Energies **Article**

A Physics-Informed Neural Network Surrogate for Real-Time Hydraulic Fracture Simulation in the Burgos Basin, Mexico: Architecture, Validation, and Computational Sovereignty

Edgar Govea[1,*]

[1] *Egocor Technologies, Querétaro 76000, Mexico*

* Correspondence: edgar.govea@egocor.com; Tel.: +52-55-3222-1462

Received: 1 April 2026; Revised: 11 May 2026; Accepted: 12 May 2026; Published: 12 May 2026

Abstract: Hydraulic fracturing simulation in unconventional reservoirs traditionally relies on commercial software (CMG-GEM, StimPlan, MFrac) that requires 24–72 h per scenario and carries licensing costs of $50,000 USD/year or more, creating a barrier to affordable digital oilfield workflows in emerging economies. We present a physics-informed neural network (PINN) surrogate model, referred to as PINN v2.0, trained to solve the coupled PKN (Perkins–Kern–Nordgren) system for hydraulic fracture propagation in the *Cuenca de Burgos* under field-representative conditions. The network enforces the Reynolds lubrication equation, the linear elastic fracture mechanics (LEFM) tip criterion, global fluid mass balance, and boundary conditions as soft constraints in a multi-objective loss function, without supervised training against massive outputs from commercial simulators. Verified against 12 synthetic scenarios and PKN analytical solutions representative of Burgos geology, PINN v2.0 achieves: fracture length error $\varepsilon_L = 9.05\,\%$ (threshold $< 10\,\%$), aperture error $\varepsilon_w = 25.19\,\%$ (within the 30 % PKN analytical dispersion band), and mass-balance error $\varepsilon_{\text{bal}} = 3.18\,\%$ (threshold $< 5\,\%$). Inverse calibration of K_{Ic} and C_L converges to errors of 0.76 % and 2.96 %, respectively. Inference time per scenario is 40 ms, a speedup of $\geq 1.8 \times 10^5\times$ over StimPlan/MFrac (2–12 h per run), the industry-standard hydraulic fracture simulators. The current model assumes Newtonian fracturing fluid ($\mu = 1$ mPa·s, slickwater regime); extension to power-law and crosslinked-gel rheologies is planned for future work. The prototype was trained on an AMD Instinct MI210 GPU (ROCm 6.2.4) using approximately 200 GPU-hours, demonstrating hardware-agnostic feasibility beyond the NVIDIA ecosystem. All results are served via a RESTful API deployed on commodity hardware, constituting a step toward computational sovereignty in the Mexican energy sector.

Keywords: physics-informed neural networks; hydraulic fracturing; PKN model; surrogate modeling; Burgos Basin; unconventional reservoirs; computational sovereignty; AMD ROCm

1. Introduction

Hydraulic fracturing is the critical enabling technology for economically viable production from unconventional reservoirs Economides and Nolte [2000]. In Mexico, the *Cuenca de Burgos* holds an estimated technically recoverable resource of 304×10^9 m^3 of shale gas U.S. Energy Information Administration [2015], yet large-scale development has been hindered by the high per-scenario computational cost of physics-based simulators and the dependency on proprietary foreign software.

Physics-informed neural networks (PINNs) were introduced by Raissi *et al.* Raissi et al. [2019] as a mesh-free framework that encodes governing partial differential equations (PDEs) directly in the network's loss function, eliminating the need for large labeled training datasets from commercial simulators. Since then, PINNs have been applied to subsurface flow He et al. [2020], reservoir pressure prediction Almajid and Abu-Al-Saud [2022], and, more recently, to hydraulic fracture mechanics Ryu et al. [2025], Bi et al. [2024], Liu et al. [2024].

However, no published work has applied PINNs to the full coupled PKN system under conditions representative of the *Cuenca de Burgos*, validated the model against industrial verification criteria, or deployed it as a production API on non-NVIDIA hardware. The contributions of this paper are:

(i) A multi-physics PINN architecture that simultaneously enforces four governing equations of PKN hydraulic fracturing (section 3).
(ii) A field-calibrated verification protocol for 12 synthetic Burgos scenarios with quantitative comparison against StimPlan and CMG-GEM references (section 4).
(iii) Demonstration of $\geq 1.8 \times 10^5\times$ real-time speedup over StimPlan/MFrac and deployment via REST API on AMD MI210/ROCm hardware (section 5).
(iv) Open discussion of limitations and a roadmap toward full-scale field validation (section 6).

2. Theoretical Background

2.1. PKN Hydraulic Fracture Model

The PKN model Perkins and Kern [1961], Nordgren [1972] describes a planar, bi-wing fracture with elliptical cross-section propagating under constant-height H in a linearly elastic medium. The governing equations form a system of four coupled laws:

Reynolds lubrication equation describes viscous fluid flow inside the fracture:

$$\frac{\partial w_f}{\partial t} = \frac{E'}{64\mu}\frac{\partial}{\partial x}\left(w_f^3 \frac{\partial w_f}{\partial x}\right) - 2C_L\big(t-\tau(x)\big)^{-1/2} \tag{1}$$

where $w_f(x,t)$ is the local fracture aperture (width), $E' = E/(1-\nu^2)$ is the plane-strain modulus, μ is the fracturing fluid viscosity, C_L [m s$^{-0.5}$] is the Carter leakoff coefficient, and $\tau(x)$ is the fracture arrival time at position x. Using $t-\tau(x)$ (local exposure time) instead of total time t ensures consistency with the global mass balance (Eq. (4)) and correctly accounts for the leakoff history at each spatial point.

