

A PRACTICAL THEORY OF FUNGIBILITY

Shorish, Jamsheed; Stephenson, Matt; Zargham, Michael

DOI:

[10.57938/59e7abc9-c99e-49e5-b2f1-a41376d0f44e](https://doi.org/10.57938/59e7abc9-c99e-49e5-b2f1-a41376d0f44e)

Published: 27/05/2021

Document Version

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Shorish, J., Stephenson, M., & Zargham, M. (2021). *A PRACTICAL THEORY OF FUNGIBILITY*. Working Paper Series / Institute for Cryptoeconomics / Interdisciplinary Research <https://doi.org/10.57938/59e7abc9-c99e-49e5-b2f1-a41376d0f44e>

A PRACTICAL THEORY OF FUNGIBILITY

J. SHORISH¹², M. STEPHENSON¹³, AND M. ZARGHAM¹²

ABSTRACT. We formalize ‘degrees of fungibility’ by differentiating goods according to both their underlying attributes and the perceived value and/or usefulness of those attributes to a value assessor. This allows us to distinguish between goods that appear to be ‘exactly the same’ from those goods that appear to be ‘nearly the same’. Such a distinction is of particular importance in the design space of digital goods, which may exist both natively in the digital space and as surrogates, i.e. as digital representations of physical goods. We provide motivating examples where digital objects are too fungible for certain desired uses, and proceed to develop a formal framework under which degrees of fungibility can be defined and characterized. We close by bridging this framework to applications in machine learning and market design.

1. INTRODUCTION

The world is a more standardized place than it used to be. Over the past centuries, nation states established standardizations like uniform measurements, official languages, and legal codifications. International institutions and markets took widely varying products and commoditized them to further price discovery and trade. These projects which so shaped our modern world were essentially about making things more interchangeable, more fungible.[1]

If the characteristic project of the past few centuries was to *fungibilize* the non-fungible, then the digital realm’s project might well be the opposite. We begin with near-perfect interchangeability in the world of bits, and from elementary particles of logic we have all manner of matter with which to build uniqueness. Which meta-data, for instance, are worth preserving?

If creating, identifying, and valuing non-fungibility is our design challenge then we will need a sufficient characterization. We hope more precision in these matters can illuminate a path towards meeting this challenge.

2. CONCEPTS AND INTUITION

To be fungible is to be interchangeable. The digital money in your bank account, for instance, is fungible — one dollar is as good as any other dollar. And since these digital dollars are interchangeable, it would seem bizarre if you went to buy something with a debit

¹BLOCKSCIENCE ²VIENNA UNIVERSITY OF ECONOMICS AND BUSINESS ³PLANCK

Date: May 27, 2021.

Thanks to David Sisson and Vernon Smith for helpful suggestions. Note that neither reviewed this manuscript and, of course, any errors are the authors’ alone.

card and you were asked which “particular” dollar in your checking account you’d like to pay with. Fungibility as an exchange standard—and the apparent ‘paradoxes’ which can result from its failure—have been recognized as an important cornerstone for understanding both micro- and macro-level behavior since at least the work of Nobel Laureate Richard Thaler [2].

In fact, for digital money in your bank account it isn’t clear what one “particular” dollar would even mean. These are just counts in a database. And thus each unit of your digital money is not only functionally interchangeable but is also indistinguishable, *salva veritate*. As such, we might say that digital money in a checking account approaches a platonic ideal of fungibility.¹²

A dollar in cash strays somewhat from this fungible “ideal”; for instance, one dollar cash can at least be distinguished from another. A given dollar has different spatio-temporal coordinates than another certainly, but also may have different serial numbers, years, markings etc. Such distinguishing aspects make the cash dollar “semi-fungible”.³

Because we don’t usually care which particular dollar we pay for something with, in practice we tend to treat a cash as fungible.⁴ However, because cash dollars are *distinguishable*, they can be treated as non-fungible if someone chose. For instance, it’s not uncommon to go into a small business and see framed the particular dollar bill from the business’s first sale. For the small business owner, that framed dollar bill is not interchangeable with another.

Now suppose that you secretly switch that business’s framed dollar bill out with another. Maybe nobody ever notices. But still, there would at least exist *the concept* of what was and was not the right bill.⁵ That this concept of “the right dollar” exists for cash but not for digital money directs our attention to non-fungibility as a design feature.

