KAN EXTENSIONS AND NONSENSICAL GENERALIZATIONS

ADAM ANDERSON

Abstract. The categorical concept of Kan extensions form a more general notion of both limits and adjoints. The general definition of Kan extensions is given and motivated by several concrete examples. After providing the necessary background on some basic categorical objects and theorems, the relationship between Kan extensions, limits, and adjoints is expressed through two theorems from [3].

1. Some Preliminary Categorical Concepts

A tremendous array of fields within mathematics draws heavily upon the ideas of limits and adjoints. While these notions are sufficiently general for most uses, there exists a more abstract concept introduced by Kan [2], which encompasses both limits and adjoints.

Although the following discussion assumes familiarity with the basic language of category theory, we begin by summarizing some terminology and a few results for reference and clarity. The notion of a limit is the first of these. Limits are easily understood through the auxiliary notion of a cone.

Definition 1.1. Given a functor $F : \mathcal{D} \to \mathcal{C}$, a *cone* on F is a pair (C, p_D) consisting of:

- an object $C \in \mathcal{C}$,
- a morphism $p_D : C \to FD$ in $\mathcal{C}$, for every object $D \in \mathcal{D}$,

such that for every morphism $d : D \to D'$ in $\mathcal{D}$, $p_{D'} = Fd \circ p_D$.

The name "cone" is used for a reason; pictorially, cones are situations in which there are morphisms that take the object C to the objects FD_i, with the following diagram commuting:

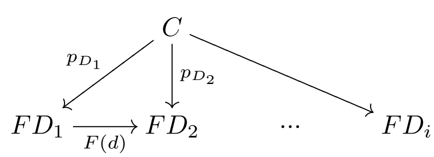

Definition 1.2. A *limit* of a functor is a universal cone. i.e. A limit of a functor $F : \mathcal{D} \to \mathcal{C}$ is a cone $(L, (p_D)_{D\in\mathcal{D}})$ on F such that, for every cone $(M, (q_D)_{D\in\mathcal{D}})$ on F, there exists a unique morphism $m : M \to L$ such that for every object $D \in \mathcal{D}$, $q_D = p_D \circ m$.

Kan extensions also generalize another important categorical structure: the adjunction.

Date: Summer 2007.

Definition 1.3. Let $\mathcal{C}$ and $\mathcal{D}$ be categories. An *adjunction* from $\mathcal{C}$ to $\mathcal{D}$ is a triple (F, G, φ), in which $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{C}$ are functors. The adjunction also has a function φ, which assigns to each pair of objects $c \in \mathcal{C}$ and $d \in \mathcal{D}$ a bijection of sets

$$\varphi = \varphi_{c,d} : \mathcal{D}(Fc, d) \simeq \mathcal{C}(c, Gd),$$

which is natural in c and d.

Remark 1.4. F is called the *left adjoint* of G, while G is called the *right adjoint* of F.

For brevity, we give the following result about adjunctions without a proof. It is standard, and the proof can be found in [1] or [3], for example.

Theorem 1.5. *An adjunction $(F, G, \varphi) : \mathcal{C} \leftrightarrow \mathcal{D}$ determines a natural transformation*[1] *$\eta : I_{\mathcal{C}} \to_N GF$, such that for each object $c \in \mathcal{C}$, the arrow η_c is universal from the object c to G. We call η the* unit *of the adjunction.*

Analogous to the unit is the *counit*. Given the adjunction in 1.5, the counit is another natural transformation $\epsilon : FG \to_N I_{\mathcal{D}}$. Not surprisingly, a theorem analogous to 1.5 holds for the counit. Another important fact is that the unit and counit transformations determine an adjunction. Together, the definitions of the unit and counit suggest that each should somehow act as a type of "inverse" of the other. Strictly speaking, of course, they are not inverses, but this intuition can be formalized in the following two equations known as the triangular identities.

