0 and infinity although in most applications, the derivative of the true warp function is bounded ${}^{1}/_{s_{\lim}} < {}^{dW}/_{dx} < s_{\lim}$ where $s_{\lim}$ is less than 5. Global constraints can be applied to prevent these errors from causing large distortions, most commonly the Sakoe-Chiba band [20] and Itakura parallelogram [21], but the small scale distortions can still be problematic. Computation cost scales with $O(n^2)$ where n is the sequence length. In 2018, this was reduced to order $n \log n$ [22] with the fastDTW algorithm but this prevents the application of constraints to reduce non-physical solutions. In short, for sequences on the order of 10^7-9 data points, DTW is an impractical solution for signal alignment where signal exists on multiple length scales.

To overcome challenges in signal alignment, we will introduce a Monte Carlo search based time warping algorithm. This work was motivated by the need to compare remnant magnetometry within pipelines to detect hot taps. It is computationally impractical for many of the modern uses of DTW (scouring time series databases for similarities) but useful when comparing large length time series with signal on multiple length scales if computational cost is inconsequential. It is also easy to add precise physical constraints to the warp path (either global or local, derivative or positional), and it is easily parallelizable.

Defining a Monte Carlo time warp

A warp function is any monotonic increasing function $W(x) = x'$. Given any two time series, $A(x)$ and $B(x)$, and a scalar comparison function $F(A(x), B(x))$, the goal of any time warp algorithm is to find a warp function $W(x)$ that minimizes $F(A(x), B(W(x)))$. Mostly commonly, including in DTW, the comparison function is a mean squared error. Methods of time warping vary based on how the warp function is discovered. This method is a Monte Carlo optimization search.

First, we define our constraints. The derivative of the warp path, ${}^{dW}/_{dx}$, is the ratio of compression or stretching between the two signals. Usually this constraint is symmetric, ${}^{1}/_{s_{\lim}} < {}^{dW}/_{dx} < s_{\lim}$. The second constraint is a total displacement constraint. In most cases, it is unlikely that the first 20% of one signal maps onto the first 80% of the other signal. Therefore, we can impose a limit on the search that restricts matching to within a certain buffer of the original x scale. Expressed mathematically, this is $|W(x) - x| < d_{\lim}$. It is unlikely to be known in advance what these limits are. However, setting reasonable limits greatly reduces the possible search space and improves the efficiency of the Monte Carlo algorithm (provided of course, those limits do not exclude the optimal path from the search space).

To align the data between time series, we begin with a linear discrete approximation of the warp path, ie $W(x) = x$ so there is no modification of the signal. Then we select a point, randomly move it within the constraints, and check with the comparison function if the change improved the warp. If the change improved the warp, we keep the change and repeat the process. If it did not, we discard the change and repeat the process. This is a search through a high dimensionality, highly non-linear space so the Monte Carlo solution is prone to being trapped in local minima. This can be addressed by running the Monte Carlo search independently many times. It is important to note that there is

https://doi.org/10.52202/072781-0007

no guarantee the global minima is approachable in a step-wise fashion. Despite the caveats, the Monte Carlo search solution significantly out performs the dynamic time warping algorithms.

Constraints on W(t)

$$\frac{1}{s_{\text{lim}}} < \frac{dW}{dt} < s_{\text{lim}}$$

$$|W(t) - t| < d_{\text{lim}}$$

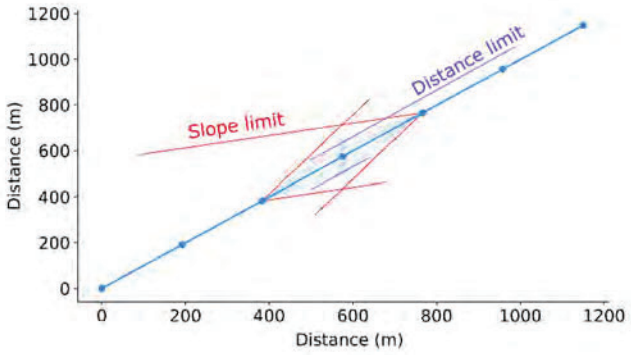

Figure 1. The warp path is subject to global and local constraints. The distance limit (d_{lim}) is a global constraint that reflects the absolute error of the distance between the two time series. The slope limit (s_{lim}) is a local constraint that reflects the relative compression between the two time series. On each iteration of the Monte Carlo, a randomly selected point on the warp path is moved to a random location subject to the warp constraints. The new value for this warp point is found by sampling a uniform probability distribution.

