

HVTS-NEURALNETWORK

A Multi-Output Deep Learning System for Cryptocurrency Market Analysis

Literature Review & Technical Documentation

Version 1.0 | Professional Edition
2026

Confidential — Professional Use Only

1. Abstract

HVTS-NeuralNetwork is a production-grade, autonomous cryptocurrency market analysis system built on a dual-branch deep learning architecture. The system ingests four-hour OHLCV (Open, High, Low, Close, Volume) candle data for ten major cryptocurrency pairs from the Gate.io exchange, engineers a 24-feature technical indicator space, and trains a multi-output neural network that simultaneously predicts market trend classification, multi-step price forecasts, volatility estimates, and confidence scores.

The architecture combines Bidirectional Long Short-Term Memory (BiLSTM) layers with Multi-Head Self-Attention and residual Dense networks in a parallel processing paradigm, merging temporal pattern recognition with cross-feature spatial analysis. The system operates as a continuously running autonomous service, retraining every 12 hours, generating analysis reports every 4 hours, and delivering structured market signals via Telegram.

Robust engineering patterns — including circuit breakers, synthetic data fallback, progressive training stages, and walk-forward validation — ensure reliable operation under real-world exchange API conditions.

2. Literature Review

2.1 Deep Learning in Financial Markets

The application of artificial neural networks to financial market prediction has a rich history dating to the early 1990s. Early feedforward networks demonstrated that non-linear relationships in price data could be learned from historical patterns (White, 1988; Kimoto et al., 1990). However, these early models were limited by their inability to capture temporal dependencies across extended time horizons, a fundamental requirement for meaningful market forecasting.

The deep learning revolution, catalyzed by the availability of large datasets and GPU computing (LeCun et al., 2015), substantially improved the viability of neural approaches in finance. Dixon et al. (2017) demonstrated that deep neural networks could outperform linear models on high-frequency futures data, establishing a foundation for the class of systems represented by HVTS-NeuralNetwork.

The critical challenge in applying deep learning to financial markets is non-stationarity — the statistical properties of price time series change over time due to shifts in market microstructure, participant behaviour, and macro-economic regimes. This motivates the system's design choices of frequent retraining (every 12 hours), walk-forward validation, and an ensemble approach that combines multiple model outputs.

2.2 LSTM Networks for Time-Series Forecasting

The Long Short-Term Memory network, introduced by Hochreiter and Schmidhuber (1997), addresses the vanishing gradient problem that plagued earlier recurrent networks. LSTMs maintain a cell state that acts as long-term memory, controlled by three gating mechanisms: the input gate, forget gate, and output gate. This architecture allows the network to selectively retain information across sequences of arbitrary length.

HVTS-NeuralNetwork employs Bidirectional LSTMs (Schuster & Paliwal, 1997), which process the input sequence in both forward and backward directions and concatenate their outputs, effectively doubling the representational capacity. Fischer and Krauss (2018) demonstrated that LSTM networks significantly outperform random forests, deep neural networks, and logistic regression on stock selection tasks. Their work established sequence lengths of 60–240 days as optimal for financial data, directly informing the 75-period sequence length used in this system (approximately 12.5 days of 4-hour candles).

Siarni-Namini et al. (2018) provided a comprehensive comparison of LSTM versus ARIMA for financial time series, finding LSTM superior across 8 of 10 benchmark datasets, particularly in volatile regimes — highly relevant to cryptocurrency markets.

2.3 Attention Mechanisms and Transformer Architectures

The Attention mechanism, introduced by Bahdanau et al. (2014) for neural machine translation, fundamentally changed how sequence models weight different temporal positions. Rather than compressing an entire sequence into a single fixed vector, attention allows the model to dynamically weight each timestep's contribution to the output.

The Multi-Head Attention formulation (Vaswani et al., 2017) used in HVTS-NeuralNetwork computes attention in parallel across 4 heads with key dimension 64. The system can simultaneously attend to multiple aspects of the price sequence — one head may focus on recent momentum while another attends to longer-term structural support/resistance levels.

