

An Interpretable Framework for Startup Success Assessment Using Entrepreneurial Ecosystem Intelligence

Fardin Sarrafnia¹, Ahmad Movahedian Attar² and Shahabaldin Gheysari³

¹M2 Control Systems & IT, University of Grenoble Alpes, Grenoble, France

²Beedie School of Business, Simon Fraser University, Vancouver, Canada

³Middle East Paidar Kesht Hoosh Company, Esfahan, Iran

Fardin.sarrafnia@etu.univ-grenoble-alpes.fr

Fardinsarrafnia98@gmail.com

ahmad_movahedian_attar@sfu.ca

shahabgheysari@gmail.com

Abstract: Entrepreneurship Support Organisations (ESOs) need metrics to evaluate ventures within entrepreneurial ecosystems, but assessing startup readiness remains difficult. Current practice relies mainly on descriptive indicators that are hard to measure and to reproduce, and little is documented about how these indicators are operationalised or how well they perform. The Open Innovation Framework for Pitch Competitions (OIF-PC) (Movahedian Attar et al., 2025) takes a systems-thinking view in which entrepreneurial support is treated as a funnel whose design parameters shape venture development. Its constructs, however, are defined at programme level and have not yet been expressed as measurable attributes. Expressing them at venture level would give programmes a more useful basis for evaluation. This study therefore operationalises the OIF-PC orientation as measurable attributes of startups and examines whether these attributes are associated with observed outcomes. A secondary dataset of 923 records describing 922 unique ventures, with 45 raw attributes covering funding, milestones, networks, and the ecosystem, was analysed using Logistic Regression, Random Forest and XGBoost. SHAP (SHapley Additive exPlanations) was used to interpret model outputs and to quantify the contribution of original and derived variables to the assessed readiness score. Access to resources is associated with more favourable outcomes. How those resources are used, and how actively a venture engages with its entrepreneurial environment, account for a comparable share of the model's attribution, with network connectivity the single most influential attribute. This indicates that ecosystem constructs of the OIF-PC type can be operationalised through venture-level variables, but it is not a validation of the framework itself. On this basis, the study introduces the Startup Success Readiness Index (SSRI), which combines four pillars (Relational Strength, Milestone Achievement, Investor Diversity and Funding Efficiency) into a weighted score using SHAP-derived attribution. The SSRI gives ESOs a transparent readiness measure that they can use to profile ventures, find capability gaps and inform programme design.

Keywords: Startup Success Assessment; Entrepreneurship Support Organisations (ESOs); Open Innovation Framework for Pitch Competitions (OIF-PC); Machine Learning; Startup Success Readiness Index (SSRI); SHAP

1. Introduction

Entrepreneurial ventures contribute to innovation, job creation and economic growth. A large share of them close early in their development, and many of the remainder never reach a sustainable scale. Support programmes have grown in number in response, with the aim of developing and scaling promising ventures.

These programmes offer entrepreneurs mentorship, funding opportunities, professional networks and strategic guidance. Identifying which ventures have strong growth potential is harder. Selection decisions are commonly based on expert opinion, scattered indicators or short-term measures that may not reflect longer-term outcomes well, leaving support organisations without a clear account of which venture characteristics matter most.

The role of entrepreneurial ecosystems has attracted increasing attention in this context. Startup performance is associated with the internal characteristics of the firm and also with the interaction between entrepreneurs, investors, mentors, institutions and support organisations. These relational aspects, however, are rarely measured at the level of the individual venture.

This study therefore asks whether venture characteristics can be measured from observable data and used to assess how closely a venture resembles those that were ultimately acquired rather than closed. It is framed as assessment rather than prediction: the dataset is cross-sectional and the outcome is already observed, so all findings are reported as associations. The study makes two contributions. First, using interpretable machine learning, it operationalises ecosystem constructs of the OIF-PC type as measurable attributes and identifies which of them are associated with observed outcomes. Second, it introduces the Startup Success Readiness

Index (SSRI), a transparent composite score derived from SHAP attribution that allows ESOs to profile ventures on four pillars and to locate capability gaps.