Theorem 1.6. *Let $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{C}$ be functors. If G is left adjoint to F with unit and counit natural transformations $\eta : 1_{\mathcal{C}} \to_N GF$ and $\epsilon : FG \to_N 1_{\mathcal{D}}$ respectively, then*

$$G\epsilon \circ \eta G = 1_G \tag{1}$$

$$\epsilon F \circ F\eta = 1_F \tag{2}$$

In pictures we have:

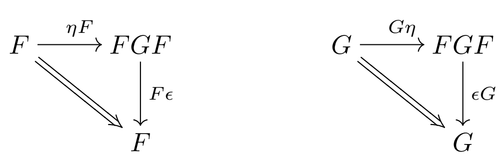

Proof. From the definition of the adjunction, we have the isomorphism:

$$\varphi = \varphi_{c,d} : \mathcal{D}(Fc, d) \simeq \mathcal{C}(c, Gd). \tag{3}$$

If we plug $1_{Gd} : Gd \to Gd$ into the right-hand side of (3), and recall that $\epsilon : FG \to_N 1_{\mathcal{D}}$, then we know that $\varphi^{-1}(1_{Gd}) = \epsilon_d$. It follows that:

$$\begin{aligned} 1_{Gd} &= \varphi(\epsilon_d) \\ &= G\epsilon \circ \eta G. \end{aligned}$$

Reasoning by analogy, we find that $\varphi(1_F c) = \eta_c$, and:

$$\begin{aligned} 1_{Fc} &= \varphi^{-1}(\eta_c) \\ &= \epsilon F \circ F\eta, \end{aligned}$$

[1]Different authors use various notations for natural transformations. I use an arrow with a subscript N: $\to_N$

which are the two triangular identities. In other words, the diagrams above commute, whence arises the name "triangular identities". □

2. Kan Extensions

With this basic machinery and terminology we can finally introduce the formal definition of the Kan extension.

Definition 2.1. Let $K : \mathcal{M} \to \mathcal{C}$ and $T : \mathcal{M} \to \mathcal{A}$ be functors. The *right Kan extension* of T along K is a pair (R, ϵ), with a functor $R : \mathcal{C} \to \mathcal{A}$ and a natural transformation $\epsilon : RK \to_N T$, which satisfies the following universal property:

If (H, δ) is another pair with a functor $H : \mathcal{C} \to \mathcal{A}$ and a natural transformation $\delta : HK \to_N T$, then there exists a unique natural transformation $\sigma : H \to_N R$ such that $\delta = \epsilon \cdot \sigma K : SK \to_N T$.

The functor R in definition 2.1 is usually denoted $R = \mathrm{Ran}_K T$. Because definition 2.1 guarantees the existence of a unique natural transformation σ, the map $\sigma \mapsto \epsilon \cdot \sigma K$ is bijective. So for a right Kan extension, we have a bijection:

$$Nat(H, Ran_K T) \simeq Nat(HK, T), \tag{4}$$

which determines the Kan extension by specifying just the functors T and K.

One can visualize Kan extensions through the diagrams:

(5)

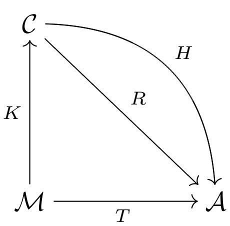

(6)

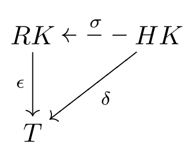

In diagrams (5) and (6), R is the functor of the right Kan extension of T along K, while ϵ is the natural transformation. ϵ maps the composition $R \circ K$ to the functor T. We now pick any other functor from $\mathcal{C}$ to $\mathcal{A}$–say H–along with a natural transformation δ, mapping the composition $H \circ K$ to T. There will then be a unique natural transformation σ that takes H to R such that $\delta = \epsilon \cdot \sigma K$. To put it more simply, the pair that comprises the Kan extension is universal. It is also important to note that only diagram (6) commutes. The functors in (5) do not commute in general, although the diagram is conceptually helpful.

The fact that we specifically define a "right" Kan extension suggests that there ought to be a "left" Kan extension. Not surprisingly, a *left Kan extension* is the dual of a right Kan extension. The notation $\mathrm{Lan}_K T$ is typically used to denote the left Kan extension of T along K. Dualizing definition 2.1 has the net effect of just reversing the direction of all the natural transformations. So instead of having the diagram:

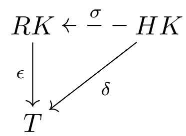

we have the diagram:

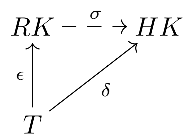

Diagram (5), however, remains unchanged.