Perhaps our example’s business owner and their desire memorialize their first transaction is insufficiently important to warrant tracking such information. But that may not always be so.[5] For instance, consider that *the* digital jersey that an e-sports star wore when they made the game-saving play does not exist as a good (and is thus not available to fans.) This may represent a failure of non-fungible design — these e-sports “jerseys” (typically called “skins”) are not designed to index uniquely important events and associations. They are thus too fungible for some uses and, given the high price and ubiquity of sports memorabilia, this is a failure one might expect to be corrected in the future.

¹Since we just mentioned Plato, and in the preceding sentence alluded to Leibniz, we might as well acknowledge that we are treading lightly along some fundamental issues. Theories of identity, reference and self-reference, and set theoretic reconciliations of the above could all be imagined in the margins here.

²On the other end of things, the platonic ideal of non-fungibility would be something that is maximally unique, completely unsubstutable. Perhaps the universe itself?

³This example builds on insights from Matt Condon & Jonathan Mann [3].

⁴And it’s no accident that fungibility is often tied to some notion of function; the word itself derives from the Latin meaning “to perform”. Cash dollars are usually regarded as interchangeable because we care about their function in getting us something else.

⁵The right bill is probably the one that exhibited an unbroken conceptual lineage, tracing “exactly one connected path over space and time”. [4]

With problems like these in mind, we proceed to develop a framework in which fungibility and standardization via commodities can be more rigorously understood and applied.

3. DEFINITIONS AND PROPERTIES

Let us consider first a set of *items* \mathcal{S} . While in principal items may be added or removed over time, so that the set \mathcal{S} depends upon time t (e.g. \mathcal{S}_t), such dynamics will be reserved for later consideration. Let items have *attributes* $x \in \mathcal{X}$.⁶ The space \mathcal{X} is akin to a feature space in a machine learning application: it is an encoding of the information about an item that may be useful in characterizing that object. In general \mathcal{X} may be arbitrary, but the choice to exclude or include a particular descriptor of an item in the space is to decide what constitutes the item itself.

Formally, an item in our model may be defined as follows:

Definition 1. An *item* is an element $i \in \mathcal{S}$ such that there exists a mapping

$$(1) \quad m : \mathcal{S} \rightarrow \mathcal{M} \subseteq \mathcal{X},$$

where \mathcal{M} is the *codomain* of $m(\cdot)$ and \mathcal{X} denotes the space of *attributes* describing any item.

We may think of the attribute space \mathcal{X} as encompassing all possible combinations of attributes that items may have, and the codomains $\mathcal{M}, \mathcal{M}'$ of different attribute maps $m(\cdot), m'(\cdot)$ (resp.) may occupy different subsets of \mathcal{X} . In addition, it is not assumed that all items are perfectly unique (formally, the map m is not assumed to be injective into its codomain \mathcal{M}). This means that it is possible to have items which are indistinguishable according to their attributes:

Definition 2. Two items $i, j \in \mathcal{S}$ are *indistinguishable* under attribute map m if and only if

$$(2) \quad m(i) = m(j) \in \mathcal{M}.$$

It is important to note that indistinguishability is defined *relative* to a specified map m . If one constructs another map m' it may be that $m'(i) \neq m'(j) \in \mathcal{M}'$ and in this case, i and j are indistinguishable under m but are distinguishable under m' .

Definition 3. A subset of items $\bar{\mathcal{S}} \subset \mathcal{S}$ with common attributes $x \in \mathcal{X}$ under m is a *commodity*:

$$(3) \quad \bar{\mathcal{S}} := \{i \in \mathcal{S} : m(i) = x\}$$

and the cardinality $|\bar{\mathcal{S}}| \gg 1$.

With these definitions in place we can begin to explore the ways in which a value assessor, such as a consumer, can place value upon (collections of) items *via their attributes*. We first proceed with a simple representation of preferences over the attribute space \mathcal{X} , using a collection of functions as the ‘building blocks’ of a utility representation.

⁶The attribute or *characteristic* space is a useful description space for multi-faceted goods—cf. e.g. [6] for the seminal treatment of hedonic goods and their associated pricing for an early example from economics.