2. Literature Review

2.1 Startup Success and Entrepreneurial Ecosystems

Research on new venture success initially concentrated on success factors, including entrepreneur traits and the financial, market and technical competences of new ventures (Song et al., 2008; Giardino et al., 2014). Evidence at this level is uneven, however. Song et al. (2008) find that only a subset of the frequently proposed firm-level factors is consistently associated with outcomes, which suggests that firm-level attributes alone leave much of the variation unexplained. More recent research has therefore shifted towards the ecosystem, examining the role of networks, institutions, entrepreneurs, incubators, policy-makers and the surrounding society in a new venture's development (Isenberg, 2010; Stam, 2015; Spigel, 2017).

A related stream has examined incubators, accelerators and pitch competitions as mechanisms for developing entrepreneurs' skills, networks and access to resources (Cohen, 2013; Pauwels et al., 2016; Hallen et al., 2020). The findings of this stream are mixed. Interventions differ considerably in design and selection practice, and the benefit associated with participation varies from programme to programme. This variation points to evaluation approaches that focus on the underlying design parameters rather than on programme participation as a single indicator.

This stream also treats open innovation, which depends on external interaction and collaboration between firms, as one of the mechanisms associated with entrepreneurial success (Chesbrough et al., 2006; West and Bogers, 2014). Despite this conceptual progress, the ecosystem perspective has rarely been translated into measurable indicators that can support explanatory analysis and the design of interventions. This gap is the starting point for the present study.

2.2 Support Programmes and Evaluation Challenges

Support services have grown quickly over the last decade and are widely used to stimulate innovation and to help new businesses develop (Cohen, 2013; Pauwels et al., 2016). Their effectiveness is nonetheless difficult to demonstrate (Hallen et al., 2020). Evaluations commonly rely on indicators such as participation rates, budgets achieved, or short-term results. These describe programme activity, but they do not give a clear picture of how support mechanisms relate to longer-term business development.

Many frameworks identify observable outcomes but provide little guidance on which ecosystem conditions or support mechanisms would need to be improved to obtain better outcomes. Evaluation therefore tends to describe what a programme did rather than indicate what its managers should change.

2.3 OIF-PC as a Structured Ecosystem Evaluation Framework

Recent work has proposed more structured ways of accounting for support activities within innovation ecosystems. One of these is the Open Innovation Framework for Pitch Competitions (OIF-PC) (Movahedian Attar et al., 2025), which sets out how support mechanisms relate to business development.

OIF-PC treats pitch competitions as one piece of the larger ecosystem rather than as isolated events. It adopts a systems-thinking perspective in which the entrepreneurial path is modelled as a funnel: from an initial idea, through mentorship and incubation, to a pitch competition and an operating company. Its design parameters include mentorship duration, mentor and team diversity, and financial support, and its output is a set of successful or discontinued ventures. On this basis, the framework proposes three indicators: an Innovation Alignment Index (IAI), a Funnel Diversity Score (FDS), and a Post-Competition Funding Index (PCFI). The three indicators are intended to reflect the quality of support activities and their relation to business progress.

Because these indicators are defined at the level of a funnel, they cannot be computed directly for an individual venture observed outside such a programme. In this study, OIF-PC is therefore used as a conceptual orientation only: it motivates which families of venture attributes are worth measuring, and the analysis does not test the framework itself.

2.4 Machine Learning for Evaluation

Machine learning is widely used in data analysis because of its ability to capture patterns, non-linear relationships, and interactions among variables (Kuhn and Johnson, 2013). Most studies applying these methods in an entrepreneurial context focus on individual firms rather than on the broader ecosystem. A further limitation is interpretability (Doshi-Velez and Kim, 2017). Such models can reach high accuracy while giving little insight into which conditions are associated with the outcome. Post-hoc explanation methods, which attribute a model's output to its individual inputs, have been developed in response, SHAP among them (Ribeiro et al., 2016; Lundberg and Lee, 2017; Molnar, 2022).

3. Methodology

3.1 Data Description

This empirical study is based on a secondary dataset of 923 records describing 922 unique ventures (one exact duplicate was removed) and 45 attributes representing different dimensions of entrepreneurship. It covers financial attributes (venture capacity and investment activity), temporal attributes (growth trajectory and funding timelines), industry and geographical attributes, and network metrics for ecosystem connectivity; the Appendix lists all attributes with short descriptions. The target variable is derived from the recorded venture status, which takes exactly two values here: acquired ($n = 596$, 64.6%) and closed ($n = 326$, 35.4%), coded as successful (1) and unsuccessful (0) respectively. No operational category is present in the data. The classes are moderately imbalanced, so evaluation measures that are not distorted by class proportion are required.