Kan extensions can intuitively be considered as "extending" the codomain of one functor into the codomain of another functor in a sort of universal manner. In order to illustrate this concept, we construct several examples. While the first of these is simple, it shows an elementary construction of the Kan extension.

Example 2.2. (Products) Consider the following diagram:

(7)

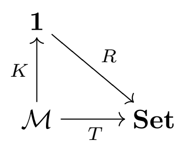

$\mathbf{1}$ is just the category with a single element and an identity arrow. **Set** is the category whose objects are small sets and whose arrows are functions between them. Finally, $\mathcal{M}$ is the *discrete category* $\{1, 2\}$. By this, we simply mean a category with two objects, whose only morphisms are identities. Since T takes two objects to the category of sets, it is equivalent to just a pair of sets. Similarly, it is obvious that R is a single set $a \in \mathbf{Set}$ and K just takes the two objects of $\mathcal{M}$ to the unique object in $\mathbf{1}$.

If we are interested in finding $\mathrm{Ran}_K T$, the right Kan extension of T along K (if it exists), we need to find a functor R that fulfills the universal property. R simply takes the unique element in $\mathbf{1}$ to some set, so our task is simply to pick the right set. Additionally, we need to find a natural transformation $\epsilon : RK \to_N T$.

No matter what, ϵ will be a cone with a vertex a and a base given by T. Now suppose we have another functor $S : \mathbf{1} \to \mathbf{Set}$, which goes to another set $b \in \mathbf{Set}$. We also have another natural transformation $\delta : SK \to_N T$. Similar to ϵ, δ will also be a cone with a vertex b and base T. The natural transformation $\sigma : S \to_N R$ is just a morphism from the vertex b to the vertex a. At this point, we want σ to be unique, and we have two cones. It is now obvious that the Kan extension here is equivalent to the product of a and b. So R goes to $a \prod b$, and ϵ is the cone with the canonical projection functions.

Example 2.3 (Representable Functors)**.** Another example of Kan extensions can be constructed using representable functors, which are defined by the following[2]:

[2]Different authors give slightly different definitions, but the relevant point is that, up to isomorphism, representable functors are covariant hom-functors.

Definition 2.4. Given a category $\mathcal{C}$ and a fixed object $c \in \mathcal{C}$, we define the *representable functor*

$$\mathcal{C}(c,-) : \mathcal{C} \to \mathbf{Set}.$$

For an object $d \in \mathcal{C}$, $\mathcal{C}(c,-)$ performs the assignment

$$\mathcal{C}(c,-)(d) = \mathcal{C}(c,d),$$

For a morphism $f : d \to e$ in $\mathcal{C}$, the assignment is

$$\mathcal{C}(c,-)(f) = f \circ g,$$

where $g : c \to d$ is a morphism in $\mathcal{C}$.

Consider small categories $\mathcal{A}$ and $\mathcal{B}$, and a functor $F : \mathcal{A} \to \mathcal{B}$. Given a fixed object $a \in \mathcal{A}$, we also consider the functors $\mathcal{A}(a,-)$ and $\mathcal{B}(Fa,-)$. Finally, let us also suppose that we have a functor $G : \mathcal{B} \to \mathbf{Set}$. The Yoneda lemma gives the existence of a bijection:

$$\psi : Nat(\mathcal{B}(Fa,-), G) \simeq GFa$$

Given a functor G and a particular object a, this determines a unique natural isomorphism $\sigma : \mathcal{B}(Fa,-) \to_N G$; if there were two distinct isomorphisms, then there could not be a bijective mapping to the object GFa. If we let $\epsilon : \mathcal{A}(a,-) \to_N \mathcal{B}(Fa, F-)$, then for a given a, the left Kan extension of $\mathcal{A}(a,-)$ along F is the pair (R, ϵ).