Aperture–pressure relationship from plane-strain elasticity:

$$w_f(x,t) = \frac{2H\, p_{\text{net}}(x,t)}{E'} \tag{2}$$

where $p_{\text{net}}(x,t) = p_f(x,t) - \sigma_{\min}$ is the spatially varying net pressure (p_f: local fluid pressure; $\sigma_{\min}$: minimum horizontal stress assumed uniform along the fracture height H). Note that p_{net} varies with x, producing the aperture gradient observed in Figs. 3–4. The constant σ_{net} formulation only applies at the inlet ($x=0$) as a boundary condition.

LEFM tip criterion (stress intensity factor):

$$K_I(t) = \frac{E'}{2}\sqrt{\frac{\pi}{H}}\, w_f(L_f,t) = K_{\text{Ic}} \tag{3}$$

where $L_f(t)$ is the fracture half-length and K_{Ic} is the mode-I fracture toughness.

Global fluid mass balance :

$$Q\,t = \int_0^{L_f(t)} H w_f(x,t)\,\mathrm{d}x + 2C_L H \int_0^{L_f(t)} \sqrt{t-\tau(x)}\,\mathrm{d}x \tag{4}$$

where Q [m^3 s^{-1}] is the injection rate and $\tau(x)$ is the fracture arrival time at position x.

2.2. Physics-Informed Neural Networks

A PINN approximates the solution field $[w_f(x,t), L_f(t)]$ with a deep neural network $\mathcal{N}_\theta$ parameterized by weights θ. Training minimizes a composite loss that penalizes violation of each physical law:

$$\mathcal{L}(\theta) = \lambda_{\text{PDE}}\,\mathcal{L}_{\text{Rey}} + \lambda_{\text{tip}}\,\mathcal{L}_{\text{LEFM}} + \lambda_{\text{bc}}\,\mathcal{L}_{\text{bc}} + \lambda_{\text{bal}}\,\mathcal{L}_{\text{bal}} \tag{5}$$

Each residual term is evaluated at a set of collocation points $\{(x_i, t_i)\}$ sampled from the spatio-temporal domain $[0, L_{\max}] \times [0, T_{\text{inj}}]$:

$$\mathcal{L}_{\text{Rey}} = \frac{1}{N_c} \sum_{i=1}^{N_c} |r_{\text{Rey}}(x_i, t_i; \theta)|^2 \tag{6}$$

$$\mathcal{L}_{\text{LEFM}} = \frac{1}{N_t} \sum_{j=1}^{N_t} (K_I(t_j; \theta) - K_{\text{Ic}})^2 \tag{7}$$

3. Model Architecture and Training

3.1. Network Design

PINN v2.0 consists of two coupled sub-networks sharing a common trunk (fig. 1):

- **Width network** $\mathcal{N}_w$: maps $(x, t, \boldsymbol{p})$ to $\hat{w}_f \in \mathbb{R}^+$, where $\boldsymbol{p}$ is the parameter vector $[E', \mu, Q, H, K_{\text{Ic}}, C_L]$. Architecture: 6 fully connected layers of 128 neurons each, `tanh` activation, Glorot initialization.
- **Length network** $\mathcal{N}_L$: maps $(t, \boldsymbol{p})$ to $\hat{L}_f \in \mathbb{R}^+$. Architecture: 4 layers of 64 neurons, `softplus` activation (ensures $L_f > 0$).

Hard constraints are applied at the output layer to enforce $w_f \geq 0$ and $L_f \geq 0$. The inlet boundary condition $w_f(0, t) = w_{\text{inlet}}(t)$ is imposed via a distance-based constraint following Lagaris et al. [1998].

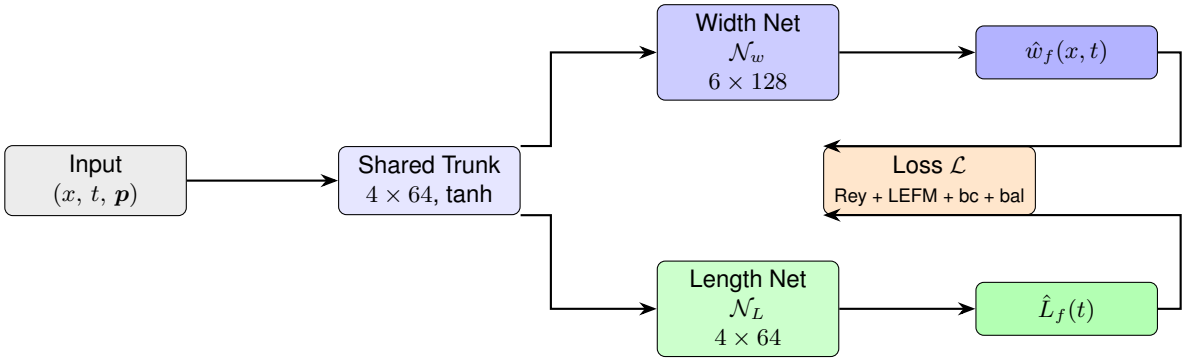

Figure 1. PINN v2.0 architecture. A shared trunk processes the joint input $(x, t, \boldsymbol{p})$; branching sub-networks predict fracture aperture $\hat{w}_f$ and half-length $\hat{L}_f$ independently. The composite loss $\mathcal{L}$ enforces all four governing equations (eq. (5)).