The sample contains only ventures whose trajectory has resolved into acquisition or closure; ventures still operating at the time of collection are absent by construction. The observed 64.6% acquisition share is therefore a property of the sampling frame, not a base rate for startups in general. The analysis concerns which of the two resolutions occurred, given that the venture resolved at all, and the framework is positioned as a readiness assessment instrument rather than a prediction model.

Missingness is concentrated in a small number of attributes. Missing closure information is fully determined by the outcome, and this attribute is therefore excluded. Missing milestone information (18.3% and 17.2%) is more common among younger ventures, which probably reflects their stage of development and reporting behaviour, and is handled as described in Section 3.2.

3.2 Feature Engineering and Preprocessing

The modelling feature set was derived from the 45 attributes of the working file under a selection rule fixed before any model was fitted. Attributes that encode the outcome were removed: labels, an exact copy of the recorded status, and closed_at, whose presence identifies closed ventures. Also removed were identifiers, high-cardinality text fields already represented by indicator variables, one duplicated column, and one external ranking whose measurement date could not be established. Two derived variables were constructed: the natural logarithm of total funding, which reduces a strong right skew, and a binary indicator of whether any milestone was recorded, which is also the missing-data indicator for the milestone dates. No composite ratio or index was created at this stage, because each candidate composite was a deterministic function of attributes already retained; composition into the four readiness dimensions is performed on model contributions rather than on raw attributes (Section 4.4). This yields 28 features: 16 assigned to the four readiness dimensions and 12 sector and location controls (Appendix A).

Preprocessing was implemented as a single pipeline fitted on the training partition only. The four attributes with missing values were median-imputed because the median is less sensitive than the mean to the long right tails of the time-to-event variables. The nine continuous and count features were standardised, which the coefficient-based model requires and which leaves the tree ensembles unchanged. The 19 binary indicators passed through untransformed. Class rebalancing was tested rather than assumed to be necessary. SMOTE (Chawla et al., 2002), SMOTE-NC and class weighting were each compared against the untreated training set in repeated cross-validation. None improved ranking performance, so no resampling was applied.

3.3 Temporal Scope, Leakage Control and Interpretive Limits

One limitation of the data is that several features, especially those related to funding and milestones, represent information that becomes available at later stages of the venture life cycle. These variables therefore describe

a venture's observed trajectory to date rather than its characteristics at founding. The framework is scoped accordingly: it assesses ventures on what has already been observed, and attributes generated by the resolution event itself are excluded from the analysis.

Leakage was controlled at two levels. The risk is that information from the held-out ventures reaches the model and inflates the reported performance. Attributes that encode the outcome were removed before modelling (Section 3.2), and every fitted transformation was contained in a pipeline estimated on the training partition alone. No quantity computed from the 185 held-out ventures, whether an imputation median or a scaling constant, was available to the model at any stage.

The results are interpreted as statistical associations rather than causal evidence: they indicate which venture characteristics correspond to observed outcomes, not that those characteristics produce them. Establishing causal claims would require a longitudinal design in which venture attributes are recorded before the outcome is realised.

3.4 Model Development

Three algorithms were compared: Logistic Regression as an interpretable linear baseline, and Random Forest (Breiman, 2001) and XGBoost (Chen and Guestrin, 2016) for their ability to represent non-linear relationships and interactions between features. A majority-class classifier was included as a null reference, so that the reported figures can be judged against what is obtainable without any model.

The 922 ventures were split once, using a fixed random seed, into 737 training and 185 held-out test ventures, stratified on the outcome so that both partitions carry the same acquired/closed proportion. Hyperparameters were tuned by exhaustive grid search (105 configurations across the three learners), evaluated by 5-fold stratified cross-validation on the training partition only; the test partition was not used at any point during model development. To verify that the selected configuration is not an artefact of tuning and reporting on the same folds, the entire tuning procedure was re-run inside a nested cross-validation, which placed the resulting optimism at 0.002–0.004 ROC-AUC. The searched grids and the selected settings are reported in the Appendix.