3. Relationship to Adjoints and Limits

With all of the necessary categorical preliminaries, we are now ready to consider the connections between Kan extensions, adjoints, and limits. The first result is produced by letting $\mathcal{C} = \mathbf{1}$ in definition 2.1. This causes the natural transformation $\alpha : T \to_N SK_1$ to become a cocone, implying that left Kan extensions and colimits are equivalent. In so doing, we generalize the result of example 2.2 to any limit, beyond just products. Also note that the basic argument that was used in example 2.2 easily gives the proof of the theorem.

Theorem 3.1. *A functor $T : \mathcal{M} \to \mathcal{A}$ has a colimit if and only if it has a left Kan extension along the unique functor $K_1 : \mathcal{M} \to \mathbf{1}$, and then Colim$(T)$ is the value of Lan$_{K_1}T$ on the object of $\mathbf{1}$.*

Proof. Suppose T has a colimit. We are looking for a functor $F : \mathbf{1} \to A$. Since $\mathbf{1}$ has just one object, F is just an object $a \in \mathcal{A}$. Next consider the natural transformation $\alpha : T \to_N SK_1$. The functor T maps objects in $\mathcal{M}$ to objects in $\mathcal{A}$, while the functor $SK : \mathcal{M} \to \mathcal{A}$ maps every object of $\mathcal{M}$ to a constant object $a \in \mathcal{A}$. Hence, α is a cocone. If a is the colimit object of T, then the cocone formed by α is universal, and the Kan extension of T along K_1 exists.

Suppose $\mathrm{Lan}_{K_1}T$ exists. $\mathrm{Lan}_{K_1}T = L$ is the functor $S : \mathbf{1} \to \mathcal{A}$ along with the universal natural transformation $\epsilon : T \to_N LK_1$. By the same reasoning as in the other direction of the proof, ϵ forms a cocone of T, which must be universal. Hence T has a colimit.

As usual, the dual of the theorem holds, substituting limits for colimits and right Kan extensions for left ones. □

In addition, there is a deep connection between the existence of Kan extensions and adjoints. The following theorem gives a criterion for the existence of adjoints in terms of Kan extensions.

Theorem 3.2 (Criteria for the existence of an adjoint)**.** *Let $\mathcal{A}$ and $\mathcal{X}$ be categories. A functor $G : \mathcal{A} \to \mathcal{X}$ has a left adjoints if and only if the right Kan extension $Ran_G 1_{\mathcal{A}} : \mathcal{X} \to \mathcal{A}$ exists and is preserved by G.*

Proof. Despite the state of confusion often generated by the proliferation of commutative diagrams, the following argument is actually simple. The general idea for the first direction is straightforward: given an adjunction, we know that we have a bijection. Incidentally, Kan extensions can also be determined by a bijection, as we saw in equation (4). By picking our categories correctly and finding the right natural transformations, we show that one bijection implies the other, yielding the right Kan extension.

Suppose G has a left adjoint $F : \mathcal{X} \to \mathcal{A}$, with a unit $\eta : 1_{\mathcal{X}} \to_N GF$ and a counit $\epsilon : FG \to_N 1_{\mathcal{A}}$. Let $S : \mathcal{X} \to \mathcal{C}$.

For all functors $H : \mathcal{A} \to \mathcal{C}$, we can construct a bijection

$$Nat(S, HF) \simeq Nat(SG, H), \tag{8}$$

which is natural in S. To do this, we use the following assignments:

$$\{\sigma : S \to_N HF\} \quad \mapsto \quad \{\ SG \xrightarrow{\sigma G} HFG \xrightarrow{H\epsilon} H\ \} \tag{9}$$

$$\{\tau : SG \to_N H\} \quad \mapsto \quad \{\ S \xrightarrow{S\eta} SGF \xrightarrow{\tau F} HF\ \} \tag{10}$$

To show the naturality of the bijection in (8), first apply the map (9) followed by the map (10). We first note that the following diagram is commutative:

$$\begin{array}{ccccc} S & \xrightarrow{\sigma} & HF & & \\ \downarrow{\scriptstyle S\eta} & & \downarrow{\scriptstyle HF\eta} & \searrow\!\!\searrow & \\ SGF & \xrightarrow[\sigma GF]{} & HFGF & \xrightarrow[H\epsilon F]{} & HF \end{array} \tag{11}$$

To show this, first consider the square on the left. It commutes because each path is just the two equivalent ways of composing $\sigma\eta$. The triangle on the right commutes because of the triangular identities for units and counits. The application of (9) followed by (10) is represented in diagram (11) by moving down from S and across the bottom of the diagram to the right. Since the diagram commutes, applying (9) and then (10) just gives the identity map $\sigma \mapsto \sigma$.