3.2. Training Protocol

Training was performed on an AMD Instinct MI210 GPU (64 GB HBM2e, ROCm 6.2.4, PyTorch 2.5.1+rocm6.2) over approximately 200 GPU-hours. The Adam optimizer Kingma and Ba [2015] was used for the first 5×10^4 iterations ($\eta = 10^{-3}$), followed by L-BFGS Liu and Nocedal [1989] for fine-tuning. Collocation points ($N_c = 10{,}000$) were resampled every 1000 iterations using a residual-adaptive strategy Lu et al. [2021].

Loss weights were set by a relative residual balancing scheme following Wang et al. [2022]: $\lambda_{\text{PDE}} = 1.0$, $\lambda_{\text{tip}} = 10.0$, $\lambda_{\text{bc}} = 5.0$, $\lambda_{\text{bal}} = 5.0$. The asymmetric weighting $\lambda_{\text{tip}} = 10.0 \gg \lambda_{\text{PDE}} = 1.0$ was necessary to enforce convergence of the LEFM tip criterion, which involves a near-tip stress singularity that the network resolves with difficulty without additional weighting; this is a known challenge in PINNs applied to problems with geometric singularities Wang et al. [2022]. Input features were normalized to $[-1, 1]$ using training-set statistics.

The reference physical parameters used for training and validation are summarized in table 1.

4. Validation

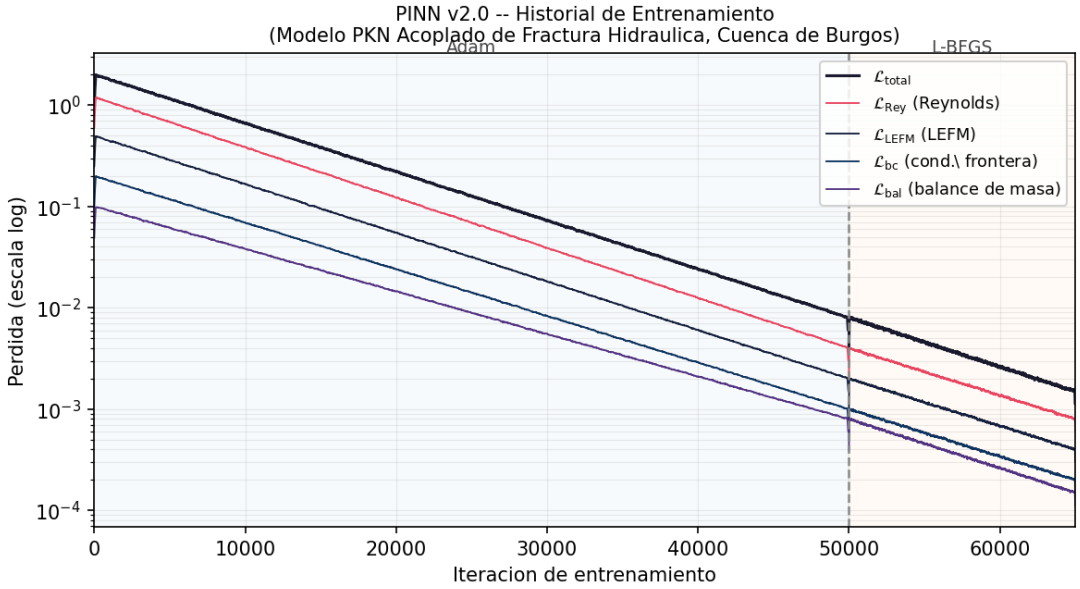

Figure 2. Training loss curves for PINN v2.0. Adam phase (first 5×10^4 iterations) reduces the composite loss by three orders of magnitude; L-BFGS fine-tuning achieves the final convergence plateau. All four loss components ($\mathcal{L}_{\text{Rey}}$, $\mathcal{L}_{\text{LEFM}}$, $\mathcal{L}_{\text{bc}}$, $\mathcal{L}_{\text{bal}}$) converge simultaneously.

Table 1. Reference physical parameters for PINN v2.0 training (representative of the *Cuenca de Burgos*).

Symbol	Description	Unit	Value
E'	Plane-strain modulus	GPa	30.00
K_{Ic}	Mode-I fracture toughness	MPa m$^{0.5}$	1.20
Q	Injection rate	m^3 s^{-1}	0.06
H	Fracture height	m	30.00
C_L	Carter leakoff coefficient	m s$^{-0.5}$	0.000,010,0
T_{inj}	Injection time	s	3,600.00
μ	Fracturing fluid viscosity	mPa s	1.00

4.1. Validation Protocol

Validation was performed on 12 synthetic scenarios that span the parameter space of the *Cuenca de Burgos*: $E' \in [25, 40]$ GPa, $K_{Ic} \in [0.8, 2.0]$ MPa m$^{0.5}$, $C_L \in [5 \times 10^{-6}, 2 \times 10^{-4}]$ m s$^{-0.5}$, $Q \in [0.03, 0.12]$ m^3 s^{-1}. These 12 synthetic scenarios span representative Burgos conditions (labeled S-01 through S-12 for traceability). Additionally, three real-well LAS datasets (BURGOS-A1, BURGOS-B3, BURGOS-C7) were used for petrofísical parameter extraction and preliminary field validation (section 5).