3.5 Evaluation Metrics

Model performance was assessed using Accuracy, Precision, Recall, F1 score, ROC-AUC and PR-AUC. Accuracy is sensitive to class imbalance and is therefore reported only against the majority-class baseline. Precision is the proportion of ventures assessed as successful that were in fact acquired, and so reflects false positives; Recall is the proportion of acquired ventures that were identified, and so reflects false negatives. The F1 score balances the two.

Because the framework orders ventures by assessed readiness rather than classifying them at a fixed cut-off, ROC-AUC was specified in advance as the primary metric: it evaluates ranking across all decision thresholds and is unaffected by the class proportion (Fawcett, 2006). PR-AUC is reported alongside it because it concentrates on the positive class and disregards true negatives (Saito and Rehmsmeier, 2015); as acquired ventures are the more frequent class here, PR-AUC must be read against the 0.646 baseline of an uninformative ranking. Threshold-dependent measures are reported at a stated operating point (Section 4.2).

4. Results

4.1 Bivariate Associations

Point-biserial correlations between each attribute and the outcome are consistently positive but modest. Network connectivity shows the strongest association (relationships, $r = 0.36$, $p < 0.001$), followed by milestone count ($r = 0.33$), total funding on the logarithmic scale ($r = 0.30$), age at last milestone ($r = 0.26$), number of funding rounds ($r = 0.21$) and average investors per round ($r = 0.19$). No coefficient reaches $|0.40|$, indicating that no single attribute accounts for the outcome on its own.

Two features of these associations shaped the modelling. First, the rank correlation exceeds the linear one for several attributes, most clearly for relationships ($p = 0.47$ against $r = 0.36$). This indicates a monotone but non-linear relationship, and it is evidence, obtained before any model was fitted, that non-linear learners should perform better than a linear baseline. Second, the temporal attributes are correlated with each other. The strongest pair is age at last funding and age at first milestone ($r = 0.72$), and the maximum variance inflation

factor is 8.0. Multicollinearity of this magnitude destabilises coefficient-based interpretation in the Logistic Regression baseline but leaves the tree ensembles, which split on one attribute at a time, unaffected.

4.2 Model Performance

Table 1 reports performance on the 185 held-out ventures. Both tree ensembles outperform the linear baseline across all measures, and all three models exceed the majority-class reference. XGBoost attains the highest ROC-AUC (0.858) and PR-AUC (0.912) and is adopted as the reference model for interpretation, since the framework ranks ventures by readiness.

The margin over Random Forest is small. In repeated cross-validation on the training partition, the two ensembles differ by 0.006 ROC-AUC, which does not reach conventional significance (paired t -test, $p = 0.060$) and is smaller than the fold-to-fold variation of either model. XGBoost is adopted because it leads on both pre-specified metrics, but the findings that follow do not depend on that choice: the attribution ranking is reproduced under Random Forest (Section 4.3). The figures in Table 1 correspond to the default 0.5 cut-off; because the operating point is a policy decision rather than a property of the model, it is reported explicitly and can be shifted, at a known cost, to favour recall on closed ventures.

Table 1: Test-set performance on 185 held-out ventures (positive class = acquired)

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC	PR-AUC
Majority-class baseline	0.649	0.649	1.000	0.787	0.500	0.649
Logistic Regression	0.714	0.782	0.775	0.778	0.797	0.861
Random Forest	0.789	0.800	0.900	0.847	0.838	0.892
XGBoost (adopted)	0.778	0.811	0.858	0.834	0.858	0.912

4.3 Attribution Analysis

For each venture, SHAP assigns to every attribute the amount by which that attribute moved the model's output away from the average output. The contributions for a venture therefore sum exactly to its assessed readiness. Averaging the absolute contributions over the training ventures gives each attribute's share of total attribution mass.

Attribution is concentrated. Relationships alone carry 24.8% of the total, followed by age at last milestone (12.7%), total funding (12.6%), milestone count (12.4%) and average investors per round (11.0%). These five attributes account for roughly three-quarters of all attribution, and the twelve sector and location controls together carry 5.4%. Each of the five is oriented in the expected direction, with higher values corresponding to higher assessed readiness (Figure 1).