In the other direction, applying (10) and then (9), we have another diagram:

$$\begin{array}{ccccc} SG & \xrightarrow{\tau} & H & & \\ \downarrow{\scriptstyle G\eta G} & & \downarrow{\scriptstyle HF\eta} & \searrow\!\!\searrow & \\ SGFG & \xrightarrow[\tau FG]{} & HFG & \xrightarrow[H\epsilon]{} & H \end{array} \tag{12}$$

which commutes by the same arguments. This implies that the application of (10) followed by (9) gives the identity $\tau \mapsto \tau$. The bijection (8) is therefore natural in S.

Let $H = 1_{\mathcal{A}}$ with $\mathcal{C} = \mathcal{A}$. In this case, the bijection (8) requires that $F = \mathrm{Ran}_G 1_{\mathcal{A}}$, because bijection (4) determines the right Kan extension. Letting $H = G$ instead, we have $GF = \mathrm{Ran}_G G$, showing that G preserves our right Kan extension.

In the other direction, the argument is similarly straightforward. We again start with the bijection defining the Kan extension. From this, we can deduce the unit and counit of the desired adjunction, and hence the existence of the adjunction.

Suppose that $R = \mathrm{Ran}_G 1_{\mathcal{A}}$ exists and that R is preserved by $G : \mathcal{A} \to \mathcal{X}$. From the definition of the right Kan extension, we have the bijection (4):

$$\phi = \phi_S : Nat(S, R) \simeq Nat(SG, 1_{\mathcal{A}}),$$

which is natural in $S : \mathcal{X} \to \mathcal{A}$. Since R is preserved by G, by composing with G we obtain another bijection:

$$\psi = \psi_H : Nat(H, GR) \simeq Nat(HG, G),$$

which is natural in $H : \mathcal{X} \to \mathcal{X}$. From these bijections, we deduce:

$$\begin{aligned} \phi_R(1) &= \epsilon : RG \to_N 1_{\mathcal{A}} \\ \psi_{GR}(1) &= G\epsilon : GRG \to_N G \end{aligned}$$

If we define $\eta : 1 \to_N GR$ by $\eta = \phi_{id}^{-1}(1 : G \to G)$, then $\psi\phi = 1$ and

$$G\epsilon \cdot \eta G = 1_G.$$

This is one of the triangle identities which define the adjunction with G left adjoint to F. To show the other, it suffices to show that $\phi(\epsilon R \cdot R\eta) = \epsilon$ using the bijection ϕ_R above. In other words, we want to show that the following diagram commutes:

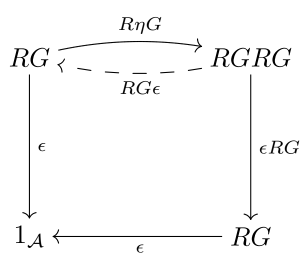

If we insert the dashed natural transformation $RG\epsilon$ and use the first triangular identity obtained above, $G\epsilon \cdot \eta G = 1$, then the square just states the equivalence of two expressions for $\epsilon\epsilon : RGRG \to 1$. Because of the first triangular identity, the square commutes, and we obtain the second of the desired triangular identities. Together they define the desired adjunction. □

References

[1] F. Borceux. (1994). *Handbook of categorical algebra 1: Basic category theory*, Cambridge: Cambridge University Press.

[2] D. Kan, *Adjoint Functors*, Transactions of the American Mathematical Society, vol. 87 (1958) pp. 294-329.

[3] S. Mac Lane. (1998). *Categories for the working mathematician*, New York: Springer.