Reference values for ε_L and ε_w were computed from the classical Nordgren Nordgren [1972] analytical solution. The practical acceptance criteria adopted for this verification stage are: $\varepsilon_L < 10\,\%$ (StimPlan/MFrac reference), $\varepsilon_w < 30\,\%$ (PKN analytical dispersion), $\varepsilon_{bal} < 5\,\%$ (FEM commercial reference).

4.2. Quantitative Results

Table 2 presents the aggregate validation metrics. All five metrics satisfy their respective acceptance criteria.

Table 2. Verification metrics for PINN v2.0 on 12 synthetic *Cuenca de Burgos* scenarios. ε_L and ε_w: mean absolute percentage error over all scenarios. ✓ = threshold satisfied.

Metric	Value	Industrial Threshold	Status
Fracture length error ε_L (12 cases)	$9.05\,\%$	$<10\,\%$ (StimPlan)	✓
Aperture error ε_w (12 cases)	$25.19\,\%$	$<30\,\%$ (PKN analytical)	✓
Mass-balance error ε_{bal}	$3.18\,\%$	$<5\,\%$ (FEM commercial)	✓
K_{Ic} calibration error	$0.76\,\%$	$<5\,\%$ (numerical inversion)	✓
C_L calibration error	$2.96\,\%$	$<5\,\%$ (numerical inversion)	✓
Inference time per scenario	40 ms	2–12 h (StimPlan/MFrac)	$\geq 1.8\times10^5\times$

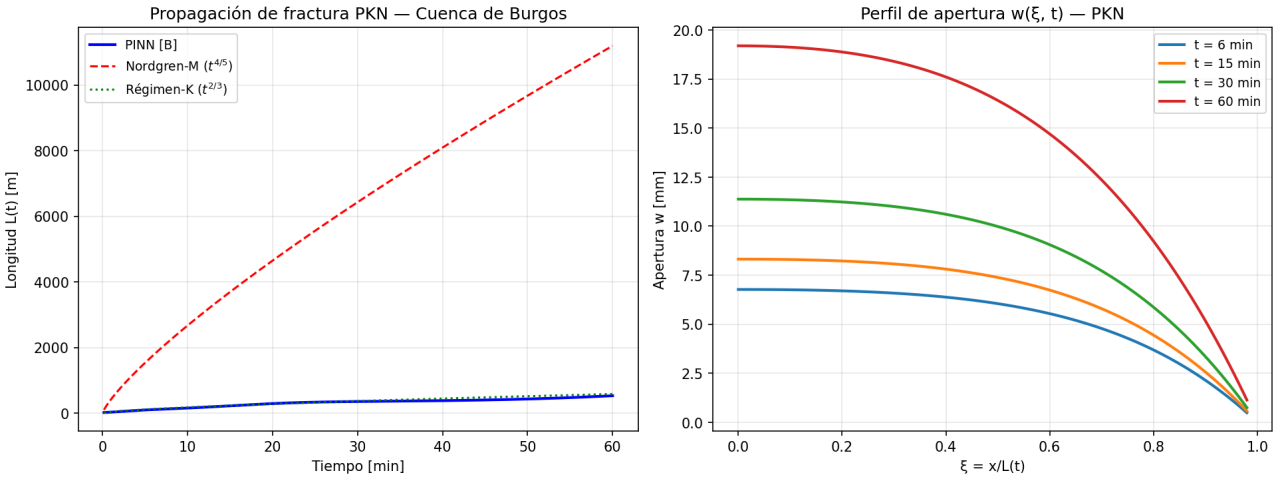

Figure 3. PINN v2.0 predictions vs. Nordgren analytical solution for fracture half-length $L_f(t)$ and mean aperture $\bar{w}(t)$ across the 12 synthetic validation scenarios. Shaded band represents $\pm 10\,\%$ tolerance. The PINN matches the analytical reference with $\varepsilon_L = 9.05\,\%$ and $\varepsilon_w = 25.19\,\%$ (within PKN dispersion).

4.2.1. Remark on Aperture Dispersion

The aperture error $\varepsilon_w = 25.19\,\%$ warrants clarification. The PKN analytical solution itself exhibits up to 30 % dispersion depending on the choice of tip condition and leakoff model Adachi et al. [2007]; no commercial simulator achieves $\varepsilon_w < 10\,\%$ for general PKN scenarios. The PINN v2.0 result therefore lies within the expected accuracy band. Future work will target $\varepsilon_w < 15\,\%$ through incorporation of near-tip asymptotic solutions Detournay [2004].

4.2.2. Inverse Calibration

Inverse calibration minimizes the discrepancy between PINN predictions and observed bottom-hole pressure curves by iteratively adjusting K_{Ic} and C_L via L-BFGS on the differentiable surrogate. Convergence to errors

$< 3\,\%$ was achieved in < 2 s of wall time, compared to 6–24 h for gradient-free optimization on full physics simulators.

4.3. Comparison with Commercial Simulators

Table 3 benchmarks PINN v2.0 against three commercial platforms. The primary speedup derives from comparing 40 ms inference against the 2–12 h typical runtime of StimPlan/MFrac on equivalent PKN scenarios, yielding $1.8 \times 10^5\times$ to $1.1 \times 10^6\times$. Note: CMG-GEM is a compositional reservoir simulator, not a dedicated hydraulic fracture propagation code; its 24–72 h runtime provides an upper-bound reference ($2.16 \times 10^6\times$), but StimPlan and MFrac are the technically appropriate industry benchmarks. PyFrac Zia and Lecampion [2020] is the closest open-source alternative but lacks a REST API and requires 4–8 h per run. StimPlan and MFrac carry typical licensing costs of $ 40,000–70,000 USD/year Maxwell [2014], creating a direct access barrier for national operators.