Aggregated by dimension, the shares are Funding Efficiency 27.9%, Milestone Achievement 26.1%, Relational Strength 24.8% and Investor Diversity 15.7%. The number of attributes in each dimension affects these totals. Funding Efficiency reaches its share through four attributes, whereas Relational Strength reaches a comparable share through one. Per attribute, Relational Strength is more than three times as influential as any other dimension (0.558 against 0.157 for Funding Efficiency), so the ordering of dimensions reflects how many attributes each one contains as much as how much each contributes. The ordering of individual attributes is stable across learners; the ordering of dimension totals is not, and per-attribute values are therefore reported alongside dimension totals throughout.




Figure 1: Attribute contributions to assessed readiness (SHAP)

4.4 Startup Success Readiness Index (SSRI)

The assessed readiness score is the model's output for a venture, a probability-scaled quantity that requires the fitted model to compute. The SSRI is a transparent composite index on a 0–100 scale that reproduces the model's reasoning in one line of arithmetic and can be computed on paper. The index is intended for programme use, where a score that cannot be inspected is of limited value.

Its construction proceeds in three steps. Each attribute is first rescaled to a common scale by min–max normalisation using bounds taken from the training partition only, and oriented so that a higher value always corresponds to higher readiness. For attribute j on venture i , with training bounds L_j and U_j and direction $\sigma_j \in \{+1, -1\}$:

$$\tilde{x}_{ij} = \text{clip}((x_{ij} - L_j) / (U_j - L_j), 0, 1), z_{ij} = \tilde{x}_{ij} \text{ if } \sigma_j = +1, \text{ otherwise } 1 - \tilde{x}_{ij}$$

Second, each attribute is weighted by its share of attribution mass. Writing s_j for the mean absolute SHAP value of attribute j and S_p for the sum over the attributes of dimension p , the within-dimension weight is $w_{j|p} = s_j / S_p$ and the dimension weight is $W_p = S_p / \sum S$. The dimension sub-score and the composite are then:

$$P_{ip} = 100 \times \sum_{j \in p} w_{j|p} \cdot z_{ij}$$

$$\text{SSRI}_i = \sum_p W_p \cdot P_{ip} = 0.2949 P_{FE} + 0.2764 P_{MA} + 0.2624 P_{RS} + 0.1662 P_{ID}$$

The construction has three consequences. First, the weights are not chosen: they are the model's own attribution shares, so an attribute the model does not use receives weight zero automatically. Five of the sixteen attributes receive exactly zero weight, and the index therefore rests on eleven attributes. Second, since $W_p \cdot w_{j|p} = s_j / \sum S$, the two-level structure is algebraically identical to weighting every attribute by its plain share of total attribution mass; the dimension structure determines how a score is explained but has no effect on its value, and reassigning attributes between dimensions leaves every venture's SSRI unchanged. Third, the twelve sector and location controls are excluded from the index by design, since they are not levers available to a programme designer, and the weights of the remaining attributes are renormalised accordingly.

Table 2: Dimension weights derived from attribution mass

Dimension	Active/total attributes	W_p
Funding Efficiency (FE)	4 / 4	29.49%
Milestone Achievement (MA)	3 / 4	27.64%
Relational Strength (RS)	1 / 1	26.24%
Investor Diversity (ID)	3 / 7	16.62%

5. Discussion

5.1 Operationalising Ecosystem Constructs

This study asked whether ecosystem constructs can be mapped onto computable metrics that can be examined empirically. The results support this in a limited sense. OIF-PC defines its indicators at the level of a support programme, not the individual venture; what this analysis shows is that the constructs behind them have venture-level counterparts that are measurable and associated with observed outcomes. The framework itself is not tested here.

The attribution results show three correspondences. Funding attributes speak to the concern behind the Post-Competition Funding Index: ventures with larger and more frequent funding histories receive higher assessed readiness. Milestone attributes speak to the concern behind the Innovation Alignment Index, since milestone achievement records the conversion of intentions into observable progress rather than innovation as such. Relationship intensity is a venture-level counterpart to the Funnel Diversity Score (FDS): FDS was proposed for competition cohorts, but the number of recorded ties reflects the breadth of a venture's interaction with ecosystem actors, and it is the single most influential attribute in the model. These correspondences are interpretive rather than empirical equivalences, and give a basis for further work rather than measurement validation.