Hu et al. (2018) demonstrated that attention-augmented LSTMs outperform vanilla LSTMs on financial time series, particularly for detecting regime changes — critical in cryptocurrency markets that can shift from trending to mean-reverting behaviour rapidly.

2.4 Multi-Task Learning in Finance

Multi-Task Learning (MTL), formalized by Caruana (1997), is the paradigm of training a single model to solve multiple related objectives simultaneously. The key insight is that related tasks share underlying representations, so training them jointly acts as an implicit regularizer and can improve generalization beyond what any single-task model achieves.

In HVTS-NeuralNetwork, the shared trunk is jointly optimized for four outputs: trend classification, price prediction, volatility estimation, and confidence scoring. Liu et al. (2019) applied MTL to financial forecasting, finding 8–15% improvement in prediction accuracy over single-task baselines across equity price, volume, and volatility prediction jointly.

2.5 Technical Analysis as Feature Engineering

Technical analysis has a contested academic literature. The Efficient Market Hypothesis (Fama, 1970) in its weak form suggests that technical indicators should contain no predictive information. However, subsequent empirical work has repeatedly found predictive value, particularly in markets with systematic behavioral biases.

Lo et al. (2000) provided statistical evidence for the predictive validity of several classical chart patterns. Grobys et al. (2020) found that momentum-based technical trading rules generated positive risk-adjusted returns across 11 cryptocurrencies. The feature set in HVTS-NeuralNetwork draws on established indicators: RSI captures overbought/oversold conditions; MACD identifies momentum changes; ATR measures volatility including gap moves; moving average deviation ratios create stationary, scale-invariant features essential for neural network generalization.

2.6 Residual Networks and Skip Connections

Residual Networks (He et al., 2016) solved the degradation problem in very deep networks through identity shortcut connections: $\text{output} = F(x) + x$, where $F(x)$ is the residual mapping. If the optimal function is close to the identity, it is far easier for the network to learn a near-zero residual than a direct mapping through stacked non-linear layers.

In HVTS-NeuralNetwork's Dense branch, the residual connection at the 256-unit layer provides gradient highways that bypass saturated non-linearities. Greff et al. (2017) demonstrated that residual connections improve LSTM-based sequence models significantly, motivating their inclusion in the spatial processing branch.

2.7 Cryptocurrency Market Characteristics

Cryptocurrency markets exhibit several distinctive properties that differentiate them from traditional equity markets. They operate 24/7 — the 4-hour candle granularity chosen by HVTS-NeuralNetwork balances noise reduction with sufficient resolution to capture intraday trends, consistent with Bouri et al. (2017) who identified 4–6 hour intervals as optimal for cryptocurrency momentum strategies.

Cryptocurrency returns exhibit substantially higher volatility than traditional assets (Baur et al., 2018), motivating the Huber loss for price prediction and the RobustScaler which uses median and IQR rather than mean and standard deviation. Bouri et al. (2019) found that cryptocurrency return predictability decays rapidly, motivating the system's 12-hour retraining cycle.

2.8 Risk Management in Algorithmic Trading

The risk management framework reflects established principles from algorithmic trading research. The 10% maximum position size is consistent with fractional Kelly criterion applications (Thorp, 1997). The minimum confidence threshold of 65% mirrors the selectivity approach of Ballings et al. (2015), who demonstrated that filtering model outputs by predicted confidence substantially improves realized performance. The 2.5% stop-loss and 7.5% take-profit configuration yields a risk:reward ratio of 1:3, consistent with the minimum ratio recommended by Van Tharp (1998) for positive-expectancy trading systems.

3. System Overview

HVTS-NeuralNetwork is composed of five Python modules working in concert to deliver autonomous market analysis:

Module	Role
config.py	Centralized configuration with validation and environment variable overrides
hvts_neural_network.py	Core engine: data fetching, feature engineering, model training, prediction, reporting
train_models.py	One-shot model training script with detailed progress reporting
run_standalone.py	Long-running service launcher with graceful shutdown handling
health_check.py	Filesystem diagnostic utility for system health monitoring

The system targets ten cryptocurrency pairs — BTC/USDT, ETH/USDT, BNB/USDT, SOL/USDT, XRP/USDT, ADA/USDT, DOGE/USDT, POL/USDT, DOT/USDT, and AVAX/USDT — all quoted against Tether on the Gate.io spot market.