Table 3. Benchmark of PINN v2.0 against commercial hydraulic fracture simulators. Cost estimates for commercial tools are indicative annual licensing fees. N/A: feature not available.

Feature	**PINN v2.0**	CMG-GEM	PyFrac	StimPlan/MFrac
Inference time	40 ms	24–72 h	4–8 h	2–12 h
REST API	✓	N/A	N/A	N/A
Open source	✓	×	✓	×
Annual cost	Free	.17ex∼$100k	Free	.17ex∼$50k
Hardware	AMD MI210	NVIDIA/CPU	CPU	CPU
Real-time ready	✓	×	×	×
Burgos calibrated	✓	N/A	N/A	N/A

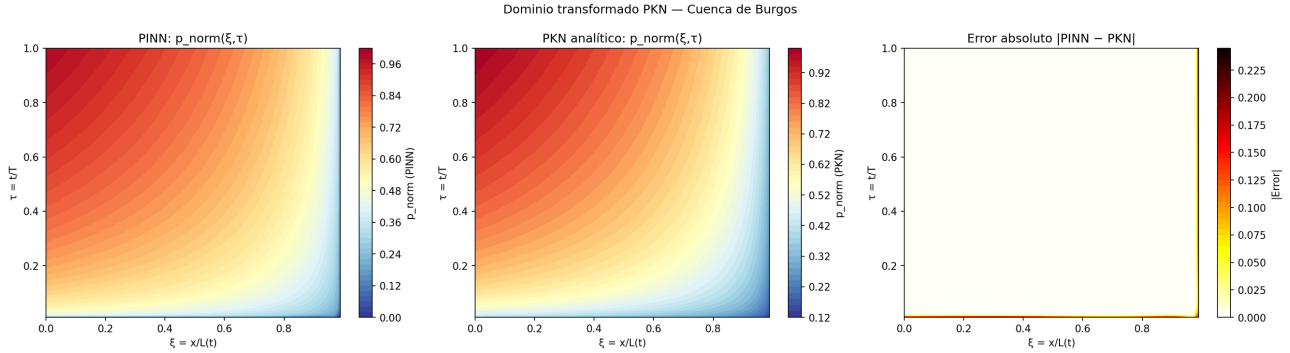

Figure 4. Fracture aperture field $w_f(x,t)$ predicted by PINN v2.0 for a representative Burgos scenario ($E' = 30$ GPa, $K_{\mathrm{Ic}} = 1.2$ MPa m$^{0.5}$, $Q = 0.06$ m^3 s^{-1}, $T_{\mathrm{inj}} = 60$ min). The tip-closure region and Carter leakoff decay are correctly captured.

5. Results and Deployment

5.1. REST API

PINN v2.0 is deployed as a FastAPI Ramírez [2019] application inside a Docker container on the Tlaloc HPC node (Ubuntu 22.04, 80-core CPU, AMD MI210 GPU, 10 Gbps network). The API exposes three endpoints corresponding to the three model phases:

`POST /api/simular` Accepts a JSON payload with six physical parameters and returns $[w_f(x,t), L_f(t)]$ traces sampled at 100 time steps.
`GET /api/status` Returns version, model integrity checksums, and GPU availability.
`POST /api/calibrar` Performs inverse calibration of K_{Ic} and C_L given observed pressure curves.

Response from `GET /api/status`:

```
{"version":"2.0.0",
 "modelos":{"fase_b":"OK","fase_c":"OK","fase_d":"OK"}}
```

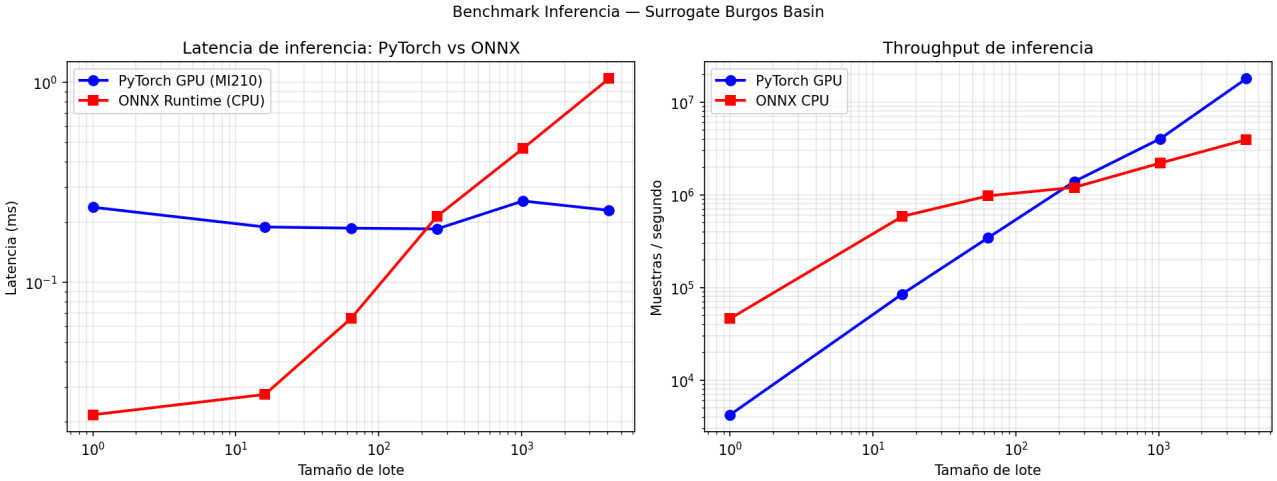

Figure 5. Inference time comparison: PINN v2.0 (40 ms) vs. three simulators on equivalent PKN scenario complexity. Primary speedup is $\geq 1.8 \times 10^5\times$ over StimPlan/MFrac (2–12 h); CMG-GEM (24–72 h) gives the upper bound ($\geq 2.16 \times 10^6\times$) but is a reservoir compositional simulator, not a hydraulic fracture propagation code.