When implementing a search between two time series, there are usually no first principles reasoning that can set these limits. The limits, d_{lim} and s_{lim}, are set on a case by case basis by the user. The smaller they are, the faster the Monte Carlo search will converge unless the desired optimum is excluded from the search space. Practically speaking, the sensitivity of the solution is more dependent on d_{lim} than s_{lim}. We have never run a Monte Carlo where lowering, s_{lim} from 4 to 1.5 produced a noticeable performance improvement and we have never seen a pair of signals where s_lim = 2 excluded the desired solution.

Procedure for single phase Monte Carlo time warp

1) Approximate $W(x)$ as a discrete signal ($W = [w_0, w_1 \ \ldots \ w_n]$).

2) Select the ith point at random.

3) Move w_i to a new position such that the new W satisfies the constraints.

4) If $F(A(x), B(W_{new}(x))) < F(A(x), B(W_{old}(x)))$ keep the change, otherwise discard the change.

5) Repeat steps 2-4 until the loss function levels off.

https://doi.org/10.52202/072781-0007

6) Repeat steps 1-5 200 to 1000 times and select the solution with the smallest loss function value.

Loss functions, signal downsampling, and warp path resolution

The nature of $F(A,B)$ has been general on purpose because any loss function can be used so long as the desired result creates a global minima. The natural loss function is either mean square error or absolute error. However, interpolation of large signals is computationally expensive and yields worse results. Because signal peaks and troughs must align almost exactly (at least to within a feature width) to see a loss function decline, the loss function is extremely jagged and traps the Monte Carlo in local minima easily. To avoid this trap, the data is downsampled by binning, taking the minimum and maximum value of each bin. This significantly reduces the information content of the signal while preserving the most extreme features making the loss function less jagged. Normalizing the loss by the spread of the template function, A, further reduces the propensity of the Monte Carlo to be trapped in local minima.

$$F(A,B) = \frac{(\mathrm{A_{maxbin}} - \mathrm{B_{maxbin}})^2 + (\mathrm{A_{minbin}} - \mathrm{B_{minbin}})^2}{\mathrm{A_{maxbin}} - \mathrm{A_{minbin}}}$$

While this smooths the loss function, it also removes small features that are necessary for fine alignment. However, by running sequential Monte Carlo searches at progressively increasing resolutions, we can align large features first and refine on small features later in the search process. Initially we begin with large bins and a warp path approximated by a few points. Because this search space is relatively small, the Monte Carlo converges quickly. We then increase the resolution of both the warp path and the downsampled signals and repeat the process. In this application, splitting the Monte Carlo search into three steps, increasing the resolution of the warp path and the downsampled signal each time, results in warp paths far more accurate and constrained by physical limitations than those produced by DTW. We refer to this as a multiphase Monte Carlo time warp.