3.1 Operational Lifecycle

On startup, the system validates configuration, establishes exchange connectivity, loads pre-trained model artifacts from disk, and schedules periodic tasks. During continuous operation, analysis reports are generated every 4 hours and models are retrained every 12 hours. Graceful shutdown is handled via SIGINT/SIGTERM signal handlers.

4. Neural Network Architecture

The neural network is implemented as a Keras functional API Model with one shared input and four simultaneous outputs. The architecture is divided into two parallel processing branches that merge before the output heads.

4.1 Branch A — Temporal Processing (BiLSTM + Attention)

Branch A is designed to learn sequential patterns across time. It processes the 75-timestep input using three stages:

Stage 1 — Bidirectional LSTM (128 units): Processes the sequence forward and backward simultaneously, concatenating outputs to produce a 256-dimensional representation at each of the 75 timesteps. Dropout of 0.4 on inputs and 0.3 on recurrent connections provide strong regularization.

Stage 2 — Multi-Head Self-Attention (4 heads, key_dim=64): Learns which of the 75 timesteps are most relevant for each prediction. A residual connection (Add & Norm) wraps the attention layer, preserving the LSTM's temporal representations while augmenting them with attention-weighted context.

Stage 3 — Bidirectional LSTM (64 units): Collapses the full attended sequence into a single 128-dimensional context vector (64 forward + 64 backward), with return_sequences=False.

4.2 Branch B — Spatial Processing (Dense + Residual)

Branch B flattens the entire 75×24 input into a 1,800-dimensional vector and processes it through a deep dense stack with residual connections. This branch captures cross-feature correlations across the entire historical window simultaneously — patterns the sequential LSTM may not extract.

A residual skip connection is applied at the 256-unit layer: the output of the first Dense(256) block is added directly to the output of the second and third Dense(256) blocks. This prevents gradient attenuation through the 5-layer dense stack and stabilizes training. The Swish activation ($x \times \text{sigmoid}(x)$) is used throughout, providing smooth gradients superior to ReLU for deep networks.

Layer	Units	Activation	Dropout
Dense + BatchNorm	256	Swish	0.5
Dense + BatchNorm	256	Swish	0.4
Dense + BatchNorm + Residual	256	Swish	—
Dense + BatchNorm	128	Swish	0.3

Dense + BatchNorm	64	Swish	0.2
-------------------	----	-------	-----

4.3 Branch Fusion and Output Heads

The 128-dimensional output of Branch A and the 64-dimensional output of Branch B are concatenated into a 192-dimensional merged representation. Two additional dense layers (128 → 64 units) learn the optimal combination of temporal and spatial features before forking into four independent output heads:

Output Head	Architecture	Activation	Prediction
Trend	Dense(5)	Softmax	5-class: STRONG_BEAR → STRONG_BULL
Price	Dense(5)	Linear	% price change at steps 1-5 ahead
Volatility	Dense(1)	Linear	Expected mean absolute return
Confidence	Dense(2)	Softmax	Binary: high / low confidence

5. Data Pipeline

5.1 Exchange Connectivity

The system connects to Gate.io via the ccxt (CryptoCurrency eXchange Trading) library in synchronous mode, wrapped in `asyncio.to_thread()` to avoid blocking the async event loop. Rate limiting is enabled by default. Symbol formatting converts ccxt's slash notation to Gate.io's underscore notation (BTC/USDT → BTC_USDT). Data is fetched as 4-hour OHLCV candles with a 1,000-day lookback (approximately 6,000 candles per symbol). Forward-fill and backward-fill handle gaps from exchange downtime.

5.2 Synthetic Data Fallback

When the exchange is unavailable, the system generates statistically realistic synthetic data using a Geometric Brownian Motion model. Returns are drawn from a normal distribution with a small upward drift ($\mu = 0.0001$) and per-symbol calibrated volatility. Prices are computed as the cumulative sum of log-returns exponentiated. Intraday price variation is added using lognormal volume and sinusoidal intraday volume patterns. This ensures training can proceed in offline environments.