5.2 Validation of the Startup Success Readiness Index (SSRI)

The index was evaluated on the 185 reserved ventures, which were used neither for model fitting nor for any weight or normalisation bound. SSRI is positively associated with the outcome (point-biserial $r = 0.432$, $p = 8.2 \times 10^{-10}$; Spearman $\rho = 0.472$) and orders the reserved ventures with ROC-AUC = 0.785. Mean SSRI is 38.8 for acquired ventures against 29.3 for closed ventures (Cohen's $d = 1.00$). When ventures are ranked by SSRI and divided into quintiles, the observed acquisition rate falls from 86.5% in the highest quintile to 29.7% in the lowest. Figure 2 shows the two distributions. They overlap substantially but differ clearly in location.

Three qualifications apply. First, the index is measurably weaker than the model it summarises: the difference in ROC-AUC is 0.073 (95% CI 0.022 to 0.126), so transparency is obtained at a measurable cost in ordering quality. Second, relationships used alone are not statistically distinguishable from the full composite (difference 0.033, 95% CI -0.027 to 0.090); the index is therefore justified by the diagnostic profile it gives across four dimensions, not by better discrimination. Third, attribution weighting is not distinguishable from equal weighting at this sample size (difference 0.030, 95% CI -0.001 to 0.060), so the weights should be read as the model's measured emphasis, not as a demonstrated improvement over a simple average.




Figure 2: Distribution of SSRI by observed outcome

5.3 Interpretation and Practical Implications

The attribution results point to how resources are used rather than to how much of them a venture holds: in this sample, capital is associated with favourable outcomes when it appears together with relational and execution attributes rather than on its own. Three implications follow for support programmes, subject to the associative nature of the evidence. First, relational activity (deliberate connection to mentors, investors and industry partners) is the dimension with the highest attribution per attribute and the most direct programme lever. Second, progress is better represented by several recorded intermediate milestones than by a single terminal assessment. Third, support directed at how capital is deployed is a plausible lever alongside support directed at raising more of it. In the funnel logic that motivated the study, these are design parameters of a support programme rather than properties of a venture.

6. Conclusion

Motivated by the funnel logic of OIF-PC, this study expressed ecosystem characteristics as measurable venture attributes and examined them with interpretable machine learning models. These attributes are associated with observed outcomes, network connectivity most strongly. The study contributes the Startup Success Readiness Index, a transparent composite that translates model attribution into a score programmes can inspect and act on. Its limits are equally clear: the evidence is associative, the sample contains only resolved ventures, and the index summarises the model at a measurable cost in ordering quality.

Ethics Declaration

No ethical approval was required for this study. It did not involve human participants, animals or sensitive personal data, and it used only publicly available information and secondary sources.

AI Declaration

Artificial intelligence (AI) tools were used to help with language editing and to improve the clarity of the text. The authors reviewed, edited and verified every suggestion before it was included.