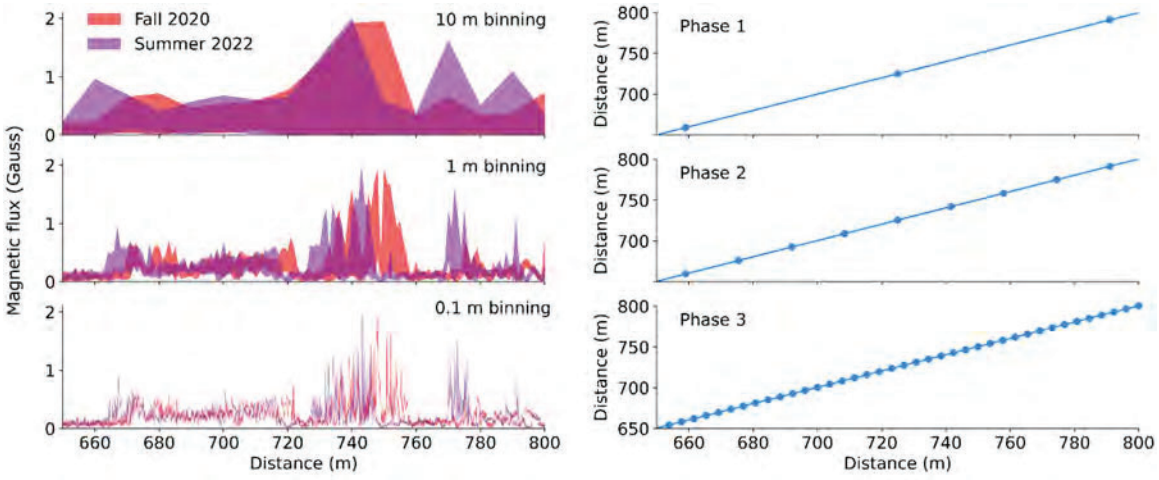

Figure 2. In each phase on the Monte Carlo search, the resolution of the warp path and the downsampled signal increases by a factor of 10. The warp path progression shown here is approximately 5 for illustration purposes (otherwise the dots on the final graph would be indistinguishable).

https://doi.org/10.52202/072781-0007

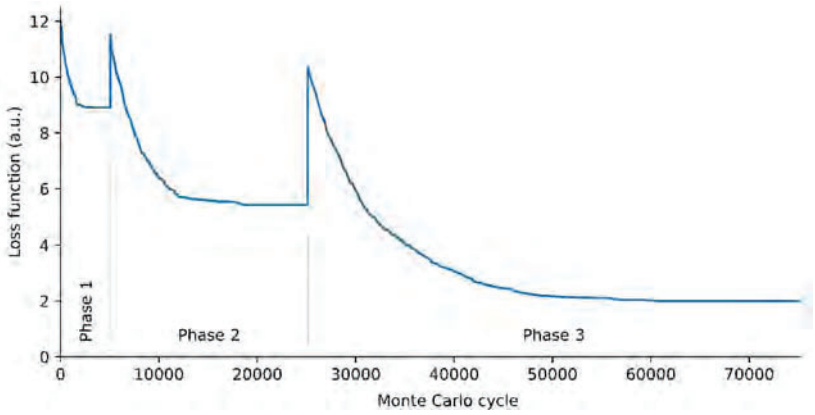

Figure 3. The loss function versus Monte Carlo cycle for the alignment of two time series. Each phase shows an exponential decay because each drop in the loss function lowers the probability that the next Monte Carlo cycle will find an improvement. Each new phase begins with a step change in the loss function because the loss function is now being computed over a larger number of points. This three phase Monte-Carlo search is typical of our alignment process.

The multi-phase strategy can still be trapped by undesirable local minima although is empirically less likely. It is therefore necessary to run the Monte Carlo search hundreds of times independently and select the best result. This gives a similar procedure for the multi-phase process:

Procedure for multi-phase Monte Carlo time warp

1) Approximate $W(x)$ as a discrete signal ($W = [w_0, w_1 \ \ldots \ w_n]$).

2) Select the ith at random.

3) Move w_i to a new position such that the new W satisfies the constraints.

4) If $F(A(x), B(W_{new}(x))) < F(A(x), B(W_{old}(x)))$ keep the change, otherwise discard the change.

5) Repeat steps 2-4 until the loss function levels off.

6) Increase the resolution of $A(x)$, $B(x)$, and $W(x)$ and repeat steps 1-5 for two or three levels of resolution.

7) Repeat steps 1-6 200 to 1000 times and select the solution with the smallest loss function value.

Case study - locating a hot tap in remnant magnetometry

In 2023, Pan American Energy ran a test case to demonstrate remnant magnetometry is an effective and economical solution for ongoing monitoring for illegal hot taps. They introduced a new hot tap in a 3 km 6 inch crude oil pipeline. Tools were deployed twice before and twice after the hot tap installation. A precise and accurate distance scale is generated by utilizing the identified girth welds

https://doi.org/10.52202/072781-0007

and fittings in the magnetic flux data with known information provided by the operator. The signals were then aligned using the Monte Carlo time warp algorithm presented in this paper. They were compared by taking the absolute value after subtracting the 0.1 m binned signals. This is reported below as absolute difference.