Proposition 1. Consider the attribute space \mathcal{X} and suppose there exists a set \mathcal{B} of functions $f : \mathcal{X} \rightarrow \mathbb{R}$, called (for reasons given shortly) the set of basis functions, such that $\forall f \in \mathcal{B}$ a (supremum or an infinity) norm is defined by:

$$(4) \quad \|f(x) - f(x')\|_\infty := \max_{f \in \mathcal{B}} |f(x) - f(x')| > 0 \quad \forall x, x' \in \mathcal{X}, x \neq x'.$$

Then \mathcal{X} is a metric space: there exists a distance measure $d : \mathcal{X} \rightarrow \mathbb{R}_+$ such that

$$(5) \quad d(x, x') > 0 \quad \forall x, x' \in \mathcal{X}, x \neq x'$$

and $d(x, x) = 0 \quad \forall x \in \mathcal{X}$.

Assumption 1. Given any two attribute points $x, x' \in \mathcal{X}$, $x \neq x'$, the set of basis functions \mathcal{B} is such that

$$(6) \quad f_k(x) = f_k(x') \Rightarrow x = x', \quad \forall f_k \in \mathcal{B}, \quad \forall k = 1, \dots, n,$$

i.e. every basis function is an injection (a 1:1 function).

The manner in which items are valued depends upon their *context*, i.e. the circumstances under which the item finds itself being considered for use. For the purpose of this work, a context is any circumstance which places a value *on* an item, or derives a value *for* an item. For example, a context may be the circumstances of a consumer of a good (and their associated preferences under those circumstances), or a specific use case (and its associated goals or deliverables).

Definition 4. A **context** is a ‘frame’ (cf. e.g. [7]) within which a value assessor (such as a consumer) is capable of determining their preferences—here we assume that an assessor has preferences directly over attributes, rather than directly over items (which provide attributes). We denote a context by $c \in \mathbb{R}^n$, i.e. a context is a vector of real numbers, and $n := |\mathcal{B}|$ is the cardinality of the set of basis functions \mathcal{B} .

Definition 5. Given a context $c \in \mathbb{R}^n$, a **utility representation** of preferences over attributes is defined by a function $u_c : \mathcal{X} \rightarrow \mathbb{R}$, such that:

$$(7) \quad u_c(x) := \sum_{k=1}^n c_k f_k(x)$$

where k is an index over the elements of the basis \mathcal{B} .⁷

Defining a utility representation as an expansion over \mathcal{B} (which justifies its designation as a set of basis functions), rather than as being drawn directly from \mathcal{B} alone, provides a richer characterization of utility. For example, even though the set of basis functions \mathcal{B} is injective it is not the case that every context ‘cares about’ every attribute. In other words,

⁷We assume $n < \infty$ for simplicity, i.e. there are a finite number of basis functions, but the approach does not preclude a countable or uncountable infinite of basis functions (where for the latter, the sum is replaced with an integral) in what follows.

there may be two attributes $x, x' \in \mathcal{X}$ where $u_c(x) = u_c(x')$, but $x \neq x'$ (compare with Assumption 1 above).⁸

A context-dependent utility representation also allows fungibility to be defined by context (and, implicitly, by the mapping m from items to attributes). This means that *changing the context under which attributes are valued may change an item's fungibility with another item*. From a utility representation point of view, then, fixing a context fixes the way that attributes can be compared. If two sets of attributes x and x' provide the same utility for a given context, their underlying items are said to be fungible:

Definition 6. *Two items $i, j \in \mathcal{S}$ are **fungible** in context $c \in \mathbb{R}^m$ if and only if*

$$(8) \quad u_c(m(i)) = u_c(m(j)),$$

where the utility function $u_c : \mathcal{X} \rightarrow \mathbb{R}$ is constructed according to Definition 5 and $m : \mathcal{S} \rightarrow \mathcal{M} \subseteq \mathcal{X}$ is the attribute map for the set of items \mathcal{S} as defined in Definition 1.