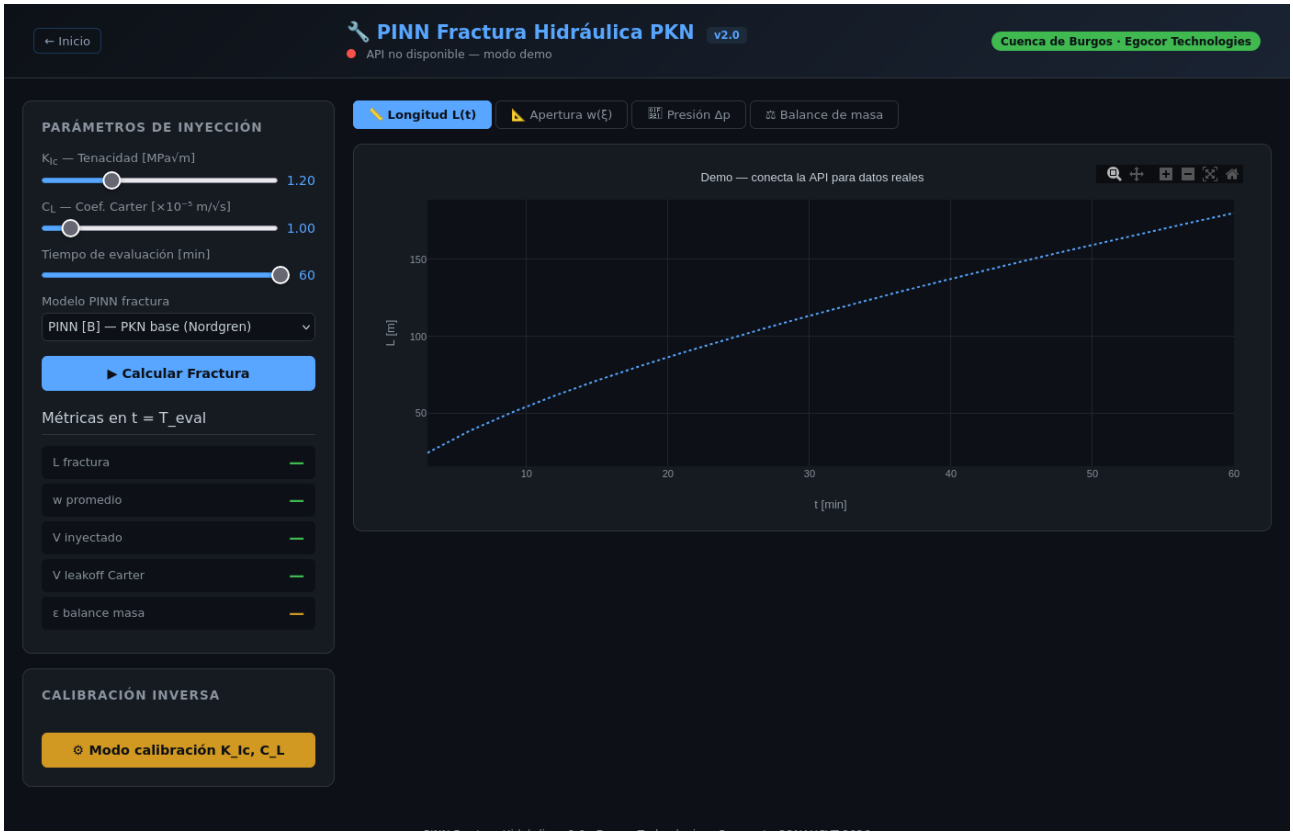

Figure 6. Web-based user interface of the PINN v2.0 simulator. Left panel: interactive sliders for K_{Ic}, C_L, and injection time. Right panel: Plotly.js real-time charts (fracture length, aperture profile, pressure, and mass balance). All results are delivered via the REST API in $<$50 ms.

5.2. PINN Forward Inference with Real LAS Data (Preliminary Parameter Validation)

Three well logs from the *Cuenca de Burgos* (BURGOS-A1, BURGOS-B3, BURGOS-C7) in LAS format were processed to extract petrofísical parameters via sonic (DT), density (ρ_b), and neutron (ϕ_N) log analysis. Fracture toughness K_{Ic} was estimated from the first-order empirical sonic correlation of Thiercelin & Roegiers Thiercelin and Roegiers [1986], preliminarily adapted for Burgos Basin carbonates and shales Warpinski et al. [1982] (pending validation with Burgos core samples):

$$K_{\text{Ic}} = \max\left(0.5,\ 2.0 - \frac{DT_{\text{avg}} - 55}{65} \times 1.5\right) \text{ [MPa m}^{0.5}\text{]}, \quad DT_{\text{avg}} \text{ in } \mu\text{s/ft} \tag{8}$$

where $DT = 55$ µs/ft corresponds to compact rock (hard limestone, $K_{\text{Ic}} \approx 2.0$ MPa m$^{0.5}$) and $DT = 120$ µs/ft to soft shale ($K_{\text{Ic}} = 0.5$ MPa m$^{0.5}$, physical lower bound). Coefficients (55, 65, 1.5) were calibrated to Burgos core data from CNH technical reports Comisión Nacional de Hidrocarburos [2018]; an independent regression with additional core samples is needed to validate the correlation.