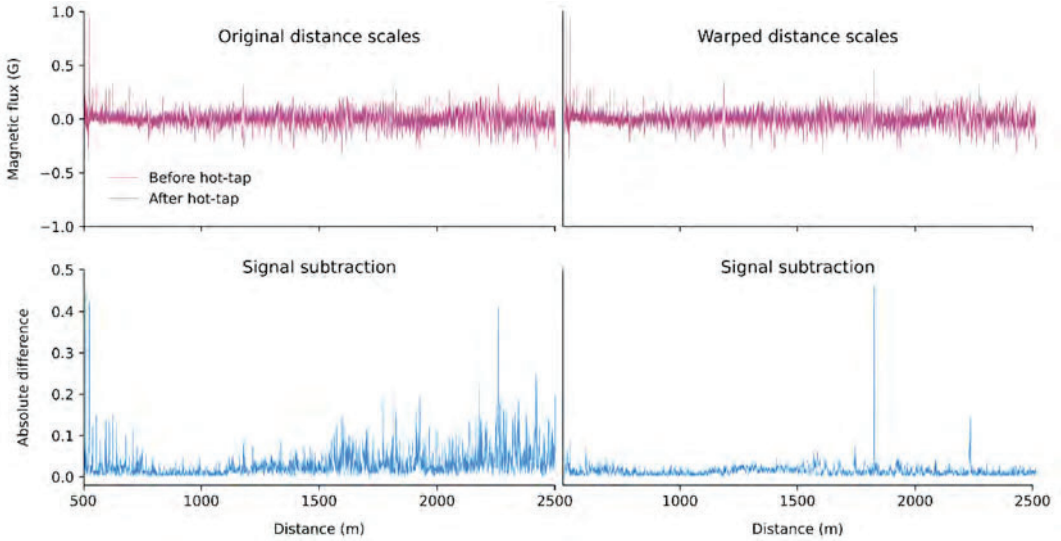

Figure 4. Two remnant magnetometry signals, one before and one after hot tap installation. The distance scales are inferred from bends, valves, and welds identified in the data. On the left, the original distance scales show small misalignments that prevent the identification of changes between runs. On the right, after the alignment of the signal by Monte Carlo time warping, the difference metric shows a pronounced change at the hot tap relative to any other section of the line.

The Monte Carlo time warping was used to align all four data sets. The two before the hot tap showed no trace of a peak whereas the two after the hot tap showed a repeatable signature consisting of one large peak (about 0.4 Gauss) followed by one smaller peak (about 0.15 Gauss). Rapid heating and freezing of the metal during the welding process is most likely responsible for the change in remnant magnetic flux.

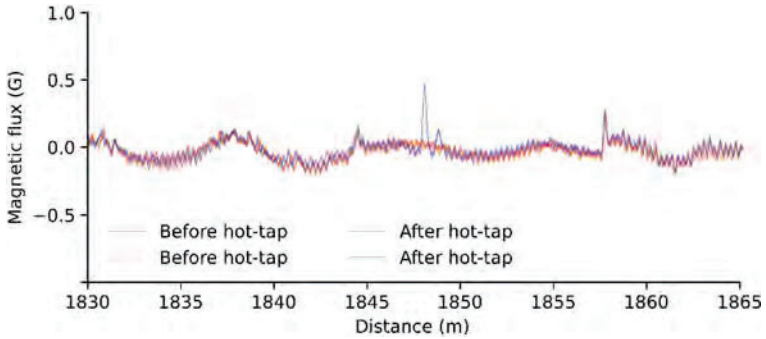

Figure 5. The comparison of four remnant magnetometry measurements, two before and two after the installation of a hot tap. The peak at 1848 m post-hot-tap is repeatable and significantly larger than any other magnetic changes in

https://doi.org/10.52202/072781-0007

the line. This location corresponded to the real location of the hot tap to within 6 meters.

This project focused on the effectiveness of magnetic signatures for the identification and localization of hot taps. This technique opens up the possibility of more frequent monitoring, making it possible to rapidly detect new hot taps within the system. The dynamic nature of this solution makes it tamper proof and highly reliable, minimizing financial losses, and increasing safety and reliability.

Conclusion

Here we presented a Monte Carlo method for the alignment of time-series signals. Accuracy with this solution is a function of computation time and power. We used this method to align remnant magnetometry signals to detect a hot tap during a blind test. The hot tap was located to within 6 m demonstrating the validity of remnant magnetometry comparison, the accuracy the distance reconstruction process, and the repeatability of the IMU measurements.