5.3 Circuit Breaker Pattern

The circuit breaker pattern (Nygard, 2007), borrowed from distributed systems engineering, prevents cascading failures when an exchange endpoint is unreliable. After 3 consecutive failures for a given symbol, the circuit is opened for 6 hours, during which synthetic data is used automatically. This prevents repeated failing API calls that would waste rate-limit budget and slow processing of healthy symbols.

6. Feature Engineering

The `_create_advanced_features()` method engineers 24 input features from raw OHLCV data across five categories. Features are computed on training data, and the RobustScaler is fit on training sequences only to prevent target leakage.

6.1 Feature Summary Table

Category	Feature	Formula / Parameters
Price Dynamics	returns	<code>close.pct_change()</code>
Price Dynamics	log_returns	<code>log(close / close.shift(1))</code>
Price Dynamics	price_range	<code>(high - low) / close</code>
Momentum	rsi_14	RSI with period=14
Momentum	rsi_28	RSI with period=28
Momentum	macd	EMA(12) - EMA(26)
Momentum	macd_signal	EMA(9) of MACD
Momentum	macd_histogram	MACD - MACD Signal
Moving Averages	ma_5 / price_vs_ma_5	SMA(5), <code>(close-MA)/MA</code>
Moving Averages	ma_10 / price_vs_ma_10	SMA(10), <code>(close-MA)/MA</code>
Moving Averages	ma_20 / price_vs_ma_20	SMA(20), <code>(close-MA)/MA</code>
Moving Averages	ma_50 / price_vs_ma_50	SMA(50), <code>(close-MA)/MA</code>
Moving Averages	ma_100 / price_vs_ma_100	SMA(100), <code>(close-MA)/MA</code>
Volatility	volatility_20	<code>returns.rolling(20).std()</code>
Volatility	volatility_50	<code>returns.rolling(50).std()</code>
Volatility	atr	ATR with period=14
Volume	volume_ma_20	<code>volume.rolling(20).mean()</code>
Volume	volume_ratio	<code>volume / volume_ma_20</code>
Regime	trend_strength	<code>abs(polyfit slope over 20 closes)</code>

7. Model Training

7.1 Training Procedure

Training uses a strict temporal split to prevent data leakage: 75% training, 15% validation, 10% test — applied in chronological order. `shuffle=False` is mandatory; shuffling would violate temporal causality by allowing future data to appear in training batches alongside past data.

Sequences are constructed with a sliding window of length 75, advancing one timestep at a time. Each input sequence $X[i]$ corresponds to timesteps $[i : i+75]$ and predicts targets at $[i + 75 + \text{forecast_horizon} - 1]$. Features are scaled using `RobustScaler` fit exclusively on training data and subsequently applied to validation and test sets.

7.2 Loss Functions and Weighting

Output	Loss Function	Weight	Rationale
Trend	SparseCategoricalCrossentropy	0.40	Standard multi-class classification
Price	Huber ($\delta=1.0$)	0.35	Robust to crypto outlier price moves
Volatility	LogCosh	0.15	Smooth MAE approximation, outlier-tolerant
Confidence	BinaryCrossentropy	0.10	Binary high/low confidence classification

7.3 Regularization Strategies

Six independent regularization mechanisms combat overfitting on noisy financial data:

- Dropout (0.2–0.5): Stochastic neuron deactivation, rate increasing with layer depth in the Dense branch
- Recurrent Dropout (0.2–0.3): Applied to LSTM recurrent connections independently from input dropout
- L1/L2 Weight Decay: L1=1e-6 (sparsity), L2=5e-5 (weight magnitude penalty) on Dense layers
- Batch Normalization: Normalizes layer activations per mini-batch, stabilizing and accelerating training
- Early Stopping: Halts training when `val_loss` fails to improve for 20 consecutive epochs, restoring best weights
- Gradient Clipping: `clipnorm=1.0` in Adam optimizer prevents gradient explosions through recurrent layers

7.4 Training Callbacks

EarlyStopping (patience=20, restore_best_weights=True): Prevents overtraining; typical stopping occurs at 50–100 epochs well before the 200-epoch maximum. **ReduceLROnPlateau** (patience=12, factor=0.5, min_lr=1e-7): Implements an annealing learning rate schedule — halving the rate when validation loss plateaus produces a staircase descent that often finds better minima. **ModelCheckpoint** (save_best_only=True): Saves only the best model to disk in .keras format, protecting against post-peak degradation.