References

- Breiman, L. (2001) "Random Forests", *Machine Learning*, Vol. 45, No. 1, pp 5–32.
- Chawla, N.V., Bowyer, K.W., Hall, L.O. and Kegelmeyer, W.P. (2002) "SMOTE: Synthetic Minority Over-sampling Technique", *Journal of Artificial Intelligence Research*, Vol. 16, pp 321–357.
- Chen, T. and Guestrin, C. (2016) "XGBoost: A Scalable Tree Boosting System", *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp 785–794.
- Chesbrough, H., Vanhaverbeke, W. and West, J. (2006) *Open Innovation: Researching a New Paradigm*, Oxford University Press, Oxford.
- Cohen, S. (2013) "What Do Accelerators Do? Insights from Incubators and Angels", *Innovations: Technology, Governance, Globalization*, Vol. 8, No. 3–4, pp 19–25.
- Doshi-Velez, F. and Kim, B. (2017) "Towards a Rigorous Science of Interpretable Machine Learning", *arXiv preprint arXiv:1702.08608*.
- Fawcett, T. (2006) "An Introduction to ROC Analysis", *Pattern Recognition Letters*, Vol. 27, No. 8, pp 861–874.
- Giardino, C., Wang, X. and Abrahamsson, P. (2014) "Why Early-Stage Software Startups Fail: A Behavioral Framework", *Lecture Notes in Business Information Processing*, Vol. 182, pp 27–41.
- Hallen, B.L., Cohen, S. and Bingham, C.B. (2020) "Do Accelerators Work? If So, How?", *Organization Science*, Vol. 31, No. 2, pp 378–414.
- Isenberg, D.J. (2010) "How to Start an Entrepreneurial Revolution", *Harvard Business Review*, Vol. 88, No. 6, pp 40–50.
- Kuhn, M. and Johnson, K. (2013) *Applied Predictive Modeling*, Springer, New York.
- Lundberg, S.M. and Lee, S.I. (2017) "A Unified Approach to Interpreting Model Predictions", *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS)*, pp 4765–4774.
- Molnar, C. (2022) *Interpretable Machine Learning*, 2nd ed., Independently published.
- Movahedian Attar, A., Gupta, A. and Taherdoost, H. (2025) "OIF-PC: An Open Innovation Framework for Pitch Competitions", *Proceedings of the XXXVI ISPIM Innovation Conference*, Bergen, Norway, 15–18 June.
- Pauwels, C., Clarysse, B., Wright, M. and Van Hove, J. (2016) "Understanding a New Generation Incubation Model: The Accelerator", *Technovation*, Vols. 50–51, pp 13–24.
- Ribeiro, M.T., Singh, S. and Guestrin, C. (2016) "Why Should I Trust You? Explaining the Predictions of Any Classifier", *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp 1135–1144.
- Saito, T. and Rehmsmeier, M. (2015) "The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets", *PLoS ONE*, Vol. 10, No. 3, e0118432.
- Song, M., Podoynitsyna, K., Van Der Bij, H. and Halman, J.I.M. (2008) "Success Factors in New Ventures: A Meta-analysis", *Journal of Product Innovation Management*, Vol. 25, No. 1, pp 7–27.
- Spigel, B. (2017) "The Relational Organization of Entrepreneurial Ecosystems", *Entrepreneurship Theory and Practice*, Vol. 41, No. 1, pp 49–72.
- Stam, E. (2015) "Entrepreneurial Ecosystems and Regional Policy: A Sympathetic Critique", *European Planning Studies*, Vol. 23, No. 9, pp 1759–1769.
- West, J. and Bogers, M. (2014) "Leveraging External Sources of Innovation: A Review of Research on Open Innovation", *Journal of Product Innovation Management*, Vol. 31, No. 4, pp 814–831.

Appendix A: Model attributes

The model uses 28 attributes: 16 assigned to the four readiness dimensions and 12 sector and location controls. Weights are the attribution shares used in the SSRI; a weight of 0.00 indicates the attribute was not used by the model.

Attribute	Dimension	Description	Weight
relationships	RS	People and organisations	26.24%
age_last_milestone_year	MA	Years from founding to most recent milestone	13.48%
log_funding_total_usd	FE	log(1 + total capital raised, USD)	13.29%
milestones	MA	Count of recorded milestone events	13.10%

Attribute	Dimension	Description	Weight
avg_participants	ID	Average number of investors per round	11.63%
age_last_funding_year	FE	Years from founding to most recent funding	6.60%
age_first_funding_year	FE	Years from founding to first funding	6.50%
has_VC	ID	A venture-capital firm participated	4.16%
funding_rounds	FE	Number of distinct funding rounds	3.10%
age_first_milestone_year	MA	Years from founding to first milestone	1.07%
has_roundD	ID	Series D stage reached	0.84%
has_angel	ID	Angel investor participated	0.00%
has_roundA / has_roundB / has_roundC	ID	Series A / B / C stage reached	0.00%
has_milestone	MA	At least one milestone recorded (imputation indicator)	0.00%
12 control indicators	—	Headquarters state (CA, NY, MA, TX) and sector (software, web, mobile, enterprise, advertising, games/video, e-commerce, biotech)	excluded by design

Hyperparameters were selected by exhaustive grid search over 105 configurations (Logistic Regression 14, Random Forest 27, XGBoost 64) under 5-fold stratified cross-validation on the training partition. Selected settings: Logistic Regression $C = 0.1$, L2 penalty; Random Forest max depth 6, sqrt features per split, 200 trees; XGBoost max depth 2, learning rate 0.03, subsample 1.0, colsample 1.0, 400 trees.