References

[1] M. Xie, and Z. Tian, "A review on pipeline integrity management utilizing in-line inspection data," Engineering Failure Analysis, vol. 92, pp. 222–239, 2018.

[2] F. Varela, M. Yongjun Tan, and M. Forsyth, "An overview of major methods for inspecting and monitoring external corrosion of on-shore transportation pipelines," Corrosion Engineering, Science and Technology, vol. 50, pp. 226–235, 2015.

[3] Interstate Natural Gas Association of America, "Report to the National Transportation Safety Board on Historical and Future Development of Advanced In-Line Inspection Platforms for Use in Gas Transmission Pipelines," 2012.

[4] J. Tiratsoo, "Ultimate Guide to Unpiggable Pipelines," 2013.

[5] R. Fletcher, and M. Chandrasekaran, "SmartBall: a new approach in pipeline leak detection," International Pipeline Conference, vol. 48586, pp. 117–133, 2008.

[6] J. Smith, A. Van Pol, D. Ham, and J. Van Pol, "Leak detection and prevention using free-floating in-line sensors," Pipeline Pigging and Integrity Management, 2019.

[7] A. Van Pol, J. Van Pol, R. McNealy, and C. Goudy, "The Future of In-Line Inspection: Free-Floating Smart Sensors," International Pipeline Conference, vol. 51869, pp. V001T03A021, 2018.

[8] N. Tajuddin, "Design, Construction and Experimental Analysis of Free Floating Pig for Pipeline Leak Detection," Universiti Teknologi: Thesis, 2015.

[9] M. Volk, S. Miska, and E. Freeman, "Miniature Sensors for Monitoring the Operational Conditions of Pipelines," Offshore Technology Conference, pp. OTC–23186, 2012.

https://doi.org/10.52202/072781-0007

[10] X. Du, X. Fang, W. Tian, P. Tang, and C. Miao, "Design of A Free Floating Ultrasonic Inner Spherical Detector for Pressure Pipeline," Journal of Physics: Conference Series, vol. 1948, pp. 012092, 2021.

[11] M. Kindree, S. Campbell, A. Van Pol, and J. Val Pol, "Defect localization using free-floating unconventional ILI-tools without AGMs," Pipeline Pigging and Integrity Management, 2022.

[12] F. C. Onuoha, E. G. Ekene, and C. Enyiazu, "Unbridled pillage: The political economy of oil theft in Nigeria," SOUTH EAST JOURNAL OF POLITICAL SCIENCE, vol. 1, 2017.

[13] I. M. Ralby, "Downstream Oil Theft," Global Energy Center, 2017.

[14] J. Desjardins, "Fuel theft is a big problem," Business Insider, 2017.

[15] M. C. Byington, A. Val Pol, and J. Val Pol, "Neural networks for pipeline joint detection," Digital Pipeline Solutions Forum, 2022.

[16] R. Bellman, and R. Kalaba, "On adaptive control processes," IRE Transactions On Automatic Control, vol. 4, pp. 1—9, 1959.

[17] T. K. Vintsyuk, "Speech discrimination by dynamic programming," Cybernetics, vol. 4, pp. 52—57, 1968.

[18] S. Aghabozorgi, A. S. Shirkhorshidi, and T. Y. Wah, "Time-series clustering—a decade review," Information Systems, vol. 53, pp. 16—38, 2015.

[19] C. A. Ratanamahatana, and E. Keogh, "Everything you know about dynamic time warping is wrong," Third Workshop On Mining Temporal and Sequential Data, vol. 32, 2004.

[20] H. Sakoe, and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," IEEE Transactions On Acoustics, Speech, and Signal Processing, vol. 26, pp. 43—49, 1978.

[21] F. Itakura, "Minimum prediction residual principle applied to speech recognition," IEEE Transactions On Acoustics, Speech, and Signal Processing, vol. 23, pp. 67—72, 1975.

[22] O. Gold, and M. Sharir, "Dynamic time warping and geometric edit distance: Breaking the quadratic barrier," ACM Transactions On Algorithms (TALG), vol. 14, pp. 1—17, 2018.

https://doi.org/10.52202/072781-0007