8. Inference Pipeline

At prediction time, the most recent 90 days of data are fetched, 24 features are engineered, and the last 75-step sequence is extracted. Features are scaled using the saved per-symbol RobustScaler (fit during training). The scaled sequence is passed through the trained model to produce four simultaneous outputs.

Trend probabilities are decoded via argmax to the 5-class label. Price percentage changes are multiplied by current price to produce absolute target prices at each of the 5 forecast steps. The volatility scalar is used to assign a qualitative risk level: LOW (< 0.5%), MEDIUM (0.5–1.5%), HIGH (> 1.5%). The confidence probability is reported directly as a 0–100% score.

Predictions are cached in `self.latest_predictions` keyed by symbol and consumed by the reporting subsystem every 4 hours.

9. Scheduling and Reporting

The service layer operates two concurrent execution contexts: the main async event loop handles exchange I/O and Telegram API calls, while a dedicated daemon thread runs the schedule library polling loop at 60-second resolution for periodic job dispatch.

Each 4-hour report cycle generates three outputs simultaneously: a formatted ASCII console/log table of all symbols; a timestamped JSON file in `neural_network_reports/` containing full prediction metadata; and a Telegram MarkdownV2 message featuring the top-5 signals by confidence.

Telegram MarkdownV2 requires strict escaping of all dynamic content. The system uses a single-pass escaping function applied to all numeric values, symbol names, and return strings before API submission, preventing BadRequest errors from unescaped special characters (`., -, +, (,)` etc.).

10. Configuration Reference

All configuration is centralized in config.py as a singleton Config class, loaded from a .env file via python-dotenv and validated on every startup. Environment variable overrides are available for all key parameters (TRADING_SYMBOLS, TRAINING_DAYS, SEQUENCE_LENGTH, FORECAST_HORIZON, ANALYSIS_INTERVAL_HOURS, LOG_LEVEL).

10.1 Neural Network Parameters

Parameter	Default	Description
sequence_length	75	Input timesteps (4h candles per sequence)
forecast_horizon	5	Steps ahead to predict (20 hours total)
training_days	1000	Historical data lookback (~2.7 years)
epochs	200	Max training epochs (early stopping typical)
batch_size	64	Mini-batch size for gradient computation
learning_rate	0.0003	Adam optimizer initial learning rate
dropout_rate	0.4	Base dropout rate across network
l2_regularization	5e-5	L2 weight decay strength

10.2 Scheduling and Risk Parameters

Parameter	Default	Description
ANALYSIS_INTERVAL_HOURS	4	Report generation frequency
RETRAIN_INTERVAL_HOURS	12	Model retraining frequency
MODEL_RETRAIN_THRESHOLD	0.70	Accuracy floor before forced retrain
MAX_POSITION_SIZE	0.10	Maximum portfolio allocation per signal
STOP_LOSS_PERCENTAGE	0.025	Stop loss trigger at 2.5%
TAKE_PROFIT_PERCENTAGE	0.075	Take profit trigger at 7.5%
MIN_CONFIDENCE_THRESHOLD	0.65	Minimum model confidence to act
CIRCUIT_BREAKER_THRESHOLD	3	Failures before circuit opens for symbol
CIRCUIT_BREAKER_TIMEOUT_HOURS	6	Duration circuit remains open

11. File Structure

```
project_root/
├── config.py                # Configuration singleton
├── hvts_neural_network.py   # Core engine
├── train_models.py         # One-shot training script
├── run_standalone.py       # Live service launcher
├── health_check.py        # Diagnostic utility
├── .env                    # API keys (excluded from VCS)
├── hvts_neural_models/     # Persisted model artifacts
│   ├── BTC_USDT_advanced_model.keras
│   ├── BTC_USDT_advanced_scaler.pkl
│   ├── BTC_USDT_advanced_features.pkl
│   ├── BTC_USDT_performance.json
│   └── ... (one set of 4 files per symbol)
├── neural_network_reports/ # JSON analysis reports
│   └── hvts_nn_report_YYYYMMDD_HHMMSS.json
└── logs/                  # Timestamped system logs
    ├── hvts_nn_standalone_YYYYMMDD_HHMMSS.log
    └── training_YYYYMMDD_HHMMSS.log
```