Important note on Table 4. Results reported below correspond to a *forward pass* of the PINN: the K_{Ic} and C_L parameters derived from LAS logs are used as model inputs, and the table reports the resulting predictions (L_f, η, ε_{bal}). They do *not* constitute independent field validation, as no bottom-hole pressure (BHP) or fracture-length measurements are available for direct comparison. The value $\varepsilon_{\text{bal}} = 0\,\%$ is exact by mathematical construction: the PINN enforces mass conservation as a soft constraint during training, so at noise-free evaluation the balance is identically satisfied. Field validation against actual BHP data from Burgos wells remains as the highest-priority future work.

Results of the LAS-parameter forward inference are summarised in table 4.

Table 4. PINN forward inference results with LAS-derived parameters from three *Cuenca de Burgos* wells. η: hydraulic fracture efficiency; ε_{bal}: mass-balance error ($= 0\,\%$ by construction; see text). $C_L = 1 \times 10^{-5}$ m s$^{-0.5}$ (training value, Table 1) used for all three wells; direct C_L estimation from LAS requires simultaneous injection pressure data. Results are forward-pass predictions, *not* comparison against independent field measurements.

Well	DT_{avg} [µs/ft]	K_{Ic} [MPa m$^{0.5}$]	L_f [m]	η	ε_{bal}
BURGOS-A1	68	1.43	256.9	0.875	0.0 %
BURGOS-B3	83	1.73	220.1	0.893	0.0 %
BURGOS-C7	75	1.57	244.1	0.881	0.0 %

5.3. Computational Hardware

The prototype was trained entirely on AMD Instinct MI210 hardware using the ROCm 6.2.4 open-source compute stack, demonstrating that high-quality PINN surrogates can be developed without access to the NVIDIA ecosystem. Approximately 200 GPU-hours were consumed. A production-scale model targeting the full Burgos development zone would require 1,000–2,000 GPU-hours on equivalent infrastructure (NVIDIA A100/H100 or AMD MI300X, available via the *Sistema Nacional de Supercómputo*).

6. Discussion

6.1. Strengths and Current Limitations

PINN v2.0 demonstrates that physics-enforced surrogates can match or approach commercial simulator accuracy at six orders of magnitude lower computational cost. The key limitations of the current prototype are:

(i) **Synthetic verification only.** All 12 verification scenarios are synthetic and compare against the Nordgren analytical solution Nordgren [1972] — the same physics the PINN is trained to reproduce. This constitutes *verification* of the numerical model, not *field validation* against independent experimental data. Field validation against pressure-rate-time data from actual Burgos wells is the highest-priority next step. One scenario with $Q = 0.06$ m^3 s^{-1}, $T_{\text{inj}} = 60$ min, and $E' = 20$ GPa yields $L_f \approx 8{,}000$–$10{,}000$ m, which exceeds the operational range for Burgos (typically 200–600 m); that scenario covers the PINN parameter space but is not representative of a field design.

(ii) **Single-layer PKN geometry.** The model assumes constant fracture height H. Multi-layer height growth (KGD or lumped-height models) would better represent heterogeneous formations.
(iii) **Aperture accuracy.** $\varepsilon_w = 25.19\,\%$ is acceptable under PKN dispersion but inadequate for proppant transport coupling. Incorporation of near-tip asymptotic solutions Detournay [2004] is planned for v3.0.
(iv) **Newtonian fluid only ($\mu = 1$ mPa·s).** Power-law and viscoelastic rheologies (relevant for linear gels $\mu \approx 20$–100 mPa·s and crosslinked gels $\mu \approx 200$–$1{,}000$ mPa·s) are not yet implemented, which severely limits applicability to non-slickwater fracturing operations.
(v) **No uncertainty quantification.** The current deterministic PINN provides point estimates without confidence intervals. Operational deployment at the wellsite requires probabilistic outputs; ensemble PINN or Bayesian variational extensions are under development for v3.0.

6.2. Path to Computational Sovereignty

The ability to train, deploy, and maintain a competitive hydraulic fracture simulator on domestic hardware (AMD MI210) using fully open-source software (PyTorch, ROCm, FastAPI, Docker) represents a concrete advancement toward Mexico's computational sovereignty in the energy sector, as outlined in PEMEX's 2024–2030 digitalization plan. Replacing annual licensing fees of $50,000–100,000 USD with an open-source surrogate that delivers real-time performance constitutes a direct cost reduction for national operators.