Although we have not defined utility directly from a preference ordering, the definition of fungibility may also be couched as *context dependent indifference*: for a particular context, fungible attributes partition the attribute space into equivalence classes in the same way that indifference sets (cf. e.g. [9]) partitions the space of goods. Ranking equivalence classes using e.g. real numbers as the ordinal field (as is usually done in utility theory) recovers the definition of fungibility given above: a value assessor viewing two items as fungible may be thought of as being ‘indifferent’ in the attribute space (via the mapping m) **for the given context**. If it should happen that two items are fungible for *every* context, then we recover a definition of perfect fungibility:

Definition 7. *Two items $i, j \in \mathcal{S}$ are **perfectly fungible** under \mathcal{B} if they are fungible in all contexts $c \in \mathbb{R}^n$.*

Theorem 1. *Given the set of basis functions \mathcal{B} and an attribute mapping m , two items are perfectly fungible if and only if they are indistinguishable under m .*

Corollary 1. *All items within a commodity $\bar{\mathcal{S}}$ are perfectly fungible with each other.*

Commodities arise naturally from standardization. It makes sense to think of a particular map m as standard if it is being used to define an invariant for the commodity $\bar{\mathcal{S}}$. In practice this means that items $i \in \bar{\mathcal{S}}$ are manufactured with the express purpose of satisfying the invariant $m(i) = m(j) \forall i, j \in \bar{\mathcal{S}}$. The standard is performative in so far as it brings the commodity $\bar{\mathcal{S}}$ into existence by defining it, and making it possible to create more items which satisfy it.

The above constructions provide a formal characterization of the variations of items and how they manifest as potential variations across different contexts under which these items might be valued.

⁸Note that via m it is also possible, for a given context c , to *pull back* u_c to \mathcal{S} . If it could be assumed that m were given and fixed, this would imply a model where products are differentiated in the item space, rather than in the attribute space—cf. e.g. [8] for a seminal model of monopolistic competition using this commodity approach (Definition 3 above provides some of the ‘machinery’ for proceeding in this direction).

4. RELATIONSHIP TO MACHINE LEARNING

The above formalism lays the groundwork for interpreting markets as online “machine learning” processes. This improves our ability to use data driven methods to better understand and forecast future market activity, as well as to design better market mechanisms. A “better” market mechanism is one capable of distilling private preference signals into public price signals which are more granular and less noisy, while retaining the ability to track changes in the underlying signals. For further reading see e.g. [10] for a landmark contribution to understanding underlying shocks using time series filtering, [11] for an overview of time series methods in economics and finance, [12] for an application of particle filtering for structural parameter estimation and more recently [13] considering formal game structures as estimation mechanisms.

Building upon Section 3, let us consider an enumerable set⁹ \mathcal{S} with $|\mathcal{S}| < \infty$ items in it, where item attributes are described by the induced metric space (\mathcal{M}, d) of an attribute map m , where $|\mathcal{M}| < \infty$. Furthermore, let us assume that there are a set of events or encounters where an item $i \in \mathcal{S}$ has its value revealed for a context $c \in \mathcal{C} \subset \mathbb{R}^n$. Further suppose that a discrete set of such encounters manifest for a set of contexts with cardinality $|\mathcal{C}| < \infty$. One might think of \mathcal{C} as a cloud of points in \mathbb{R}^n where each point is a parameterization of a utility function u_c .

At this point we can look at what kinds of data sets could arise naturally from contexts such as purchasing decisions over items. We can define a set of encounters as $\mathcal{E} = \{(i, c) : i \in \mathcal{S}, c \in \mathcal{C}\}$, which can be interpreted as a bipartite graph from vertices which are items i to vertices which are contexts c . Each encounter $e = (i, c) \in \mathcal{E}$ implies a unique value assignment

$$(9) \quad y_{(i,c)} = u_c(m(i)).$$

In practice, the only direct observation is that item i was purchased for (say) a price p , where we define the price as an observation based upon the valuation $y_{(i,c)}$ (so that e.g. $p = p(y_{(i,c)})$ for some function $p(\cdot)$ over valuations). We may have some model of the attributes of $m(i) \in \mathcal{X}$, but we do not know with certainty what attributes drove the purchase of i . Furthermore, we may have a set of basis functions over our feature space \mathcal{X} but we do not know with certainty what the vector $c \in \mathbb{R}^n$ precipitating that purchase was. This may seem like a lost cause, we know so little! However, this construction is adjacent to the class of collaborative filtering [15] algorithms commonly used in recommendation engines, albeit with a focus on regression, rather than binary classification. Collaborative filtering in particular works well when there are observable features for the contexts such as data about a platform’s users, in addition to the attributes of items [16]. The method involves sparse matrix factorizations that can discover which users will prefer which products, using a form of pattern recognition [17].