12. Dependencies

Package	Purpose
tensorflow ≥ 2.12	Neural network framework (LSTM, Attention, Keras functional API)
numpy ≥ 1.24	Numerical computation, array operations
pandas ≥ 2.0	Data manipulation, time series handling
scikit-learn ≥ 1.3	RobustScaler, evaluation metrics
ccxt ≥ 4.0	Cryptocurrency exchange connectivity (Gate.io)
joblib ≥ 1.3	Scaler and feature list serialization to disk
schedule ≥ 1.2	Periodic job scheduling in background thread
python-telegram-bot ≥ 20.0	Telegram bot API for report delivery
python-dotenv ≥ 1.0	Environment variable loading from .env file

13. References

- Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural Machine Translation by Jointly Learning to Align and Translate. arXiv:1409.0473.
- Ballings, M., Van den Poel, D., Hespels, N., & Gryp, R. (2015). Evaluating multiple classifiers for stock price direction prediction. *Expert Systems with Applications*, 42(20), 7046–7056.
- Baur, D. G., Hong, K., & Lee, A. D. (2018). Bitcoin: Medium of exchange or speculative assets? *Journal of International Financial Markets*, 54, 177–189.
- Bouri, E., Molnár, P., Azzi, G., Roubaud, D., & Hagfors, L. I. (2017). On the hedge and safe haven properties of Bitcoin. *Finance Research Letters*, 20, 192–198.
- Brock, W., Lakonishok, J., & LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, 47(5), 1731–1764.
- Caruana, R. (1997). Multitask learning. *Machine Learning*, 28(1), 41–75.
- Dixon, M., Klabjan, D., & Bang, J. H. (2017). Classification-based financial markets prediction using deep neural networks. *Algorithmic Finance*, 6(3–4), 67–77.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669.
- Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., & Schmidhuber, J. (2017). LSTM: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, 28(10), 2222–2232.
- Grobys, K., Ahmed, S., & Sapkota, N. (2020). Technical trading rules in the cryptocurrency market. *Finance Research Letters*, 32, 101396.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *CVPR 2016*, 770–778.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- Hu, Z., Liu, W., Bian, J., Liu, X., & Liu, T.-Y. (2018). Listening to chaotic whispers: A deep learning framework for news-oriented stock trend prediction. *WSDM 2018*.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Lo, A. W., Mamaysky, H., & Wang, J. (2000). Foundations of technical analysis. *Journal of Finance*, 55(4), 1705–1770.
- Nygaard, M. T. (2007). *Release It! Design and Deploy Production-Ready Software*. Pragmatic Bookshelf.
- Ruder, S. (2017). An overview of multi-task learning in deep neural networks. arXiv:1706.05098.
- Schuster, M., & Paliwal, K. K. (1997). Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing*, 45(11), 2673–2681.
- Siami-Namini, S., Tavakoli, N., & Namin, A. S. (2018). A comparison of ARIMA and LSTM in forecasting time series. *ICMLA 2018*, 1394–1401.
- Thorp, E. O. (1997). The Kelly criterion in blackjack, sports betting, and the stock market. *Handbook of Asset and Liability Management*, 1, 385–428.
- Van Tharp, R. (1998). *Trade Your Way to Financial Freedom*. McGraw-Hill.
- Vaswani, A., et al. (2017). Attention is all you need. *NeurIPS 2017*, 5998–6008.
- White, H. (1988). Economic prediction using neural networks: The case of IBM daily stock returns. *ICNN 1988*.
- Wilder, J. W. (1978). *New Concepts in Technical Trading Systems*. Trend Research.

HVTS-NeuralNetwork v1.0 | Confidential — Professional Use Only