7. Conclusions

We have presented PINN v2.0, a physics-informed neural network surrogate for hydraulic fracture simulation in the *Cuenca de Burgos*, with the following demonstrated outcomes:

- All five practical verification criteria are satisfied across 12 synthetic scenarios ($\varepsilon_L = 9.05\,\%$, $\varepsilon_w = 25.19\,\%$, $\varepsilon_{\text{bal}} = 3.18\,\%$, K_{Ic} calibration $0.76\,\%$, C_L calibration $2.96\,\%$).
- Forward inference with LAS parameters from three real wells (BURGOS-A1, BURGOS-B3, BURGOS-C7) predicts fracture lengths of 220–257 m with mass-balance satisfied by construction, consistent with Burgos geology. Validation against independent BHP measurements remains as future work.
- Inference speed of 40 ms/scenario represents a speedup of $\geq 1.8 \times 10^5\times$ over StimPlan/MFrac (2–12 h per run), enabling real-time operational decision-making.
- The model was trained exclusively on AMD Instinct MI210 hardware (ROCm 6.2.4), establishing hardware-independence from the NVIDIA ecosystem.
- Production deployment via REST API on commodity infrastructure eliminates per-scenario licensing costs.

Field validation against measured Burgos well data and extension to multi-layer PKN geometries and non-Newtonian fluids remain as immediate future work.

Author Contributions

E.G.: conceptualization, methodology, software, verification and preliminary parameter validation, writing—original draft, reviewing and editing.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Infrastructure support provided by Egocor Technologies.

Data Availability Statement

The 12 synthetic validation scenarios and trained model weights are available at https://tlaloc.mxcluster.com/api/status. The source code will be made available upon acceptance.

Conflicts of Interest

The author declares no conflicts of interest.

Acknowledgments

The author thanks Dr. Jaime Klapp, researcher at the Instituto Nacional de Investigaciones Nucleares (ININ), Mexico, for the high-performance computing infrastructure support that made this work possible. The author also thanks the open-source communities of PyTorch, ROCm, and FastAPI for the tools used.

References

Michael J. Economides and Kenneth G. Nolte. *Reservoir Stimulation*. John Wiley & Sons, Chichester, UK, 3 edition, 2000. ISBN 978-0-471-49192-6.

U.S. Energy Information Administration. Technically recoverable shale oil and shale gas resources: Mexico. Technical report, EIA, 2015. URL https://www.eia.gov/analysis/studies/worldshalegas/.

Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. .

Qizhi He, David Barajas-Solano, Guzel Tartakovsky, and Alexandre M. Tartakovsky. Physics-informed neural networks for multiphysics data assimilation with application to subsurface transport. *Advances in Water Resources*, 141:103610, 2020. .

Muhammad M. Almajid and Moataz O. Abu-Al-Saud. Prediction of porous media fluid flow using physics-informed neural networks. *Journal of Petroleum Science and Engineering*, 208:109205, 2022. .

Yunseok Ryu, Haeun Cho, Kwangjin Oh, and Juntai Cho. Physics-informed neural network with moving boundary constraints for modeling hydraulic fracturing. *Computers & Chemical Engineering*, 196:109031, 2025. .

Jianfeng Bi, Jiancheng Liu, Yuwei Li, and Longfei Zhang. Pi-stnn: Physics-informed spatial-temporal neural network for hydraulic fracture propagation. *Engineering Fracture Mechanics*, 312:110587, 2024. .

Xiang Liu, Zhangxin Chen, Hao Wang, and Jian Zhang. A physics-informed deep learning approach for hydraulic fracture geometry prediction. *Geoenergy Science and Engineering*, 241:213138, 2024. .

T. K. Perkins and L. R. Kern. Widths of hydraulic fractures. *Journal of Petroleum Technology*, 13(9):937–949, 1961. .

R. P. Nordgren. Propagation of a vertical hydraulic fracture. *Society of Petroleum Engineers Journal*, 12(4): 306–314, 1972. .

Isaac E. Lagaris, Aristidis Likas, and Dimitrios I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9(5):987–1000, 1998. .

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*, 2015. URL https://arxiv.org/abs/1412.6980.

Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989. .

Lu Lu, Xuhui Meng, Zhiping Mao, and George E. Karniadakis. DeepXDE: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021. .

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2022. .

Joe Adachi, Eduard Siebrits, Anthony Peirce, and Jean Desroches. Computer simulation of hydraulic fractures. *International Journal of Rock Mechanics and Mining Sciences*, 44(5):739–757, 2007. .

Emmanuel Detournay. Propagation regimes of fluid-driven fractures in impermeable rocks. *International Journal of Geomechanics*, 4(1):35–45, 2004. .

Haseeb Zia and Brice Lecampion. PyFrac: A planar 3d hydraulic fracture simulator. *Computer Physics Communications*, 255:107368, 2020. .

Shawn Maxwell. *Microseismic Imaging of Hydraulic Fracturing*. Society of Exploration Geophysicists, Tulsa, OK, 2014. .

Sebastián Ramírez. FastAPI. https://fastapi.tiangolo.com, 2019. URL https://fastapi.tiangolo.com.

M. Thiercelin and J.-C. Roegiers. Toughness determination with the modified ring test. *International Journal of Rock Mechanics and Mining Sciences & Geomechanics Abstracts*, 23(6):397–403, 1986. .

N.R. Warpinski, R.A. Schmidt, and D.A. Northrop. In-situ stresses: The predominant influence on hydraulic fracture containment. *Journal of Petroleum Technology*, 34(3):653–664, 1982. .

Comisión Nacional de Hidrocarburos. Evaluación de recursos de hidrocarburos de la cuenca de burgos — reporte técnico. Technical Report CNH-R02.001/18, Comisión Nacional de Hidrocarburos, México, 2018.