However, even if we have no context data our formalism can still be used to inform a choice of machine learning methods. For example, we may choose to organize event data

⁹Finiteness of the sets used in this representation is to ensure the *computability* of the associated machine learning (e.g. convex optimization [14]) problem.

as a vector \vec{p} of length $r = |\mathcal{E}|$:

$$(10) \quad [\vec{p}]_e = p_e = p(y_{(i,c)})$$

where the index e maps to a pair (i, c) . In this form it is clear that we have a bunch of records of encounters resulting in feature label pairs **(features, labels)** where $x = m(i)$ are the features and p_e is the label for each event $e = (i, c)$. This is a canonical structure for supervised machine learning. If we have some knowledge of what kinds of (basis) functions over the attributes are likely to make up the utility functions, then we can apply non-linear kernel methods, (cf. e.g. [18]). Specifically, we are interested in machine learning methods where the model is updated after each new observation rather than trained in batches, as in [19].

5. NEXT STEPS

Our mathematical representation of *fungibility* expresses the importance and interdependence of an items descriptions, and the context in which that item is evaluated. This approach is novel in that it suggests that markets may be shaped by the way items are *described*, not merely by what attributes these items have; this point holds doubly for goods which are natively digital. Furthermore, the definitions we have constructed provide a basis from which some observed economic phenomena such as Gresham’s law [20] and context-dependent preference switching [7] can be evaluated. Our next steps include establishing formal results to this effect.

In addition to giving the mapping of items to attributes first class status in our theory, we have dismantled our utility functions into linear combinations of basis functions. In doing so, we make more explicit the inherent similarity “AI” and markets in their capacity to synthesize data into estimates [21]. In the blockchain world, market designs such as constant function market makers are increasingly interpreted as price sensors [22].

The practical successes of constant function market makers [23] for price discovery during the Cambrian explosion of tokens on the Ethereum network [24] demonstrate *market mechanisms* that rely upon signal processing properties which (like online learning algorithms) are derived using iterative convex optimization [25] (or as a generalization of convex programming, including quasi-convex optimization and generalized geometric programs). It stands to reason that automated market makers (AMMs) designed according to the same principle can *enable* price discovery, even amongst items that are highly variable and only partially fungible in practice. We will call these markets **Automated Regression Markets** (ARMs) and endeavor to derive analytically and demonstrate numerically the criteria under which such markets can achieve price discovery, even for items whose particular collection of attributes has not previously been labeled by earlier purchases within that market.

6. APPENDIX: PROOF OF THEOREM 1

[\Rightarrow] Consider an attribute map $m : \mathcal{S} \rightarrow \mathcal{M} \subseteq \mathcal{X}$ and two items $i, j \in \mathcal{S}$ that are identical under m , i.e.

$$m(i) = m(j).$$

Then $\forall c \in \mathbb{R}^n$,

$$u_c(m(i)) := \sum_k c_k f_k(m(i)) = \sum_k c_k f_k(m(j)) =: u_c(m(j)),$$

i.e. items i and j are perfectly fungible.

[\Leftarrow] Suppose that two items $i, j \in \mathcal{S}$ are perfectly fungible, i.e. under an attribute map m , $\forall c \in \mathbb{R}^n$,

$$u_c(m(i)) := \sum_k c_k f_k(m(i)) = \sum_k c_k f_k(m(j)) =: u_c(m(j)).$$

Since perfect fungibility holds for all c , select a context $\bar{c}^s := (0, \dots, 1, 0, \dots, 0)$, i.e. a context whose s th element is 1 and zero otherwise. Applying this context to the definition of perfect fungibility implies

$$u_{\bar{c}^s}(m(i)) = f_s(m(i)) = f_s(m(j)) = u_{\bar{c}^s}(m(j)), \quad s \in \{1, \dots, n\}.$$

Since every basis function is an injection (cf. Assumption 1), this immediately gives

$$m(i) = m(j),$$

i.e. items i and j are indistinguishable under m . □

REFERENCES

- [1] S. Krislov, *How nations choose product standards and standards change nations*. University of Pittsburgh Pre, 1997.
- [2] R. H. Thaler, “Anomalies: Saving, fungibility, and mental accounts,” *Journal of Economic Perspectives*, vol. 4, no. 1, pp. 193–205, 1990.
- [3] M. G. Condon and J. Mann, “Sharing ideas about idea sharing with planck.” Digitally Rare.
- [4] M. Fernández *et al.*, “A deep dive on erc-721,” URL: <https://medium.com/ecf-review/more-than-cats-and-pandas-a-deep-dive-on-erc-721-in-preparation-for-ecfs-nonfungible-summit-bebb824443a0>, 2018.
- [5] R. Watkins and M. Molesworth, “Attachment to digital virtual possessions in videogames,” in *Research in consumer behavior*, Emerald Group Publishing Limited, 2012.
- [6] S. Rosen, “Hedonic prices and implicit markets: Product differentiation in pure competition,” *Journal of Political Economy*, vol. 82, no. 1, pp. 34–55, 1974.
- [7] A. Tversky and I. Simonson, “Context-dependent preferences,” *Management science*, vol. 39, no. 10, pp. 1179–1189, 1993.
- [8] A. K. Dixit and J. E. Stiglitz, “Monopolistic competition and optimum product diversity,” *The American Economic Review*, vol. 67, no. 3, pp. 297–308, 1977.
- [9] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford University Press, 1995.
- [10] F. E. Kydland and E. C. Prescott, “Time to build and aggregate fluctuations,” *Econometrica*, vol. 50, no. 6, pp. 1345–1370, 1982.
- [11] J. D. Hamilton, *Time Series Analysis*. Princeton University Press, 1994.

- [12] J. Fernández-Villaverde and J. F. Rubio-Ramírez, “Estimating macroeconomic models: A likelihood approach,” *The Review of Economic Studies*, vol. 74, no. 4, pp. 1059–1087, 2007.
- [13] M. Zargham, K. Paruch, and J. Shorish, “Economic games as estimators,” in *Mathematical Research for Blockchain Economy*, pp. 125–142, Springer, 2020.
- [14] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [15] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, “Item-based collaborative filtering recommendation algorithms,” in *Proceedings of the 10th International Conference on World Wide Web, WWW '01*, (New York, NY, USA), p. 285–295, Association for Computing Machinery, 2001.
- [16] X. Su and T. M. Khoshgoftaar, “A survey of collaborative filtering techniques,” *Advances in artificial intelligence*, vol. 2009, 2009.
- [17] A. Mnih and R. R. Salakhutdinov, “Probabilistic matrix factorization,” *Advances in neural information processing systems*, vol. 20, pp. 1257–1264, 2007.
- [18] J. Shawe-Taylor, N. Cristianini, *et al.*, *Kernel methods for pattern analysis*. Cambridge university press, 2004.
- [19] W. Liu, J. C. Principe, and S. Haykin, *Kernel adaptive filtering: a comprehensive introduction*, vol. 57. John Wiley & Sons, 2011.
- [20] R. Mundell, “Uses and abuses of gresham’s law in the history of money,” *Zagreb Journal of Economics*, vol. 2, no. 2, pp. 3–38, 1998.
- [21] M. I. Jordan, “Dr. ai or: How i learned to stop worrying and love economics,” *Harvard Data Science Review*, vol. 1, no. 1, 2019.
- [22] F. Martinelli and N. Mushegian, “A non-custodial portfolio manager, liquidity provider, and price sensor,” *URL: <https://balancer.finance/whitepaper>*, 2019.
- [23] G. Angeris and T. Chitra, “Improved price oracles: Constant function market makers,” in *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies, AFT '20*, (New York, NY, USA), p. 80–91, Association for Computing Machinery, 2020.
- [24] A. J. Morales, S. Somin, Y. Altshuler, and A. Pentland, “User behavior and token adoption on erc20,” *arXiv preprint arXiv:2005.12218*, 2020.
- [25] S. Shalev-Shwartz *et al.*, “Online learning and online convex optimization,” *Foundations and trends in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2011.